

A Framework for Practical Anonymous Credentials from Lattices

Jonathan Bootle jbt@zurich.ibm.com IBM Research Europe - Zurich, Switzerland	Vadim Lyubashevsky vad@zurich.ibm.com IBM Research Europe - Zurich, Switzerland
Ngoc Khanh Nguyen khanh.nguyen@epfl.ch EPFL, Switzerland	Alessandro Sorniotti aso@zurich.ibm.com IBM Research Europe - Zurich, Switzerland

Abstract. We present a framework for building practical anonymous credential schemes based on the hardness of lattice problems. The running time of the prover and verifier is independent of the number of users and linear in the number of attributes. The scheme is also compact in practice, with the proofs being as small as a few dozen kilobytes for arbitrarily large (say up to 2^{128}) numbers of users with each user having several attributes. The security of our scheme is based on a new family of lattice assumptions which roughly states that given short pre-images of random elements in some set S , it is hard to create a pre-image for a fresh element in such a set. We show that if the set admits efficient zero-knowledge proofs of knowledge of a commitment to a set element and its pre-image, then this yields practically-efficient privacy-preserving primitives such as blind signatures, anonymous credentials, and group signatures. We propose a candidate instantiation of a function from this family which allows for such proofs and thus yields practical lattice-based primitives.

1 Introduction

With the recent announcement that the American National Institute of Standards and Technology (NIST) will be standardizing three lattice-based encryption and digital signature schemes, together with the NSA releasing their new CNSA 2.0 crypto suite which mandates that lattice-based constructions be the main cryptographic tools for communication security beginning from 2030, it is looking increasingly likely that lattices will form the future foundation of public key cryptography. While we already have several very efficient almost-standardized schemes for encryption and digital signatures [17, 33, 69], the state of more advanced quantum-safe cryptography is more murky with a lot of constructions being noticeably less efficient than their classical counterparts. The last several years, however, have seen tremendous progress in constructing efficient zero-knowledge proofs for lattice relations [15, 74, 35, 7, 34, 57] and this led to rather compact and practical constructions of schemes like ring signatures, group signatures and confidential payment protocols [36, 59, 61, 57]. A set of important primitives that are currently lacking in truly efficient instantiations fall under the global umbrella of *anonymous credentials*.

At a very high level, in an anonymous credential scheme [26, 27, 29, 51] an issuer signs a set of credentials for a user. The user can then present a subset of this credential set and give a zero-knowledge proof that they were indeed signed by the issuer. For the scheme to be secure, it should be impossible for a user to present a credential that the issuer never signed. For anonymity, when a user presents his credential, all his other credentials should remain hidden – even from the issuer. Furthermore, multiple credential presentations by the same user should remain unlinkable. Manifestations of such schemes have been adopted by the *self-sovereign identity* community, which has developed open-source projects [44, 43, 72, 62] and standard drafts [50, 71] catering for a market whose size is estimated in the billions of US dollars [66]. While there are some very efficient instantiations of these schemes based on non-quantum-safe assumptions, to the best of our knowledge there are no truly practical quantum-safe instantiations that are available.¹

¹ It is of course possible to construct these schemes using succinct zero-knowledge proofs such as STARKs, but the cost of these constructions appears to be in the hundreds of kilobytes range.

Some recent lattice-based constructions of related primitives that previously didn't have any efficient quantum-safe instantiation based the security of the scheme on novel, but very plausible, lattice assumptions. For example, the recent blind signature in [3] was based on the new “one-more-ISIS” assumption and publicly-verifiable SNARKs were recently instantiated using a new k-R-ISIS assumption in [5]. The blind signature in [3] is one of the most concretely-efficient quantum-safe constructions to date, while the work in [5] is the first construction of its kind that gets within the vicinity of practicality. It is not too surprising that the range of assumptions needed to be expanded in order to construct more efficient lattice-based primitives. Analogously to the non-post-quantum setting, simple primitives like digital signatures can be efficiently based on the hardness of the basic discrete logarithm problem, while the most efficient constructions of other more advanced schemes crucially rely on much more esoteric assumptions. And so, analogously to the classical setting, it makes sense to explore other assumptions for constructing advanced lattice-based primitives as well.

In this paper we propose a new family of lattice problems some of whose members admit a practically-efficient zero-knowledge proof for proving knowledge of a solution. We then show how to apply this zero-knowledge proof to create fairly simple constructions of various efficient privacy-preserving lattice-based primitives, such as blind signatures and anonymous credentials, based on the presumed hardness of the new problem. The specific assumption we use in our paper is a particular instantiation from this family of assumptions that yields efficient zero-knowledge proofs. The family of problems is very natural and can be seen as a generalization of the underlying problem upon which the classic GPV signature scheme [41] is based.

Lattice-Based hash-and-Sign Signatures and Extensions. Privacy-based primitives are often constructed as some combination of a hash-and-sign digital signature scheme and a zero-knowledge proof system. Hash-and-sign lattice-based digital signatures based on the hardness of the standard SIS problem, first constructed in [41], are abstractly based on the hardness of the following problem over some distribution of input matrices A , which is parametrized by a global function f (there are also specific parameters n, m, β, β' which are determined based on the security parameter):

Definition 1.1 (The ISIS $_f$ Problem (informal)). *We are given a matrix $A \in \mathbb{Z}_p^{n \times m}$ (chosen from some distribution), a function $f : [N] \rightarrow \mathbb{Z}_p^n$, and access to an oracle who chooses a random input $x \in [N]$ and outputs it together with a vector $\|\vec{s}\| \leq \beta$ satisfying $A\vec{s} = f(x)$. The game is won by coming up with a fresh tuple $(x', \vec{s}') \in [N] \times \mathbb{Z}^m$ where $\|\vec{s}'\| \leq \beta'$ and $A\vec{s}' = f(x')$.²*

The hardness of the above problem depends on the choice of f , the distribution of A and how the oracle chooses the vector \vec{s} .³ When the function f is instantiated as a cryptographic hash function (e.g. SHA) which is modelled as a random oracle, A is chosen uniformly-random, and \vec{s} is chosen from a distribution with enough entropy, then the ISIS $_f$ problem is equivalent to the well-known SIS and ISIS problems. If f is modelled as a random oracle, then $f(x)$ is uniformly-random in \mathbb{Z}_p^n , and so there is no advantage gained by seeing a pre-image \vec{s} of a random element in \mathbb{Z}_p^n because one could, in principle, do the pre-image sampling in reverse. That is, one could instead generate a random \vec{s} and if the distribution of \vec{s} has enough entropy, then it is the pre-image of a uniformly-random $\vec{t} = A\vec{s}$ (since $(A, \vec{t} = A\vec{s})$ is uniformly-random by the leftover hash lemma).

Modelling f as a random oracle and furthermore being able to create a matrix A together with a trapdoor which allows for pre-image sampling, gives rise to lattice-based hash-and-sign signatures based on the hardness of ISIS $_f$ [41]. The idea is simply to sign a message (digest) x with the pre-image \vec{s} . This allows us to sign random messages x , but when f is modelled as a random oracle, it doesn't matter whether x is chosen at random or adaptively because $f(x)$ is uniformly-random regardless – thus allowing the signing of arbitrary messages is as secure as only being able to sign random ones.

When f is not a cryptographic hash function, then we don't immediately obtain a signature scheme based on ISIS $_f$ by treating x as the message to be signed precisely because the ISIS $_f$ problem requires x to be

² Note that x' can be equal to one of the x , as long as $\vec{s}' \neq \vec{s}$.

³ We will also want the set $[N]$ to be exponentially large so that there is a negligible chance of obtaining pre-images on the same $f(x)$.

random, whereas in an attack on a signature scheme the x is chosen by the adversary. One could, of course, propose a stronger version of ISIS_f in which the x is chosen by the attacking algorithm, but we feel that this may require the function f to be needlessly “complicated” (e.g. something that’s close to a cryptographic hash function like SHA) in order for the problem to remain hard.⁴ And we do not want the function circuit to be too complex because building anonymous credentials from a signature scheme will additionally require proving knowledge of the signature. Even though recent techniques for proving arbitrarily large circuits have proof sizes of around 60 KB [12], this is still larger than proofs of simple linear or quadratic relations which can be as small as a dozen KB for common lattice relations [57]. Furthermore, the more complex proofs of arbitrary circuits will be less efficient, in terms of computation time, than proofs of simpler relations.

The main result of this work is to show how one can build an anonymous credential scheme based upon the above-sketched construction of a signature scheme for random messages. We also propose a simple function f which admits a very efficient zero-knowledge proof of the relation $A\vec{s}' = f(x')$, and for which we believe that the ISIS_f problem is hard. We now discuss our choice of f and will sketch the anonymous credential scheme construction and proof in Section 1.1.

Arguably the simplest way to instantiate $f(x)$ is to let it be a linear function. For example, if we use the natural correspondence between binary vectors of length $\log N$ with the set $[N]$, then if $f(\vec{v}) = B\vec{v}$, where $\vec{v} \in \{0, 1\}^{\log N}$, one can use the practically-efficient zero-knowledge proofs from [57] to commit to \vec{s} and \vec{v} and show that they have small norms and satisfy $A\vec{s} = B\vec{v}$. At first glance, it might seem that the linearity of the function interacts dangerously with the inherent linearity of lattices. For example, if $\vec{v}_1 + \vec{v}_2 = \vec{v}_3$ and we have $A\vec{s}_i = f(\vec{v}_i)$ for $i = 1, 2$, then $A(\vec{s}_1 + \vec{s}_2) = f(\vec{v}_3)$. In our definition of the problem, however, the \vec{v}_i are chosen at random and it is very unlikely that the sum of two randomly-chosen binary vectors will sum to another binary one (i.e. the probability is $(3/4)^{\log N}$). And if we sum more than 2 binary vectors, then the probability that the resulting pre-image \vec{s} which is a sum of the pre-images, will still be as small becomes negligible.⁵ So the hardness of ISIS_f for the above-defined function relies in large part on the fact that the domain of f is not closed under addition, even though f is a linear function over a larger domain. And so we conjecture that instantiating f in this way (i.e. choosing f as a random linear function over some ring and its domain as binary vectors) leads to a hard instance of ISIS_f . Furthermore, proving $A\vec{s} = B\vec{v}$ for some short \vec{s} and binary \vec{v} can be done very efficiently using the proofs from [57]. In particular, we do not see any better algorithm for solving the ISIS_f problem when f is defined as above vs. when f is a cryptographic hash function as in [41].

1.1 Blind Signatures and Anonymous Credentials from ISIS_f

We will now give a high level sketch of our main construction – an efficient anonymous credential scheme based on the hardness of the ISIS_f problem when one is able to give an efficient zero-knowledge proof of a solution to ISIS_f . We will first sketch the construction and proof of a blind signature scheme and then explain how essentially the same protocol can be turned into an anonymous credential scheme.

The Blind Signature Scheme. Recall that in a blind signature scheme, the user asks a signer to sign a message \vec{m} and can then prove to everyone that he has a signature for \vec{m} . The blindness property of the scheme states that the signer cannot know during which interaction he signed \vec{m} , and the soundness property states that a user who interacted with the signer k times can only produce k valid signed messages.

The signer creates a matrix A and a trapdoor which allows him to sample short $\vec{s} \sim D_\sigma$, for some standard deviation σ , satisfying $A\vec{s} = \vec{t}$ for any $\vec{t} \in \mathbb{Z}_p^n$. In the concrete instantiation, the matrix A , along with the

⁴ Still, investigating this stronger assumption with appropriate functions f could be an interesting research direction. The one-more-ISIS assumption in [3] does allow the adversary access to an oracle to obtain pre-images of arbitrary vectors, but requires him to find pre-images for a random set of vectors to win the game.

⁵ For example, the expected squared norm of a 512-dimensional Gaussian of standard deviation σ is $512\sigma^2$. The probability that the sum of three such Gaussians have squared norm less than this is less than 2^{-160} .

trapdoor and sampling algorithm, is as in the NTRU-based signature scheme [30, 69].⁶ The signer also creates random matrices B_1 and B_2 (their dimensions will be clear from their usage). The public key is A, B_1, B_2 and the function f , while the secret key is the trapdoor for A .

If the user wants to sign a (digest of a) message \vec{m} , he sends to the signer an Ajtai commitment of his message committed under randomness \vec{r} as $\vec{c} = B_1\vec{m} + B_2\vec{r}$, along with a zero-knowledge proof (using [57]) that \vec{m} and \vec{r} have small norms and the linear relation $\vec{c} = B_1\vec{m} + B_2\vec{r}$ is satisfied. The signer checks the proof, creates a uniformly-random “tag” $x \in [N]$ and then uses the trapdoor for A to create a short pre-image \vec{s} satisfying

$$A\vec{s} = f(x) + B_1\vec{m} + B_2\vec{r}. \quad (1)$$

The pre-image \vec{s} and the tag x are sent to the user. In order to prove that he has a signature of a message \vec{m} , the user reveals \vec{m} and creates a zero-knowledge proof of knowledge of \vec{s}, x satisfying (1). In the case that f is a linear function, the exact same proof from [57] for proving $\vec{c} = B_1\vec{m} + B_2\vec{r}$ can be used here as well.

The anonymity of the blind signature (i.e. that the signer cannot figure out during which interaction \vec{m} was signed) is ensured by the fact that \vec{m} is transferred to the signer in a computationally-hiding commitment scheme and all the proofs are zero-knowledge. The more interesting part of the proof is showing that a user cannot produce more signed messages than the number of queries that he makes to the signer. We will show that if there is an adversary who can forge in the blind signature scheme, then there is an algorithm who can solve the ISIS_f problem.

Given the matrix A and a function f in the ISIS_f problem, the reduction chooses matrices with small coefficients R_1 and R_2 , and then creates $B_1 = AR_1$ and $B_2 = AR_2$. By the LWE assumption, these look indistinguishable from uniform. The public key for the blind signature is thus A, B_1, B_2 , along with the function f . In his first move, the adversary sends an Ajtai commitment to a message \vec{m} using randomness \vec{r} along with a zero-knowledge proof that the commitment is valid. In particular, the zero-knowledge proof ensures that the commitment is of the form $B_1\vec{m} + B_2\vec{r}$, and the reduction can extract the \vec{m} and \vec{r} . By the construction of B_1 and B_2 , we have $B_1\vec{m} + B_2\vec{r} = A(R_1\vec{m} + R_2\vec{r})$. The reduction then calls the oracle in the ISIS_f definition and receives some \vec{s}, x satisfying $A\vec{s} = f(x)$. Notice that the reduction can now create $\vec{s}' = \vec{s} + R_1\vec{m} + R_2\vec{r}$ that satisfies

$$A\vec{s}' = f(x) + B_1\vec{m} + B_2\vec{r}. \quad (2)$$

The tuple (\vec{s}', x) could be a valid signature except for the fact that the distribution of \vec{s}' is not a discrete Gaussian (as in the real scheme), but is a Gaussian perturbed by $R_1\vec{m} + R_2\vec{r}$. To get \vec{s}' to be a discrete Gaussian, the reduction can use rejection sampling as in [53, Theorem 4.6] which converts shifted Gaussians into zero-centered ones. If we accept, then we can send \vec{s}' and x to the adversary as the signature of \vec{m} . If there is a rejection, then we can again query the oracle to obtain another \vec{s} and x and create another potential vector \vec{s}' , and so on until the rejection sampling procedure accepts.

The above shows that the view of the adversary is identical in both the simulation and the real signature scheme, and thus the adversary should be successful in creating a signature for a new message \vec{m} . In other words, he is able to prove knowledge of an $x \in [N]$ and vectors \vec{m}, \vec{r} satisfying (1), and these can be extracted by the extractor. Since $B_1 = AR_1$ and $B_2 = AR_2$, we have the potential solution to the ISIS_f problem being

$$A(\vec{s} - R_1\vec{m} - R_2\vec{r}) = f(x). \quad (3)$$

There are three possibilities for the x that is used in the forgery

1. x has not been queried to the oracle
2. x has been queried to the oracle, but was not seen by the adversary (i.e. it was discarded by the reduction during rejection sampling)
3. x has been queried to the oracle and seen by the adversary

⁶ This implies that the scheme would be instantiated over some polynomial ring rather than \mathbb{Z}_p , as in the description in this section. But since \mathbb{Z}_p is a sub-ring of the polynomial ring, all operations in the polynomial ring can also be described as operations over vectors over \mathbb{Z}_p , and so the description in this section is actually more general than we will need to be.

In case 1, Equation (3) directly gives a solution to the ISIS_f problem. In case 2, the fact that x was already queried means that the reduction has a vector \vec{s}' satisfying $A\vec{s}' = f(x)$. If $\vec{s}' \neq \vec{s} - R_1\vec{m} - R_2\vec{r}$, then we again have a solution to ISIS_f . Because the adversary never saw \vec{s}' and the entropy of pre-image sampling is (exponentially) large, there is only a negligible chance that $\vec{s}' = \vec{s} - R_1\vec{m} - R_2\vec{r}$.

In case 3, the adversary knows the value $\vec{s} = \vec{s}' + R_1\vec{m}' + R_2\vec{r}'$ where \vec{s}' was such that $A\vec{s}' = f(x)$, and it produces a forgery $\vec{s}, \vec{m} \neq \vec{m}', \vec{r}$ satisfying $A(\vec{s} - R_1\vec{m} - R_2\vec{r}) = f(x)$. If $\vec{s}' = \vec{s} - R_1\vec{m}' - R_2\vec{r}' \neq \vec{s} - R_1\vec{m} - R_2\vec{r}$, then we have a solution to the ISIS_f problem. In order for this to not be a solution, the adversary would need to create $\vec{s}, \vec{m} \neq \vec{m}', \vec{r}$ such that

$$(\vec{s} - \vec{s}') + R_1(\vec{m} - \vec{m}') + R_2(\vec{r} - \vec{r}') = 0. \quad (4)$$

To prove that even an all-powerful adversary will fail to create such an equality, we need to rely on the entropy of the matrix R_1 . In particular, we can set the distribution of R_1 such that for every column of R_1 , there exists another R'_1 which differs from R_1 in only that column, satisfying $AR_1 = AR'_1$ (c.f. [55, Lemma 4.4]). If for every R_1 , there is such a set of R'_1 , then for the R'_1 which differs in the column corresponding to the coordinate where \vec{m} and \vec{m}' differ, we will have $R_1(\vec{m} - \vec{m}') \neq R'_1(\vec{m} - \vec{m}')$. An important point is also that when receiving the vector \vec{s} , which is a result of rejection sampling that involves $R_1\vec{m}$, the value of R_1 is hidden since the output of the rejection sampling hides the specific $R_1\vec{m}$. In short, the success probability of the reduction breaking the ISIS_f problem is at least half the probability that an adversary can forge the blind signature.

Adapting the above to anonymous credentials and other privacy primitives. The framework for creating an anonymous credential scheme is virtually identical to the one above for a blind signature scheme. The first part of the scheme is the issuance of a credential for a set of attributes \vec{m} . This is done in the exact same manner as signing in the blind signature scheme. The user submits a commitment to a vector of credentials \vec{m} along with a zero-knowledge proof that the commitment is validly formed, and the issuer creates a random tag x along with a pre-image vector \vec{s} satisfying (1). The one difference in an anonymous credential scheme over a blind signature is that the user may not wish to reveal the entire attribute vector \vec{m} , whereas in the blind signature scheme, the whole message is revealed. In this case, he can simply reveal the sub-vector \vec{m}' of \vec{m} that he wishes and then prove knowledge of the remaining part of \vec{m} in the zero-knowledge proof. The security proof for the anonymous credential scheme is virtually identical to the one for the blind signature.

One can also easily adapt the framework to create a group signature scheme. The only change versus an anonymous credential scheme would be that the user would additionally create a lattice-based encryption of x along with the zero-knowledge proof, and additionally prove (again using the zero-knowledge proof from [57]) that the ciphertext is a valid encryption of x . The size of the signature will be larger than that in [56], but the advantage is that signing and verification time does not scale linearly in the group size.

1.2 Related Work

Prior works building privacy-preserving primitives such as group signatures, blind signatures, and anonymous credentials [68, 61, 67], circumvented the need to prove knowledge of a random oracle pre-image by using the standard-model digital signature framework of [1] instead of the hash-and-sign based approach. But these standard-model signatures are a factor 5X - 10X longer because instantiating standard-model signatures requires larger parameters and does not allow instantiating the trapdoored matrix A via the NTRU assumption as we do.

Other constructions of efficient blind signatures [3, 13] which do not start from the signature in [1], and do in fact use the NTRU trapdoor, unfortunately cannot be extended to anonymous credential schemes while utilizing the compact proofs from [57] for simple lattice relations. The technical reason is that in these blind signature schemes, the user gets the signer to blindly sign $H(\mu)$ for some message μ and cryptographic hash function H . In the proof, he reveals the whole μ , and proves knowledge of a simple relation that includes $H(\mu)$. Because the user will only ever give one proof that includes μ , revealing μ is perfectly fine. In an anonymous credential scheme, however, one may want to get a set of credentials (i.e. $\mu = \mu' || \mu''$) blindly

signed, and then reveal only the μ' part of it. This will require proving knowledge of a μ'' that satisfies some relation involving $H(\mu' || \mu'')$, which results in longer and slower proofs.⁷

1.3 Concurrent Work

We discuss concurrent works [45, 47], of which [45] was recently updated with applications to anonymous credentials and appears at CRYPTO'23.

In common with prior work [68, 61, 67], both [45] and [47] avoid proving knowledge of random oracle preimages. Both works follow the paradigm of ‘signatures with efficient protocols’ and design lattice-based digital signature schemes tailored towards efficient zero-knowledge proofs of a message and signature. The second work, [47] defines a new cryptographic primitive called commit-transferable signatures (CTS) in order to instantiate this paradigm more efficiently. Similarly to our work, they rely on existing frameworks such as [58] (with [45] additionally relying on [74])) to produce efficient zero-knowledge proofs to use inside their constructions.

The biggest difference between our work and [45, 47] lies in the use of cryptographic assumptions. The anonymous credential constructions of [45, 47] rely solely on the SIS and LWE assumptions and their ring variants. In contrast, our work relies on the new ISIS_f assumption which we introduce. The main benefit of this new assumption over the approaches in [45, 47] is that it allows us to use zero-knowledge proofs of [58] to prove significantly fewer and simpler relations. This leads to a large efficiency improvement in our anonymous credential constructions, with [45, 47] offering credential sizes of roughly 640KB and 500KB respectively, and our work offering sizes of roughly 120KB (under ISIS_f). Further, while [45, 47] rely on trapdoor sampling techniques for standard lattices from [64], we rely on NTRU lattices and NTRU-based trapdoors [30, 69]. Finally, as noted in [47], we use minimalistic definitions of anonymous credentials which do not incorporate pseudonyms or user secret keys. Including these aspects would require additional commitments and zero-knowledge proofs which would increase the complexity of our anonymous credentials, although we do not believe they would add too much extra overhead.

1.4 Discussion and Open Problems

The most pertinent open problem is to analyze the hardness of the ISIS_f problem when instantiated with a random linear f , or perhaps other classes of slightly higher-degree functions f which still yield efficient zero-knowledge proofs.

Another open problem is to come up with a possibly stronger assumption that would result in tighter reductions for our schemes. In particular, we feel that breaking our blind signature and anonymous credential scheme is more difficult than what the security reductions imply. For example, we need to set parameters so as to allow rejection sampling to work and for the matrix R_1 to be chosen such that with high probability there is another R'_1 satisfying $AR_1 = AR'_1$.

If one looks closely at the proof sketch in Section 1.1, what we have essentially done is give a reduction from ISIS_f to an “interactive” version of ISIS_f in which the adversary has more control over the pre-images that he receives. In the ISIS_f problem, the adversary is given an oracle that outputs \vec{s}, x satisfying $A\vec{s} = f(x)$, whereas in the “interactive” version, the adversary can choose \vec{m} and \vec{r} as in (1) and then receives a random x and \vec{s} satisfying (1). His goal is then to satisfy (1) for a fresh \vec{m} . We have shown that if the random matrices B_1 and B_2 are constructed as AR_1 and AR_2 , and other parameters are set appropriately, then using rejection sampling we are able to transform \vec{s}, x that satisfy $A\vec{s} = f(x)$ into \vec{s}', x that satisfy (1). Furthermore, picking an R_1 from a wide-enough distribution to guarantee the existence of another R'_1 satisfying $AR_1 = AR'_1$ was required to assure that the adversary could not create a forgery that satisfied (4). It is very much possible, however, that setting the scheme parameters to allow for these particularities of the proof to go through is not really necessary for security. Since asymptotically, the interactive ISIS_f problem is as hard as ISIS_f , it

⁷ The blind signature protocol of [13] does require a zero-knowledge proof involving $H(\cdot)$, but this proof is only done in the intermediate interaction between the user and the signer, and so its relative inefficiency does not affect the signature size.

Attributes	Assumption			
	ISIS _f	interactive ISIS _f	Strong-RSA [23]	qSDH [20]
8	122KB	26KB	1319B	608B
16	133KB	29KB	1910B	865B

Table 1: Output sizes for the anonymous credential schemes with 8 and 16 attributes.

makes sense to examine the interactive version on its own with a more favorable parameter setting. That is, one could make the assumption that the interactive version of the ISIS_f problem is hard in the parameter range that does not yield a reduction from ISIS_f.⁸ In Table 1, we give example parameters of instantiations of an anonymous credential scheme based on the ISIS_f assumption and parameters required if one assumes hardness of the interactive version. The scheme based on ISIS_f is noticeably more compact than the previously most efficient anonymous credential scheme [45], which has output sizes of around 650KB. But as we also see, there is a noticeable advantage in setting the scheme parameters smaller and assuming that the problem (interactive ISIS_f) is still hard. We believe that analyzing the security of this version of the problem, and possibly building other more efficient schemes based on it, is a promising research direction when it comes to building practical lattice-based schemes.

One of the main appeals of lattice cryptography, in addition to its versatility, is that its underlying operations are very fast when instantiated over polynomial rings. It is therefore quite conceivable that lattice-based constructions will be the fastest option out of all the post-quantum alternatives. Since most anonymous credential schemes involve real-time interaction (e.g. credit card usage), speed is a very important consideration in their real world deployment. Unlike the many efforts to construct efficient software for discrete log, pairing-based, and PCP/IOP-based proof systems, there has not been much concentrated effort to develop software for efficient lattice-based primitives. Part of the problem has been that there were not many lattice-based protocols which were compact enough to be considered for practical deployment. We hope that this paper provides some motivation for creating implementations of lattice-based privacy primitives that are real-world ready.

2 Preliminaries

Let λ be a security parameter which is provided in unary to all involved algorithms. For $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$. Let \mathbb{Z}_p denote the ring of integers modulo p . We write $\vec{v} \in \mathbb{Z}_p^m$ to denote vectors over a ring \mathbb{Z}_p . Matrices over \mathbb{Z}_p will be written as regular capital letters. By default, all vectors are column vectors.

We define $\mathcal{U}(S)$ to be the uniform distribution on the finite set S . We write $x \leftarrow D$ when x is sampled according to the distribution D . Sometimes, we abuse the notation and write $x \leftarrow S$ to denote $x \leftarrow \mathcal{U}(S)$.

Lattices. Let $B = \{\vec{b}_1, \dots, \vec{b}_n\} \subseteq \mathbb{R}^n$ be a set of linearly independent vectors. The n -dimensional full-rank lattice generated by B is defined as follows:

$$\Lambda = \Lambda(B) := \left\{ \sum_{i=1}^n c_i \vec{b}_i : c_1, \dots, c_n \in \mathbb{Z} \right\}.$$

We denote $\tilde{B} = (\vec{b}_i)_{i \in [n]}$ to be the Gram-Schmidt orthogonalization of B . Further, define the Gram-Schmidt norm of B as $\|\tilde{B}\| := \max_{i \in [n]} \|\vec{b}_i\|$.

⁸ This is similar to how the SIS and LWE problems were first introduced – solving their random instances was shown to be as hard as solving lattice problems in the worst case [4, 65, 70], but now the SIS and LWE problems are used with parameters which do not satisfy these original reductions because these problems have since been very well-studied on their own.

Cyclotomic Rings. For a power of two $d \geq 4$ and a positive integer p , let $\mathcal{K} = \mathbb{Q}[X]/(X^d + 1)$ denote the $2d$ -th cyclotomic field and $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$ be the corresponding ring of integers. Lower-case letters denote elements in \mathcal{R} and bold lower-case (resp. upper-case) letters represent column vectors (resp. matrices) with coefficients in \mathcal{R} . For a modulus $q \in \mathbb{N}$, we define $\mathcal{R}_q := \mathcal{R}/(q) = \mathbb{Z}_q[X]/(X^d + 1)$. Further, we define \mathcal{R}_q^\times to be the set of polynomials in \mathcal{R}_q which are invertible over \mathcal{R}_q . For $\eta \in \mathbb{N}$, define $S_\eta := \{x \in \mathcal{R} : \|x\|_\infty \leq \eta\}$. Recall that, for any $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, the q -ary lattice $\Lambda_q^\perp(\mathbf{A})$ is defined as $\Lambda_q^\perp(\mathbf{A}) := \{\mathbf{x} \in \mathcal{R}^m : \mathbf{A}\mathbf{x} = \mathbf{0} \pmod{q}\}$.

Coefficient representation and multiplication matrices. For a polynomial vector $\mathbf{x} \in \mathcal{R}^l$, define $\text{Coeffs}(\mathbf{x}) \in \mathbb{Z}^{ld}$ to be coefficient vector of \mathbf{x} . Similarly, $\text{Coeffs}^{-1}(\vec{x}) \in \mathcal{R}^l$ is the polynomial vector with coefficients \vec{x} .

For a polynomial $f = f_0 + f_1X + \dots + f_{d-1}X^{d-1} \in \mathcal{R}$, we define the *multiplication matrix* $\text{rot}(f) \in \mathbb{Z}^{d \times d}$ as:

$$\text{rot}(f) = \begin{bmatrix} f_0 & -f_{d-1} & \dots & -f_1 \\ f_1 & f_0 & \dots & -f_2 \\ \vdots & \vdots & \ddots & \vdots \\ f_{d-1} & \dots & f_1 & f_0 \end{bmatrix}.$$

In particular, we will use the property that for $f, g \in \mathcal{R}$, $\text{rot}(fg) = \text{rot}(f)\text{rot}(g)$. We extend this definition to matrices over \mathcal{R} . Namely, for a matrix $\mathbf{F} = (f_{i,j}) \in \mathcal{R}_q^{n \times m}$, we define

$$\text{rot}(\mathbf{F}) = \begin{bmatrix} \text{rot}(f_{1,1}) & \text{rot}(f_{1,2}) & \dots & \text{rot}(f_{1,m}) \\ \vdots & \vdots & \ddots & \vdots \\ \text{rot}(f_{n,1}) & \text{rot}(f_{n,2}) & \dots & \text{rot}(f_{n,m}) \end{bmatrix} \in \mathbb{Z}^{nd \times md}.$$

Then, for any polynomial vector $\mathbf{x} \in \mathcal{R}_q^m$ we have the following property over \mathbb{Z}_q :

$$\text{Coeffs}(\mathbf{F}\mathbf{x}) = \text{rot}(\mathbf{F})\text{Coeffs}(\mathbf{x}).$$

Discrete Gaussian distribution. We recall the discrete Gaussian distribution used for the rejection sampling and trapdoor sampling.

Definition 2.1. *The n -dimensional Gaussian function $\rho_{\mathfrak{s}, \vec{c}} : \mathbb{R} \rightarrow (0, 1]$ is defined by*

$$\rho_{\mathfrak{s}, \vec{c}}(\vec{x}) := \exp\left(-\frac{\|\vec{x} - \vec{c}\|^2}{2\mathfrak{s}^2}\right).$$

For any coset $\Lambda + \vec{t}$ of a full-rank lattice $\Lambda \subset \mathbb{R}^n$, $\rho_{\mathfrak{s}, \vec{c}}(\Lambda + \vec{t}) := \sum_{\vec{x} \in \Lambda + \vec{t}} \rho_{\mathfrak{s}, \vec{c}}(\vec{x})$. Then, the discrete Gaussian distribution over a coset of a lattice $\Lambda + \vec{t}$ centred around $\vec{c} \in \mathbb{R}^n$ with standard deviation $\mathfrak{s} > 0$ is given by

$$\forall \vec{x} \in \Lambda + \vec{t}, D_{\Lambda + \vec{t}, \mathfrak{s}, \vec{c}}(\vec{x}) := \frac{\rho_{\mathfrak{s}, \vec{c}}(\vec{x})}{\rho_{\mathfrak{s}, \vec{c}}(\Lambda + \vec{t})}.$$

We write $D_{\mathfrak{s}, \vec{c}}^n$ when $\Lambda = \mathbb{Z}^n$. Similarly, we ignore the subscript \vec{c} when the distribution is centred around $\vec{0} \in \mathbb{Z}^n$.

Smoothing parameter. We recall the definition of a smoothing parameter [65]. Namely, for any n -dimensional lattice Λ and a real number $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest $s > 0$ such that $\rho_{\frac{1}{s\sqrt{2\pi}}}(\Lambda^* \setminus \{0\}) \leq \epsilon$. We also consider the scaled version $\eta'_\epsilon(\Lambda) = \frac{1}{\sqrt{2\pi}}\eta_\epsilon(\Lambda)$. Further, for any $\epsilon > 0$ we define $\eta_{\min}(\epsilon) \in \mathbb{R}$ to be such that $\Pr_{\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}}[\eta'_\epsilon(\Lambda_q^\perp(\mathbf{A})) > \eta_{\min}(\epsilon)] \leq 2^{-d}$. Note that η_{\min} can be computed as in [18, Lemma 2.5].

Tail Bounds We recall the following tail bounds which are widely used in the literature. The first one focuses on discrete Gaussians over integers and follows from [9, Lemma 1.5(i)] and was adapted in [53, Lemma 4.4]. The next one by Micciancio and Regev [65, Lemma 4.4] is a tail bound on discrete Gaussians over any full-rank lattices.

Lemma 2.2 ([53]). *Let $\vec{x} \leftarrow D_{\mathfrak{s}}^n$ and $t > 1$. Then*

$$\Pr [\|\vec{x}\| > t \cdot \mathfrak{s}\sqrt{n}] < \left(te^{\frac{1-t^2}{2}} \right)^n.$$

Lemma 2.3 ([65]). *Let Λ be an n -dimensional full-rank lattice, $\vec{c} \in \mathbb{R}^n$, $0 < \epsilon < 1$ and $\mathfrak{s} \geq \eta'_\epsilon(\Lambda)$. Let $\vec{x} \leftarrow D_{\Lambda, \mathfrak{s}, \vec{c}}$. Then $\Pr [\|\vec{x} - \vec{c}\| \geq \mathfrak{s}\sqrt{n}] < \frac{1+\epsilon}{1-\epsilon} \cdot 2^{-n}$.*

Preimage sampling. Let $\mathbf{A} \in \mathcal{R}_q^{n \times m}$. Then, we denote $\mathbf{A}_{\mathfrak{s}}^{-1}(\mathbf{u})$ to be the random variable $\mathbf{x} \leftarrow D_{\mathfrak{s}}^{md}$ conditioned on $\mathbf{A}\mathbf{x} = \mathbf{u}$ over \mathcal{R}_q .

Rejection sampling. Rejection sampling [52, 53] is a widely used technique to ensure the zero-knowledge property of many lattice-based (non-)interactive proofs.

Lemma 2.4 (Rejection Sampling [53]). *Let $V \subseteq \mathcal{R}^\ell$ be a set of polynomials with norm at most T and $\rho: V \rightarrow [0, 1]$ be a probability distribution. Fix the standard deviation $\mathfrak{s} = \alpha T$ for $\alpha = O(\sqrt{\lambda})$. Let*

$$M = \exp \left(\sqrt{\frac{2(\lambda + 1)}{\log e}} \cdot \frac{1}{\alpha} + \frac{1}{2\alpha^2} \right) = O(1).$$

Now, sample $\mathbf{v} \leftarrow \rho$ and $\mathbf{y} \leftarrow D_{\mathfrak{s}}^\ell$, set $\mathbf{z} = \mathbf{y} + \mathbf{v}$, and run $b \leftarrow \text{Rej}_1(\mathbf{z}, \mathbf{v}, \mathfrak{s}, M)$ as defined in Figure 1. Then, the probability that $b = 1$ is at least $(1 - 2^{-\lambda})/M$ and the distribution of (\mathbf{v}, \mathbf{z}) , conditioned on $b = 1$, is within statistical distance of $2^{-\lambda}$ of the product distribution $\rho \times D_{\mathfrak{s}}^\ell$.

```

Rej( $\vec{z}, \vec{v}, \mathfrak{s}, M$ )
1:  $u \leftarrow [0, 1]$ 
2: if  $u > \frac{1}{M} \cdot \exp\left(\frac{-2(\vec{z}, \vec{v}) + \|\vec{v}\|^2}{2\mathfrak{s}^2}\right)$  then
3:   return 0 (i.e. reject)
4: else
5:   return 1 (i.e. accept)

```

Fig. 1: Standard rejection sampling algorithm [53].

Recently, Boschini et al. [18] proposed a generalized rejection sampling method for ellipsoidal Gaussians over any lattice, which was later used in the context of multi-signatures. Here, we use the (simplified) result from [18, Theorem B.1] to apply rejection sampling on q -ary lattices⁹.

Lemma 2.5 (Generalized Rejection Sampling [18]). *Take any $\alpha, T > 0$ and $\epsilon \leq 1/2$. Let $\mathbf{v} \in \mathcal{R}_q^m$ be such that $\|\mathbf{v}\| \leq T$, $\mathbf{A} \in \mathcal{R}_q^{n \times m}$, $\mathbf{w} \in \mathcal{R}_q^n$ and $\mathbf{t} := \mathbf{A}\mathbf{v} \in \mathcal{R}_q^n$. Also, pick $\mathfrak{s} \geq \max(\alpha T, \eta'_\epsilon(\Lambda_q^\perp(\mathbf{A})))$. Then, for any*

$$t > 0, \quad M := \exp\left(\frac{1}{2\alpha^2} + \frac{t}{\alpha}\right), \quad \epsilon := 2\left(\frac{1+\epsilon}{1-\epsilon}\right) \exp\left(-2t^2 \cdot \frac{\pi-1}{\pi}\right),$$

the statistical distance between distributions RejSamp and SimRS defined in Figure 2 is at most $\frac{\epsilon}{2M} + \frac{2\epsilon}{M}$. Moreover, the probability that RejSamp outputs something is at least $\frac{1-\epsilon}{M} \left(1 - \frac{4\epsilon}{(1+\epsilon)^2}\right)$.

⁹ An almost identical application was described in [18, Section B.4] in the context of proving statistical honest-verifier zero-knowledge of the Fiat-Shamir with aborts protocol [53].

$\text{RejSamp}(\mathbf{A}, \mathbf{v}, \mathbf{t}, \mathbf{w})$	$\text{SimRS}(\mathbf{A}, \mathbf{t}, \mathbf{w})$
1: if $\mathbf{A}\mathbf{v} \neq \mathbf{t}$ then return \perp	1: if $\mathbf{A}\mathbf{v} \neq \mathbf{t}$ then return \perp
2: $\mathbf{y} \leftarrow \mathbf{A}_s^{-1}(\mathbf{w})$	2: $\mathbf{z} \leftarrow \mathbf{A}_s^{-1}(\mathbf{t} + \mathbf{w})$
3: $\mathbf{z} = \mathbf{y} + \mathbf{v}$	3: return $(\mathbf{A}, \mathbf{t}, \mathbf{w}, \mathbf{z})$ with prob. $\frac{1}{M}$
4: return $(\mathbf{A}, \mathbf{t}, \mathbf{w}, \mathbf{z})$ with prob. $\min\left(\frac{D_s^{md}(\mathbf{z})}{M \cdot D_{s,\mathbf{v}}^{md}(\mathbf{z})}, 1\right)$	

Fig. 2: Rejection sampling on q -ary lattices.

2.1 NTRU Lattices

Using terminology from above, let d be a power of two, q a positive integer and $f, g \in \mathcal{R}$ such that f is invertible over \mathcal{R}_q . Let $h = g/f \in \mathcal{R}_q$. The NTRU lattice associated to h and q is defined as

$$A_{h,q} := \{(u, v) \in \mathcal{R}^2 : u + vh = 0 \pmod{q}\}.$$

Then, $A_{h,q}$ is a $2d$ -dimensional full-rank lattice generated by the rows of

$$A_{h,q} := \begin{bmatrix} -\text{rot}(h) & I_d \\ q \cdot I_d & 0 \end{bmatrix} \in \mathbb{Z}^{2d \times 2d}.$$

We recall that there is an efficient algorithm NTRU.TrapGen , which given modulus q and the ring dimension d , outputs $h \in \mathcal{R}_q$ and a short basis of $A_{h,q}$. This is the core part of the key generation of the Falcon signature scheme [69].

Lemma 2.6 ([30, 69]). *There is an efficient algorithm $\text{NTRU.TrapGen}(q, d)$ which outputs $h \in \mathcal{R}_q$ and a basis \mathbf{B} of $A_{h,q}$ such that $\|\tilde{\mathbf{B}}\| \leq 1.17\sqrt{q}$.*

The short basis can now be used for preimage sampling using the well-known GPV framework [41] and its concrete instantiation in [30]. Namely, for any $c \in \mathcal{R}$, one can efficiently sample $(u, v) \in \mathcal{R}^2$ from a discrete Gaussian distribution conditioned on $u + vh = c \pmod{q}$. We use the extended result following [73, Lemma 2.7], i.e. given additionally $\mathbf{a} \in \mathcal{R}^m$, we can efficiently sample $(u, v, \mathbf{w}) \in \mathcal{R}^{m+2}$ from a discrete Gaussian distribution conditioned on $u + vh + \mathbf{a}^T \mathbf{m} = c \pmod{q}$.

Lemma 2.7 ([31, 30, 41]). *Let $n \in \mathbb{N}$ and $\epsilon = 2^{-\lambda}/(4d)$. There is a PPT algorithm GSampler , which takes $(h, \mathbf{B}) \leftarrow \text{NTRU.TrapGen}(q, d)$, $\mathbf{a} \in \mathcal{R}_q^m$, standard deviation $\mathfrak{s} > 0$ and a target vector $c \in \mathcal{R}_q$ as input, and outputs a triple $(u, v, \mathbf{w}) \in \mathcal{R}_q^{m+2}$ such that*

$$\Delta\left([h \ \mathbf{a}^T \ 1]_s^{-1}(c), \text{GSampler}(h, \mathbf{a}, \mathbf{B}, \mathfrak{s}, c)\right) \leq 2^{-\lambda}$$

as long as

$$\mathfrak{s} \geq 1.17\sqrt{q} \cdot \eta'_\epsilon(\mathbb{Z}) \quad \text{where} \quad \eta'_\epsilon(\mathbb{Z}) \approx \frac{1}{\pi} \cdot \sqrt{\frac{1}{2} \ln\left(2 + \frac{2}{\epsilon}\right)}.$$

2.2 Module-SIS and Module-LWE Problems

The security of our schemes relies on the well-known computational lattice problems Module-LWE (MLWE) and Module-SIS (MSIS) [48]. Both problems are defined over \mathcal{R}_q .

Definition 2.8 (MSIS $_{n,m,\mathcal{B}}$). *Given $A \leftarrow \mathcal{R}_q^{n \times m}$, the Module-SIS problem with parameters $n, m > 0$ and $0 < \mathcal{B} < q$ asks to find $\mathbf{z} \in \mathcal{R}_q^m$ such that $\mathbf{A}\mathbf{z} = \mathbf{0}$ over \mathcal{R}_q and $0 < \|\mathbf{z}\| \leq \mathcal{B}$. An algorithm \mathcal{A} is said to have advantage ϵ in solving MSIS $_{n,m,\mathcal{B},q}$ if*

$$\text{Adv}_{n,m,\mathcal{B},q}^{\text{MSIS}}(\mathcal{A}) := \Pr\left[0 < \|\mathbf{z}\|_\infty \leq \mathcal{B} \wedge \mathbf{A}\mathbf{z} = \mathbf{0} \mid A \leftarrow \mathcal{R}_q^{n \times m}; \mathbf{z} \leftarrow \mathcal{A}(A)\right] \geq \epsilon.$$

If the modulus is clear from the context, then we drop q from the subscript.

As for Module-LWE, we consider its “knapsack version” which is equivalent to its standard definition up to an additive negligible factor for typically chosen parameters [36, Appendix C]. In the paper we will use both versions and assume they are equivalent.

Definition 2.9 (MLWE_{*n,m,χ,q*}). *The Module-LWE problem with parameters $m \geq n > 0$ and an error distribution χ over \mathcal{R} asks the adversary \mathcal{A} to distinguish between the following two cases: 1) (A, \mathbf{As}) for $A \leftarrow \mathcal{R}_q^{n \times m}$, a secret vector $\mathbf{s} \leftarrow \chi^m$, and 2) $(A, \mathbf{b}) \leftarrow \mathcal{R}_q^{n \times m} \times \mathcal{R}_q^n$. Then, \mathcal{A} is said to have advantage ϵ in solving MLWE_{*n,m,χ,q*} if*

$$\begin{aligned} \text{Adv}_{n,m,\chi,q}^{\text{MLWE}}(\mathcal{A}) := & \left| \Pr \left[b = 1 \mid A \leftarrow \mathcal{R}_q^{n \times m}; \mathbf{s} \leftarrow \chi^m; b \leftarrow \mathcal{A}(A, \mathbf{As}) \right] \right. \\ & \left. - \Pr \left[b = 1 \mid A \leftarrow \mathcal{R}_q^{n \times m}; \mathbf{b} \leftarrow \mathcal{R}_q^n; b \leftarrow \mathcal{A}(A, \mathbf{b}) \right] \right| \geq \epsilon. \end{aligned} \quad (5)$$

For simplicity, if χ is a uniform distribution over S_η then we simply write MLWE_{*n,m,η,q*}. Also, we drop the subscript q if the modulus is clear from the context.

2.3 Non-Interactive Zero-Knowledge Proofs in the ROM

In this paper, we consider binary relations $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$. Then define $L_R := \{x \in \{0, 1\}^* : \exists w, (x, w) \in R\}$ to be the language corresponding to R . We refer to x as a statement and w as a witness.

We recall the (slightly adapted) definitions of non-interactive zero-knowledge proofs (NIZK) in the random oracle model from [67].

Definition 2.10 (NIZK). *A non-interactive zero-knowledge proof system Π_{NIZK} for the relation R consists of three PPT algorithms (Setup, Prove, Verify) which are defined as follows:*

- $\text{Setup}(1^\lambda) \rightarrow \text{crs}$: the setup algorithm which outputs the common reference string $\text{crs} \in \{0, 1\}^{\ell(\lambda)}$,
- $\text{Prove}^H(\text{crs}, x, w) \rightarrow \pi / \perp$: the prover algorithm takes as input the common reference string $\text{crs} \in \{0, 1\}^\ell$, statement x and witness w , either outputs a proof π or an abort symbol \perp ,
- $\text{Verify}^H(\text{crs}, x, \pi) \rightarrow 0/1$: the verifier algorithm takes as input the common reference string $\text{crs} \in \{0, 1\}^\ell$, statement x and a proof π and outputs a bit b where $b = 1$ means accept and $b = 0$ means reject.

Unless stated otherwise, we assume that crs can be generated as the output of another random oracle. In other words, crs is a common random string. Hence, our protocols do not require a trusted setup.

We recall the key properties of NIZK used in this paper: (i) correctness, (ii) zero-knowledge, and (iii) multi-proof extractability (i.e. straight-line extractability).

Definition 2.11 (Correctness). *We say that a NIZK proof system Π_{NIZK} is correct if for all $\text{crs} \in \{0, 1\}^\ell$ and $(\text{crs}, x, w) \in R$, the probability that $\text{Prove}^H(\text{crs}, x, w)$ outputs \perp is $\text{negl}(\lambda)$, and:*

$$\Pr \left[\pi \leftarrow \text{Prove}^H(\text{crs}, x, w) : \text{Verify}(\text{crs}, x, w) = 1 \mid \pi \neq \perp \right] = 1 - \text{negl}(\lambda).$$

Definition 2.12 (Zero-Knowledge). *We say that a NIZK proof system Π_{NIZK} is zero-knowledge if there exists a simulator $\mathcal{S} = (\mathcal{S}_0, \mathcal{S}_1)$ which consists of two PPT algorithms with a shared state such that for any PPT adversary \mathcal{A} we have*

$$\text{Adv}_{\Pi_{\text{NIZK}}}^{\text{ZK}}(\mathcal{A}) := \left| \Pr[1 \leftarrow \mathcal{A}^{\text{H,Prove}}(\text{crs})] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{S}_0, \mathcal{S}_1}(\text{crs})] \right| = \text{negl}(\lambda)$$

where Prove and \mathcal{S} are prover and simulator oracles which, given (x, w) , output \perp if $(\text{crs}, x, w) \notin R$ and otherwise return $\text{Prove}^H(\text{crs}, x, w)$ and $\mathcal{S}_1(\text{crs}, x)$ respectively. The probability is also taken over the randomness of generating the common reference string $\text{crs} \leftarrow \text{Setup}(1^\lambda)$.

Finally, we consider multi-proof extractability which corresponds to straight-line extractability. Here, the adversary can pick the statements adaptively. In order to perform extraction in this stronger setting, the common reference string is simulated and the corresponding trapdoor is provided to the extractor.

Definition 2.13 (Multi-Proof Extractability). *The following hold:*

CRS Simulatability. For any PPT adversary \mathcal{A} , we have:

$$\text{Adv}_{\Pi_{\text{NIZK}}}^{\text{crs}}(\mathcal{A}) := \left| \Pr[\text{crs} \leftarrow \text{Setup}(1^\lambda) : 1 \leftarrow \mathcal{A}^{\text{H}}(\text{crs})] - \Pr[(\text{c}\tilde{\text{r}}\text{s}, \tau) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda) : 1 \leftarrow \mathcal{A}^{\text{H}}(\text{c}\tilde{\text{r}}\text{s})] \right|$$

is negligible.

Straight-Line Extractability. There exist constants e_1, e_2, c such that for any $Q_{\text{H}}, Q_s \in \text{poly}(\lambda)$ and any PPT adversary \mathcal{A} that makes at most Q_{H} random oracle queries with

$$\Pr \left[\begin{array}{l} (\text{c}\tilde{\text{r}}\text{s}, \tau) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda), \\ \{(x_i, \pi_i)\}_{i \in [Q_s]} \leftarrow \mathcal{A}^{\text{H}}(\text{c}\tilde{\text{r}}\text{s}) \end{array} : \forall i \in [Q_s], \text{Verify}^{\text{H}}(\text{c}\tilde{\text{r}}\text{s}, x_i, \pi_i) = 1 \right] \geq \varepsilon(\lambda),$$

where $\varepsilon(\lambda)$ is non-negligible, we have

$$\Pr \left[\begin{array}{l} (\text{c}\tilde{\text{r}}\text{s}, \tau) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda), \{(x_i, \pi_i)\}_{i \in [Q_s]} \leftarrow \mathcal{A}^{\text{H}}(\text{c}\tilde{\text{r}}\text{s}), \\ \{w_i \leftarrow \text{Multi-Extract}(Q_{\text{H}}, Q_s, 1/\varepsilon, \text{c}\tilde{\text{r}}\text{s}, \tau, x_i, \pi_i)\}_{i \in [Q_s]} : \\ \forall i \in [Q_s], (\text{c}\tilde{\text{r}}\text{s}, x_i, w_i) \in R \wedge \text{Verify}^{\text{H}}(\text{c}\tilde{\text{r}}\text{s}, x_i, \pi_i) \end{array} \right] \geq \frac{1}{2} \cdot \varepsilon(\lambda) - \text{negl}(\lambda)$$

where the runtime of the extractor is upper-bounded by $Q_{\text{H}}^{e_1} \cdot Q_s^{e_2} \cdot \frac{1}{\varepsilon(\lambda)^c} \cdot \text{poly}(\lambda)$.

For proving various soundness properties, we will deal with *expected* PPT knowledge extractors. Since security of our protocols rely on computation assumptions, it is essential for us to transform these extractors into *strict* PPT algorithms. Below, we show a standard way to achieve this goal.

Lemma 2.14 (Expected-Time to Strict-Time Transformation). Take any efficient binary relation R and any statement x . Let \mathcal{A} be a probabilistic algorithm which runs in expected time at most T and

$$\Pr_{\text{crs} \leftarrow \{0,1\}^\ell} [\mathcal{A}(\text{crs}, x) \rightarrow w : (x, w) \in R] = \varepsilon.$$

Then, there is an algorithm \mathcal{B} with an oracle access to \mathcal{A} , which runs in time at most $\frac{2T}{\varepsilon}$ and

$$\Pr_{\text{crs} \leftarrow \{0,1\}^\ell} [\mathcal{B}^{\mathcal{A}}(\text{crs}, x) \rightarrow 1] \geq \varepsilon/2.$$

Proof. The algorithm \mathcal{B} is pretty intuitive: given input (crs, x) , it runs $\mathcal{A}(\text{crs}, x)$ and checks whether \mathcal{A} did output a valid witness w . However, if the total runtime is more than $T_{\text{max}} := 2T/\varepsilon$, then \mathcal{B} aborts.

Denote X to be the binary random variable, which determines whether the output w from \mathcal{A} is valid. Similarly, let Y be the runtime of $\mathcal{A}(\text{crs}, x)$. Then, we are interested in

$$\Pr[X = 1 \wedge Y \leq T_{\text{max}}] \geq \Pr[X = 1] - \Pr[Y > T_{\text{max}}].$$

Note that $\Pr[X = 1] = \varepsilon$, and also by Markov inequality we have $\Pr[Y > T_{\text{max}}] \leq \varepsilon/2$. Hence, this concludes the proof. \square

3 The ISIS_f Assumption

This section introduces a new family of lattice-based assumptions called ISIS_f (where ISIS stands for Inhomogenous Shortest Integer Solution). With a similar flavour to the recently proposed lattice assumptions [5, 2], the adversary is given short preimages of random (but not necessarily uniformly random) images and its task is to come up with either a short preimage of a new image, or a new preimage (i.e. not the one received earlier) to one of the images which was sent. More formally, let $\mathbf{n}, \mathbf{m}, k \in \mathbb{N}$ where the first two variables correspond to the matrix dimensions and the last one represents the number of samples. We consider a uniformly random $\mathbf{n} \times \mathbf{m}$ matrix \mathbf{A} over \mathcal{R}_q , and an efficiently computable function $f : [N] \rightarrow \mathbb{Z}_q^n$. The adversary is given $x_1, \dots, x_k \leftarrow [N]$ as well as vectors $\mathbf{s}_1, \dots, \mathbf{s}_k \in \mathbb{Z}_q^m$, where each $\mathbf{s}_i \leftarrow \mathbf{A}_s^{-1}(f(x_i))$. Then, it needs to come up with a new pair (x^*, \mathbf{s}^*) (in particular, different from the previous ones) such that $\mathbf{A}\mathbf{s}^* = f(x^*)$ and \mathbf{s}^* is short.




Fig. 3: The ISIS_f experiment.

Definition 3.1 (The ISIS_f Problem). Let $\text{pp} := (q, d, n, m, N, k, \mathfrak{s}, \mathcal{B})$ be a tuple of functions of the security parameter λ . Consider any efficiently computable function $f : [N] \rightarrow \mathcal{R}_q^n$. The ISIS_f assumption is defined by the experiment in Figure 3. For an adversary \mathcal{A} , we define

$$\text{Adv}_{\text{pp}}^{\text{ISIS}_f}(\mathcal{A}) = \Pr[\text{Exp}_{\text{pp}}^{\text{ISIS}_f}(\mathcal{A}) \rightarrow 1].$$

The $\text{ISIS}_f^{\text{pp}}$ assumption states that for every PPT adversary \mathcal{A} , $\text{Adv}_{\text{pp}}^{\text{ISIS}_f}(\mathcal{A})$ is negligible.

3.1 Concrete Instantiations of f

3.1.1 Random Function

An intuitive choice for $f = \text{RF}$ is a random function, modelled in the security analysis as a random oracle. In this case, if $N = \omega(\text{poly}(\lambda))$ then following the GPV signature proof [41, Proposition 6.1], ISIS_f can be tightly reduced to the plain Module-SIS assumption. Unfortunately, the issue with picking f to be random is that in our constructions we need to prove knowledge of a preimage of f . This becomes slightly awkward since then we would require proof systems that can actually prove statements related to the random oracle f .

3.1.2 Binary Encoding

In our constructions we will use the following family of functions f . Let $t \in \mathbb{N}$ and take any matrix $B \in \mathbb{Z}_q^{nd \times t}$. Then, the function is defined as:

$$f(x) := \text{Coeffs}^{-1}(B \cdot \text{enc}(x)) \in \mathcal{R}_q^n$$

where $\text{enc}(x) \in \{0, 1\}^t$ is a binary decomposition of $x - 1$. Hence, naturally $N = 2^t$. Clearly, proving knowledge of a preimage x of f , i.e. $f(x) = \mathbf{y}$, is equivalent to proving knowledge of a binary vector $\vec{u} \in \{0, 1\}^t$ such that $B\vec{u} = \text{Coeffs}(\mathbf{y})$. This, in turn, can be proven using the [57] framework. The main purpose of the matrix B is to provide flexibility when setting N and has no significant impact on the security (apart from a few naive special cases). We denote the corresponding ISIS_f problem as ISIS_{bin} . In the following, we propose a few standard attacks on ISIS_{bin} .

Relations to finding a short vector in the lattice. Regardless of the choice of f , one of the most naive attacks would be to simply try any x and then use the lattice attacks to find a short vector \mathbf{s}^* such that

$$\mathbf{A}\mathbf{s}^* = f(x)$$

which has a flavour of the Inhomogenous-MSIS problem. Now, if B is a zero (resp. identity) matrix then this corresponds to the plain (resp. Hermite Normal Form) Module-SIS [33].

We propose another attack, which this time actually makes use of the k pairs $(x_i, \mathbf{s}_i)_{i \in [k]}$. Denote $\vec{u}_i = \text{enc}(x_i)$ and $\vec{s}_i = \text{Coeffs}(\mathbf{s}_i)$ for $i \in [k]$. Then, by definition of ISIS_{bin} we have ¹⁰:

$$[\text{rot}(\mathbf{A}) - B]S = 0 \quad \text{where} \quad S := \begin{bmatrix} \vec{s}_1 & \cdots & \vec{s}_k \\ \vec{u}_1 & \cdots & \vec{u}_k \end{bmatrix} \in \mathbb{Z}_q^{(md+t) \times k}. \quad (6)$$

Here, the attack is to consider the lattice Λ generated by columns of S and find a short vector $(\vec{s}^*, \vec{u}^*) \in \Lambda$, such that $\|\vec{s}^*\| \leq \mathcal{B}$, $\vec{u}^* \in \{0, 1\}^t$ and $(\vec{s}^*, \vec{u}^*) \notin S$. We briefly explain why this allows us to solve ISIS_{bin} . Suppose we found a vector $(v_1, \dots, v_k) \in \mathbb{Z}_q^k$ such that $\vec{s}^* = \sum_{i=1}^k v_i \vec{s}_i$ and $\vec{u}^* = \sum_{i=1}^k v_i \vec{u}_i$, where \vec{s}^* and \vec{u}^* satisfy all the conditions above. Then, for $\mathbf{s}^* = \text{Coeffs}^{-1}(\vec{s}^*)$ we have

$$\mathbf{A}\mathbf{s}^* = \sum_{i=1}^k v_i \mathbf{s}^* = \sum_{i=1}^k v_i \text{Coeffs}^{-1}(B\vec{u}_i) = \text{Coeffs}^{-1} \left(\sum_{i=1}^k v_i \cdot B\vec{u}_i \right) = \text{Coeffs}^{-1}(B\vec{u}^*).$$

Finally, by the conditions above, this gives us a valid ISIS_{bin} solution.

A naive solution to obtain (\vec{s}^*, \vec{u}^*) would be to hope that S has full-rank $md + t$ for large enough instances k . Then, the adversary could pick any valid (\vec{s}^*, \vec{u}^*) and compute the linear combination $v_1, \dots, v_k \in \mathbb{Z}_q$ using linear algebra. The hardness of this problem lies in the fact that by (6) and the rank-nullity theorem, we have that the rank of S is at most nd . Since in all our applications $n \leq m$, this implies that S will never have rank at least $md + t$.

Relations to integer linear programming. Let us denote $\vec{u}_i = \text{enc}(x_i)$ for $i \in [k]$. The attack involves finding a short non-zero vector $\vec{v} \in \mathbb{Z}_q^k$ such that $\sum_{i=1}^k v_i \vec{u}_i = \vec{0} \pmod{q}$. Note that if $\vec{u}_1, \dots, \vec{u}_k$ as well as \vec{v} are sufficiently short w.r.t. q then the equation above holds over integers. Now, we can use techniques from integer linear programming (ILP) [16, 42] to find such a vector \vec{v} . In particular, Herold and May [42, Theorem 1] showed that under some mild assumptions, one can efficiently find such a binary vector $\vec{v} \in \{0, 1\}^k$, as long as $k \geq 2t$. Let $i \in [k]$ such that $v_i \neq 0$. Then, if we denote $\mathbf{s}^* := -(\sum_{j \neq i} v_j \mathbf{s}_j)$, then

$$\mathbf{A}\mathbf{s}^* = -\sum_{j \neq i} v_j \mathbf{A}\mathbf{s}_j = -\sum_{j \neq i} v_j \text{Coeffs}^{-1}(B\vec{u}_j) = v_i \text{Coeffs}^{-1}(B\vec{u}_i) = \text{Coeffs}^{-1}(B\vec{u}_i).$$

If \mathbf{s} is larger than the norm of the shortest norm of $\Lambda_q^\perp(\mathbf{A})$, then by the unpredictability argument (c.f. Lemma 3.8), with non-negligible probability we have $\mathbf{s}^* \neq \mathbf{s}_i$. Now, using the standard tail bounds (c.f. Section 2), we can upper-bound the norm of \mathbf{s}^* by $(k-1) \cdot \mathbf{s} \sqrt{md}$. Hence, if \mathcal{B} is larger than that value then $(\mathbf{s}^*, f^{-1}(\mathbf{b}_i))$ is a valid ISIS_{bin} solution¹¹. Therefore, in order to prevent such attacks, we need to set $\mathcal{B} \approx \mathbf{s} \sqrt{md}$ to be close to the tail bound of the discrete Gaussian distribution with standard deviation \mathbf{s} .

3.2 Interactive Version

We are ready to introduce the interactive version of the ISIS_f problem which will be used as an underlying cryptographic assumption for our anonymous credentials. The main difference is that we allow the adversary to make *preimage queries* with respect to adaptively chosen messages, i.e. it has access to the preimage oracle \mathcal{O}_{pre} . Here, \mathcal{O}_{pre} on input $(\mathbf{m}_i, \mathbf{r}_i)$ first samples $x_i \leftarrow [N]$ and then returns x_i along with

$$\mathbf{s}_i \leftarrow \mathbf{A}_s^{-1} \left(f(x_i) + \mathbf{C} \begin{bmatrix} \mathbf{m}_i \\ \mathbf{r}_i \end{bmatrix} \right)$$

where \mathbf{C} is a uniformly random public matrix. The adversary's goal is to come up with a new tuple $(x^*, \mathbf{s}^*, \mathbf{m}^*, \mathbf{r}^*)$ such that $\mathbf{s}^*, \mathbf{m}^*, \mathbf{r}^*$ are short vectors,

$$\mathbf{A}\mathbf{s}^* = f(x^*) + \mathbf{C} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} \quad (7)$$

¹⁰ We recall the skew-circulant matrices and the function rot in Section 2.

¹¹ Obviously, one could try to use ILP for smaller k to decrease the bound on \mathbf{s}^* .

$\text{Exp}_{\text{pp}}^{\text{Int-ISIS}_f}(\mathcal{A})$	$\mathcal{O}_{\text{pre}}(\mathbf{m}, \mathbf{r})$
1: $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$	1: if $\ (\mathbf{m}, \mathbf{r})\ < \mathcal{B}_m$
2: $\mathbf{C} \leftarrow \mathcal{R}_q^{n \times (\ell_m + \ell_r)}$	2: then return \perp
3: $\mathcal{M} = \emptyset$	3: $x \leftarrow [N]$
4: $(x^*, \mathbf{s}^*, \mathbf{m}^*, \mathbf{r}^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{pre}}}(\mathbf{A}, \mathbf{C})$	4: $\mathbf{s} \leftarrow \mathbf{A}_s^{-1} \left(f(x) + \mathbf{C} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix} \right)$
5: if $(\mathbf{m}^* \notin \mathcal{M}) \wedge \left(\mathbf{A} \mathbf{s}^* = f(x^*) + \mathbf{C} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} \right) \wedge (0 < \ \mathbf{s}^*\ \leq \mathcal{B}_s) \wedge (\ (\mathbf{m}^*, \mathbf{r}^*)\ \leq \mathcal{B}_m)$	5: $\mathcal{M} \leftarrow \mathcal{M} \cup \{\mathbf{m}^*\}$
6: then return 1	6: return (x, \mathbf{s})
7: else return 0	

Fig. 4: Interactive version of the ISIS_f problem.

and $\mathbf{m}^* \notin \{\mathbf{m}_1, \dots, \mathbf{m}_k\}$ where k is the number of oracle queries. In order to avoid trivial attacks, we require the inputs $(\mathbf{m}_i, \mathbf{r}_i)$ to the preimage oracle \mathcal{O}_{pre} to be short vectors as well.

Definition 3.2 (The Interactive ISIS_f Problem). Define public parameters $\text{pp} := (q, d, n, m, \ell_m, \ell_r, N, \mathfrak{s}, \mathcal{B}_s, \mathcal{B}_m)$ as a tuple of functions of the security parameter λ . Consider any efficiently computable function $f : [N] \rightarrow \mathcal{R}_q^n$. The Int-ISIS_f assumption is defined by the experiment in Figure 4. For an adversary \mathcal{A} , we define

$$\text{Adv}_{\text{pp}}^{\text{Int-ISIS}_f}(\mathcal{A}) = \Pr[\text{Exp}_{\text{pp}}^{\text{Int-ISIS}_f}(\mathcal{A}) \rightarrow 1].$$

The $\text{Int-ISIS}_f^{\text{pp}}$ assumption states that for every PPT adversary \mathcal{A} , $\text{Adv}_{\text{pp}}^{\text{Int-ISIS}_f}(\mathcal{A})$ is negligible.

The main result of this section is an efficient reduction from the interactive ISIS_f to the one in Definition 3.2.

Theorem 3.3 ($\text{Int-ISIS}_f \implies \text{ISIS}_f$). Let $\text{pp} := (q, d, n, m, \ell_m, \ell_r, N, \mathfrak{s}, \mathcal{B}_s, \mathcal{B}_m)$ be public parameters such that $q/2 > \mathcal{B}_m \geq 1, m = n \log q + r$, and

$$M := \exp\left(1 + \frac{1}{2\alpha^2}\right) \quad \text{and} \quad \varepsilon := 2 \left(\frac{1+\epsilon}{1-\epsilon}\right) \exp\left(-2\alpha^2 \cdot \frac{\pi-1}{\pi}\right)$$

where $\epsilon \leq 1/2, \alpha \geq 1$ and $r > 0$. Suppose $\mathfrak{s} \geq \max\left(\eta_{\min}(\epsilon), \alpha \mathcal{B}_m d \sqrt{(\ell_m + \ell_r)m}\right)$. Then, for every adversary \mathcal{A} which makes at most Q queries to \mathcal{O}_{pre} , there is an adversary \mathcal{A}' which runs in time essentially identical to \mathcal{A} and

$$\begin{aligned} \text{Adv}_{\text{pp}'}^{\text{ISIS}_f}(\mathcal{A}') &\geq \frac{1}{6Q} \text{Adv}_{\text{pp}}^{\text{Int-ISIS}_f}(\mathcal{A}) - \frac{\ell_m + \ell_r}{12Q} \sqrt{\left(1 + 2^{-r}\right)^d - 1} - \frac{2^{-\lambda}}{6} - \frac{T_{\max}^2 Q}{12N} \\ &\quad - \frac{\ell_m}{3} \cdot 2^{-rd} - \left(Q - \frac{2}{3}\right) T_{\max} \left(\frac{\varepsilon}{2M} - \frac{2\epsilon}{M}\right) - \frac{2^{-d+2}}{3} \end{aligned}$$

where $\text{pp}' := (q, d, n, m, N, T_{\max} Q, \mathfrak{s}, \mathcal{B} = \mathcal{B}_s + \mathcal{B}_m d \sqrt{(\ell_m + \ell_r)m})$ and T_{\max} satisfies $(1 - \frac{1}{M})^{T_{\max}} \leq 2^{-\lambda}$.

Proof. Assume without loss of generality that \mathcal{A} makes exactly Q queries. We will prove the statement using a hybrid argument. Namely, in each game we will either modify the execution of the \mathcal{O}_{pre} oracle, or change the execution of the challenger. In the following, define E_i to be the event that \mathcal{A} wins the Int-ISIS_f experiment.

Game₁: This is the standard Interactive ISIS_f game. Hence, $\Pr[E_1] = \varepsilon$. Let us denote the preimage oracle $\mathcal{O}_0 := \mathcal{O}_{\text{pre}}$.

Game₂: Here, instead of sampling \mathbf{C} uniformly random, the challenger first picks a uniformly random matrix of binary polynomials $\mathbf{D} \in \mathcal{R}_q^{m \times (\ell_m + \ell_r)}$ and sets $\mathbf{C} = \mathbf{AD}$.

$\mathcal{O}_1(\mathbf{m}, \mathbf{r})$	$\mathcal{O}_2(\mathbf{m}, \mathbf{r})$	$\mathcal{O}_3(\mathbf{m}, \mathbf{r})$
1: if $\ (\mathbf{m}, \mathbf{r})\ > \mathcal{B}_m$ 2: then return \perp 3: $\text{ctr} = 0$ 4: while $\text{ctr} \geq 0$ do 5: $x \leftarrow [N]$ 6: $\mathbf{s} \leftarrow \mathbf{A}_s^{-1} \left(f(x) + \mathbf{C} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix} \right)$ 7: $u \leftarrow [0, 1]$ 8: if $u \leq \frac{1}{M}$: 9: $\mathcal{M} \leftarrow \mathcal{M} \cup \{\mathbf{m}\}$ 10: return (x, \mathbf{s}) 11: $\text{ctr} = \text{ctr} + 1$ 12: return \perp	1: if $\ (\mathbf{m}, \mathbf{r})\ > \mathcal{B}_m$ 2: then return \perp 3: $\text{ctr} = 0$ 4: $\rho = \perp$ 5: while $\text{ctr} < T_{\max}$ do 6: $x \leftarrow [N]$ 7: $\mathbf{s} \leftarrow \mathbf{A}_s^{-1} \left(f(x) + \mathbf{C} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix} \right)$ 8: $u \leftarrow [0, 1]$ 9: if $u \leq \frac{1}{M} \wedge \rho = \perp$: 10: $\mathcal{M} \leftarrow \mathcal{M} \cup \{\mathbf{m}\}$ 11: $\rho = (x, \mathbf{s})$ 12: $\text{ctr} = \text{ctr} + 1$ 13: return ρ	1: if $\ (\mathbf{m}, \mathbf{r})\ > \mathcal{B}_m$ 2: then return \perp 3: $\mathbf{v} = \mathbf{D} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix}$ 4: $\text{ctr} = 0$ 5: $\rho = \perp$ 6: while $\text{ctr} < T_{\max}$ do 7: $x \leftarrow [N]$ 8: $\bar{\mathbf{s}} \leftarrow \mathbf{A}_s^{-1}(f(x))$ 9: $\mathbf{s} = \bar{\mathbf{s}} + \mathbf{v}$ 10: $u \leftarrow [0, 1]$ 11: if $u \leq \min \left(\frac{D_s^{\text{nd}}(\mathbf{s})}{M \cdot D_{s, \mathbf{v}}^{\text{nd}}(\mathbf{s})}, 1 \right) \wedge \rho = \perp$: 12: $\mathcal{M} \leftarrow \mathcal{M} \cup \{\mathbf{m}\}$ 13: $\rho = (x, \mathbf{s})$ 14: $\text{ctr} = \text{ctr} + 1$ 15: return ρ

Fig. 5: Preimage oracles \mathcal{O}_{pre} used in the hybrid argument.

Lemma 3.4. $\Pr[E_2] \geq \Pr[E_1] - \frac{\ell_m + \ell_r}{2} \sqrt{(1 + 2^{-r})^d - 1}$.

Proof. The result follows directly from the ring version of the leftover hash lemma [19, Lemma 7]. \square

Game₃: In this game, we run \mathcal{O}_{pre} as \mathcal{O}_1 defined in Figure 5.

Lemma 3.5. $\Pr[E_3] = \Pr[E_2]$.

Proof. Define CTR to be the random variable which is the value of ctr when \mathcal{O}_1 actually outputs (x, \mathbf{s}) . Then, for any (\mathbf{m}, \mathbf{r}) such that $\|(\mathbf{m}, \mathbf{r})\| \leq \mathcal{B}_m$, and any possible output (x, \mathbf{s}) we have:

$$\begin{aligned}
\Pr[(x, \mathbf{s}) \leftarrow \mathcal{O}_1(\mathbf{m}, \mathbf{r})] &= \sum_{i=0}^{\infty} \Pr[(x, \mathbf{s}) \leftarrow \mathcal{O}_1(\mathbf{m}, \mathbf{r}) \wedge \text{CTR} = i] \\
&= \sum_{i=0}^{\infty} \frac{1}{M} \left(1 - \frac{1}{M}\right)^i \Pr[(x, \mathbf{s}) \leftarrow \mathcal{O}_0(\mathbf{m}, \mathbf{r})] \\
&= \Pr[(x, \mathbf{s}) \leftarrow \mathcal{O}_0(\mathbf{m}, \mathbf{r})].
\end{aligned}$$

When $\|(\mathbf{m}, \mathbf{r})\| > \mathcal{B}_m$, then both \mathcal{O}_1 and \mathcal{O}_0 output \perp in the first step. Hence, the two oracles behave identically and thus $\Pr[E_3] = \Pr[E_2]$. \square

Game₄: In this game, we run \mathcal{O}_{pre} as \mathcal{O}_2 defined in Figure 5.

Lemma 3.6. $\Pr[E_4] \geq \Pr[E_3] - Q2^{-\lambda}$.

Proof. The only change is that we abort when ctr reaches T_{\max} . Hence, the statistical distance between $\mathcal{O}_2(\mathbf{m}, \mathbf{r})$ and $\mathcal{O}_1(\mathbf{m}, \mathbf{r})$ is at most $(1 - \frac{1}{M})^{T_{\max}} \leq 2^{-\lambda}$. Hence, $\Pr[E_4] \geq \Pr[E_3] - Q(1 - \frac{1}{M})^{T_{\max}} \geq \Pr[E_3] - Q2^{-\lambda}$. \square

Game₅: It differs from the previous game in a sense that we abort whenever \mathcal{O}_2 picks an index x , which was already sampled, even if it was not sent to the adversary. Recall that we still use the preimage oracle \mathcal{O}_2 .

Lemma 3.7. $\Pr[E_5] \geq \Pr[E_4] - \frac{T_{\max}^2 Q^2}{2N}$.

Proof. The inequality follows directly from the birthday bound. \square

Game₆: We still use the preimage oracle \mathcal{O}_2 . The only change is in the challenger execution. First, the challenger additionally keeps track of all generated variables. Namely, define

$$\mathcal{X} := \left\{ \left(x, \mathbf{s} - \mathbf{D} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix} \right) : x \text{ and } \mathbf{s} \text{ were generated while querying } \mathcal{O}_2(\mathbf{m}, \mathbf{r}) \right\}. \quad (8)$$

In particular, \mathcal{X} contains pairs which were not sent to adversary. Further,

$$\mathcal{Y} := \left\{ \left(x, \mathbf{s} - \mathbf{D} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix} \right) : x \text{ and } \mathbf{s} \text{ were sent to } \mathcal{A} \text{ by } \mathcal{O}_2(\mathbf{m}, \mathbf{r}) \right\}. \quad (9)$$

Then, we have the following simple relations: $|\mathcal{X}| = T_{\max}Q$, $|\mathcal{Y}| = |\mathcal{M}| \leq Q$ and $\mathcal{Y} \subseteq \mathcal{X}$. In this game, the challenger aborts when the output vectors from \mathcal{A} satisfy the following:

$$\left(x^*, \mathbf{s}^* - \mathbf{D} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} \right) \in \mathcal{X} \setminus \mathcal{Y}.$$

Lemma 3.8. $\Pr[\mathbf{E}_6] \geq \frac{1}{2} \cdot \Pr[\mathbf{E}_5]$.

Proof. Denote $\mathcal{X}_0 := \{x \in [N] : \exists \mathbf{s}, (x, \mathbf{s}) \in \mathcal{X}\}$ and similarly for \mathcal{Y}_0 . It is easy to see that

$$\Pr[\mathbf{E}_6 \wedge x^* \in \mathcal{Y}_0] = \Pr[\mathbf{E}_5 \wedge x^* \in \mathcal{Y}_0] \quad \text{and} \quad \Pr[\mathbf{E}_6 \wedge x^* \notin \mathcal{X}_0] = \Pr[\mathbf{E}_5 \wedge x^* \notin \mathcal{X}_0].$$

In the following, we show $\Pr[\mathbf{E}_6 \wedge x^* \in \mathcal{X}_0 \setminus \mathcal{Y}_0] \geq \frac{1}{2} \Pr[\mathbf{E}_5 \wedge x^* \in \mathcal{X}_0 \setminus \mathcal{Y}_0]$ which concludes the proof. Suppose that \mathcal{A} wins **Game₅** and the output index is in $\mathcal{X}_0 \setminus \mathcal{Y}_0$, i.e. there exists a unique \mathbf{z} such that $(x^*, \mathbf{z}) \in \mathcal{X} \setminus \mathcal{Y}$ (by conditions of **Game₅**). Because this pair does not belong to \mathcal{Y} , it was rejected during the execution of \mathcal{O}_2 on some input $(\mathbf{m}', \mathbf{r}')$. Define

$$\mathbf{z}^* := \mathbf{s}^* - \mathbf{D} \begin{bmatrix} \mathbf{m}^* - \mathbf{m}' \\ \mathbf{r}^* - \mathbf{r}' \end{bmatrix}.$$

By definition of $\mathcal{X} \setminus \mathcal{Y}$, $\mathbf{s}' := \mathbf{z} + \mathbf{D} \begin{bmatrix} \mathbf{m}' \\ \mathbf{r}' \end{bmatrix}$ is perfectly hidden among the preimages of $f(x^*) + \mathbf{C} \begin{bmatrix} \mathbf{m}' \\ \mathbf{r}' \end{bmatrix}$ under \mathbf{A} since the adversary did not actually see \mathbf{s}' . Let \mathbf{u} be a polynomial vector with ternary coefficients in $\Lambda_q^\perp(\mathbf{A})$, which exists since $m > n \log q$. Note that $\|\mathbf{u}\| \leq \sqrt{md}$ and $\mathbf{s} \geq \alpha \mathcal{B}_m d \sqrt{(\ell_m + \ell_r)m} \geq \sqrt{md} \geq \|\mathbf{u}\|$. Therefore, $\mathbf{s}^* - \mathbf{D} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} = \mathbf{z} \iff \mathbf{s}' = \mathbf{z}^*$, and also

$$\begin{aligned} \Pr \left[\mathbf{s}' = \mathbf{z}^* \mid \mathbf{E}_5 \wedge x^* \in \mathcal{X}_0 \setminus \mathcal{Y}_0 \right] &\leq \frac{\exp\left(\frac{-\|\mathbf{z}^*\|^2}{-2s^2}\right)}{\exp\left(\frac{-\|\mathbf{z}^*\|^2}{2s^2}\right) + \exp\left(\frac{-\|\mathbf{z}^* + \mathbf{u}\|^2}{2s^2}\right) + \exp\left(\frac{-\|\mathbf{z}^* + \mathbf{u}\|^2}{2s^2}\right)} \\ &\leq \frac{1}{1 + \left(\exp\left(-\frac{\langle \mathbf{z}^*, \mathbf{u} \rangle}{s^2}\right) + \exp\left(\frac{\langle \mathbf{z}^*, \mathbf{u} \rangle}{s^2}\right)\right) \exp\left(-\frac{\|\mathbf{u}\|^2}{2s^2}\right)} \\ &\leq \frac{1}{1 + 2 \exp\left(-\frac{\|\mathbf{u}\|^2}{2s^2}\right)} \\ &\leq \frac{1}{1 + 2 \exp\left(-\frac{1}{2}\right)} \leq \frac{1}{2}. \end{aligned}$$

By taking the complement, we deduce that

$$\Pr \left[\mathbf{s}^* - \mathbf{D} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} \neq \mathbf{z} \mid \mathbf{E}_5 \wedge x^* \in \mathcal{X}_0 \setminus \mathcal{Y}_0 \right] \geq \frac{1}{2}.$$

Finally, note that if $\mathbf{E}_5 \wedge x^* \in \mathcal{X}_0 \setminus \mathcal{Y}_0$ hold and additionally $\mathbf{s}^* - \mathbf{D} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} \neq \mathbf{z}$, then we must have

$$\left(x^*, \mathbf{s}^* - \mathbf{D} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} \right) \notin \mathcal{X} \setminus \mathcal{Y}$$

Game₉

1. Generate $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$.
2. Sample $\mathbf{D} \leftarrow \{0, 1\}^{md \times (\ell_m + \ell_r)d}$ and set $\mathbf{C} = \mathbf{AD}$.
3. Set $\mathcal{M} = \emptyset$.
4. Sample $j^* \leftarrow [Q]$.
5. Run $(x^*, \mathbf{s}^*, \mathbf{m}^*, \mathbf{r}^*) \leftarrow \mathcal{A}^{\mathcal{O}'}$ (\mathbf{A}, \mathbf{C}), where \mathcal{O}' is defined below.
6. Define $\tilde{\mathbf{s}} = \mathbf{s}^* - \mathbf{D} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix}$
7. If all the conditions below hold, then return 1 and 0 otherwise:
 - $\mathbf{m}^* \notin \mathcal{M}$ and $\|(\mathbf{m}^*, \mathbf{r}^*)\| \leq \mathcal{B}_m$ and $\|\mathbf{s}^*\| \leq \mathcal{B}_s$ (from the Int-ISIS_f game),
 - $\mathbf{A}\tilde{\mathbf{s}} = f(x^*)$ (from the Int-ISIS_f game),
 - $\forall (x, \mathbf{s}) \neq (x', \mathbf{s}') \in \mathcal{X}, x \neq x'$ (from Game_5),
 - $(x^*, \tilde{\mathbf{s}}) \notin \mathcal{X} \setminus \mathcal{Y}$ (from Game_6),
 - if $x^* \in \mathcal{Y}_0$, then x^* must have been generated in the j^* -th oracle query (from Game_7),
 - $(x^*, \tilde{\mathbf{s}}) \notin \mathcal{Y}$ (from Game_9).

Fig. 6: Description of Game_9 . Here, \mathcal{O}' behaves as \mathcal{O}_2 , apart from the j^* -th query where it behaves like \mathcal{O}_3 (defined in Figure 5.)

since $(x^*, \mathbf{z}) \in \mathcal{X} \setminus \mathcal{Y}$ and there is at most one pair in \mathcal{X} of the form (x^*, \cdot) by the condition introduced in Game_5 . This implies that

$$\Pr[\mathbf{E}_6 \wedge x^* \in \mathcal{X}_0 \setminus \mathcal{Y}_0] \geq \frac{1}{2} \Pr[\mathbf{E}_5 \wedge x^* \in \mathcal{X}_0 \setminus \mathcal{Y}_0]$$

which concludes the proof. \square

Game₇: At the beginning, the challenger samples an index $j^* \leftarrow [Q]$. Then at the end, if $x^* \in \mathcal{Y}_0$ and x^* was *not* sampled in the j^* -th query then the challenger aborts. We still use the preimage oracle \mathcal{O}_2 .

Lemma 3.9. $\Pr[\mathbf{E}_7] \geq \frac{1}{Q} \Pr[\mathbf{E}_6]$.

Proof. Clearly, if $x^* \notin \mathcal{Y}_0$ then the game behaves identically as before. However, if $x^* \in \mathcal{Y}$ then with probability at least $1/Q$ we have that x^* was sampled during the j^* -th oracle query. \square

Game₈: In this game, only for the j^* -th oracle query, instead of executing \mathcal{O}_2 we run \mathcal{O}_3 as in Figure 5.

Lemma 3.10. $\Pr[\mathbf{E}_8] \geq \Pr[\mathbf{E}_7] - T_{\max} \left(\frac{\epsilon}{2M} + \frac{2\epsilon}{M} \right) - 2^d$.

Proof. First, we exclude the case that $\eta'_\epsilon(\Lambda_q^\perp(\mathbf{A})) > \eta_{\min}$ which occurs with probability 2^{-d} . Then, by our parameter selection and Lemma 2.5, the statistical distance between $\mathcal{O}_3(\mathbf{m}, \mathbf{r})$ and $\mathcal{O}_2(\mathbf{m}, \mathbf{r})$ is at most $\frac{\epsilon}{2M} + \frac{2\epsilon}{M}$. Hence, the statement follows by the hybrid argument. \square

Game₉: For presentation purposes, we summarise the security game in Figure 6. The difference from the previous game is that the adversary loses whenever

$$\left(x^*, \mathbf{s}^* - \mathbf{D} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} \right) \in \mathcal{Y}.$$

Lemma 3.11. $\Pr[\mathbf{E}_9] \geq \frac{1}{3} \Pr[\mathbf{E}_8] - \frac{\ell_m \cdot 2^{-rd}}{3}$.

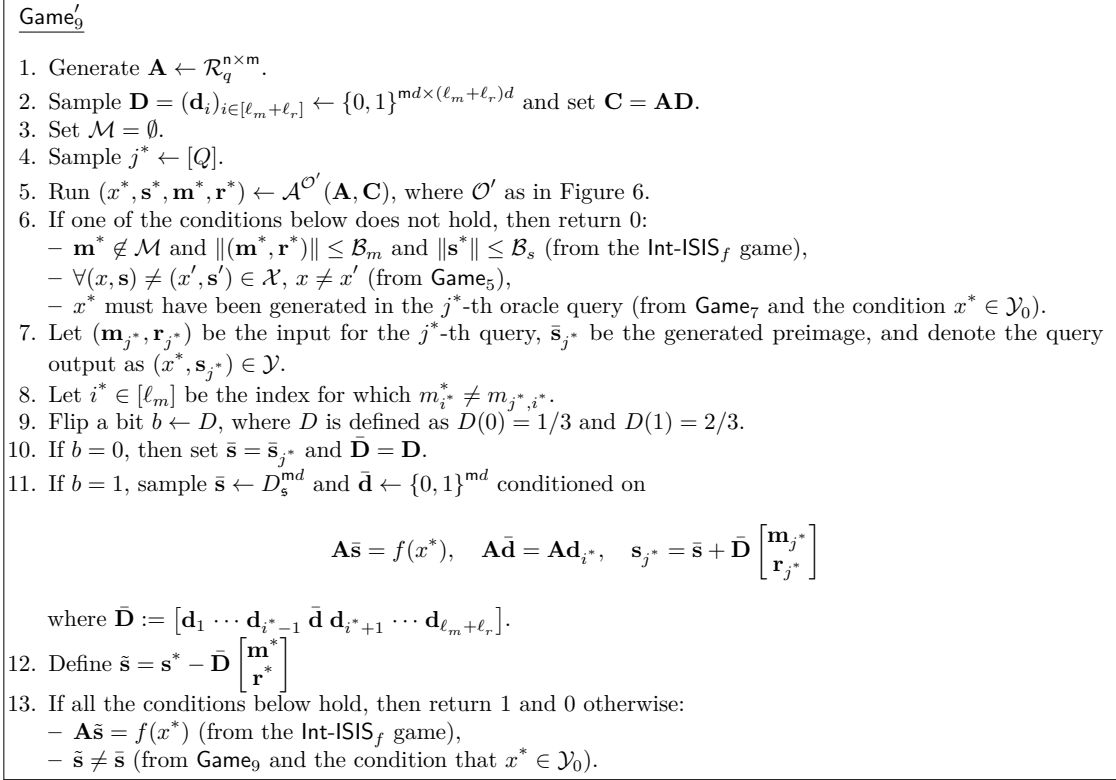


Fig. 7: Description of the alternative game Game'₉.

Proof. We follow the proof strategy from [54, Theorem 3.2]. Note that conditioned on $x^* \notin \mathcal{Y}_0$ the two security games are identical. Hence, in the following we show

$$\Pr[\mathbf{E}_9 \wedge x^* \in \mathcal{Y}_0] \geq \frac{1}{3} \Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0] - \frac{\ell_m \cdot 2^{-rd}}{3}.$$

Let us consider the following alternative game Game'₉ defined in Figure 7 and denote \mathbf{E}'_9 to be the event that the adversary wins Game'₉. We claim that $\Pr[\mathbf{E}'_9] = \Pr[\mathbf{E}_9 \wedge x^* \in \mathcal{Y}_0]$. Indeed, the only change is that the challenger *resamples* the vectors $\bar{\mathbf{s}}_{j^*}$ (which is the preimage generated in the j^* -th query) and \mathbf{D} according to their original distributions conditioned on what the adversary \mathcal{A} already knows (e.g. it knows the value \mathbf{s}_j). Hence, from now on we focus on Game'₉.

Let us denote $\tilde{\mathbf{s}}_{j^*} := \mathbf{s}^* - \mathbf{D} \begin{bmatrix} \mathbf{m}_{j^*}^* \\ \mathbf{r}_{j^*}^* \end{bmatrix}$. In the following, we will use the fact that when $b = 0$, we have $\tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}$.

Next, we partition the success probability into two parts: $\Pr[\mathbf{E}'_9] = \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} \neq \mathbf{s}_j] + \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} = \mathbf{s}_j]$. Now, one observes that

$$\begin{aligned} \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} \neq \mathbf{s}_j] &\geq \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} \neq \bar{\mathbf{s}}_{j^*} \wedge b = 0] \\ &\geq \frac{1}{3} \cdot \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} \neq \bar{\mathbf{s}}_{j^*} | b = 0] \\ &\geq \frac{1}{3} \Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \tilde{\mathbf{s}} \neq \bar{\mathbf{s}} \wedge \tilde{\mathbf{s}}_{j^*} \neq \bar{\mathbf{s}}_{j^*} | b = 0] \\ &\geq \frac{1}{3} \Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \tilde{\mathbf{s}}_{j^*} \neq \bar{\mathbf{s}}_{j^*}]. \end{aligned}$$

For the second part, let us denote

$$S(\mathbf{D}, i^*) := \left\{ [\mathbf{d}_1 \cdots \mathbf{d}_{i^*-1} \quad \mathbf{e} \quad \mathbf{d}_{i^*+1} \cdots \mathbf{d}_{\ell_m + \ell_r}] : \mathbf{A}\mathbf{e} = \mathbf{A}\mathbf{d}_{i^*} \wedge \mathbf{e} \in \{0, 1\}^{(\ell_m + \ell_r)d} \right\}.$$

Clearly, $\mathbf{D} \in S(\mathbf{D}, i^*)$. Further,

$$\begin{aligned} \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} = \mathbf{s}_j] &\geq \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*} \wedge |S(\mathbf{D}, i^*)| \geq 2] \\ &\geq \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*} \wedge |S(\mathbf{D}, i^*)| \geq 2 \wedge b = 1]. \end{aligned}$$

For presentation purposes, define the event

$$E := (\mathbf{E}_8) \wedge (x^* \in \mathcal{Y}_0) \wedge (\tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*}) \wedge (|S(\mathbf{D}, i^*)| \geq 2) \wedge (b = 1).$$

Then,

$$\begin{aligned} \Pr[E] &\geq \frac{2}{3} \Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*} \wedge |S(\mathbf{D}, i^*)| \geq 2] \\ &\geq \frac{2}{3} (\Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*}] - \Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*} \wedge |S(\mathbf{D}, i^*)| = 1]) \\ &\geq \frac{2}{3} \left(\Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*}] - \ell_m \left(\frac{q^n}{2^m} \right)^d \right) \end{aligned}$$

where for the last inequality we used the fact from [53, Lemma 5.2] to deduce that

$$\begin{aligned} \Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*} \wedge |S(\mathbf{D}, i^*)| = 1] &\leq \Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*} \wedge \exists i, |S(\mathbf{D}, i)| = 1] \\ &\leq \Pr[\exists i, |S(\mathbf{D}, i)| = 1] \\ &\leq \sum_{i=1}^{\ell_m} \Pr[|S(\mathbf{D}, i)| = 1] \\ &\leq \ell_m \left(\frac{q^n}{2^m} \right)^d = \ell_m \cdot 2^{-rd}. \end{aligned}$$

Thus,

$$\begin{aligned} \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*}] &\geq \Pr[\mathbf{E}'_9 \wedge \tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*} \wedge |S(\mathbf{D}, i^*)| \geq 2 \wedge b = 1] \\ &\geq \Pr[\tilde{\mathbf{s}} \neq \bar{\mathbf{s}} | E] \cdot \Pr[E]. \end{aligned}$$

We claim that

$$\Pr[\tilde{\mathbf{s}} = \bar{\mathbf{s}} | E] \leq \Pr[\bar{\mathbf{s}} = \bar{\mathbf{s}}_{j^*} \wedge \bar{\mathbf{D}} = \mathbf{D} | E].$$

Indeed, suppose $\tilde{\mathbf{s}} = \bar{\mathbf{s}}$ which by definition implies

$$\mathbf{s}^* = \bar{\mathbf{s}} + \bar{\mathbf{D}} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix}.$$

We know from E that $\tilde{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*}$, i.e.

$$\bar{\mathbf{s}}_{j^*} + \mathbf{D} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} = \mathbf{s}^* = \bar{\mathbf{s}} + \bar{\mathbf{D}} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix}.$$

Also, we know from Step 11 of Figure 7 that:

$$\bar{\mathbf{s}}_{j^*} + \mathbf{D} \begin{bmatrix} \mathbf{m}_{j^*}^* \\ \mathbf{r}_{j^*}^* \end{bmatrix} = \mathbf{s}_{j^*} = \bar{\mathbf{s}} + \bar{\mathbf{D}} \begin{bmatrix} \mathbf{m}_{j^*}^* \\ \mathbf{r}_{j^*}^* \end{bmatrix}.$$

By subtracting the two equations, we end up with

$$(\mathbf{D} - \bar{\mathbf{D}}) \begin{bmatrix} \mathbf{m}^* - \mathbf{m}_{j^*}^* \\ \mathbf{r}^* - \mathbf{r}_{j^*}^* \end{bmatrix} = \mathbf{0}.$$

Since both matrices $\mathbf{D}, \bar{\mathbf{D}}$ belong to $S(\mathbf{D}, i^*)$, we have $(m_{i^*}^* - m_{j^*, i^*})(\mathbf{d}_{j, i^*} - \bar{\mathbf{d}}_{i^*}) = \mathbf{0}$ where $\mathbf{d}_{i^*}, \bar{\mathbf{d}}_{i^*}$ are the i^* -th columns of $\mathbf{D}, \bar{\mathbf{D}}$ respectively. Since $2\mathcal{B}_m < q$, we deduce that this equation holds over integers. Further, we know that $m_{i^*}^* \neq m_{j^*, i^*}$ which implies that $\mathbf{d}_{i^*} = \bar{\mathbf{d}}_{i^*}$ and thus $\mathbf{D} = \bar{\mathbf{D}}$. Consequently, $\bar{\mathbf{s}} = \bar{\mathbf{s}}_{j^*}$ and the claim follows. Hence, we have

$$\Pr[\bar{\mathbf{s}} = \bar{\mathbf{s}}|E] \leq \Pr[\bar{\mathbf{s}} = \bar{\mathbf{s}}_{j^*} \wedge \bar{\mathbf{D}} = \mathbf{D}|E] \leq \Pr[\bar{\mathbf{D}} = \mathbf{D}|E] \leq \frac{1}{2}.$$

This is because E includes the event that $S(\mathbf{D}, i^*)$ contains at least two elements. Therefore, we conclude that

$$\begin{aligned} \Pr[\mathbf{E}'_9 \wedge \bar{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*}] &\geq \frac{1}{2} \cdot \frac{2}{3} \left(\Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \bar{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*}] - 2^{-rd} \right) \\ &\geq \frac{1}{3} \Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0 \wedge \bar{\mathbf{s}}_{j^*} = \bar{\mathbf{s}}_{j^*}] - \frac{\ell_m \cdot 2^{-rd}}{3} \end{aligned}$$

and thus

$$\Pr[\mathbf{E}_9 \wedge x^* \in \mathcal{Y}_0] \geq \Pr[\mathbf{E}'_9] \geq \frac{1}{3} \Pr[\mathbf{E}_8 \wedge x^* \in \mathcal{Y}_0] - \frac{\ell_m \cdot 2^{-rd}}{3}$$

which finishes the proof. \square

Game₁₀: In this game, throughout the experiment \mathcal{A} queries \mathcal{O}_{pre} defined as \mathcal{O}_3 in Figure 5 (not only for the j^* -th query).

Lemma 3.12. $\Pr[\mathbf{E}_{10}] \geq \Pr[\mathbf{E}_9] - (Q-1)T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} \right) - 2^{-d}$.

Proof. It follows identically as in Lemma 3.10. \square

Finally, our reduction runs as a challenger in **Game₁₀**. Namely, it first gets $T_{\max}Q$ pairs $(x_i, \mathbf{s}_i)_{i \in [T_{\max}Q]}$ from the ISIS_f challenger, i.e. $\mathbf{As} = f(x_i)$, and then uses them to simulate the query outputs from \mathcal{O}_3 . Note that if \mathcal{A} wins **Game₁₀** then the set \mathcal{X} contains exactly these pairs, and further the output $(x^*, \bar{\mathbf{s}})$ does not belong to \mathcal{X} (from **Game₆** and **Game₉**). Also,

$$\|\bar{\mathbf{s}}\| \leq \|\mathbf{s}^*\| + \left\| \mathbf{D} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix} \right\| \leq \mathcal{B}_s + \mathcal{B}_m d \sqrt{(\ell_m + \ell_r)\mathbf{m}} = \mathcal{B}.$$

Therefore, the reduction outputs a valid ISIS_f solution. Now, using all the previous lemmas, the probability that the reduction wins the ISIS_f game can be lower-bounded as

$$\begin{aligned} \Pr[\mathbf{E}_{10}] &\geq \Pr[\mathbf{E}_9] - (Q-1)T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} \right) - 2^{-d} \\ &\geq \frac{1}{3} \Pr[\mathbf{E}_8] - \frac{\ell_m \cdot 2^{-rd}}{3} - (Q-1)T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} \right) - 2^{-d} \\ &\geq \frac{1}{3} \Pr[\mathbf{E}_7] - \frac{\ell_m \cdot 2^{-rd}}{3} - \left(Q - \frac{2}{3} \right) T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} \right) - \frac{2^{-d+2}}{3} \\ &\geq \frac{1}{3Q} \Pr[\mathbf{E}_6] - \frac{\ell_m \cdot 2^{-rd}}{3} - \left(Q - \frac{2}{3} \right) T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} \right) - \frac{2^{-d+2}}{3} \\ &\geq \frac{1}{6Q} \Pr[\mathbf{E}_5] - \frac{\ell_m \cdot 2^{-rd}}{3} - \left(Q - \frac{2}{3} \right) T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} \right) - \frac{2^{-d+2}}{3} \\ &\geq \frac{1}{6Q} \Pr[\mathbf{E}_4] - \frac{T_{\max}^2 Q}{12N} - \frac{\ell_m \cdot 2^{-rd}}{3} - \left(Q - \frac{2}{3} \right) T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} \right) - \frac{2^{-d+2}}{3} \\ &\geq \frac{1}{6Q} \Pr[\mathbf{E}_3] - \frac{2^{-\lambda}}{6} - \frac{T_{\max}^2 Q}{12N} - \frac{\ell_m \cdot 2^{-rd}}{3} - \left(Q - \frac{2}{3} \right) T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} \right) - \frac{2^{-d+2}}{3} \\ &\geq \frac{1}{6Q} \Pr[\mathbf{E}_2] - \frac{2^{-\lambda}}{6} - \frac{T_{\max}^2 Q}{12N} - \frac{\ell_m \cdot 2^{-rd}}{3} - \left(Q - \frac{2}{3} \right) T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} \right) - \frac{2^{-d+2}}{3} \\ &\geq \frac{1}{6Q} \Pr[\mathbf{E}_1] - \text{AddLoss} \end{aligned}$$

where

$$\begin{aligned} \text{AddLoss} := & \frac{\ell_m + \ell_r}{12Q} \sqrt{(1 + 2^{-r})^d - 1} + \frac{2^{-\lambda}}{6} + \frac{T_{\max}^2 Q}{12N} + \frac{\ell_m \cdot 2^{-rd}}{3} \\ & + \left(Q - \frac{2}{3}\right) T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M}\right) + \frac{2^{-d+2}}{3} \end{aligned}$$

which concludes the proof. \square

Asymptotic parameter selection. Let $d = O(\lambda)$, $\alpha = O(\sqrt{\lambda})$ and $\epsilon = \text{negl}(\lambda)$. Then, $M = O(1)$, $\varepsilon = \text{negl}(\lambda)$ and $T_{\max} = O(\lambda)$.

Suppose we want Int-ISIS_f to be hard against efficient adversaries, and thus $Q = \text{poly}(\lambda)$. This implies that we need N to be exponential in λ . Next, to use the leftover hash lemma, we set $r \geq \omega(\log d)$. With these parameters, if a PPT adversary \mathcal{A} solves Int-ISIS_f with probability δ , then there is a PPT adversary which solves ISIS_f with probability $\delta/(12Q) - \text{negl}(\lambda)$.

3.3 Applications to Exotic Signatures

We briefly explain how to build digital signatures [41, 69], group signatures [11, 68] and blind signatures [39, 3, 67] from the (interactive) ISIS_f assumption.

Signature schemes. We can directly build standard model signatures using the hash-and-sign GPV framework. Namely, the secret key is a trapdoor for the public matrix \mathbf{A} . Let \mathbf{C} be another uniformly random matrix (which corresponds to the \mathbf{C} in the interactive ISIS_f problem). Then, to sign a (short) message \mathbf{m} , we sample a uniformly random $x \leftarrow [N]$ and compute a short preimage \mathbf{s} which satisfies $\mathbf{A}\mathbf{s} = f(x) + \mathbf{C}\mathbf{m}$. Thus, the signature is a pair (x, \mathbf{s}) . Unforgeability follows directly from the Int-ISIS_f assumption.

Group signatures. Let L be the size of the group and for $i \in [L]$, let \mathbf{m}_i to be the polynomial vector whose coefficient vector is the binary representation of i . We follow the standard sign-and-encrypt approach. Namely, the setup authority generates matrices (\mathbf{A}, \mathbf{C}) , along with the trapdoor for preimage sampling. Then, for each user $i \in [L]$, it generates a uniformly random $x_i \leftarrow [N]$ and computes a short preimage \mathbf{s}_i that satisfies $\mathbf{A}\mathbf{s}_i = f(x_i) + \mathbf{C}\mathbf{m}_i$. Hence, the signature for user i consists of a zero-knowledge proof of knowledge π that $i \in [L], x_i \in [N]$, and short \mathbf{s}_i such that $\mathbf{A}\mathbf{s}_i = f(x_i) + \mathbf{C}\mathbf{m}_i$ ¹². Since we obtain π by applying the Fiat-Shamir transformation to an interactive proof, the message to sign is included as an input to the hash function.

Blind signatures. We follow the Fischlin framework for constructing round-optimal blind signatures. The public key are the matrices (\mathbf{A}, \mathbf{C}) , while the secret key is a trapdoor for \mathbf{A} . Suppose the user wants to obtain a signature on a short message vector \mathbf{m} . It first computes the Ajtai commitment [4] on \mathbf{m} , i.e. $\mathbf{u} := \mathbf{C} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix}$, where \mathbf{r} is a fresh, short randomness vector, and sends \mathbf{u} to the signer along with a proof π_u of well-formedness of \mathbf{u} . The signer, who possess the trapdoor for matrix \mathbf{A} , samples $x \leftarrow [N]$ and returns to the user a short vector \mathbf{s} which satisfies $\mathbf{A}\mathbf{s} = f(x) + \mathbf{u}$. Finally, the user outputs as a signature a zero-knowledge proof of knowledge π_s of $x \in [N]$, short randomness \mathbf{r} and a short vector \mathbf{s} such that

$$\mathbf{A}\mathbf{s} = f(x) + \mathbf{C} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix},$$

where \mathbf{m} is a part of the statement. This is also the approach we take for constructing anonymous credentials in Section 8 (see Section 4 for more background on anonymous credentials).

¹² The signature should also contain a verifiable encryption of i using the opener's public key.

4 Background on Anonymous Credentials

We recall the definitions of the anonymous credentials. The entities involved in the scheme are the following: (i) the issuer \mathcal{I} certifies attributes of holders, (ii) the holder \mathcal{H} receives credentials about their attributes and uses them to authenticate to third parties, and (iii) the verifier \mathcal{V} engages in a protocol with a holder \mathcal{H} , and learns a subset of the attributes of the holder.

Definition 4.1 (Anonymous Credentials). *An anonymous credentials scheme is a triple $(\text{Init}, \text{Issue}_{\mathcal{I}, \mathcal{H}}, \text{Verify}_{\mathcal{H}, \mathcal{V}})$ of algorithms defined as follows:*

- $(\text{ipk}, \text{isk}) \leftarrow \text{Init}(1^\lambda, 1^l)$: *the initialisation algorithm takes as input the security parameter 1^λ and the nonzero size l (in unary) of the array of attributes the issuer will certify, and generates the issuer public key ipk and the issuer secret key isk .*
- $\text{cred}/\perp \leftarrow \text{Issue}_{\mathcal{I}, \mathcal{H}}((\text{isk}, \text{attrs}', \text{idx}), (\text{ipk}, \text{attrs}, \text{idx}))$: *this is the issuing protocol between issuer \mathcal{I} and holder \mathcal{H} . The holder has as input the issuer public key ipk , an array $\text{attrs} := [\mathbf{a}_i]_{i \in [l]}$ of attributes $\mathbf{a}_i \in \{0, 1\}^*$ and an array $\text{idx} \in [l]^{l'}$, $1 \leq l' \leq l$ of l' pairwise distinct indices for the attributes in attrs that will be revealed to the issuer. The issuer has as input the issuer secret key isk , an array of disclosed attributes $\text{attrs}' := [\mathbf{a}'_i]_{i \in \text{idx}}$ and indices idx . The holder outputs a credential cred over attributes attrs if $\forall i \in \text{idx}, \text{attrs}'_i = \text{attrs}_i$ or \perp otherwise.*
- $(\perp, 1/0) \leftarrow \text{Verify}_{\mathcal{H}, \mathcal{V}}((\text{cred}, \text{attrs}, \text{idx}), (\text{ipk}, \text{attrs}', \text{idx}))$: *this is a protocol between a holder \mathcal{H} and a verifier \mathcal{V} . The holder \mathcal{H} has as input a credential cred , an array $\text{attrs} := [\mathbf{a}_i]_{i \in [l]}$ of attributes $\mathbf{a}_i \in \{0, 1\}^*$ and an array $\text{idx} \in [l]^{l'}$, $1 \leq l' \leq l$ of l' pairwise distinct indices for the attributes in attrs that will be revealed to the verifier. The verifier \mathcal{V} has as input the issuer public key ipk of issuer \mathcal{I} , an array of disclosed attributes $\text{attrs}' := [\mathbf{a}'_i]_{i \in \text{idx}}$ and indices idx . The verifier outputs 1 if cred is a valid credential from issuer \mathcal{I} over attributes attrs such that $\forall i \in \text{idx}, \text{attrs}'_i = \text{attrs}_i$, and 0 otherwise.*

Related work. Anonymous credential systems offer to individuals the means to achieve control over the way information about themselves is exchanged for purposes such as authentication to third parties and payments. Camenisch and Lysyanskaya [24] present one of the first formal definitions and practical constructions for anonymous credential systems. The protocols that realise these schemes enable holders to receive credentials over their attributes from issuers, and later present subsets of these attributes to verifiers. These interactions preserve the privacy of holders upon issuance (since the holder can hide a subset of the certified attributes) and upon verification (since verifications are unlinkable and support the disclosure of subsets of the certified attributes). Subsequent constructions based on CL signature scheme [23], BBS+ signature scheme [14, 8] and associated efficient proof system [20], or on signatures based on SXDH assumption [49] are integrated into prominent open-source self-sovereign identity projects such as Hyperledger Indy [44], Hyperledger Aries [43] or the MATTR stack [63], or are proposed as an IETF standard draft [50]. These more recent constructions are based on a combination of signature schemes supporting blind signing (to achieve issuer anonymity) and multi-message (to efficiently accumulate multiple attributes into a single signature); commitment schemes and efficient zero-knowledge proof systems (to achieve unlinkable presentations and prove properties about pairs of attributes).

Additional aspects. The definitions above are minimalistic: while they capture the essential features of anonymous credential systems (blind issuance for a subset of attributes, unlinkability of verifications and selective disclosure of attributes), they exclude several aspects that feature in prominent related works, whose absence we review and justify here. The first aspect relates to pseudonyms [26, 27, 29, 51], which are optional, verifier-specific identifiers that a holder can present upon verification. Pseudonyms enable selective tracking of holders, since verifiers can link verifications that used the same pseudonym. Pseudonyms at different verifiers cannot be linked together, and verifier-specific pseudonyms cannot be forged arbitrarily by holders to fraudulently present different personas. Pseudonyms are typically realised by adding a user initialisation and registration phase during which the user generates a private key and has this key blindly signed by issuers; pseudonyms can then be derived by committing the private key in a non-hiding way against a verifier-specific

base and proving equality between the signed private key and the committed one. While neither user private keys nor pseudonyms are captured in our definitions, we remark that these can be easily integrated on top of any scheme that instantiates our definition: user secret key may simply be one of the undisclosed attributes upon issuance; and verification may be extended to include commitments to undisclosed attributes alongside proofs of equality between committed and signed attributes to achieve pseudonyms. A closely related aspect is that of all-or-nothing transferability, that requires that sharing a credential amongst colluding holders require sharing private key material, which represents a strong incentive not to engage in such practices. This is once again typically achieved by including private keys in the issuance, whose knowledge must be proved upon verification. Such keys might further be stored in HSMs or similar hardware facilities to make the sharing of these keys impossible. Another aspect relates to enforceable one-show credentials [24]. This property has meanwhile been decoupled from identity system proper, and has been further developed in payment-systems under the wider notion of double-spending resistance. Finally, the topic of revocation of anonymous credentials has received a lot of attention. The approaches proposed in the literature [25, 21, 22] can be typically integrated on top of a basic scheme, provided that it supports blind issuance and proof of knowledge of equality of committed attributes. These can be simply integrated on top of our construction, as it shall become apparent later.

Security Definitions. Informally, we require the scheme to be secure against any coalition comprising a malicious issuer, malicious holders and verifiers who attempt to learn more information about a victim holder (e.g. learn an undisclosed attribute upon verification) or link the same victim holder across multiple verifications. For example, if a holder repeatedly authenticates disclosing only their vaccination status, no coalition of malicious verifiers/holders should learn any information about any of the other certified attributes, nor should they be able to learn whether the same holder or two different ones disclosed twice that same vaccination status. We also require credentials to be unforgeable, constraining the ability of any coalition of malicious holders to successfully pass verification against an honest verifier for an array of attributes different from those that were certified by an honest issuer.

In what follows we present definitions of *correctness*, *anonymity* and *one-more unforgeability*.

Definition 4.2 (Correctness). An anonymous credential scheme is correct if for any $\lambda \in \mathbb{N}$, $l \in \mathbb{N}$, $l' \in [l]$, $l'' \in [l]$, any $\text{idx}' \in [l]^{l'}$ such that its elements are pairwise distinct, any $\text{idx}'' \in [l]^{l''}$ such that its elements are pairwise distinct, any array $\text{attrs} := [a_0, \dots, a_{l-1}]$, $\text{attrs}' := [a'_0, \dots, a'_{l'-1}]$ and $\text{attrs}'' := [a''_0, \dots, a''_{l''-1}]$ of attributes such that $\forall i \in \text{idx}', \text{attrs}_i = \text{attrs}'_i$ and $\forall i \in \text{idx}'', \text{attrs}_i = \text{attrs}''_i$:

$$\Pr \left[(\text{ipk}, \text{isk}) \leftarrow \text{Init}(1^\lambda, l) : \perp \leftarrow \text{Issue}_{\mathcal{I}, \mathcal{H}}((\text{isk}, \text{attrs}', \text{idx}'), (\text{ipk}, \text{attrs}, \text{idx}')) \right] \leq \text{negl}(\lambda) \quad (10)$$

and

$$\Pr \left[\begin{array}{c} (\text{ipk}, \text{isk}) \leftarrow \text{Init}(1^\lambda, l), \\ \text{cred} \leftarrow \text{Issue}_{\mathcal{I}, \mathcal{H}}((\text{isk}, \text{attrs}', \text{idx}'), (\text{ipk}, \text{attrs}, \text{idx}')) \end{array} : \begin{array}{c} \langle \perp, 0 \rangle \leftarrow \\ \text{Verify}_{\mathcal{H}, \mathcal{V}}(\langle \text{cred}, \text{attrs}, \text{idx}'' \rangle, \\ (\text{ipk}, \text{attrs}'', \text{idx}'')) \end{array} \right] \leq \text{negl}(\lambda). \quad (11)$$

Definition 4.3 (Anonymity Game). The anonymity game is played between adversary \mathcal{A} and challenger \mathcal{C} .

1. *Setup.* The adversary receives the security parameter 1^λ and outputs the issuer public key ipk .
2. *Challenge.* \mathcal{A} chooses two arrays of attributes attrs_0 and attrs_1 such that $|\text{attrs}_0| = l$ and $|\text{attrs}_1| = l$, and an array of indices $\text{idx}^* \in [l]^{l'}$, $1 \leq l' \leq l$ such that $\forall i \in \text{idx}^*, \text{attrs}_{0,i} = \text{attrs}_{1,i}$. The challenger \mathcal{C} picks $b \in \{0, 1\}$. Then, \mathcal{A} and \mathcal{C} run the $\text{Issue}_{\mathcal{I}, \mathcal{H}}(\langle \cdot, \cdot, \cdot \rangle, (\text{ipk}, \text{attrs}_b, \text{idx}^*))$ protocol, where the challenger plays the role of the holder on input attribute array attrs_b and revealed indices idx^* and the adversary plays the role of the issuer. If the protocol succeeds, the challenger stores the generated cred_b . Next, the adversary and the challenger run the $\text{Verify}_{\mathcal{H}, \mathcal{V}}(\langle \text{cred}_b, \text{attrs}_b, \text{idx}^* \rangle, \langle \cdot, \cdot, \cdot \rangle)$ protocol, where the challenger plays the role of the holder on input the credential cred_b generated previously, attribute array attrs_b and revealed indices idx^* and the adversary plays the role of the verifier.

3. *Response.* \mathcal{A} outputs b' and wins the game if $b' = b$.

Definition 4.4 (Anonymity). A credential scheme has the *Anonymity* property if no PPT adversary \mathcal{A} can win the Anonymity game with advantage greater than $\frac{1}{2} + \text{negl}(\lambda)$.

Definition 4.5 (One-More Unforgeability Game). The one-more unforgeability game is played between adversary \mathcal{A} and challenger \mathcal{C} .

1. *Setup.* The challenger runs the `Init` algorithm, sends `ipk` to the adversary and stores `isk` and `QIssue`.
2. *Queries.* \mathcal{A} has access to the following oracle:
 - $\mathcal{O}_{\text{Issue}}(\mathcal{H}, \text{attrs}, \text{idx})$: adversary \mathcal{A} requests the creation of a credential for holder \mathcal{H} over disclosed attributes `attrs` at indices `idx`. Then, \mathcal{C} and \mathcal{A} run the `IssueT, H`((`isk`, `attrs`, `idx`), (\cdot, \cdot, \cdot)) protocol where the challenger plays the role of the issuer \mathcal{T} on input the issuer secret key `isk`, the requested attributes `attrs` and indices `idx`, and the adversary plays the role of the holder. If it succeeds, the issuer increments `QIssue`.
3. *Response.* We say that \mathcal{A} wins the one-more unforgeability game if it outputs $\mathcal{L} := \{\langle \text{attrs}_i, \text{idx}_i \rangle\}$ of size `QIssue` + 1 such that the following conditions hold:
 - for all $i \in [Q_{\text{Issue}} + 1]$, $\langle \perp, 1 \rangle \leftarrow \text{Verify}_{\mathcal{H}, \mathcal{V}}(\langle \cdot, \cdot, \cdot \rangle, (\text{ipk}, \text{attrs}_i, \text{idx}_i))$ where the challenger plays the role of the verifier on input the issuer public key `ipk` and each of the pairs of attributes and indices $\langle \text{attrs}_i, \text{idx}_i \rangle \in \mathcal{L}$, and the adversary plays the role of the holder,
 - for each distinct pairs $\langle \text{attrs}_i, \text{idx}_i \rangle \in \mathcal{L}$ and $\langle \text{attrs}_j, \text{idx}_j \rangle \in \mathcal{L}$, there exists at least one index i^* in both `idxi` and `idxj` such that $\text{attrs}_{i, i^*} \neq \text{attrs}_{j, i^*}$.

Definition 4.6 (One-More Unforgeability). A credential scheme is one-more unforgeable if no PPT adversary \mathcal{A} can win the one-more unforgeability game with greater than negligible advantage.

5 Proof of Knowledge of a Preimage of SIS_{bin}

In the construction of anonymous credentials, we will need to provide a proof of knowledge of the ISIS_{bin} secret, as in Equation 7. For the interactive ISIS_{bin} case, it means proving knowledge of a vector $\vec{u} \in \mathbb{Z}_{\hat{q}}^n$, as well as short vectors $\vec{s} \in \mathbb{Z}_{\hat{q}}^m$ and $\vec{r} \in \mathbb{Z}_{\hat{q}}^{\ell_r}$ such that

$$\begin{aligned} \vec{u} &\in \{0, 1\}^t, \quad \|\vec{s}\| \leq \mathcal{B}_s \quad \text{and} \quad \|\vec{r}\| \leq \mathcal{B}_r, \\ P\vec{s} &= B\vec{u} + C \begin{bmatrix} \vec{m} \\ \vec{r} \end{bmatrix} \pmod{\hat{q}}, \end{aligned} \tag{12}$$

where

$$P \in \mathbb{Z}_{\hat{q}}^{n \times m}, \quad C \in \mathbb{Z}_{\hat{q}}^{n \times (\ell_m + \ell_r)}, \quad \vec{m} \in \mathbb{Z}_{\hat{q}}^{\ell_m}, \quad B \in \mathbb{Z}_{\hat{q}}^{n \times t}, \quad \text{Bounds} := (\mathcal{B}_s, \mathcal{B}_r)$$

are parts of the statement, along with some auxiliary information `aux`. Clearly, as long as $q|\hat{q}$, we can map equations such as (7) to this setting with standard transformation from \mathcal{R}_q to \mathbb{Z}_q^d using multiplication matrices (c.f. Section 2) and by lifting from \mathbb{Z}_q to $\mathbb{Z}_{\hat{q}}$.

In order to prove (12), we make use of the LNP framework, developed recently by Lyubashevsky et al. [57], for proving various lattice-related statements. For completeness, we provide the full protocol for proving (12) and refer to [57] for more details.

5.1 Background

Notation. In the following, we introduce a few variables which might be also present in the other sections. We highlight that these are only defined for this section, and thus they have no relation with variables outside this section, unless specified otherwise.

Denote \mathbf{H} to be the random oracle used for the NIZK. Let \hat{d} be a power-of-two and $\hat{\mathcal{R}} := \mathbb{Z}[X]/(X^{\hat{d}} + 1)$. For a modulus $\hat{q} \in \mathbb{N}$, define $\hat{\mathcal{R}}_{\hat{q}} := \hat{\mathcal{R}}/(\hat{q})$. In this work, $\hat{q} = \hat{p}_1 \hat{p}_2$ is a product of two primes \hat{p}_i of the

form $\hat{p}_i = 5 \pmod{8}$ and $\hat{p}_1 < \hat{p}_2$. For a polynomial $x \in \hat{\mathcal{R}}_{\hat{q}}$, we define \tilde{x} to be the constant coefficient of x . Furthermore, we assume that \hat{d} divides the security parameter λ .

Let $\sigma : \hat{\mathcal{R}}_{\hat{q}} \rightarrow \hat{\mathcal{R}}_{\hat{q}}$ be the ring automorphism defined by $\sigma(X) := X^{-1}$. For a polynomial vector (and similarly for a matrix) $\mathbf{x} = (x_1, \dots, x_n) \in \hat{\mathcal{R}}_{\hat{q}}^n$, define $\sigma(\mathbf{x}) := (\sigma(x_1), \dots, \sigma(x_n))$. Further, let $\tau \in \mathbb{N}$ be such that $1/\hat{p}_1^\tau = \text{negl}(\lambda)$, i.e. a parameter used for soundness amplification.

ABDLOP Commitment. It was introduced in [57] and extends both the Ajtai [4] and BDLOP [10] constructions. Let $n, m_1, m_2, \ell \in \mathbb{N}$ and $m_2 \geq n + \ell$. The commitment key consists of uniformly random matrices

$$\text{crs} := (\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}) \in \hat{\mathcal{R}}_{\hat{q}}^{n \times m_1} \times \hat{\mathcal{R}}_{\hat{q}}^{n \times m_2} \times \hat{\mathcal{R}}_{\hat{q}}^{\ell \times m_2}.$$

Now, to commit to a message vector $\mathbf{s}_1 \in \hat{\mathcal{R}}_{\hat{q}}^{m_1}$ with small coefficients as well as a “full-fledged” polynomial vector $\mathbf{m} \in \hat{\mathcal{R}}_{\hat{q}}^\ell$, we sample a randomness vector $\mathbf{s}_2 \leftarrow \chi^{m_2}$, where χ is a probability distribution over $\hat{\mathcal{R}}_{\hat{q}}$, and compute:

$$\text{ABDLOP.Com}(\text{crs}, \mathbf{s}_1, \mathbf{m}; \mathbf{s}_2) := \begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_B \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \end{bmatrix} \mathbf{s}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \end{bmatrix} \mathbf{s}_2 + \begin{bmatrix} \mathbf{0} \\ \mathbf{m} \end{bmatrix}.$$

Note that when $\ell = 0$ (resp. $m_1 = 0$) then this construction ends up being the Ajtai (resp. BDLOP) commitment scheme. In particular, the commitment size does not depend on the length m_1 of \mathbf{s}_1 (but it does on ℓ). Therefore, we will commit to long vectors with small coefficients in the “Ajtai” part \mathbf{s}_1 and commit to a few *garbage* polynomials used for the proofs in the “BDLOP” part \mathbf{m} . Hence, we call \mathbf{t}_A (resp. \mathbf{t}_B) the “top part” (resp. “bottom part”) of the commitment. An opening of the commitment is a triple $(\mathbf{s}_1, \mathbf{m}, \mathbf{s}_2)$ ¹³.

Relaxed openings. As folklore in lattice-based cryptography, we also consider relaxed openings of a commitment which involve so-called *relaxation factors*. First, we fix $\xi, \nu > 0$ and a power-of-two k and define the following set \mathcal{C} as¹⁴:

$$\mathcal{C} := \left\{ c \in \hat{\mathcal{R}}_{\hat{q}} : \|c\|_\infty \leq \xi \wedge \sigma_{-1}(c) = c \wedge \sqrt[2k]{\|c^{2k}\|_1} \leq \nu \right\}. \quad (13)$$

Roughly speaking, the first condition $\sigma_{-1}(c) = c$ is needed to prove quadratic equations in the committed messages which might additionally involve automorphisms. On the other hand, the second condition allows us to use [57, Lemma 2.15] and deduce that if $\|\mathbf{r}\| \leq \mathcal{B}$ and $c \in \mathcal{C}$ then $\|\mathbf{c}\mathbf{r}\| \leq \nu\mathcal{B}$.

We define the set of relaxation factors as

$$\bar{\mathcal{C}} := \{c - c' : c, c' \in \mathcal{C} \text{ and } c \neq c'\},$$

i.e. set of differences of any two distinct elements in \mathcal{C} . We will choose the constant ν such that (experimentally) the probability for $c \leftarrow S_\xi$ to satisfy $\sqrt[2k]{\|c^{2k}\|_1} \leq \nu$ is at least 99%. In Figure 8 we show example parameters (ξ, ν, k) and the lower-bound on the size of \mathcal{C} .

\hat{d}	ξ	ν	k	$ \mathcal{C} $
64	8	140	32	2^{129}
128	2	59	32	2^{147}

Fig. 8: Example parameters to instantiate \mathcal{C} for a modulus \hat{q} such that its smallest prime divisor p_1 is greater than 16.

For security of our protocols, we need the invertibility property of the set \mathcal{C} , i.e. the difference of any two distinct elements of \mathcal{C} is invertible over $\hat{\mathcal{R}}_{\hat{q}}$. Hence, we apply [57, Lemma 2.6] and thus we only need the

¹³ Message \mathbf{m} does not need to be included in the opening since it can be deterministically computed from \mathbf{t}_B and \mathbf{s}_2 .

¹⁴ Looking ahead, this set will be a challenge space and will play a big role in our protocols.

condition $\xi < \hat{p}_1/2$ where \hat{p}_1 is the smallest prime dividing q . Secondly, to achieve negligible soundness error, we will need $|\mathcal{C}|$ to be exponentially large.

Now, we are ready to define the ABDLOP relaxed openings.

Definition 5.1. *A relaxed opening of the ABDLOP commitment $(\mathbf{t}_A, \mathbf{t}_B)$ w.r.t. a commitment key crs is a tuple $(\mathbf{s}_1, \mathbf{m}, \mathbf{s}_2, c) \in \hat{\mathcal{R}}_q^{m_1} \times \hat{\mathcal{R}}_q^\ell \times \hat{\mathcal{R}}_q^{m_2} \times \bar{\mathcal{C}}$ which satisfies:*

$$\begin{aligned} \text{ABDLOP.Com}(\text{crs}, \mathbf{s}_1, \mathbf{m}; \mathbf{s}_2) &= (\mathbf{t}_A, \mathbf{t}_B) \\ \|\mathbf{cs}_1\| &\leq \mathcal{B}_1 \quad \text{and} \quad \|\mathbf{cs}_2\| \leq \mathcal{B}_2. \end{aligned}$$

Security properties. As shown in [57, Lemma 3.1], the ABDLOP commitment is binding with respect to relaxed openings under the Module-SIS assumption. As for hiding, it follows from the fact that under the Module-LWE assumption $\begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \end{bmatrix} \mathbf{s}_2$ looks pseudorandom.

Lemma 5.2 (Binding). *Let $\xi < p_1/2$ where p_1 is the smallest prime dividing \hat{q} . Then, the ABDLOP commitment is computationally binding with respect to relaxed openings under the $\text{MSIS}_{n, m_1+m_2, \mathcal{B}}$ assumption where $\mathcal{B} := 4\nu\sqrt{\mathcal{B}_1^2 + \mathcal{B}_2^2}$.*

Lemma 5.3 (Hiding). *The ABDLOP commitment is computationally hiding under the $\text{MLWE}_{n+\ell, m_2, \chi}$ assumption.*

Approximate Range Proofs. We use the results developed in [40, 57] to upper-bound the operator norm of a matrix $\mathfrak{R} \leftarrow \text{Bin}_\xi^{256 \times m}$. Here, Bin denotes the binomial distribution with a positive integer parameter ξ , which is the distribution $\sum_{i=1}^\xi (a_i - b_i)$, where $a_i, b_i \leftarrow \{0, 1\}$. The variance of this distribution is $\xi/2$ and it holds that $\text{Bin}_{\xi_1} \pm \text{Bin}_{\xi_2} = \text{Bin}_{\xi_1 + \xi_2}$.

The bound on the operator norm of \mathfrak{R} is used for the *approximate range proof* part, where we want to show that \vec{s} has relatively small coefficients w.r.t. to the proof system modulus \hat{q} . Namely, let us define functions $\omega_{\min}(\lambda)$ and $\omega_{\max}(\lambda)$ such that for any $\vec{w} \in \mathbb{Z}^m$ and any $\xi \in \mathbb{N}$:

$$\begin{aligned} \Pr_{\mathfrak{R} \leftarrow \text{Bin}_\xi^{256 \times m}} [\|\mathfrak{R}\vec{w}\|^2 > \xi \cdot \|\vec{w}\|^2 \cdot \omega_{\min}(\lambda)^2] &\leq 2^{-2\lambda}, \\ \Pr_{\mathfrak{R} \leftarrow \text{Bin}_\xi^{256 \times m}} [\|\mathfrak{R}\vec{w}\|^2 < \xi \cdot \|\vec{w}\|^2 \cdot \omega_{\max}(\lambda)^2] &\leq 2^{-\lambda}. \end{aligned}$$

It was shown in [40] (under a few natural heuristic assumptions) that $\omega_{\min}(128) = \sqrt{13}$ and $\omega_{\max}(128) = \sqrt{337}$. In the soundness argument, we will use the following result from [57, Lemma 2.9] in a slightly more general form (if one substitutes $\lambda = 128$ then it is identical to the aforementioned lemma):

Lemma 5.4. *Fix $m, P \in \mathbb{N}$ and a bound $b \leq P/(8\sqrt{2} \cdot \omega_{\min}(\lambda) \cdot m)$, and let $\vec{w} \in [\pm P/2]^m$ with $\|\vec{w}\| \geq b$, and let \vec{y} be an arbitrary vector in $[\pm P/2]^m$. Then*

$$\Pr_{R \leftarrow \text{Bin}_1^{256 \times m}} \left[\|R\vec{w} + \vec{y} \bmod P\| < \frac{b}{\sqrt{2}} \cdot \omega_{\min}(\lambda) \right] < 2^{-\lambda}.$$

5.2 The Protocol

We provide a brief summary of the Fiat-Shamir transformed [38] protocol from [57] to prove (12). We assume that variables n, m, t, ℓ_r are divisible by \hat{d} . First, note that by the sum-of-four-squares theorem, $\|\vec{s}\| \leq \mathcal{B}_s$ is equivalent to the existence of four integers $a_1, a_2, a_3, a_4 \in \mathbb{Z}$ such that $\mathcal{B}_s^2 - \|\vec{s}\|^2 = a_1^2 + a_2^2 + a_3^2 + a_4^2$. Thus, we can define a vector $\mathbf{a} := (a_1, a_2, a_3, a_4, 0, \dots, 0) \in \mathbb{Z}_{\hat{q}}^d$ and observe that:

$$\left\| \begin{bmatrix} \vec{s} \\ \mathbf{a} \end{bmatrix} \right\| = \mathcal{B}_s \quad \text{and} \quad [P \ 0] \begin{bmatrix} \vec{s} \\ \mathbf{a} \end{bmatrix} = P\vec{s}.$$

Hence, we slightly abuse the notation and define $\mathbf{m} := \mathbf{m} + \hat{d}, \vec{s} := (\vec{s}, \vec{a}) \in \mathbb{Z}_{\hat{q}}^{\mathbf{m}}$ and $P := [P \ 0] \in \hat{\mathcal{R}}_{\hat{q}}^{n \times \mathbf{m}}$. Similarly, we proceed for the vector \vec{r} . Thus, we now need to prove that

$$\|\vec{s}\| = \mathcal{B}_s, \quad \|\vec{r}\| = \mathcal{B}_r, \quad \vec{u} \in \{0, 1\}^n, \quad P\vec{s} = B\vec{u} + C \begin{bmatrix} \vec{m} \\ \vec{r} \end{bmatrix}. \quad (14)$$

In this proof system, the common *random* string simply consists of uniformly random matrices which are only used for the ABDLOP commitment scheme [57]:

$$\text{crs} := (\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_y, \mathbf{B}_g, \mathbf{b}) \in \hat{\mathcal{R}}_{\hat{q}}^{n \times m_1} \times \hat{\mathcal{R}}_{\hat{q}}^{n \times m_2} \times \hat{\mathcal{R}}_{\hat{q}}^{256/\hat{d} \times m_2} \times \hat{\mathcal{R}}_{\hat{q}}^{\tau \times m_2} \times \hat{\mathcal{R}}_{\hat{q}}^{m_2}. \quad (15)$$

We define the statement as the tuple:

$$\mathbf{x} := (P, C, \vec{m}, B, \text{Bounds} := (\mathcal{B}_s, \mathcal{B}_r), \text{aux}).$$

First round. Let $\mathbf{s} = \text{Coeffs}^{-1}(\vec{s}) \in \hat{\mathcal{R}}_{\hat{q}}^{m/d}$ and similarly for \mathbf{r}, \mathbf{u} . To begin with, we commit to the witness (\mathbf{s}, \mathbf{r}) . To be consistent with notation from [57] let $\mathbf{s}_1 := (\mathbf{s}, \mathbf{r}, \mathbf{u})$ and $m_1 = \mathbf{m} + \ell_r + \mathbf{t}$. We commit to \mathbf{s}_1 using the ABDLOP commitment (which in this case is just the Ajtai commitment), i.e. we sample the randomness vector $\mathbf{s}_2 \leftarrow \chi^{m_2}$, where χ is a probability distribution on ternary polynomials in $\hat{\mathcal{R}}_{\hat{q}}$, and compute:

$$\mathbf{t}_A := \mathbf{A}_1 \mathbf{s}_1 + \mathbf{A}_2 \mathbf{s}_2 \quad \text{where} \quad (\mathbf{A}_1, \mathbf{A}_2) \in \hat{\mathcal{R}}_{\hat{q}}^{n \times m_1} \times \hat{\mathcal{R}}_{\hat{q}}^{n \times m_2}.$$

Then, we sample short masking vectors $\mathbf{y}_i \leftarrow D_{\hat{s}_i}^{m_i}$ for $i = 1, 2$, and compute $\mathbf{w} := \mathbf{A}_1 \mathbf{y}_1 + \mathbf{A}_2 \mathbf{y}_2$. Finally, we focus on the parts used to prove (14). Namely, we generate the polynomial vector $\mathbf{y}_3 \leftarrow D_{\hat{s}_3}^{256}$ and commit to it as follows

$$\mathbf{t}_y := \mathbf{B}_y \mathbf{s}_2 + \mathbf{y}_3 \quad \text{where} \quad \mathbf{B}_y \in \hat{\mathcal{R}}_{\hat{q}}^{256/\hat{d} \times m_2}.$$

Finally, we sample a polynomial vector $\mathbf{g} \leftarrow \{x \in \hat{\mathcal{R}}_{\hat{q}} : \hat{x} = 0\}^{\tau}$ and commit to it:

$$\mathbf{t}_g := \mathbf{B}_g \mathbf{s}_2 + \mathbf{g} \quad \text{where} \quad \mathbf{B}_g \leftarrow \hat{\mathcal{R}}_{\hat{q}}^{\tau \times m_2}.$$

Hence, the first message and the corresponding challenge are

$$a_1 := (\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w}) \quad \text{and} \quad (\mathfrak{R}_0, \mathfrak{R}_1) = \text{H}(1, \text{crs}, \mathbf{x}, a_1) \in \{0, 1\}^{256 \times m_1 \hat{d}} \times \{0, 1\}^{256 \times m_1 \hat{d}}.$$

Second round. Given the first challenge $(\mathfrak{R}_0, \mathfrak{R}_1)$, we compute the response

$$\vec{z}_3 := \vec{y}_3 + \mathfrak{R}_1 \vec{s}_1 \in \mathbb{Z}^{256},$$

where $\vec{y}_3 := \text{Coeffs}(\mathbf{y}_3)$, $\vec{s}_1 := \text{Coeffs}(\mathbf{s}_1)$, and $\mathfrak{R} := \mathfrak{R}_0 - \mathfrak{R}_1$, and apply rejection sampling on \vec{z}_3 . The second message and the corresponding challenge are:

$$a_2 := \vec{z}_3 \quad \text{and} \quad (\gamma_{i,j})_{i \in [\tau], j \in [256+n+3]} = \text{H}(2, \text{crs}, \mathbf{x}, a_1, a_2) \in \mathbb{Z}_{\hat{q}}^{\tau \times (256+n+3)}.$$

Third round. The prover's goal now is to prove the following relations over $\mathbb{Z}_{\hat{q}}$:

$$\begin{cases} \vec{z}_3 = \vec{y}_3 + \mathfrak{R}_1 \vec{s}_1 \\ P\vec{s} = B\vec{u} + C \begin{bmatrix} \vec{m} \\ \vec{r} \end{bmatrix} \\ \langle \vec{s}, \vec{s} \rangle = \mathcal{B}_s^2 \\ \langle \vec{r}, \vec{r} \rangle = \mathcal{B}_r^2 \\ \langle \vec{u}, \vec{u} - \vec{1} \rangle = 0 \end{cases} \quad (16)$$

where the terms in blue are secret and committed. We will use the following observation from [57, Lemma 2.4]: for any integer vectors $\vec{x}_0, \vec{x}_1 \in \mathbb{Z}_q^n$, the constant coefficient of

$$\text{Coeffs}^{-1}(\vec{x}_0)^T \sigma \left(\text{Coeffs}^{-1}(\vec{x}_1) \right) \in \hat{\mathcal{R}}_q$$

is equal to $\langle \vec{x}_0, \vec{x}_1 \rangle \in \mathbb{Z}_q$. This implies that proving the first equation is equivalent to proving that for $i \in [256]$, the constant coefficient of

$$\sigma(\mathbf{r}_i)^T \mathbf{s}_1 + \sigma(\mathbf{e}_i)^T \mathbf{y}_3 - z_{3,i} \text{ is equal to zero,}$$

where \mathbf{r}_i is the polynomial vector such that its coefficient vector is the i -th row of \mathfrak{R} , and \mathbf{e}_i is the polynomial vector so that $\text{Coeffs}(\mathbf{e}_i)$ is a unit vector¹⁵ with its i -th coefficient being 1. Similarly, the second equation is equivalent to: for all $i \in [n]$, the constant coefficient of:

$$\sigma(\mathbf{p}_i)^T \mathbf{s} - \sigma(\beta_i)^T \mathbf{u} - m_{C,i} - \sigma(\mathbf{c}_{r,i})^T \mathbf{r} \text{ is equal to zero,}$$

where:

- \mathbf{p}_i is the polynomial vector with its coefficients being the i -th row of P ,
- β_i is the polynomial vector with its coefficients being the i -th row of B ,
- we write $C = [C_m \ C_r]$ such that $C \begin{bmatrix} \vec{m} \\ \vec{r} \end{bmatrix} = C_m \vec{m} + C_r \vec{r}$, and define $\vec{m}_C := C_m \vec{m}$,
- $\mathbf{c}_{r,i}$ is the polynomial vector with its coefficients being the i -th row of C_r .

Clearly, the last three equations are equivalent to

$$\sigma(\mathbf{s})^T \mathbf{s} - \mathcal{B}_s, \quad \sigma(\mathbf{r})^T \mathbf{r} - \mathcal{B}_r, \quad \sigma \left(\mathbf{u} - \text{Coeffs}^{-1}(1^{\dagger}) \right)^T \mathbf{u}$$

having the constant coefficient equal to zero. Hence, we compute for $i \in [\tau]$:

$$\begin{aligned} h_i := & g_i + \sum_{j=1}^{256} \gamma_{i,j} \cdot \left(\sigma(\mathbf{r}_j)^T \mathbf{s}_1 + \sigma(\mathbf{e}_j)^T \mathbf{y}_3 - z_j \right) \\ & + \sum_{j=1}^n \gamma_{i,256+j} \cdot \left(\sigma(\mathbf{p}_j)^T \mathbf{s} - \sigma(\beta_j)^T \mathbf{u} - m_{C,j} - \sigma(\mathbf{c}_{r,j})^T \mathbf{r} \right) \\ & + \gamma_{i,256+n+1} \cdot \left(\sigma(\mathbf{s})^T \mathbf{s} - \mathcal{B}_s \right) + \gamma_{i,256+n+2} \cdot \left(\sigma(\mathbf{r})^T \mathbf{r} - \mathcal{B}_r \right) \\ & + \gamma_{i,256+n+3} \cdot \sigma \left(\mathbf{u} - \text{Coeffs}^{-1}(1^{\dagger}) \right)^T \mathbf{u} \in \hat{\mathcal{R}}_q. \end{aligned} \tag{17}$$

Observe that by definition of \mathbf{g} , the constant coefficients of h_1, \dots, h_τ are all zeroes. The third message and the corresponding challenge are:

$$a_3 := (h_1, \dots, h_\tau) \quad \text{and} \quad \boldsymbol{\mu} = (\mu_i)_{i \in [\tau]} = \text{H}(3, \text{crs}, \mathbf{x}, a_1, a_2, a_3) \in \hat{\mathcal{R}}_q^\tau.$$

¹⁵ Concretely, it is a binary vector with exactly one 1 entry.

Fourth round. The goal of the prover is now to prove the τ quadratic equations (with automorphisms) over $\mathcal{R}_{\hat{q}}$ from (17). We now simply linear-combine them with the challenges μ_i , i.e. we prove a single equation:

$$\begin{aligned}
0 = \sum_{i=1}^{\tau} \mu_i & \left(\sum_{j=1}^{256} \gamma_{i,j} \cdot \left(\sigma(\mathbf{r}_j)^T \mathbf{s}_1 + \sigma(\mathbf{e}_j)^T \mathbf{y}_3 - z_{3,j} \right) \right. \\
& + \sum_{j=1}^n \gamma_{i,256+j} \cdot \left(\sigma(\mathbf{p}_j)^T \mathbf{s} - \sigma(\beta_j)^T \mathbf{u} - m_{C,j} - \sigma(\mathbf{c}_{r,j})^T \mathbf{r} \right) \\
& + \gamma_{i,256+n+1} \cdot \left(\sigma(\mathbf{s})^T \mathbf{s} - \mathcal{B}_s \right) + \gamma_{i,256+n+2} \cdot \left(\sigma(\mathbf{r})^T \mathbf{r} - \mathcal{B}_r \right) \\
& \left. + \gamma_{i,256+n+3} \cdot \sigma \left(\mathbf{u} - \text{Coeffs}^{-1}(1^t) \right)^T \mathbf{u} + g_i - h_i \right). \tag{18}
\end{aligned}$$

Let us define

$$\mathbf{B} := \begin{bmatrix} \mathbf{B}_y \\ \mathbf{B}_g \end{bmatrix}, \quad \mathbf{t}_B := \begin{bmatrix} \mathbf{t}_y \\ \mathbf{t}_g \end{bmatrix}, \quad \mathbf{m} := \begin{bmatrix} \mathbf{y}_3 \\ \mathbf{g} \end{bmatrix} \in \hat{\mathcal{R}}_{\hat{q}}^{256/\hat{d}+\tau} \quad \text{and} \quad \hat{\mathbf{s}} := \begin{bmatrix} \mathbf{s}_1 \\ \sigma(\mathbf{s}_1) \\ \mathbf{m} \\ \sigma(\mathbf{m}) \end{bmatrix}. \tag{19}$$

Recall that $\mathbf{s}_1 = (\mathbf{s}, \mathbf{r}, \mathbf{u})$ and let us write $\mathbf{r}_j := (\mathbf{r}_{s,j}, \mathbf{r}_{r,j}, \mathbf{r}_{u,j})$ such that

$$\sigma(\mathbf{r}_j)^T \mathbf{s}_1 = \sigma(\mathbf{r}_{s,j})^T \mathbf{s} + \sigma(\mathbf{r}_{r,j})^T \mathbf{r} + \sigma(\mathbf{r}_{u,j})^T \mathbf{u}.$$

Then, the quadratic equation above can be written equivalently as

$$\hat{\mathbf{s}}^T \mathbf{D}_2 \hat{\mathbf{s}} + \mathbf{d}_1^T \hat{\mathbf{s}} + d_0 = 0$$

where

$$\begin{aligned}
\mathbf{D}_2 & := \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \sum_{i=1}^{\tau} \mu_i \gamma_{i,256+n+1} \cdot \mathbf{I}_{m/\hat{d}} & \sum_{i=1}^{\tau} \mu_i \gamma_{i,256+n+2} \cdot \mathbf{I}_{\ell_r/\hat{d}} & \sum_{i=1}^{\tau} \mu_i \gamma_{i,256+n+3} \cdot \mathbf{I}_{n/\hat{d}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \\
\mathbf{d}_1 & := \begin{bmatrix} \sum_{i=1}^{\tau} \mu_i \cdot \left(\sum_{j=1}^{256} \gamma_{i,j} \sigma(\mathbf{r}_{s,j}) + \sum_{j'=1}^n \gamma_{i,256+j'} \sigma(\mathbf{p}_{j'}) \right) \\ \sum_{i=1}^{\tau} \mu_i \cdot \left(\sum_{j=1}^{256} \gamma_{i,j} \sigma(\mathbf{r}_{r,j}) - \sum_{j'=1}^n \gamma_{i,256+j'} \sigma(\mathbf{c}_{r,j'}) \right) \\ \sum_{i=1}^{\tau} \mu_i \cdot \left(\sum_{j=1}^{256} \gamma_{i,j} \sigma(\mathbf{r}_{u,j}) - \sum_{j'=1}^n \gamma_{i,256+j'} \sigma(\beta_{j'}) - \gamma_{i,256+n+3} \text{Coeffs}^{-1}(1^t) \right) \\ \mathbf{0} \\ \sum_{i=1}^{\tau} \sum_{j=1}^{256} \mu_i \gamma_{i,j} \sigma(\mathbf{e}_j) \\ \boldsymbol{\mu} \\ \mathbf{0} \end{bmatrix} \tag{20} \\
d_0 & := - \sum_{i=1}^{\tau} \mu_i \cdot \left(\sum_{j=1}^{256} z_{3,j} + \sum_{j'=1}^n m_{C,j'} + \gamma_{i,256+n+1} \mathcal{B}_s + \gamma_{i,256+n+2} \mathcal{B}_r + h_i \right).
\end{aligned}$$

Eventually, we run the sub-protocol for proving a single quadratic equation with automorphisms. First, we calculate the garbage term $f_1 = \hat{\mathbf{s}}^T \mathbf{R}_2 \mathbf{y} + \mathbf{y}^T \mathbf{R}_2 \hat{\mathbf{s}} + \mathbf{r}_1^T \mathbf{y}$, where \mathbf{y} is defined as

$$\mathbf{y} := \begin{bmatrix} \mathbf{y}_1 \\ \sigma(\mathbf{y}_1) \\ -\mathbf{B}\mathbf{y}_2 \\ -\sigma(\mathbf{B}\mathbf{y}_2) \end{bmatrix} \in \hat{\mathcal{R}}_{\hat{q}}^{2(m_1+256/d+\tau)}, \tag{21}$$

and the commitment $t = \mathbf{b}^T \mathbf{s}_2 + f_1$ to f_1 . Then, we set $f_0 = \mathbf{y}^T \mathbf{R}_2 \mathbf{y} + \mathbf{b}^T \mathbf{y}_2$. Hence, the fourth message and the corresponding challenge are (we defined \mathcal{C} in Section 5.1):

$$a_4 := (t, f_0) \quad \text{and} \quad c := \mathbf{H}(4, \text{crs}, x, a_1, a_2, a_3, a_4) \in \mathcal{C}.$$

Final round. Given a challenge c , we compute $\mathbf{z}_i = c\mathbf{s}_i + \mathbf{y}_i$ for $i = 1, 2$ and apply rejection sampling on $\mathbf{z}_1, \mathbf{z}_2$. So, the prover's last message is $a_5 := (\mathbf{z}_1, \mathbf{z}_2)$ and thus the proof consists of

$$\pi := (a_1, a_2, a_3, a_4, a_5).$$

Verification algorithm. The verifier is given a proof π and recomputes the corresponding challenges as well as $\mathbf{D}_2, \mathbf{d}_1, d_0$. Let us define

$$\mathbf{z} := \begin{bmatrix} \mathbf{z}_1 \\ \sigma(\mathbf{z}_1) \\ c\mathbf{t}_B - \mathbf{B}\mathbf{z}_2 \\ \sigma(c\mathbf{t}_B - \mathbf{B}\mathbf{z}_2) \end{bmatrix}. \quad (22)$$

Finally, the verifier outputs 1 if *all* the following relations hold:

$$\begin{cases} \|\mathbf{z}_1\| \stackrel{?}{\leq} \mathcal{B}_1, \|\mathbf{z}_2\| \stackrel{?}{\leq} \mathcal{B}_2, \|\bar{\mathbf{z}}_3\| \stackrel{?}{\leq} \mathcal{B}_3, \\ \tilde{h}_i \stackrel{?}{=} 0 \text{ for } i \in [\tau], \\ \mathbf{A}_1 \mathbf{z}_1 + \mathbf{A}_2 \mathbf{z}_2 \stackrel{?}{=} \mathbf{w} + c\mathbf{t}_A, \\ \mathbf{z}^T \mathbf{D}_2 \mathbf{z} + c\mathbf{d}_1^T \mathbf{z} + c^2 d_0 - (ct - \mathbf{b}^T \mathbf{z}_2) \stackrel{?}{=} f_0 \end{cases}$$

and 0 otherwise.

For completeness, we provide a description of the NIZK proof system $\Pi_{\text{NIZK}}^{\text{ISIS}} = (\text{Prove}_{\text{ISIS}}^{\text{H}}, \text{Verify}_{\text{ISIS}}^{\text{H}})$ for proving (12) in Figures 9 and 10. Similarly as in many prior works, e.g. [53], in practice we do not include \mathbf{w}, f_0 in the proof and send c instead, since these components can be computed deterministically given other parts of π and c . However, we keep the standard Fiat-Shamir transformation for simpler security analysis.

5.3 Security Analysis

In this subsection, we prove the key properties of $\Pi_{\text{NIZK}}^{\text{ISIS}}$, i.e. correctness and zero-knowledge. As for knowledge soundness, we will prove it directly when arguing unforgeability of our anonymous credentials in Section 8.

5.3.1 Correctness

Lemma 5.5 (Correctness). *Let $\alpha_1, \alpha_2, \alpha_3, \bar{\alpha}_1, \bar{\alpha}_2 \in O(\sqrt{\lambda})$ and $\mathbf{N}, \hat{d} \in O(\lambda)$. Further, suppose χ is a probability distribution over S_1 . Recall ν, ω_{\max} from Section 5.1, and define the following variables:*

$$\begin{aligned} \mathfrak{s}_1 &:= \alpha_1 \nu \sqrt{\mathcal{B}_s^2 + \mathcal{B}_r^2 + \mathfrak{t}}, & \mathfrak{s}_2 &:= \alpha_2 \nu \sqrt{m_2 \hat{d}}, & \mathfrak{s}_3 &:= \alpha_3 \omega_{\max}(\lambda) \sqrt{\mathcal{B}_s^2 + \mathcal{B}_r^2 + \mathfrak{t}}, \\ \mathcal{B}_1 &:= \mathfrak{s}_1 \sqrt{2m_1 \hat{d}}, & \mathcal{B}_2 &:= \mathfrak{s}_2 \sqrt{2m_2 \hat{d}}, & \mathcal{B}_3 &:= 1.7 \mathfrak{s}_3 \sqrt{256} \end{aligned}$$

and the corresponding rejection rates:

$$M_i := \exp \left(\sqrt{\frac{2(\lambda + 1)}{\log e}} \cdot \frac{1}{\alpha_i} + \frac{1}{2\alpha_i^2} \right) \text{ for } i \in \{1, 2, 3\}.$$

Then, $\Pi_{\text{NIZK}}^{\text{ISIS}} := (\text{Prove}_{\text{ISIS}}^{\text{H}}, \text{Verify}_{\text{ISIS}}^{\text{H}})$ is correct.

```

ProveISIS(crs, x, w)
Input: crs = (A1, A2, By, Bg, b), x = (P, C, m̄, B, Bounds, aux), w = (s̄, r̄, ū)
Output: proof π
1: rst = 0                                ▷ boolean which indicates if a restart is necessary
2: idx = 0                                ▷ keeps track how many restarts occurred
3: s1 = (s, r, u)                        ▷ s, r, u are the polynomial vectors with coeff. s̄, r̄, ū
4: while rst = 0 ∧ idx < N do
5:   idx = idx + 1
6:   s2 ← χm2                            ▷ generate commitment randomness
7:   tA = A1s1 + A2s2                    ▷ Ajtai commitment to s1
8:   y1 ← Ds1m1d̄, y2 ← Ds2m2d̄, y3 ← Ds3256    ▷ sample masking vectors
9:   g ← {x ∈ R̂q̂ : x̄ = 0}τ                ▷ sample random polynomials with const. coeff. zero
10:  w = A1y1 + A2y2                    ▷ used to prove knowledge of a commitment opening
11:  ty = Bys2 + y3                            ▷ commitment to y3
12:  tg = Bgs2 + g                            ▷ commitment to g
13:  a1 = (tA, ty, tg, w)                            ▷ first message
14:  (R0, R1) = H(1, crs, x, a1) ∈ {0, 1}256 × m1d̄ × {0, 1}256 × m1d̄    ▷ first challenge
15:  R = R0 - R1
16:  s̄1 = Coeffs(s1)
17:  ȳ3 = Coeffs(y3)
18:  z̄3 = ȳ3 + R s̄1 ∈ Z256                ▷ masked opening of R s̄1
19:  b3 ← Rej(z̄3, R s̄1, s3, M3)
20:  a2 = z̄3                                ▷ second message
21:  (γi,j)i ∈ [τ], j ∈ [256+n+3]} = H(2, crs, x, a1, a2) ∈ Zq̂τ × (256+n+3)    ▷ second challenge
22:  for i ∈ [τ] do
23:    compute hi as in (17)
24:    a3 = h = (h1, ..., hτ)                            ▷ third message
25:    (μi)i ∈ [τ]} = H(3, crs, x, a1, a2, a3) ∈ R̂q̂τ    ▷ third challenge
26:    compute B, tB, m, ŝ, y as in (19) and (21)
27:    compute D2, d1, d0 as in (20)
28:    f1 = sT R2 y + yT R2 s + r1T y
29:    f0 = yT R2 y + bT y2
30:    t = bT s2 + f1                                ▷ commitment to f1
31:    a4 = (t, f0)                                ▷ fourth message
32:    c = H(4, crs, x, a1, a2, a3, a4) ∈ C        ▷ fourth challenge
33:    for i ∈ {1, 2} do
34:      zi = yi + c si                            ▷ maked opening of c si
35:      bi ← Rej(zi, c si, si, Mi)                ▷ rejection sampling
36:      π = (tA, ty, tg, w, z̄3, h, t, f0, z̄3, h, z1, z2)    ▷ candidate proof π
37:      rst = b1 b2 b3                            ▷ if all rejection steps passed then rst = 1
38:  if rst = 1 then return π
39:  else return ⊥

```

Fig. 9: Prove_{ISIS}^H algorithm.

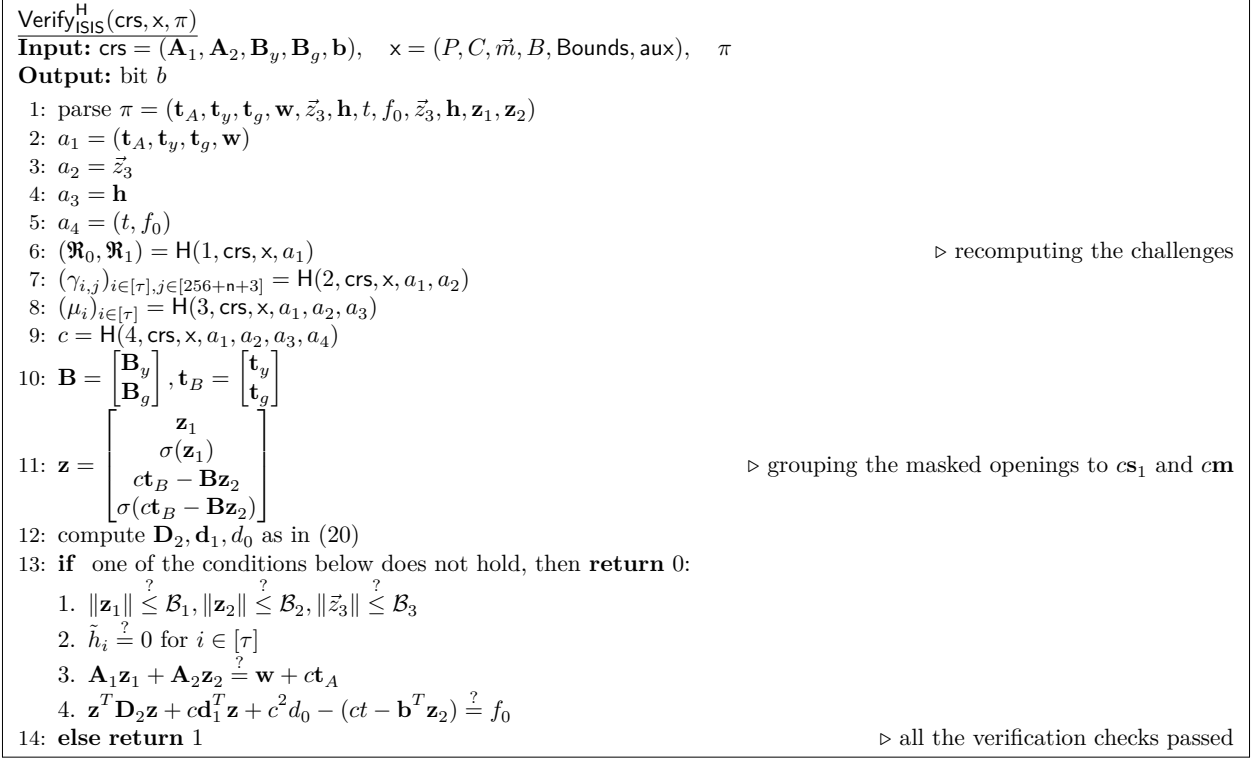


Fig. 10: $\text{Verify}_{\text{ISIS}}^{\text{H}}$ algorithm.

Proof. First, one observes that both algorithms, and especially $\text{Prove}_{\text{ISIS}}^{\text{H}}$, are *strict* polynomial time since we bounded the number of restarts by $O(\lambda)$. Next, we argue that the prover algorithm outputs \perp with negligible probability. First, using Lemma 2.4 and the following inequalities (which hold with an overwhelming probability):

$$\|\mathbf{cs}_1\| \leq \nu \sqrt{\mathcal{B}_s^2 + \mathcal{B}_r^2 + \mathbf{t}}, \quad \|\mathbf{cs}_2\| \leq \nu \sqrt{m_2 \hat{d}}, \quad \|\mathfrak{R}\vec{s}_1\| \leq \omega_{\max}(\lambda) \cdot \sqrt{\mathcal{B}_s^2 + \mathcal{B}_r^2 + \mathbf{t}}.$$

Hence, we deduce that in a single run the probability that $b_1 b_2 b_3 = 1$ is at least $(M_1 M_2 M_3)^{-1} - \text{negl}(\lambda)$ by Lemma 2.4, and thus the probability that $\text{Prove}_{\text{ISIS}}^{\text{H}}$ outputs \perp , i.e. makes exactly \mathbf{N} unsuccessful attempts, is at most

$$\left(1 - \frac{1}{M_1 M_2 M_3} + \text{negl}(\lambda)\right)^{\mathbf{N}} = \text{negl}(\lambda)$$

since $M_1, M_2, M_3 \in O(1)$ and $\mathbf{N} = O(\lambda)$.

In terms of verification equations, by Lemma 2.2 for $t = \sqrt{2}$ and the union bound, the probability that $\|\mathbf{z}_1\| \leq \mathfrak{s}_1 \sqrt{2m_1 \hat{d}}$ and $\|\mathbf{z}_2\| \leq \mathfrak{s}_2 \sqrt{2m_2 \hat{d}}$ is overwhelming. We argue similarly for $\|\vec{z}_3\| > 1.7\mathfrak{s}_3 \sqrt{256}$ where we apply Lemma 2.2 for $t = 1.7$ instead¹⁶. The rest of the verification checks hold by careful inspection of the protocol (we refer to [57] for more details). \square

¹⁶ The reason is to ensure $\approx 2^{-128}$ correctness error when setting the parameters.

5.3.2 Zero-Knowledge

Lemma 5.6 (Zero-Knowledge). *Let $m_2 \geq 256/\hat{d} + \tau + 1$, $\alpha_1, \alpha_2, \alpha_3 \in O(\sqrt{\lambda})$ and $\mathbf{N}, d \in O(\lambda)$. Define the following variables:*

$$\begin{aligned} \mathfrak{s}_1 &:= \alpha_1 \nu \sqrt{\mathcal{B}_s^2 + \mathcal{B}_r^2 + \mathfrak{t}}, & \mathfrak{s}_2 &:= \alpha_2 \nu \sqrt{m_2 \hat{d}}, & \mathfrak{s}_3 &:= \alpha_3 \omega_{\max}(\lambda) \sqrt{\mathcal{B}_s^2 + \mathcal{B}_r^2 + \mathfrak{t}}, \\ \mathcal{B}_1 &:= \mathfrak{s}_1 \sqrt{2m_1 \hat{d}}, & \mathcal{B}_2 &:= \mathfrak{s}_2 \sqrt{2m_2 \hat{d}}, & \mathcal{B}_3 &:= 1.7 \mathfrak{s}_3 \sqrt{256} \end{aligned}$$

and the corresponding rejection rates:

$$M_i := \exp \left(\sqrt{\frac{2(\lambda+1)}{\log e}} \cdot \frac{1}{\alpha_i} + \frac{1}{2\alpha_i^2} \right) \text{ for } i \in \{1, 2, 3\}.$$

Then, under the $\text{MLWE}_{n+256/\hat{d}+\tau+1, m_2, \chi, \hat{q}}$ assumption $\Pi_{\text{NIZK}}^{\text{ISIS}}$ is zero-knowledge.

Proof. We prove the statement using a hybrid argument. Let \mathcal{A} be a PPT algorithm, which makes at most Q_0 queries to \mathbf{H} and Q_1 queries to $\text{Prove}_{\text{ISIS}}$, outputs 1 with probability ϵ . In each i -th game, we will introduce PPT algorithms $\mathcal{S}_0, \mathcal{S}_1$, which have a shared state, and denote ϵ_i to be the probability that $1 \leftarrow \mathcal{A}^{\mathcal{S}_0, \mathcal{S}_1}(\text{crs})$.

Game₁: Here, \mathcal{S}_0 simulates the random oracle queries using lazy sampling and \mathcal{S}_1 runs $\text{Prove}_{\text{ISIS}}(\text{crs}, \mathbf{x}, \mathbf{w})$ truthfully. By construction, $\epsilon_1 = \epsilon$.

Game₂: When \mathcal{S}_1 is queried, before executing Line 6 it first samples $c \leftarrow \mathcal{C}$. Then in Line 32, instead of calling the random oracle \mathcal{S}_0 , it uses the challenge c generated at the very beginning. If $(4, \text{crs}, \mathbf{x}, a_1, a_2, a_3, a_4)$ was already queried earlier to \mathcal{S}_0 , then \mathcal{S}_1 aborts. Otherwise, it continues with the $\text{Prove}_{\text{ISIS}}(\text{crs}, \mathbf{x}, \mathbf{w})$ algorithm.

We argue the indistinguishability following the proof of [53, Lemma 5.3]. Namely, the only difference between **Game₁** and **Game₂** is that we directly program c in Line 32 without checking whether $\mathbf{a} := (4, \text{crs}, \mathbf{x}, a_1, a_2, a_3, a_4)$ was already set. Since \mathcal{A} calls \mathcal{S}_0 at most Q_0 times and \mathcal{S}_1 at most Q_1 times, at most $Q_0 + \mathbf{N}Q_1$ of such tuples \mathbf{a} will ever be set. Recall that \mathbf{a} , and especially a_1 , contains $\mathbf{w} = \mathbf{A}_1 \mathbf{y}_1 + \mathbf{A}_2 \mathbf{y}_2$. Therefore, at most $Q_0 + \mathbf{N}Q_1$ values of \mathbf{w} will be set. Also, with an overwhelming probability over the standard parameter choices [36, Appendix C], matrix \mathbf{A}_2 can be written in the Hermite Normal Form as $\mathbf{A}_2 = \hat{\mathbf{A}}_2^{-1} [\mathbf{A}'_2 \mathbf{I}_n]$ and thus

$$\mathbf{y}_{2,2} = \hat{\mathbf{A}}_2 (\mathbf{w} - \mathbf{A}_1 \mathbf{y}_1) - \mathbf{A}'_2 \mathbf{y}_{2,1}.$$

where $\mathbf{y}_2 := (\mathbf{y}_{2,1}, \mathbf{y}_{2,2}) \in \hat{\mathcal{R}}_{\hat{q}}^{m_2 - (n+256/\hat{d}+\tau+1)} \times \hat{\mathcal{R}}_{\hat{q}}^{n+256/\hat{d}+\tau+1}$. Hence, by [53, Lemma 4.4], for any fixed $\mathbf{w} \in \hat{\mathcal{R}}_{\hat{q}}^n$ we have

$$\begin{aligned} \Pr \left[\mathbf{A}_1 \mathbf{y}_1 + \mathbf{A}_2 \mathbf{y}_2 = \mathbf{w} : \mathbf{y}_1 \leftarrow D_{\mathfrak{s}_2}^{m_1 \hat{d}}, \mathbf{y}_2 \leftarrow D_{\mathfrak{s}_2}^{m_2 \hat{d}} \right] &\leq \max_{\mathbf{w}'_1 \in \hat{\mathcal{R}}_{\hat{q}}^n} \Pr \left[\mathbf{y}_{2,2} = \mathbf{w}'_1 : \mathbf{y}_{2,2} \leftarrow D_{\mathfrak{s}_2}^{n\hat{d}} \right] \\ &\leq 2^{-n\hat{d}}. \end{aligned}$$

So, if \mathcal{S}_1 is accessed Q_1 times and the probability of getting a collision each time is at most $(Q_0 + \mathbf{N}Q_1)2^{-n\hat{d}}$, then the probability that a collision occurs after Q_1 queries is at most $\mathbf{N}Q_1(Q_0 + \mathbf{N}Q_1)2^{-n\hat{d}}$. Since $\mathbf{N}, d \in O(\lambda)$ and Q_0, Q_1 are polynomials in λ , we get

$$|\epsilon_2 - \epsilon_1| \leq \mathbf{N}Q_1(Q_0 + \mathbf{N}Q_1)2^{-n\hat{d}} + \text{negl}(\lambda) = \text{negl}(\lambda).$$

Game₃: Here, \mathcal{S}_1 directly computes $\mathbf{z}_i = \mathbf{y}_i + c\mathfrak{s}_i$ for $i = 1, 2$ and $\mathbf{w} := \mathbf{A}_1 \mathbf{z}_1 + \mathbf{A}_2 \mathbf{z}_2 - c\mathfrak{t}_A$. Further, \mathcal{S}_1 later calculates $f_0 := \mathbf{z}^T \mathbf{D}_2 \mathbf{z} + c\mathbf{d}_1^T \mathbf{z} + c^2 d_0 - \mathfrak{f}$, where $\mathfrak{f} := c\mathfrak{t} - \mathbf{b}^T \mathbf{z}_2$, instead. By the verification equations and simple rearrangement argument, we have $\epsilon_3 = \epsilon_2$.

Game₅: Now, \mathcal{S}_1 directly samples each $\mathbf{z}_2 \leftarrow D_{\mathfrak{s}_2}^{m_2 \hat{d}}$ and sets $b_2 = 1$ with probability $1/M_2$. By Lemma 2.4, we have $|\epsilon_5 - \epsilon_4| \leq N \cdot 2^{-\lambda} = \text{negl}(\lambda)$.

Game₆: This time, \mathcal{S}_1 samples uniformly random $\nu \leftarrow \hat{\mathcal{R}}_q^{n+256/\hat{d}+\tau+1}$ and computes the commitment the following way:

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_y \\ \mathbf{t}_g \\ t \end{bmatrix} := \nu + \begin{bmatrix} \mathbf{A}\mathbf{s}_1 \\ \mathbf{y}_3 \\ \mathbf{g} \\ f_1 \end{bmatrix}.$$

One can naturally construct a PPT adversary $\mathcal{A}^{\text{MLWE}}$ which breaks Module-LWE with probability $|\epsilon_5 - \epsilon_4|/Q_1$. Hence,

$$|\epsilon_6 - \epsilon_5| \leq N \cdot Q_1 \cdot \text{Adv}_{n+256/\hat{d}+\tau+1, m_2, \chi}^{\text{MLWE}}(\mathcal{A}^{\text{MLWE}}) = \text{negl}(\lambda).$$

Game₇: Here, \mathcal{S}_1 samples directly $(\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, t) \leftarrow \hat{\mathcal{R}}_q^{n+256/\hat{d}+\tau+1}$. Clearly, we have $\epsilon_7 = \epsilon_6$.

Game₈: Now, \mathcal{S}_1 generates $\mathbf{z}_1 \leftarrow D_{\mathfrak{s}_1}^{m_1 \hat{d}}$ and sets $b_1 = 1$ with probability $1/M_1$. By Lemma 2.4, we have $|\epsilon_8 - \epsilon_7| \leq N \cdot 2^{-\lambda} = \text{negl}(\lambda)$.

Game₉: Here, \mathcal{S}_1 picks $\mathbf{h} \leftarrow \{x \in \hat{\mathcal{R}}_q : \tilde{x} = 0\}^\tau$ directly. Since earlier \mathbf{g} was chosen uniformly at random from $\{x \in \hat{\mathcal{R}}_q : \tilde{x} = 0\}^\tau$, we get $\epsilon_9 = \epsilon_8$.

Game₁₀: This time \mathcal{S}_1 samples $\vec{z}_3 \leftarrow D_{\mathfrak{s}_3}^{256}$ and sets $b_3 = 1$ with probability $1/M_3$. By Lemma 2.4, $|\epsilon_{10} - \epsilon_9| \leq N \cdot 2^{-\lambda} = \text{negl}(\lambda)$.

Now, note that in the last security game \mathcal{S}_1 does not use the witness w (recall that \mathcal{S}_0 only simulates the random oracle H using lazy sampling). Hence, by the hybrid argument, $|\epsilon_{10} - \epsilon| = \text{negl}(\lambda)$ which concludes the proof. \square

6 Efficient Multi-Proof Extractable Non-Interactive Zero-Knowledge Proofs

This section focuses on constructing multi-proof extractable NIZK for proving knowledge of a short vector $\vec{s} \in \hat{\mathcal{R}}_q^m$ such that $\|\vec{s}\| \leq \mathcal{B}$ and $P\vec{s} = \vec{u}$ over \mathbb{Z}_q where $\mathbf{P} \in \hat{\mathcal{R}}_q^{n \times m}$ and $\vec{u} \in \mathbb{Z}_q^n$ are public. We define the corresponding relation

$$R^{\text{Com}} := \{x := (P, \vec{u}), w := \vec{s} \mid P \in \mathbb{Z}_q^{n \times m}, \vec{u} \in \hat{\mathcal{R}}_q^n, \vec{s} \in \hat{\mathcal{R}}_q^m : P\vec{s} = \vec{u} \wedge \|\vec{s}\| \leq \mathcal{B}\}. \quad (23)$$

Throughout the section, we follow the footsteps of del Pino and Katsumata [67][Section 4.3] who added multi-proof extractability to the exact proof system by Bootle et al. [15]. Namely, we adapt the recently proposed lattice-based framework for proving exact statements by Lyubashevsky et al. [57] to prove relation R^{Com} . However, as shown in [57, Appendix B], the protocol only satisfies standard knowledge soundness property. Hence, we further apply the Katsumata transform [46] to achieve multi-proof extractability¹⁷

Conceptually, the protocol will be almost identical to the in Section 5. Hence, we use the same notation as in Section 5.1. As before, we assume that variables n, m are divisible by \hat{d} . By the sum-of-four-squares theorem,

¹⁷ We note that in the random oracle model one could apply the standard “encryption-to-the-sky” paradigm to obtain straight-line extractability (e.g. [3, 13]). We provide the Katsumata transform instead since it has more potential to be further used to prove QROM security, as seen in [67]. We leave the quantum security analysis of the protocol, and of our constructions for future work.

$\|\vec{s}\| \leq \mathcal{B}_s$ is equivalent to the existence of four integers $a_1, a_2, a_3, a_4 \in \mathbb{Z}$ such that $\mathcal{B}^2 - \|\vec{s}\|^2 = a_1^2 + a_2^2 + a_3^2 + a_4^2$. Thus, we can define a vector $a := (a_1, a_2, a_3, a_4, 0, \dots, 0) \in \mathbb{Z}_q^{\hat{d}}$ and observe that:

$$\left\| \begin{bmatrix} \vec{s} \\ \vec{a} \end{bmatrix} \right\| = \mathcal{B}_s \quad \text{and} \quad [P \ 0] \begin{bmatrix} \vec{s} \\ \vec{a} \end{bmatrix} = P\vec{s}.$$

Hence, we slightly abuse the notation and define $\mathbf{m} := \mathbf{m} + 1$, $\vec{s} := (\vec{s}, \vec{a}) \in \mathbb{Z}_q^{\mathbf{m}}$ and $P := [P \ 0] \in \hat{\mathcal{R}}_q^{n \times \mathbf{m}}$. Therefore,

$$\|\vec{s}\| = \mathcal{B} \quad \text{and} \quad P\vec{s} = \vec{u}. \quad (24)$$

In this proof system, the common *random* string consists of uniformly random matrices which are only used for the ABDLOP commitment scheme:

$$\text{crs} := (\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_y, \mathbf{B}_g, \mathbf{b}) \in \hat{\mathcal{R}}_q^{n \times m_1} \times \hat{\mathcal{R}}_q^{n \times m_2} \times \hat{\mathcal{R}}_q^{256/\hat{d} \times m_2} \times \hat{\mathcal{R}}_q^{\tau \times m_2} \times \hat{\mathcal{R}}_q^{m_2}. \quad (25)$$

We define the statement as the tuple $\mathbf{x} := (P, \vec{u}, \mathcal{B})$.

First round. Let $\mathbf{s} = \text{Coeffs}^{-1}(\vec{s})$ and similarly for \mathbf{u} . We start by committing to the witness $\mathbf{s}_1 := \mathbf{s}$ and let $m_1 = \mathbf{m}$. We commit to \mathbf{s}_1 using the ABDLOP commitment, i.e. we sample the randomness vector $\mathbf{s}_2 \leftarrow \chi^{m_2}$, where χ is a probability distribution on ternary polynomials in $\hat{\mathcal{R}}_q$, and compute:

$$\mathbf{t}_A := \mathbf{A}_1 \mathbf{s}_1 + \mathbf{A}_2 \mathbf{s}_2 \quad \text{where} \quad (\mathbf{A}_1, \mathbf{A}_2) \in \hat{\mathcal{R}}_q^{n \times m_1} \times \hat{\mathcal{R}}_q^{n \times m_2}.$$

Then, we sample short masking vectors $\mathbf{y}_i \leftarrow D_{\mathbf{s}_i}^{m_i}$ for $i = 1, 2$, and compute $\mathbf{w} := \mathbf{A}_1 \mathbf{y}_1 + \mathbf{A}_2 \mathbf{y}_2$. Further, we generate the polynomial vector $\mathbf{y}_3 \leftarrow D_{\mathbf{s}_3}^{256}$ and commit to it as follows

$$\mathbf{t}_y := \mathbf{B}_y \mathbf{s}_2 + \mathbf{y}_3 \quad \text{where} \quad \mathbf{B}_y \in \hat{\mathcal{R}}_q^{\frac{256}{\hat{d}} \times m_2}.$$

Finally, we sample a polynomial vector $\mathbf{g} \leftarrow \{x \in \hat{\mathcal{R}}_q : \tilde{x} = 0\}^\tau$ and commit to it:

$$\mathbf{t}_g := \mathbf{B}_g \mathbf{s}_2 + \mathbf{g} \quad \text{where} \quad \mathbf{B}_g \leftarrow \hat{\mathcal{R}}_q^{\tau \times m_2}.$$

Hence, the first message and the corresponding challenge are

$$a_1 := (\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w}) \quad \text{and} \quad (\mathfrak{R}_0, \mathfrak{R}_1) = \text{H}(1, \text{crs}, \mathbf{x}, a_1) \in \{0, 1\}^{256 \times m_1 \hat{d}} \times \{0, 1\}^{256 \times m_1 \hat{d}}.$$

Second round. Given the first challenge $(\mathfrak{R}_0, \mathfrak{R}_1)$, we compute the response $\vec{z}_3 := \vec{y}_3 + \mathfrak{R} \vec{s}_1 \in \mathbb{Z}^{256}$, where $\mathfrak{R} := \mathfrak{R}_0 - \mathfrak{R}_1$ and also $\vec{y}_3 := \text{Coeffs}(\mathbf{y}_3)$, $\vec{s}_1 = \text{Coeffs}(\mathbf{s}_1)$. Then, we apply rejection sampling on \vec{z}_3 . The second message and the corresponding challenge are:

$$a_2 := \vec{z}_3 \quad \text{and} \quad (\gamma_{i,j})_{i \in [\tau], j \in [256+n+1]} = \text{H}(2, \text{crs}, \mathbf{x}, a_1, a_2) \in \mathbb{Z}_q^{\tau \times (256+n+1)}.$$

Third round. The prover's goal now is to prove the following relations over \mathbb{Z}_q :

$$\begin{cases} \vec{z}_3 = \vec{y}_3 + \mathfrak{R} \vec{s}_1 \\ P \vec{s}_1 = \vec{u} \\ \langle \vec{s}_1, \vec{s}_1 \rangle = \mathcal{B}^2 \end{cases}$$

where the terms in blue are secret and committed. As before, we deduce that proving the first equation is equivalent to proving that for $i \in [256]$, the constant coefficient of

$$\sigma(\mathbf{r}_i)^T \mathbf{s}_1 + \sigma(\mathbf{e}_i)^T \mathbf{y}_3 - z_{3,i} \text{ is equal to zero,}$$

where \mathbf{r}_i is the polynomial vector such that its coefficient vector is the i -th row of \mathfrak{R} , and \mathbf{e}_i is the polynomial vector so that $\text{Coeffs}(\mathbf{e}_i)$ is a unit vector with its i -th coefficient being 1. Similarly, the second equation is equivalent to: for all $i \in [n]$, the constant coefficient of:

$$\sigma(\mathbf{p}_i)^T \mathbf{s}_1 - \sigma(\mathbf{e}'_i)^T \mathbf{u} \text{ is equal to zero,}$$

where \mathbf{p}_i is the polynomial vector with its coefficient vector being the i -th row of P , and \mathbf{e}'_i is the polynomial vector so that $\text{Coeffs}(\mathbf{e}'_i)$ is a unit vector with its i -th coefficient being 1. Obviously, the last equation is equivalent to $\sigma(\mathbf{s}_1)^T \mathbf{s}_1 - \mathcal{B}$ having the constant coefficient equal to zero. Hence, we compute

$$\begin{aligned} h_i &:= g_i + \sum_{j=1}^{256} \gamma_{i,j} \cdot \left(\sigma(\mathbf{r}_j)^T \mathbf{s}_1 + \sigma(\mathbf{e}_j)^T \mathbf{y}_3 - z_j \right) \\ &\quad + \sum_{j=1}^n \gamma_{i,256+j} \cdot \left(\sigma(\mathbf{p}_j)^T \mathbf{s}_1 - \sigma(\mathbf{e}'_j)^T \mathbf{u} \right) \\ &\quad + \gamma_{i,256+n+1} \cdot \left(\sigma(\mathbf{s}_1)^T \mathbf{s}_1 - \mathcal{B} \right) \in \hat{\mathcal{R}}_{\hat{q}}. \end{aligned} \tag{26}$$

for $i = 1, 2, \dots, \tau$. By definition of \mathbf{g} , the constant coefficients of h_1, \dots, h_τ are all zeroes. The third message and the corresponding challenge are:

$$a_3 := (h_1, \dots, h_\tau) \quad \text{and} \quad \boldsymbol{\mu} = (\mu_i)_{i \in [\tau]} = \mathbf{H}(\mathbf{3}, \text{crs}, \mathbf{x}, a_1, a_2, a_3) \in \hat{\mathcal{R}}_{\hat{q}}^\tau.$$

Fourth round. The goal of the prover is now to prove the τ quadratic equations (with automorphisms) over $\mathcal{R}_{\hat{q}}$ from (26). We now linear-combine them with the challenges μ_i :

$$\begin{aligned} 0 &= \sum_{i=1}^{\tau} \mu_i \left(\sum_{j=1}^{256} \gamma_{i,j} \cdot \left(\sigma(\mathbf{r}_j)^T \mathbf{s}_1 + \sigma(\mathbf{e}_j)^T \mathbf{y}_3 - z_j \right) + \sum_{j=1}^n \gamma_{i,256+j} \cdot \left(\sigma(\mathbf{p}_j)^T \mathbf{s}_1 - \sigma(\mathbf{e}'_j)^T \mathbf{u} \right) \right. \\ &\quad \left. + \gamma_{i,256+n+1} \cdot \left(\sigma(\mathbf{s}_1)^T \mathbf{s}_1 - \mathcal{B} \right) + g_i - h_i \right). \end{aligned} \tag{27}$$

Let us define

$$\mathbf{B} := \begin{bmatrix} \mathbf{B}_y \\ \mathbf{B}_g \end{bmatrix}, \quad \mathbf{t}_B := \begin{bmatrix} \mathbf{t}_y \\ \mathbf{t}_g \end{bmatrix}, \quad \mathbf{m} := \begin{bmatrix} \mathbf{y}_3 \\ \mathbf{g} \end{bmatrix} \in \hat{\mathcal{R}}_{\hat{q}}^{256/d+\tau} \quad \text{and} \quad \hat{\mathbf{s}} := \begin{bmatrix} \mathbf{s}_1 \\ \sigma(\mathbf{s}_1) \\ \mathbf{m} \\ \sigma(\mathbf{m}) \end{bmatrix}. \tag{28}$$

Then, the quadratic equation above can be written equivalently as

$$\hat{\mathbf{s}}^T \mathbf{D}_2 \hat{\mathbf{s}} + \mathbf{d}_1^T \hat{\mathbf{s}} + d_0 = 0$$

where

$$\begin{aligned}
\mathbf{D}_2 &:= \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \sum_{i=1}^{\tau} \mu_i \gamma_{i,256+n+1} \cdot \mathbf{I}_{m_1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \\
\mathbf{d}_1 &:= \begin{bmatrix} \sum_{i=1}^{\tau} \mu_i \cdot \left(\sum_{j=1}^{256} \gamma_{i,j} \sigma(\mathbf{r}_j) + \sum_{j'=1}^n \gamma_{i,256+j'} \sigma(\mathbf{p}_{j'}) \right) \\ \mathbf{0} \\ \sum_{i=1}^{\tau} \sum_{j=1}^{256} \mu_i \gamma_{i,j} \sigma(\mathbf{e}_j) \\ \boldsymbol{\mu} \\ \mathbf{0} \end{bmatrix} \\
d_0 &:= - \sum_{i=1}^{\tau} \mu_i \cdot \left(\sum_{j=1}^{256} z_j + \sum_{j'=1}^n \sigma(\mathbf{e}'_{j'}) \mathbf{u} + \gamma_{i,256+n+1} \mathbf{B} + h_i \right).
\end{aligned} \tag{29}$$

Next, we run the sub-protocol for proving a single quadratic equation with automorphisms identically as before. First, we calculate the garbage term $f_1 = \hat{\mathbf{s}}^T \mathbf{R}_2 \mathbf{y} + \mathbf{y}^T \mathbf{R}_2 \hat{\mathbf{s}} + \mathbf{r}_1^T \mathbf{y}$, where \mathbf{y} is defined as

$$\mathbf{y} := \begin{bmatrix} \mathbf{y}_1 \\ \sigma(\mathbf{y}_1) \\ -\mathbf{B}\mathbf{y}_2 \\ -\sigma(\mathbf{B}\mathbf{y}_2) \end{bmatrix} \in \hat{\mathcal{R}}_q^{2(m_1+256/d+\tau)}, \tag{30}$$

and the commitment $t = \mathbf{b}^T \mathbf{s}_2 + f_1$ to f_1 . Then, we set $f_0 = \mathbf{y}^T \mathbf{R}_2 \mathbf{y} + \mathbf{b}^T \mathbf{y}_2$. Hence, the fourth message and the corresponding challenge are:

$$a_4 := (t, f_0) \quad \text{and} \quad c := \mathbf{H}(4, \text{crs}, x, a_1, a_2, a_3, a_4) \in \mathcal{C}.$$

Final round. Given a challenge c , we compute $\mathbf{z}_i = c\mathbf{s}_i + \mathbf{y}_i$ for $i = 1, 2$ and apply rejection sampling on $\mathbf{z}_1, \mathbf{z}_2$. So, the prover's last message is $a_5 := (\mathbf{z}_1, \mathbf{z}_2)$ and thus the proof consists of

$$\pi := (a_1, a_2, a_3, a_4, a_5).$$

Verification algorithm. The verifier is given a proof π and recomputes the corresponding challenges as well as $\mathbf{D}_2, \mathbf{d}_1, d_0$. Denote

$$\mathbf{z} := \begin{bmatrix} \mathbf{z}_1 \\ \sigma(\mathbf{z}_1) \\ c\mathbf{t}_B - \mathbf{B}\mathbf{z}_2 \\ \sigma(c\mathbf{t}_B - \mathbf{B}\mathbf{z}_2) \end{bmatrix}. \tag{31}$$

Finally, the verifier outputs 1 if *all* the following relations hold:

$$\begin{cases} \|\mathbf{z}_1\| \stackrel{?}{\leq} \mathcal{B}_1, \|\mathbf{z}_2\| \stackrel{?}{\leq} \mathcal{B}_2, \|\vec{z}_3\| \stackrel{?}{\leq} \mathcal{B}_3, \\ \tilde{h}_i \stackrel{?}{=} 0 \text{ for } i \in [\tau], \\ \mathbf{A}_1 \mathbf{z}_1 + \mathbf{A}_2 \mathbf{z}_2 \stackrel{?}{=} \mathbf{w} + c\mathbf{t}_A, \\ \mathbf{z}^T \mathbf{D}_2 \mathbf{z} + c\mathbf{d}_1^T \mathbf{z} + c^2 d_0 - (c\mathbf{t} - \mathbf{b}^T \mathbf{z}_2) \stackrel{?}{=} f_0 \end{cases}$$

and 0 otherwise.

6.1 Katsumata Transform

In order to make the protocol multi-proof extractable, we apply the Katsumata transform [46]. Recall that the protocol from [57], as well as many prior linear-sized lattice-based proofs [7, 34, 15, 60, 74], follow the

LHC.Com _i (crs ^{LHC} , s _i , y _i)	LHC.Open _i (crs ^{LHC} , (com, c), st)
1: $(\bar{\mathbf{e}}_{i,1}, \bar{\mathbf{e}}_{i,2}, \bar{\mathbf{e}}_{i,3}) \leftarrow S_{\eta_i}^{m_i} \times S_{\eta_i}^{m_i} \times S_{\eta_i}^{m_i}$	1: for $j \in \{1, 2, 3\}$ do
2: $\bar{\mathbf{t}}_{i,1} = \bar{p} \cdot (\bar{\mathbf{A}}\bar{\mathbf{e}}_{i,1} + \bar{\mathbf{e}}_{i,2})$	2: $\bar{\mathbf{z}}_{i,j} = \bar{\mathbf{f}}_{i,j} + c \cdot \bar{\mathbf{e}}_{i,j}$
3: $\bar{\mathbf{t}}_{i,2} = \bar{p} \cdot (\bar{\mathbf{B}}\bar{\mathbf{e}}_{i,1} + \bar{\mathbf{e}}_{i,3}) + \mathbf{s}_i$	3: $b \leftarrow \text{Rej}([\bar{\mathbf{z}}_{i,1} \parallel \bar{\mathbf{z}}_{i,2} \parallel \bar{\mathbf{z}}_{i,3}],$
4: $(\bar{\mathbf{f}}_{i,1}, \bar{\mathbf{f}}_{i,2}, \bar{\mathbf{f}}_{i,3}) \leftarrow D_{\bar{s}_i}^{m_i} \times D_{\bar{s}_i}^{m_i} \times D_{\bar{s}_i}^{m_i}$	4: $c \cdot [\bar{\mathbf{e}}_{i,1} \parallel \bar{\mathbf{e}}_{i,2} \parallel \bar{\mathbf{e}}_{i,3}], \bar{s}_i, \bar{M}_i)$
5: $\bar{\mathbf{w}}_{i,1} = \bar{p} \cdot (\bar{\mathbf{A}}\bar{\mathbf{f}}_{i,1} + \bar{\mathbf{f}}_{i,2})$	5: if $b = 0$ then return $\text{op} = \perp$
6: $\bar{\mathbf{w}}_{i,2} = \bar{p} \cdot (\bar{\mathbf{B}}\bar{\mathbf{f}}_{i,1} + \bar{\mathbf{f}}_{i,3}) + \mathbf{y}_i$	6: else return $\text{op} = [\bar{\mathbf{z}}_{i,1} \parallel \bar{\mathbf{z}}_{i,2} \parallel \bar{\mathbf{z}}_{i,3}]$
7: $\text{com} = (\bar{\mathbf{t}}_{i,1}, \bar{\mathbf{t}}_{i,2}, \bar{\mathbf{w}}_{i,1}, \bar{\mathbf{w}}_{i,2})$	<u>LHC.Verify_i(crs^{LHC}, (com, c), (z, op ≠ ⊥))</u>
8: $\text{st} = (\bar{\mathbf{e}}_{i,1}, \bar{\mathbf{e}}_{i,2}, \bar{\mathbf{e}}_{i,3}, \bar{\mathbf{f}}_{i,1}, \bar{\mathbf{f}}_{i,2}, \bar{\mathbf{f}}_{i,3})$	1: $(\bar{\mathbf{t}}_{i,1}, \bar{\mathbf{t}}_{i,2}, \bar{\mathbf{w}}_{i,1}, \bar{\mathbf{w}}_{i,2}) \leftarrow \text{com}$
9: return (com, st)	2: $[\bar{\mathbf{z}}_{i,1} \parallel \bar{\mathbf{z}}_{i,2} \parallel \bar{\mathbf{z}}_{i,3}] \leftarrow \text{op}$
	3: if $\ \bar{\mathbf{z}}_{i,1}\ > \bar{s}_i \sqrt{2n_i d}$ then return 0
	4: for $j \in \{2, 3\}$ do
	5: if $\ \bar{\mathbf{z}}_{i,j}\ > \bar{s}_i \sqrt{2m_i \hat{d}}$ then return 0
	6: $\bar{\mathbf{z}}_A = c\bar{\mathbf{t}}_{i,1} + \bar{\mathbf{w}}_{i,1} - \bar{p} \cdot (\bar{\mathbf{A}}\bar{\mathbf{z}}_{i,1} + \bar{\mathbf{z}}_{i,2})$
	7: $\bar{\mathbf{z}}_B = c\bar{\mathbf{t}}_{i,2} + \bar{\mathbf{w}}_{i,2} - \bar{p} \cdot (\bar{\mathbf{A}}\bar{\mathbf{z}}_{i,1} + \bar{\mathbf{z}}_{i,3})$
	8: if $\bar{\mathbf{z}}_A \neq \mathbf{0} \vee \bar{\mathbf{z}}_B \neq \mathbf{z}$ then return 0
	9: else return 1

Fig. 11: Extractable linear homomorphic commitment from MLWE. Parts highlighted in gray correspond to the Katsumata transform.

commit-and-prove approach. Namely, we first commit to \mathbf{s}_1 using randomness \mathbf{s}_2 and in the final round we send the masked openings $\mathbf{z}_i = \mathbf{y}_i + c\mathbf{s}_i$ of \mathbf{s}_i to prove knowledge of a commitment opening. Now, to make this protocol straight-extractable, we use an *extractable linear homomorphic commitment* (LHC), which can be seen as a linear homomorphic encryption scheme with pseudo-random public keys, and additionally encrypt both \mathbf{s}_i and \mathbf{y}_i .

In this section, we recall the (simplified) LHC construction from Module-LWE [46, Section 3.4]¹⁸. Since we need to commit to both $(\mathbf{s}_1, \mathbf{y}_1)$ and $(\mathbf{s}_2, \mathbf{y}_2)$, we will need two instantiations. Fix $i \in \{1, 2\}$. The common random string, i.e. the commitment key $\text{crs}_i^{\text{LHC}} = (\bar{\mathbf{A}}_i, \bar{\mathbf{B}}_i)$ contains two uniformly random matrices in $\mathcal{R}_{\hat{q}}^{m_i \times n_i} \times \mathcal{R}_{\hat{q}}^{m_i \times n_i}$. Next, we have three algorithms: LHC.Com_i, LHC.Open_i and LHC.Verify_i which are described in Figure 11. In the following, let $\bar{p} < \hat{q}$ be an odd integer co-prime to \hat{q} .

We present the NIZK proof system $\Pi_{\text{NIZK}}^{\text{Com}} := (\text{Prove}_{\text{Com}}^{\text{H}}, \text{Verify}_{\text{Com}}^{\text{H}})$ for relation R^{Com} in Figures 12 and 13. Intuitively, we take the protocol from above and additionally let the prover run $(\text{com}_i, \text{st}_i) \leftarrow \text{LHC.Com}_i(\text{crs}_i^{\text{LHC}}, \mathbf{s}_i, \mathbf{y}_i)$ for $i = 1, 2$. Hence, the first message now includes com_1 and com_2 . Then, in the final round, the prover runs

$$\text{op}_i \leftarrow \text{LHC.Open}_i(\text{crs}_i^{\text{LHC}}, (\text{com}_i, c), \text{st}_i)$$

and appends op_1 and op_2 to the final message (unless one of them is \perp , then it restarts). Finally, the verifier additionally checks for $i = 1, 2$ that

$$\text{LHC.Verify}(\text{crs}_i^{\text{LHC}}, (\text{com}_i, c, \text{op}_i \neq \perp)) \stackrel{?}{=} 1.$$

6.2 Security Analysis

In this subsection, we prove the key properties of $\Pi_{\text{NIZK}}^{\text{Com}}$: correctness, zero-knowledge and multi-proof extractability.

¹⁸ The reason why we do not apply the NTRU-type construction, as in [67], is that for efficiency we will pick relatively small ring dimension \hat{d} . In this case, the construction would rely on a *Module-NTRU* type assumption, which is still relatively uncommon and not well-studied.

```

ProveComH(crs, x, w)
Input: crs = (A1, A2, By, Bg, b, Ā1, Ā2, B̄1, B̄2), x = (P, u, B), w = s̄
Output: proof π
1: rst = 0                                ▷ boolean which indicates if a restart is necessary
2: idx = 0                                ▷ keeps track how many restarts occurred
3: s1 = s                                ▷ s is the polynomial vectors with coeff. s̄
4: while rst = 0 ∧ idx < N do
5:   idx = idx + 1
6:   s2 ← χm2                            ▷ generate commitment randomness
7:   tA = A1s1 + A2s2                    ▷ Ajtai commitment to s1
8:   y1 ← Ds1m1d̄, y2 ← Ds2m2d̄, y3 ← Ds3256    ▷ sample masking vectors
9:   g ← {x ∈ R̄q̄ : x̄ = 0}T                ▷ sample random polynomials with const. coeff. zero
10:  w = A1y1 + A2y2                    ▷ used to prove knowledge of a commitment opening
11:  ty = Bys2 + y3                    ▷ commitment to y3
12:  tg = Bgs2 + g                      ▷ commitment to g
13:  (com1, st1) ← LHC.Com1((Ā1, B̄1), s1, y1)    ▷ LHC to s1 and y1
14:  (com2, st2) ← LHC.Com2((Ā2, B̄2), s2, y2)    ▷ LHC to s2 and y2
15:  a1 = (tA, ty, tg, w, v, com1, com2)    ▷ first message
16:  (R̄0, R̄1) = H(1, crs, x, a1) ∈ {0, 1}256×m1d̄ × {0, 1}256×m1d̄    ▷ first challenge
17:  R̄ = R̄0 - R̄1
18:  s̄1 = Coeffs(s1)
19:  ȳ3 = Coeffs(y3)
20:  z̄3 = ȳ3 + R̄s̄1 ∈ Z256                ▷ masked opening of R̄s̄1
21:  b3 ← Rej(z̄3, R̄s̄1, s3, M3)
22:  a2 = z̄3                                ▷ second message
23:  (γi,j)i∈[τ], j∈[256+n+1]} = H(2, crs, x, a1, a2) ∈ Zq̄τ×(256+n+1)}    ▷ second challenge
24:  for i ∈ [τ] do
25:    compute hi as in (26)
26:  a3 = h = (h1, ..., hτ)                ▷ third message
27:  (μi)i∈[τ]} = H(3, crs, x, a1, a2, a3) ∈ R̄q̄τ    ▷ third challenge
28:  compute B, tB, m, s̄, y as in (28) and (30)
29:  compute D2, d1, d0 as in (29)
30:  f1 = sTR2y + yTR2s + r1Ty
31:  f0 = yTR2y + bTy2
32:  t = bTs2 + f1                            ▷ commitment to f1
33:  a4 = (t, f0)                            ▷ fourth message
34:  c = H(4, crs, x, a1, a2, a3, a4) ∈ C    ▷ fourth challenge
35:  for i ∈ {1, 2} do
36:    zi = yi + csi                            ▷ maked opening of csi
37:    bi ← Rej(zi, csi, si, Mi)                ▷ rejection sampling
38:    opi ← LHC.Openi((Āi, B̄i), (comi, c), sti)    ▷ opening of the LHC
39:    if opi = ⊥ then b̄i = 0
40:    else b̄i = 1
41:  π = (tA, ty, tg, w, com1, com2, z̄3, h, t, f0, z̄3, h, z1, z2, op1, op2)
42:  rst = b1b2b3b̄1b̄2                            ▷ if all rejection steps passed then rst = 1
43: if rst = 1 then return π
44: else return ⊥

```

Fig. 12: Prove_{Com}^H algorithm.

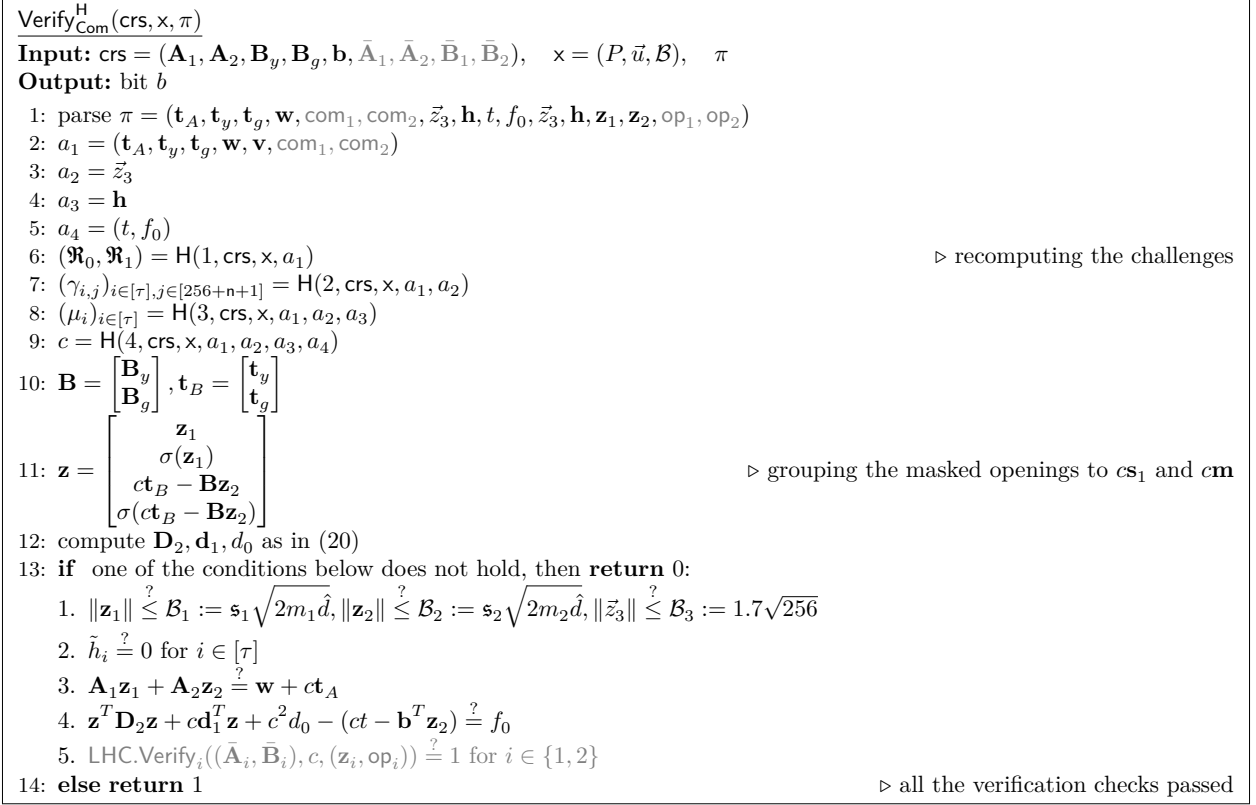


Fig. 13: $\text{Verify}_{\text{Com}}^{\text{H}}$ algorithm.

6.2.1 Correctness

Lemma 6.1 (Correctness). *Let $\alpha_1, \alpha_2, \alpha_3, \bar{\alpha}_1, \bar{\alpha}_2 \in O(\sqrt{\lambda})$, $N, d \in O(\lambda)$. Recall ν, ω_{\max} from Section 5.1 and define the following standard deviations:*

$$\begin{aligned} \mathfrak{s}_1 &:= \alpha_1 \nu \mathcal{B}, & \mathfrak{s}_2 &:= \alpha_2 \nu \sqrt{m_2 \hat{d}}, & \mathfrak{s}_3 &:= \alpha_3 \mathcal{B} \omega_{\max}(\lambda) \\ \bar{\mathfrak{s}}_1 &:= \bar{\alpha}_1 \eta_1 \nu \sqrt{(n_1 + 2m_1) \hat{d}}, & \bar{\mathfrak{s}}_2 &:= \bar{\alpha}_2 \eta_2 \nu \sqrt{(n_2 + 2m_2) \hat{d}} \end{aligned}$$

and the corresponding rejection rates for $i \in \{1, 2, 3\}$ and $j \in \{1, 2\}$:

$$M_i := \exp\left(\sqrt{\frac{2(\lambda+1)}{\log e}} \cdot \frac{1}{\alpha_i} + \frac{1}{2\alpha_i^2}\right), \quad \bar{M}_j := \exp\left(\sqrt{\frac{2(\lambda+1)}{\log e}} \cdot \frac{1}{\bar{\alpha}_j} + \frac{1}{2\bar{\alpha}_j^2}\right).$$

Then, $\Pi_{\text{NIZK}}^m := (\text{Prove}_{\text{ISIS}}^{\text{H}}, \text{Verify}_{\text{ISIS}}^{\text{H}})$ for relation R^{Com} is correct.

Proof. The proof is almost identical to Lemma 5.5. Let us argue that the prover algorithm outputs \perp with negligible probability. By Lemma 2.4 and the following inequalities (which hold with an overwhelming probability):

$$\begin{aligned} \|\mathbf{cs}_1\| &\leq \mathcal{B}\nu, & \|\mathbf{cs}_2\| &\leq \nu \sqrt{m_2 \hat{d}}, & \|\mathfrak{R}\bar{\mathfrak{s}}_1\| &\leq \mathcal{B} \cdot \omega_{\max}, \\ \|c \cdot [\bar{\mathbf{e}}_{i,1} \parallel \bar{\mathbf{e}}_{i,2} \parallel \bar{\mathbf{e}}_{i,3}]\| &\leq \eta_i \nu \sqrt{(n_i + 2m_i) \hat{d}} & \text{for } i &= 1, 2, \end{aligned}$$

we deduce that in a single run the probability that $b_1 b_2 b_3 \bar{b}_1 \bar{b}_2 = 1$ is at least $(M_1 M_2 M_3 \bar{M}_1 \bar{M}_2)^{-1} - \text{negl}(\lambda)$, and thus the probability that $\text{Prove}_{\text{Com}}^{\text{H}}$ outputs \perp , i.e. makes exactly N unsuccessful attempts, is at most

$$\left(1 - \frac{1}{M_1 M_2 M_3 \bar{M}_1 \bar{M}_2} + \text{negl}(\lambda)\right)^N = \text{negl}(\lambda)$$

since $M_1, M_2, M_3, \bar{M}_1, \bar{M}_2 \in O(1)$ and $N = O(\lambda)$.

In terms of verification equations, by Lemma 2.2 for $t = \sqrt{2}$ and the union bound, the probability that one of the following conditions holds: (i) $\|\mathbf{z}_1\| > \mathfrak{s}_1 \sqrt{2m_1 \hat{d}}$, (ii) $\|\mathbf{z}_2\| > \mathfrak{s}_2 \sqrt{2m_2 \hat{d}}$, (iii) $\|\bar{\mathbf{z}}_{i,1}\| > \bar{\mathfrak{s}}_i \sqrt{2n_i \hat{d}}$ and (iv) $\|\bar{\mathbf{z}}_{i,j}\| > \bar{\mathfrak{s}}_i \sqrt{2m_i \hat{d}}$ for $i \in \{1, 2\}, j \in \{2, 3\}$, is negligible. We argue similarly for $\|\bar{\mathbf{z}}_3\| > 1.7\mathfrak{s}_3 \sqrt{256}$ where we apply Lemma 2.2 for $t = 1.7$ instead, as in Lemma 5.5. The rest of the verification checks hold by careful inspection of the protocol. \square

$\text{LHC.ZKSim}_i(\text{crs}^{\text{LHC}}, c, \mathbf{z}_i)$
1: $(\bar{\mathbf{z}}_{i,1}, \bar{\mathbf{z}}_{i,2}, \bar{\mathbf{z}}_{i,3}) \leftarrow D_{\bar{\mathfrak{s}}_i}^{n_i} \times D_{\bar{\mathfrak{s}}_i}^{m_i} \times D_{\bar{\mathfrak{s}}_i}^{m_i}$
2: $(\bar{\mathbf{t}}_{i,1}, \bar{\mathbf{t}}_{i,2}) \leftarrow \mathcal{R}_{\hat{q}}^{m_i} \times \mathcal{R}_{\hat{q}}^{m_i}$
3: $\mathbf{w}_1 = -c\bar{\mathbf{t}}_{i,1} + \bar{p} \cdot (\bar{\mathbf{A}}\bar{\mathbf{z}}_1 + \bar{\mathbf{z}}_2)$
4: $\mathbf{w}_2 = -c\bar{\mathbf{t}}_{i,2} + \bar{p} \cdot (\bar{\mathbf{A}}\bar{\mathbf{z}}_1 + \bar{\mathbf{z}}_3) + \mathbf{z}_i$
5: $\text{com} = (\bar{\mathbf{t}}_{i,1}, \bar{\mathbf{t}}_{i,2}, \bar{\mathbf{w}}_{i,1}, \bar{\mathbf{w}}_{i,2})$
6: $u \leftarrow [0, 1)$
7: if $u > 1/\bar{M}_i$ then $\text{op} = \perp$
8: else $\text{op} = [\bar{\mathbf{z}}_{i,1} \parallel \bar{\mathbf{z}}_{i,2} \parallel \bar{\mathbf{z}}_{i,3}]$
9: return (com, op)

Fig. 14: Simulator for LHC.

6.2.2 Zero-Knowledge

Theorem 6.2 (Zero-Knowledge). *As before, let $\alpha_1, \alpha_2, \alpha_3, \bar{\alpha}_1, \bar{\alpha}_2 \in O(\sqrt{\lambda})$ and $N, n, d \in O(\lambda)$. Define the following standard deviations:*

$$\begin{aligned} \mathfrak{s}_1 &:= \alpha_1 \mathcal{B}\nu, & \mathfrak{s}_2 &:= \alpha_2 \nu \sqrt{m_2 \hat{d}}, & \mathfrak{s}_3 &:= \alpha_3 \mathcal{B}\omega_{\max}(\lambda) \\ \bar{\mathfrak{s}}_1 &:= \bar{\alpha}_1 \eta_1 \nu \sqrt{(n_1 + 2m_1) \hat{d}}, & \bar{\mathfrak{s}}_2 &:= \bar{\alpha}_2 \eta_2 \nu \sqrt{(n_2 + 2m_2) \hat{d}} \end{aligned}$$

and the corresponding rejection rates for $i \in \{1, 2, 3\}$ and $j \in \{1, 2\}$:

$$M_i := \exp\left(\sqrt{\frac{2(\lambda+1)}{\log e}} \cdot \frac{1}{\alpha_i} + \frac{1}{2\alpha_i^2}\right), \quad \bar{M}_j := \exp\left(\sqrt{\frac{2(\lambda+1)}{\log e}} \cdot \frac{1}{\bar{\alpha}_j} + \frac{1}{2\bar{\alpha}_j^2}\right).$$

Then, $\Pi_{\text{NIZK}}^{\text{Com}}$ for relation R^{Com} is zero-knowledge under the $\text{MLWE}_{2m_1, 2m_1+n_1, \eta_1}$, $\text{MLWE}_{2m_2, 2m_2+n_2, \eta_2, \hat{q}}$ and $\text{MLWE}_{n+256/\hat{d}+\tau+1, m_2, \chi, \hat{q}}$ assumptions.

Proof. The proof is almost identical to the one in Lemma 5.6 and we show it using a hybrid argument. Let \mathcal{A} be a PPT algorithm, which makes at most Q_0 queries to H and Q_1 queries to $\text{Prove}_{\text{ISIS}}$, outputs 1 with probability ϵ . In each i -th game, we will introduce PPT algorithms $\mathcal{S}_0, \mathcal{S}_1$, which have a shared state, and

denote ε_i to be the probability that $1 \leftarrow \mathcal{A}^{\mathcal{S}_0, \mathcal{S}_1}(\text{crs})$.

Game₁: Here, \mathcal{S}_0 simulates the random oracle queries using lazy sampling and \mathcal{S}_1 runs $\text{Prove}_{\text{Com}}(\text{crs}, \mathbf{x}, \mathbf{w})$ truthfully. By construction, $\varepsilon_1 = \varepsilon$.

Game₂: When \mathcal{S}_1 is queried, before executing Line 6 it first samples $c \leftarrow \mathcal{C}$. Then in Line 34, instead of calling the random oracle \mathcal{S}_0 , it uses the challenge c generated at the very beginning. If $(4, \text{crs}, \mathbf{x}, a_1, a_2, a_3, a_4)$ was already queried earlier to \mathcal{S}_0 , then \mathcal{S}_1 aborts. Otherwise, it continues with the $\text{Prove}_{\text{Com}}(\text{crs}, \mathbf{x}, \mathbf{w})$ algorithm. Similarly as in Lemma 5.6, we argue that

$$|\varepsilon_2 - \varepsilon_1| \leq \text{NQ}_1(Q_0 + \text{NQ}_1)2^{-n\hat{d}} + \text{negl}(\lambda) = \text{negl}(\lambda).$$

Game₃: In this game, immediately after sampling $c \leftarrow \mathcal{C}$ and generating $\mathbf{y}_1, \mathbf{y}_2$, the algorithm directly computes $\mathbf{z}_i = \mathbf{y}_i + c\mathbf{s}_i$ and runs

$$(\text{com}_i, \text{op}_i) \leftarrow \text{LHC.ZKSim}_i((\bar{\mathbf{A}}_i, \bar{\mathbf{B}}_i), c, \mathbf{z}_i)$$

defined in Figure 14 for $i = 1, 2$ instead. Then, by the zero-knowledge property of the extractable linear homomorphic commitment [46, Lemma 3.12], there exist PPT adversaries $\bar{\mathcal{A}}_1, \bar{\mathcal{A}}_2$ such that

$$|\varepsilon_3 - \varepsilon_2| \leq \mathbf{N} \cdot \left(\text{Adv}_{2m_1, 2m_1+n_1, \eta_1}^{\text{MLWE}}(\bar{\mathcal{A}}_1) + \text{Adv}_{2m_2, 2m_2+n_2, \eta_2}^{\text{MLWE}}(\bar{\mathcal{A}}_2) + \text{negl}(\lambda) \right) = \text{negl}(\lambda).$$

Game₄: Here, \mathcal{S}_1 directly computes $\mathbf{z}_i = \mathbf{y}_i + c\mathbf{s}_i$ for $i = 1, 2$ and $\mathbf{w} := \mathbf{A}_1\mathbf{z}_1 + \mathbf{A}_2\mathbf{z}_2 - ct_{\mathbf{A}}$. Further, \mathcal{S}_1 later calculates $f_0 := \mathbf{z}^T \mathbf{D}_2 \mathbf{z} + c\mathbf{d}_1^T \mathbf{z} + c^2 d_0 - \mathbf{f}$, where $\mathbf{f} := ct - \mathbf{b}^T \mathbf{z}_2$, instead. By the verification equations and simple rearrangement argument, we have $\varepsilon_4 = \varepsilon_3$.

Game₅: Now, \mathcal{S}_1 directly samples each $\mathbf{z}_2 \leftarrow D_{\mathbf{s}_2}^{m_2 \hat{d}}$ and sets $b_2 = 1$ with probability $1/M_2$. By Lemma 2.4, we have $|\varepsilon_5 - \varepsilon_4| \leq \mathbf{N} \cdot 2^{-\lambda} = \text{negl}(\lambda)$.

Game₆: This time, \mathcal{S}_1 samples uniformly random $\boldsymbol{\nu} \leftarrow \hat{\mathcal{R}}_{\hat{q}}^{n+256/\hat{d}+\tau+1}$ and computes the commitment the following way:

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_y \\ \mathbf{t}_g \\ t \end{bmatrix} := \boldsymbol{\nu} + \begin{bmatrix} \mathbf{A}\mathbf{s}_1 \\ \mathbf{y}_3 \\ \mathbf{g} \\ f_1 \end{bmatrix}.$$

One can then construct a PPT adversary $\mathcal{A}^{\text{MLWE}}$ which breaks Module-LWE with probability $|\varepsilon_5 - \varepsilon_4|/Q_1$. Hence,

$$|\varepsilon_6 - \varepsilon_5| \leq \mathbf{N} \cdot Q_1 \cdot \text{Adv}_{n+256/\hat{d}+\tau+1, m_2, \chi}^{\text{MLWE}}(\mathcal{A}^{\text{MLWE}}) = \text{negl}(\lambda).$$

Game₇: Here, \mathcal{S}_1 samples directly $(\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, t) \leftarrow \hat{\mathcal{R}}_{\hat{q}}^{n+256/\hat{d}+\tau+1}$. Thus, $\varepsilon_7 = \varepsilon_6$.

Game₈: Now, \mathcal{S}_1 generates $\mathbf{z}_1 \leftarrow D_{\mathbf{s}_1}^{m_1 \hat{d}}$ and sets $b_1 = 1$ with probability $1/M_1$. By Lemma 2.4, we have $|\varepsilon_8 - \varepsilon_7| \leq \mathbf{N} \cdot 2^{-\lambda} = \text{negl}(\lambda)$.

Game₉: Here, \mathcal{S}_1 picks $\mathbf{h} \leftarrow \{x \in \hat{\mathcal{R}}_{\hat{q}} : \tilde{x} = 0\}^\tau$ directly. Since earlier \mathbf{g} was chosen uniformly at random from $\{x \in \hat{\mathcal{R}}_{\hat{q}} : \tilde{x} = 0\}^\tau$, we get $\varepsilon_9 = \varepsilon_8$.

Game₁₀: This time \mathcal{S}_1 samples $\vec{z}_3 \leftarrow D_{\mathbf{s}_3}^{256}$ and sets $b_3 = 1$ with probability $1/M_3$. By Lemma 2.4, $|\varepsilon_{10} - \varepsilon_9| \leq \mathbf{N} \cdot 2^{-\lambda} = \text{negl}(\lambda)$.

We observe that in **Game₁₀** \mathcal{S}_1 does not use the witness w (recall that \mathcal{S}_0 only simulates the random oracle \mathbf{H} using lazy sampling). Hence, by the hybrid argument, $|\varepsilon_{10} - \varepsilon| = \text{negl}(\lambda)$ which concludes the proof. \square

6.2.3 Multi-Proof Extractability

We concentrate on proving multi-proof extractability. First, we provide basic intuition. Assume there is a PPT adversary which makes at most Q_H random oracle queries and outputs Q_s statement-proof pairs $(x_i, \pi_i)_{i \in [Q_s]}$ with non-negligible probability ε . Intuitively, we will use the property of *extractable* linear homomorphic commitments to extract a witness candidate w_i for each x_i from π_i . Then, for a fixed $i \in [Q_s]$, we will prove that the probability that w_i is *not* a valid witness for x_i must be negligible. Otherwise, we could run the knowledge extractor \mathcal{E} from [57, Appendix B] to obtain a witness w'_i for x_i . By nature of the [57] protocol, both w_i and w'_i are parts of two relaxed opening (c.f. Definition 5.1) to the same ABDLOP commitment. Hence, by Lemma 5.2 we conclude that $w_i \neq w'_i$ and the knowledge extractor finds a valid Module-SIS solution.

Theorem 6.3 (Multi-Proof Extractability). *Let $B := 8\nu\sqrt{\mathcal{B}_1^2 + \mathcal{B}_2^2}$ and*

$$\hat{q} > \max\left(\mathcal{B}, 16m_1\hat{d}\mathcal{B}_3, \frac{2}{\omega_{\min}(\lambda)^2}\mathcal{B}_3^2\right).$$

Then, the NIZK proof system $\Pi_{\text{NIZK}}^{\text{Com}}$ for relation R^{Com} is multi-proof extractable under the $\text{MLWE}_{m_1, m_1+n_1, \eta_1}$, $\text{MLWE}_{m_2, m_2+n_2, \eta_2}$ and the $\text{MSIS}_{n, m_1+m_2, B}$ assumptions.

Proof. We start with crs indistinguishability. Namely, define the simulator \mathcal{S}_{crs} which first samples

$$\begin{aligned} \text{crs} &:= (\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_y, \mathbf{B}_g, \mathbf{b}) \leftarrow \hat{\mathcal{R}}_{\hat{q}}^{n \times m_1} \times \hat{\mathcal{R}}_{\hat{q}}^{n \times m_1} \times \hat{\mathcal{R}}_{\hat{q}}^{n \times m_2} \times \hat{\mathcal{R}}_{\hat{q}}^{256/\hat{d} \times m_2} \times \hat{\mathcal{R}}_{\hat{q}}^{\tau \times m_2}, \\ (\bar{\mathbf{A}}_1, \bar{\mathbf{A}}_2) &\leftarrow \hat{\mathcal{R}}_{\hat{q}}^{m_1 \times n_1} \times \hat{\mathcal{R}}_{\hat{q}}^{m_2 \times n_2}. \end{aligned}$$

It also generates $\bar{\mathbf{D}}_{i,1} \leftarrow S_{\eta_i}^{m_i \times m_i}$, $\bar{\mathbf{D}}_{i,2} \leftarrow S_{\eta_i}^{m_i \times n_i}$ and sets $\bar{\mathbf{B}}_i := \bar{\mathbf{D}}_{i,1}\bar{\mathbf{A}}_i + \bar{\mathbf{D}}_{i,2}$. Finally, it outputs

$$\text{crs} = (\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_y, \mathbf{B}_g, \mathbf{b}, \bar{\mathbf{A}}_1, \bar{\mathbf{A}}_2, \bar{\mathbf{B}}_1, \bar{\mathbf{B}}_2)$$

and the trapdoor $\text{td} = (\bar{\mathbf{D}}_{1,1}, \bar{\mathbf{D}}_{1,2}, \bar{\mathbf{D}}_{2,1}, \bar{\mathbf{D}}_{2,2})$. Then, crs is indistinguishable from random crs based on the $\text{MLWE}_{m_1, n_1, \eta_1}$ and $\text{MLWE}_{m_2, n_2, \eta_2}$ assumptions.

We now turn to proving straight-line extractability. Suppose there is a PPT adversary which makes at most Q_H random oracle queries and outputs Q_s statement-proof pairs $(x_i, \pi_i)_{i \in [Q_s]}$ with non-negligible probability $\varepsilon(\lambda)$. We start by formally defining the Multi-Extract algorithm in Figure 15. To provide more intuition on what the algorithm does, consider the interactive proof implicitly defined in Figure 12 by $\text{Prove}_{\text{Com}}$. Then, a valid transcript is of the form

$$\text{tr} = ((\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \text{com}_1, \text{com}_2), (\mathfrak{R}_0, \mathfrak{R}_1), \vec{z}_3, (\gamma_{i,j}), \mathbf{h}, (\mu_i), (t, f_0), c, (\mathbf{z}_1, \mathbf{z}_2, \text{op}_1, \text{op}_2))$$

where the challenges are highlighted in gray. One of the main observations is that if there also exists a valid transcript

$$\text{tr}^* = ((\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \text{com}_1, \text{com}_2), (\mathfrak{R}_0, \mathfrak{R}_1), \vec{z}_3, (\gamma_{i,j}), \mathbf{h}, (\mu_i), (t, f_0), c, (\mathbf{z}_1^*, \mathbf{z}_2^*, \text{op}_1^*, \text{op}_2^*))$$

which has the same partial transcript as tr up to the last response, then we must have $\mathbf{z}_1 = \mathbf{z}_1^*$ and $\mathbf{z}_2 = \mathbf{z}_2^*$ assuming that the underlying encryption scheme (LHC) has no decryption error. Now, for a fixed crs , statement x and a valid transcript $\text{tr} := (\text{pref_tr}, (\mathbf{z}_1, \mathbf{z}_2, \text{op}_1, \text{op}_2))$, where pref_tr is the partial transcript which excludes the last message, define the following set:

$$\Gamma(x, \text{tr}) := \left\{ c^* \in \mathcal{C} : \exists (\mathbf{z}_1^*, \mathbf{z}_2^*, \text{op}_1^*, \text{op}_2^*) \text{ s.t. } \begin{array}{l} (\text{pref_tr}, (\mathbf{z}_1^*, \mathbf{z}_2^*, \text{op}_1^*, \text{op}_2^*)) \\ \text{is a valid transcript} \end{array} \right\}.$$

Multi-Extract($Q_H, Q_s, 1/\varepsilon, \tilde{c}rs, \tau, x, \pi$)

Input: $Q_H, Q_s, 1/\varepsilon, \tilde{c}rs := (\mathbf{A}_1, \mathbf{A}_2, \mathbf{B}_y, \mathbf{B}_g, \mathbf{b}, \bar{\mathbf{A}}_1, \bar{\mathbf{A}}_2, \bar{\mathbf{B}}_1, \bar{\mathbf{B}}_2)$,
 $td := (\bar{\mathbf{D}}_{1,1}, \bar{\mathbf{D}}_{1,2}, \bar{\mathbf{D}}_{2,1}, \bar{\mathbf{D}}_{2,2}), x := (P, \tilde{u}), \pi$

Output: witness w , or a triple $(\perp, (x, \pi, \text{aux}))$, or (\perp, \perp) (unsuccessful)

- 1: parse $\pi = (\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w}, \text{com}_1, \text{com}_2, \tilde{z}_3, \mathbf{h}, t, f_0, \tilde{z}_3, \mathbf{h}, \mathbf{z}_1, \mathbf{z}_2, \text{op}_1, \text{op}_2)$
- 2: $\mathbf{B} = \begin{bmatrix} \mathbf{B}_y \\ \mathbf{B}_g \end{bmatrix}, \mathbf{t}_B = \begin{bmatrix} \mathbf{t}_y \\ \mathbf{t}_g \end{bmatrix}$
- 3: compute $\mathbf{D}_2, \mathbf{d}_1, d_0$ as in (20)
- 4: **if** one of the conditions below does not hold, then **return** 0:
- 5: $\text{com}_1 = (\bar{\mathbf{t}}_{1,1}, \bar{\mathbf{t}}_{1,2}, \bar{\mathbf{w}}_{1,1}, \bar{\mathbf{w}}_{1,2})$
- 6: $\text{com}_2 = (\bar{\mathbf{t}}_{2,1}, \bar{\mathbf{t}}_{2,2}, \bar{\mathbf{w}}_{2,1}, \bar{\mathbf{w}}_{2,2})$
- 7: $\text{BadChall} = \{c\}$
- 8: $t = 0$
- 9: **while** $t \leq T_{\max} := \frac{\lambda \cdot 2Q_H}{\varepsilon}$ **do**
- 10: $b = 1$
- 11: $c' \leftarrow \mathcal{C} \setminus \text{BadChall}$
- 12: **for** $i \in \{1, 2\}$ **do** $\mathbf{z}'_i = (c' \bar{\mathbf{t}}_{i,2} + \bar{\mathbf{w}}_{i,2}) - \bar{\mathbf{D}}_{i,1}(c' \bar{\mathbf{t}}_{i,1} + \bar{\mathbf{w}}_{i,1}) \bmod \bar{p}$
- 13: $\mathbf{z}' = \begin{bmatrix} \mathbf{z}'_1 \\ \sigma(\mathbf{z}'_1) \\ c' \mathbf{t}_B - \mathbf{B} \mathbf{z}'_2 \\ \sigma(c' \mathbf{t}_B - \mathbf{B} \mathbf{z}'_2) \end{bmatrix}$
- 14: $\bar{\mathbf{f}}' := c' t - \mathbf{b}^T \mathbf{z}'_2$
- 15: $f'_0 = \mathbf{z}'^T \mathbf{D}_2 \mathbf{z}' + c' \mathbf{d}_1^T \mathbf{z}' + c'^2 d_0 - \bar{\mathbf{f}}'$
- 16: **if** one of the following conditions does not hold, **then** $b = 0$:
 - $\|\mathbf{z}'_1\| > \mathcal{B}_1 := 2s_1 \sqrt{2m_1 \hat{d}}$
 - $\|\mathbf{z}'_2\| > \mathcal{B}_2 := 2s_2 \sqrt{2m_2 \hat{d}}$
 - $\mathbf{A}_1 \mathbf{z}'_1 + \mathbf{A}_2 \mathbf{z}'_2 = \mathbf{w} + c' \mathbf{t}_A$
 - $f'_0 \neq f_0$
- 17: **if** $b = 1$ **then**
- 18: $c^* = c' - c$
- 19: **for** $i \in \{1, 2\}$ **do**
- 20: $\mathbf{s}_i^* = \frac{\mathbf{z}'_i - \mathbf{z}_i}{c' - c}$
- 21: $\mathbf{w} = \bar{s}_1^* \in \mathbb{Z}_{\hat{q}}^{m\hat{d}}$
- 22: **if** $(x, w) \in R^{\text{Com}}$ **then return** w
- 23: **else**
- 24: $\text{aux} = (\mathbf{s}_1^*, \mathbf{s}_2^*, c^*)$
- 25: **return** $(\perp, (x, \pi, \text{aux}))$
- 26: **else** $\text{BadChall} \leftarrow \text{BadChall} \cup \{c'\}$
- 27: $t = t + 1$
- 28: **return** (\perp, \perp)

Fig. 15: Multi-Extract algorithm.

At a high level, Multi-Extract is first given such an accepting transcript tr (technically it can manually produce one given π). It also defines a set of bad challenges as $\text{BadChall} := \{c\}$. Then, it samples $c' \leftarrow \mathcal{C} \setminus \text{BadChall}$, decrypts $(\bar{\mathbf{w}}_{i,1} + c' \bar{\mathbf{t}}_{i,1}, \bar{\mathbf{w}}_{i,2} + c' \bar{\mathbf{t}}_{i,2})$ to obtain some \mathbf{z}'_i , for $i = 1, 2$, and checks whether

$$\begin{aligned} \|\mathbf{z}'_1\| &\stackrel{?}{\leq} \mathcal{B}_1 := \mathfrak{s}_1 \sqrt{2m_1 \hat{d}}, & \|\mathbf{z}'_2\| &\stackrel{?}{\leq} \mathcal{B}_2 := \mathfrak{s}_2 \sqrt{2m_2 \hat{d}}, \\ \mathbf{A}_1 \mathbf{z}'_1 + \mathbf{A}_2 \mathbf{z}'_2 &\stackrel{?}{=} \mathbf{w} + c' \mathbf{t}_A, & \mathbf{z}'^T \mathbf{D}_2 \mathbf{z}' + c' \mathbf{d}_1^T \mathbf{z}' + c'^2 d_0 - (c' t - \mathbf{b}^T \mathbf{z}'_2) &\stackrel{?}{=} f_0, \end{aligned}$$

where

$$\mathbf{z}' := \begin{bmatrix} \mathbf{z}'_1 \\ \sigma(\mathbf{z}'_1) \\ c' \mathbf{t}_B - \mathbf{B} \mathbf{z}'_2 \\ \sigma(c' \mathbf{t}_B - \mathbf{B} \mathbf{z}'_2) \end{bmatrix}.$$

If not then the algorithm adds c' to BadChall and restarts the procedure.

Let $\text{GoodChall}(x, \text{tr})$ be the set of all challenges $c' \in \mathcal{C}$ (including c) for which the procedure above succeeds. The first step in analysing the success probability of Multi-Extract will be to lower-bound the size of $\text{GoodChall}(x, \text{tr})$. To begin with, one observes that $\Gamma(x, \text{tr}) \subseteq \text{GoodChall}(x, \text{tr})$. Hence, we focus on showing a lower bound on $\Gamma(x, \text{tr})$ instead. To this end, we present the following lemma. Since it is almost identical to [67, Lemma 4.7], we omit the proof.

Lemma 6.4 ([67]). *Let \mathcal{A} be a PPT adversary that makes at most Q_H random oracle queries and*

$$\Pr \left[\begin{array}{l} (\text{cfs}, \text{td}) \leftarrow \mathcal{S}_{\text{crs}}, \\ \{(x_k, \pi_k)\}_{k \in [Q_s]} \leftarrow \mathcal{A}^H(\text{cfs}) : \forall k \in [Q_s], \text{Verify}_{\text{Com}}^H(\text{cfs}, x_k, \pi_k) = 1 \end{array} \right] \geq \varepsilon(\lambda).$$

Then, we have

$$\Pr \left[\begin{array}{l} (\text{cfs}, \text{td}) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda), \\ \{(x_k, \pi_k)\}_{k \in [Q_s]} \leftarrow \mathcal{A}^H(1^\lambda, \text{cfs}) : \forall k \in [Q_s], \text{Verify}_{\text{Com}}^H(\text{cfs}, x_k, \pi_k) = 1 \\ \wedge |\Gamma(x_k, \text{tr}_k)| \geq \frac{\varepsilon}{2Q_H} \cdot |\mathcal{C}| \end{array} \right] \geq \frac{\varepsilon(\lambda)}{2} - \text{negl}(\lambda)$$

where the transcript tr_k is constructed naturally from π_k using $\text{Verify}_{\text{Com}}^H$.

We use the lemma above to give an upper-bound on the following probability:

$$\varepsilon_{\text{fail}} := \Pr \left[\begin{array}{l} (\text{cfs}, \text{td}) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda), \\ \{(x_k, \pi_k)\}_{k \in [Q_s]} \leftarrow \mathcal{A}^H(1^\lambda, \text{cfs}) : \exists i \in [Q_s], (\perp, \perp) \leftarrow \text{Multi-Extract}(\text{input}_i) \end{array} \right]$$

where $\text{input}_i := (Q_H, Q_s, 1/\varepsilon, \text{cfs}, \tau, x_i, \pi_i)$. Consider a fixed $i \in [Q_s]$. By Lemma 6.4, with probability at least $\varepsilon/2 - \text{negl}(\lambda)$, for each $k \in [Q_s]$ we have $|\Gamma(x_k, \text{tr}_k)| \geq \frac{\varepsilon}{2Q_H} \cdot |\mathcal{C}|$. Assuming this event occurs, then in each trial of the loop in Line 9 of Multi-Extract, the probability that we pick $c' \in \text{GoodChall}(x_i, \text{tr}_i)$ is at least $\varepsilon/(2Q_H)$. Hence, by repeating the procedure $T_{\text{max}} = \lambda \cdot (2Q_H/\varepsilon)$ times, the algorithm fails on each trial with negligible probability. Hence, by the union bound over i and the fact that $Q_s = \text{poly}(\lambda)$, we deduce that

$$\varepsilon_{\text{fail}} \leq \frac{\varepsilon}{2} + \text{negl}(\lambda).$$

Next, we focus on upper-bounding

$$\varepsilon_{\text{bind}} := \Pr \left[\begin{array}{l} (\text{cfs}, \text{td}) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda), \\ \{(x_k, \pi_k)\}_{k \in [Q_s]} \leftarrow \mathcal{A}^H(1^\lambda, \text{cfs}) : \exists i \in [Q_s], (\perp, (x, \pi, \text{aux})) \leftarrow \text{Multi-Extract}(\text{input}_i) \end{array} \right]$$

where $\text{input}_i := (Q_H, Q_s, 1/\varepsilon, \text{cfs}, \tau, x_i, \pi_i)$. To this end, we show Lemma 6.5 and conclude that $\varepsilon_{\text{bind}} = \text{negl}(\lambda)$. Finally, by simple calculation:

$$\begin{aligned} \Pr \left[\begin{array}{l} (\text{cfs}, \text{td}) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda), \\ \{(x_k, \pi_k)\}_{k \in [Q_s]} \leftarrow \mathcal{A}^H(1^\lambda, \text{cfs}) : \forall k \in [Q_s], \text{Verify}_{\text{ISIS}}^H(\text{cfs}, x_k, \pi_k) = 1, \\ \mathbf{w}_k \leftarrow \text{Multi-Extract}(\text{input}_k) \\ (x_k, \mathbf{w}_k) \in R^{\text{Com}} \end{array} \right] &= \varepsilon - \varepsilon_{\text{fail}} - \varepsilon_{\text{sp}} \\ &\leq \frac{\varepsilon}{2} - \text{negl}(\lambda) \end{aligned}$$

which concludes the proof. \square

Lemma 6.5. *Let \mathcal{A} be a PPT adversary that makes at most $Q_H = \text{poly}(\lambda)$ random oracle queries and*

$$\Pr \left[\begin{array}{l} (\text{c}\bar{\text{r}}\text{s}, \text{td}) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda), \\ \{(x_k, \pi_k)\}_{k \in [Q_s]} \leftarrow \mathcal{A}^H(\text{c}\bar{\text{r}}\text{s}) : \forall k \in [Q_s], \text{Verify}_{\text{ISIS}}^H(\text{c}\bar{\text{r}}\text{s}, x_k, \pi_k) = 1 \end{array} \right] \geq \varepsilon(\lambda),$$

where $\varepsilon(\lambda)$ is non-negligible. Then, we have

$$\Pr \left[\begin{array}{l} (\text{c}\bar{\text{r}}\text{s}, \text{td}) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda), \\ \{(x_k, \pi_k)\}_{k \in [Q_s]} \leftarrow \mathcal{A}^H(1^\lambda, \text{c}\bar{\text{r}}\text{s}) : \exists i, (\perp, (x, \pi, \text{aux})) \leftarrow \text{Multi-Extract}(\text{input}_i) \end{array} \right] = \text{negl}(\lambda)$$

where $\text{input}_i := (Q_H, Q_s, 1/\varepsilon, \text{c}\bar{\text{r}}\text{s}, \tau, x_i, \pi_i)$.

Proof. Suppose that the probability above is non-negligible and denote it as $\delta(\lambda)$. We consider a PPT algorithm \mathcal{P}^* , which given $(\text{c}\bar{\text{r}}\text{s}, \text{td})$, it essentially executes \mathcal{A} to obtain Q_s statement-proof pairs $(x_k, \pi_k)_{k \in [Q_s]}$, checks whether the proofs are valid, and then runs each $\text{Multi-Extract}(\text{input}_k)$ to check whether it outputs $(\perp, (x, \pi, \text{aux}))$. If not, then \mathcal{P}^* aborts. Otherwise, \mathcal{P}^* outputs the triple (x, π, aux) . Let us also recall the meaning of the output of \mathcal{P}^* . Here, $x = (P, \bar{u})$ is a statement, π is the corresponding proof from Section 6 written as:

$$\pi = (\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w}, \text{com}_1, \text{com}_2, \vec{z}_3, \mathbf{h}, t, f_0, \vec{z}_3, \mathbf{h}, \mathbf{z}_1, \mathbf{z}_2, \text{op}_1, \text{op}_2) \quad (32)$$

and $\text{aux} = (\mathbf{s}_1^*, \mathbf{s}_2^*, c^*)$ is the relaxed opening of \mathbf{t}_A (c.f. Definition 5.1) such that $(x, \bar{s}_1^*) \notin R^{\text{Com}}$. Clearly, the probability that \mathcal{P}^* outputs something is δ . Further, \mathcal{P}^* runs in time $Q_s \cdot Q_H \cdot \text{poly}(\lambda)/\varepsilon$ and makes at most Q_H random oracle queries.

Now, we can use the knowledge extractor \mathcal{E} from [57, Appendix B]. Roughly speaking, it runs \mathcal{P}^* once to obtain a valid triple (x, π, aux) by simulating the random oracle queries. Denote π as in (32). Then, for fixed random coins, fixed common reference string $\text{c}\bar{\text{r}}\text{s}$ and the fixed statement x (where the latter two are included in the hash input of the Fiat-Shamir transformation in π), it re-runs \mathcal{P}^* with different random oracle outputs in order to obtain additional valid triples of the form (x, π', aux') where

$$\pi' = (\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w}, \text{com}_1, \text{com}_2, \vec{z}'_3, \mathbf{h}', t', f'_0, \vec{z}'_3, \mathbf{h}', \mathbf{z}'_1, \mathbf{z}'_2, \text{op}'_1, \text{op}'_2). \quad (33)$$

Clearly, for $\text{aux} := (\mathbf{s}_1^*, \mathbf{s}_2^*, c^*)$ and $\text{aux}' = (\mathbf{s}'_1, \mathbf{s}'_2, c')$ we might have $\mathbf{s}_1^* \neq \mathbf{s}'_1$. However, in this case \mathcal{E} would find two distinct relaxed openings to \mathbf{t}_A and by Lemma 5.2 it would compute a valid solution to the Module-SIS problem. Hence, from now on we assume that every valid triple (x, π_i, aux_i) extracted from \mathcal{P}^* by \mathcal{E} satisfies $\text{aux}_i := (\mathbf{s}_1^*, \mathbf{s}_2^*, c_i)$

From [57, Theorem B.7], we deduce that given access to \mathcal{P}^* , \mathcal{E} runs in expected $Q_s \cdot Q_H^2 \cdot \text{poly}(\lambda)/\varepsilon$ time and with probability at least

$$\delta' := \delta - (Q_H + 1) \left(\frac{2}{|\mathcal{C}|} + \hat{p}_1^{-\hat{d}/2} + \hat{p}_1^{-\tau} + 2^{-\lambda} \right) \quad (34)$$

outputs one of the following:

1. a triple $(x, \text{open}, \text{aux})$, where

$$x = (P, \bar{u}, \mathcal{B}), \quad \text{open} := (\mathbf{s}_1, \mathbf{s}_2, c) \quad \text{and} \quad \text{aux} := (\mathbf{s}_1^*, \mathbf{s}_2^*, c^*)$$

such that $(x, \bar{s}_1) \in R^{\text{Com}}$, $(x, \bar{s}_1^*) \notin R^{\text{Com}}$ and both open and aux are relaxed openings of the same ABDLOP commitment \mathbf{t}_A ,

2. a non-zero vector $\bar{\mathbf{s}} \in \hat{\mathcal{R}}_q^{m_1+m_2}$ such that $\|\bar{\mathbf{s}}\| \leq B$ and

$$[\mathbf{A}_1 \ \mathbf{A}_2] \bar{\mathbf{s}} = \mathbf{0}.$$

$\text{Exp}_{\text{pp}}^{\text{NTRU-ISIS}_f}(\mathcal{A})$
1: $(a_1, \mathbf{B}) \leftarrow \text{NTRU.TrapGen}(q, d)$
2: $\mathbf{a}_2 \leftarrow \mathcal{R}_q^m$
3: $\mathbf{A} = [a_1 \ \mathbf{a}_2^T \ 1]$
4: for $i \in [k]$:
5: $x_i \leftarrow [N]$
6: $\mathbf{s}_i \leftarrow \text{GSampler}(a_1, \mathbf{a}_2, \mathbf{B}, \mathbf{s}, f(x_i))$
7: $(x^*, \mathbf{s}^*) \leftarrow \mathcal{A}(\mathbf{A}, (x_i, \mathbf{s}_i)_{i \in [k]})$
8: if $(\mathbf{A}\mathbf{s}^* = f(x^*)) \wedge (\ \mathbf{s}^*\ \leq \mathcal{B}) \wedge ((x^*, \mathbf{s}^*) \notin \{(x_i, \mathbf{s}_i) : i \in [k]\})$
9: then return 1
10: else return 0

Fig. 16: The NTRU-ISIS_f experiment.

Clearly, in the second case \mathcal{E} directly gets a valid Module-SIS solution so we only focus on the former one. Note that $(x, \bar{s}_1^*) \notin R^{\text{Com}}$ thus we must have $\bar{s}_1^* \neq \bar{s}_1$. Consequently, \mathcal{E} has found two distinct relaxed openings to the same ABDLOP commitment and by Lemma 5.2, it can compute a valid Module-SIS solution. Hence, in either case \mathcal{E} solves $\text{MSIS}_{n, m_1 + m_2, B}$.

To finish the proof, we apply the expected-time to strict-time transformation (c.f. Lemma 2.14) on \mathcal{E} . This gives us an algorithm \mathcal{B} which solves $\text{MSIS}_{n, m_1 + m_2, B}$ with probability at least $\delta'/2$ and runs in strict time

$$Q_s \cdot Q_H^2 \cdot \frac{\text{poly}(\lambda)}{\varepsilon \cdot \delta'}.$$

Therefore, if δ is non-negligible then \mathcal{B} runs in strict polynomial time and solves $\text{MSIS}_{n, m_1 + m_2, B}$ with non-negligible probability. This leads to a contradiction. \square

7 The NTRU-ISIS_f Assumption

Note that due to the preimage sampling procedure, the ISIS_f assumption is not falsifiable. In this section, we introduce the falsifiable version of the problem, called NTRU-ISIS_f, where the challenger additionally has the NTRU trapdoor to produce short preimages.

Definition 7.1 (The NTRU-ISIS_f Problem). Let $\text{pp} := (q, d, m, N, k, \mathbf{s}, \mathcal{B})$ be a tuple of functions of the security parameter λ . Consider any efficiently computable function $f : [N] \rightarrow \mathcal{R}_q$. The NTRU-ISIS_f assumption is defined by the experiment in Figure 16. For an adversary \mathcal{A} , we define

$$\text{Adv}_{\text{pp}}^{\text{NTRU-ISIS}_f}(\mathcal{A}) = \Pr[\text{Exp}_{\text{pp}}^{\text{NTRU-ISIS}_f}(\mathcal{A}) \rightarrow 1].$$

The NTRU-ISIS_f^{pp} assumption states that for every PPT adversary \mathcal{A} , the advantage $\text{Adv}_{\text{pp}}^{\text{NTRU-ISIS}_f}(\mathcal{A})$ is negligible.

Similarly as in Section 3, we consider the interactive version of the NTRU-ISIS_f problem.

Definition 7.2 (The Interactive NTRU-ISIS_f Problem). Define public parameters $\text{pp} := (q, d, m, \ell_m, \ell_r, N, \mathbf{s}, \mathcal{B}_s, \mathcal{B}_m)$ as a tuple of functions of the security parameter λ . Consider any efficiently computable function $f : [N] \rightarrow \mathcal{R}_q$. The Int-NTRU-ISIS_f assumption is defined by the experiment in Figure 17. For an adversary \mathcal{A} , we define

$$\text{Adv}_{\text{pp}}^{\text{Int-NTRU-ISIS}_f}(\mathcal{A}) = \Pr[\text{Exp}_{\text{pp}}^{\text{Int-NTRU-ISIS}_f}(\mathcal{A}) \rightarrow 1].$$

The Int-NTRU-ISIS_f^{pp} assumption states that for every PPT adversary \mathcal{A} , the advantage $\text{Adv}_{\text{pp}}^{\text{Int-NTRU-ISIS}_f}(\mathcal{A})$ is negligible.

$\text{Exp}_{\text{pp}}^{\text{Int-NTRU-ISIS}_f}(\mathcal{A})$	$\mathcal{O}_{\text{pre}}(\mathbf{m}, \mathbf{r})$
1: $(a_1, \mathbf{B}) \leftarrow \text{NTRU.TrapGen}(q, d)$	1: if $\ (\mathbf{m}, \mathbf{r})\ < \mathcal{B}_m$
2: $\mathbf{a}_2 \leftarrow \mathcal{R}_q^m$	2: then return \perp
3: $\mathbf{A} = [a_1 \mathbf{a}_2^T \ 1]$	3: $x \leftarrow [N]$
4: $\mathbf{C} \leftarrow \mathcal{R}_q^{1 \times (\ell_m + \ell_r)}$	4: $\mathbf{s} \leftarrow \text{GSampler}\left(a_1, \mathbf{a}_2, \mathbf{B}, \mathfrak{s}, f(x) + \mathbf{C} \begin{bmatrix} \mathbf{m} \\ \mathbf{r} \end{bmatrix}\right)$
5: $\mathcal{M} = \emptyset$	5: $\mathcal{M} \leftarrow \mathcal{M} \cup \{\mathbf{m}\}$
6: $(x^*, \mathfrak{s}^*, \mathbf{m}^*, \mathbf{r}^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{pre}}}(\mathbf{A}, \mathbf{C})$	6: return (x, \mathfrak{s})
7: if $(\mathbf{m}^* \notin \mathcal{M}) \wedge \left(\mathbf{A}\mathfrak{s}^* = f(x^*) + \mathbf{C} \begin{bmatrix} \mathbf{m}^* \\ \mathbf{r}^* \end{bmatrix}\right) \wedge (\ \mathfrak{s}^*\ \leq \mathcal{B}_s) \wedge$ $(\ (\mathbf{m}^*, \mathbf{r}^*)\ \leq \mathcal{B}_m)$	
8: then return 1	
9: else return 0	

Fig. 17: Interactive version of the NTRU-ISIS_f problem.

Using the proof strategy from Theorem 3.3, one can efficiently reduce Int-NTRU-ISIS_f to NTRU-ISIS_f. In the following, we only provide a proof sketch and highlight the parts where it differs from the proof of Theorem 3.3.

Theorem 7.3 (Int-NTRU-ISIS_f \implies NTRU-ISIS_f). *Define the public parameters $\text{pp} := (q, d, m, \ell_m, \ell_r, N, \mathfrak{s}, \mathcal{B}_s, \mathcal{B}_m)$ such that $\mathcal{B}_m \geq 1, (2\kappa + 1)^{m+2} > q$, and*

$$M := \exp\left(1 + \frac{1}{2\alpha^2}\right) \quad \text{and} \quad \varepsilon := 2\left(\frac{1+\epsilon}{1-\epsilon}\right) \exp\left(-2\alpha^2 \cdot \frac{\pi-1}{\pi}\right)$$

where $\epsilon = 2^{-\delta}/(4d), \alpha \geq 1$ and $\kappa \in \mathbb{Z}_+$. Suppose

$$\mathfrak{s} \geq \max\left(1.17\sqrt{q} \cdot \eta'_\epsilon(\mathbb{Z}), \alpha\mathcal{B}_m\kappa d\sqrt{(\ell_m + \ell_r)(m+2)}\right).$$

Then, for every adversary \mathcal{A} which makes at most Q queries to \mathcal{O}_{pre} , there are adversaries \mathcal{A}' and \mathcal{B} , which run in time essentially identical to \mathcal{A} and

$$\begin{aligned} \text{Adv}_{\text{pp}'}^{\text{NTRU-ISIS}_f}(\mathcal{A}') &\geq \frac{1}{6Q} \text{Adv}_{\text{pp}}^{\text{Int-NTRU-ISIS}_f}(\mathcal{A}) - \frac{\ell_m + \ell_r}{6Q} \text{Adv}_{1, m+1, \kappa}^{\text{MLWE}}(\mathcal{B}) - \frac{2^{-\lambda}}{6} - \frac{T_{\text{max}}^2 Q}{12N} \\ &\quad - \frac{\ell_m}{3} \cdot \left(\frac{q}{(2\kappa + 1)^{m+2}}\right)^d - \left(Q - \frac{2}{3}\right) T_{\text{max}} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} + 2^{-\delta+1}\right) \end{aligned}$$

where $\text{pp}' := (q, d, m, N, T_{\text{max}}Q, \mathfrak{s}, \mathcal{B} = \mathcal{B}_s + \mathcal{B}_m\kappa d\sqrt{(\ell_m + \ell_r)(m+2)})$ and T_{max} satisfies $(1 - \frac{1}{M})^{T_{\text{max}}} \leq 2^{-\lambda}$.

Proof. We closely follow the security games from the proof of Theorem 3.3 and explain the main differences. Namely, Game₁ is the interactive NTRU-ISIS_f. In Game₂, we substitute the uniformly random \mathbf{C} with $\mathbf{C} = \mathbf{A}\mathbf{D}$ where $\mathbf{D} \leftarrow \mathcal{S}_\kappa^{(m+2) \times (\ell_m + \ell_r)}$. Observe that now we do not use the leftover hash lemma anymore. Indeed, to see that the latter is pseudorandom under the $\text{MLWE}_{1, m+1, \kappa}$ assumption, denote

$$\mathbf{D} := \begin{bmatrix} d_{1,1} & \cdots & d_{\ell_m + \ell_r, 1} \\ \mathbf{d}_{1,2} & \cdots & \mathbf{d}_{\ell_m + \ell_r, 2} \\ d_{1,3} & \cdots & d_{\ell_m + \ell_r, 3} \end{bmatrix}.$$

Then,

$$\mathbf{C} = [a_1 \ \mathbf{a}_2^T \ 1] \mathbf{D} = [a_1 d_{1,1} \ \cdots \ a_1 d_{\ell_m + \ell_r, 1}] + \mathbf{u}$$

where $\mathbf{u} := \mathbf{a}_2^T [\mathbf{d}_{2,1} \ \cdots \ \mathbf{d}_{2, \ell_m + \ell_r}] + [d_{1,3} \ \cdots \ d_{\ell_m + \ell_r, 3}]$. Hence, under the Module-LWE assumption, \mathbf{u} is pseudorandom, and so is \mathbf{C} .

One can then define security games from 3 to 7 as before. For Game_6 , we note that now there exists a polynomial vector $\mathbf{u} \in \Lambda_q^\perp(\mathbf{A})$ with coefficients between -2κ and 2κ , since $q^d < (2\kappa + 1)^{(m+2)d}$. Hence, $\|\mathbf{u}\| \leq 2\kappa\sqrt{(m+2)d}$ and

$$\mathfrak{s} \geq \alpha\mathcal{B}_m\kappa d\sqrt{(\ell_m + \ell_r)(m+2)} \geq \kappa d\sqrt{(m+2)} \geq 2\kappa\sqrt{(m+2)d} \geq \|\mathbf{u}\|.$$

Thus, we can argue similarly as before that $\Pr[\mathbf{E}_6] \geq \frac{1}{2} \Pr[\mathbf{E}_5]$.

Clearly, Game_7 stays unchanged. As for Game_8 , thanks to the structure of NTRU lattices, we do not need to exclude the case when \mathfrak{s} is smaller than the smoothing parameter of $\Lambda_q^\perp(\mathbf{A})$. Now, in order to use Lemma 2.5, we need the fact that applying GSampler is statistically close to sampling preimages from the discrete Gaussian distribution, i.e. Lemma 2.7. Hence,

$$\Pr[\mathbf{E}_8] \geq \Pr[\mathbf{E}_7] - T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} + 2^{-\delta+1} \right).$$

Next, to analyse Game_9 , we need an upper-bound on the probability that $|S(\mathbf{D}, i)| = 1$ for fixed $i \in [\ell_m]$. Using the same techniques as in [53, Lemma 5.2], we can show that this happens with probability at most $\left(\frac{q}{(2\kappa+1)^{m+2}}\right)^d$. Hence,

$$\Pr[\mathbf{E}_9] \geq \frac{1}{3} \Pr[\mathbf{E}_8] - \frac{\ell_m}{3} \cdot \left(\frac{q}{(2\kappa+1)^{m+2}}\right)^d.$$

We can define Game_{10} as in the previous proof and deduce that

$$\Pr[\mathbf{E}_{10}] \geq \Pr[\mathbf{E}_9] - (Q-1)T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\epsilon}{M} + 2^{-\delta+1} \right).$$

Finally, the rest of the proof from Theorem 3.3 remains unchanged. \square

8 Anonymous Credentials from NTRU-ISIS_f

This section focuses on building an efficient anonymous credentials scheme based on the newly defined assumption NTRU-ISIS_f where f is a binary encoding function (see Section 3.1.2). Namely, let $N = 2^t$. Then, the function $f : [N] \rightarrow \mathcal{R}_q$ is defined as

$$f(x) := \text{Coeffs}^{-1}(B \cdot \text{enc}(x)) \in \mathcal{R}_q$$

where $\text{enc}(x)$ is a binary decomposition of an integer $0 \leq x-1 < 2^t$ and $B \in \mathbb{Z}_q^{d \times t}$ is a fixed matrix.

For convenience, we introduce the following notation which will be heavily used throughout this section. For a matrix $C \in \mathbb{Z}_q^{n \times lm}$ denoted as

$$C := [C_1 \ C_2 \ \dots \ C_l] \quad \text{where each } C_i \in \mathbb{Z}_q^{n \times m},$$

and a (possibly empty) subset $S = (s_1, \dots, s_k) \subseteq [l]$, where $s_1 < \dots < s_k$, we define:

$$C_S := [C_{s_1} \ C_{s_2} \ \dots \ C_{s_k}] \in \mathbb{Z}_q^{n \times km}.$$

We introduce an identical notation to vectors and matrices over \mathbb{Z}_q and \mathcal{R}_q . Further, we define $\bar{S} := [l] \setminus S$ to be the complement of S with respect to $[l]$.

Parameter	Explanation
d	degree of the cyclotomic ring \mathcal{R}
\hat{d}	degree of the cyclotomic ring $\hat{\mathcal{R}}$ used in $\Pi_{\text{NIZK}}^{\text{ISIS}}$ and $\Pi_{\text{NIZK}}^{\text{Com}}$
q	anonymous credentials modulus
\hat{q}_1	modulus for the proof system $\Pi_{\text{NIZK}}^{\text{Com}}$ divisible by q
\hat{q}_2	modulus for the proof system $\Pi_{\text{NIZK}}^{\text{ISIS}}$ divisible by q
N	domain size of the function f
t	$\log N$, length of the binary vector \vec{u}
h	length of the string output by H_M
l	number of attributes, we assume $d lh$
m	length of \mathbf{a}_2
ℓ_r	length of the randomness \mathbf{r} generated by the holder
ℓ_m	lh/d
ψ	infinity norm of \mathbf{r} and the output of H_M
\mathfrak{s}	preimages are sampled from $D_s^{(m+2)d}$

Table 2: Overview of parameters and notation.

8.1 Construction

We start with the summary of notation and parameters in Table 2, and define the building blocks, which were developed throughout the paper and will be used in our main construction:

- Function $f : [N] \rightarrow \mathcal{R}_q$ described in more detail in Section 3.1.2,
- The NIZK proof system $\Pi_{\text{NIZK}}^{\text{ISIS}} := (\text{Prove}_s^{\text{H}_{\text{ISIS}}}, \text{Verify}_s^{\text{H}_{\text{ISIS}}})$ (with a common random string crs_{ISIS}) described in Section 5 for the relation (12), where the proof system modulus is \hat{q}_1 ,
- The NIZK proof system $\Pi_{\text{NIZK}}^{\text{Com}} := (\text{Prove}_{\text{Com}}^{\text{H}_{\text{Com}}}, \text{Verify}_{\text{Com}}^{\text{H}_{\text{Com}}})$ (with a common random string crs_{Com}) described in Section 6 for the relation (23), where the proof system modulus is \hat{q}_2 ,
- Four hash functions $H_{\text{crs}}, H_{\text{Com}}, H_{\text{ISIS}}, H_M$ modelled as random oracles in the security proof. The first one is a special hash function which we will only query on input 0 to obtain the crs for the underlying primitives, i.e.

$$H_{\text{crs}}(0) = (\text{crs}_{\text{ISIS}}, \text{crs}_{\text{Com}})$$

where crs_{ISIS} and crs_{Com} are used by $\Pi_{\text{NIZK}}^{\text{ISIS}}$ and $\Pi_{\text{NIZK}}^{\text{Com}}$ respectively. Next, $H_{\text{Com}}, H_{\text{ISIS}}$ are used by the NIZK proof systems $\Pi_{\text{NIZK}}^{\text{ISIS}}$ and $\Pi_{\text{NIZK}}^{\text{Com}}$. Finally, we apply H_M to hash attributes into a vector of length h with coefficients between $-\psi$ and ψ . As usual, one can use a single hash function instead of four by applying the domain separation technique.

- NTRU trapdoor sampling functions ($\text{NTRU.TrapGen}, \text{GSampler}$) defined in Section 2.1.

We present the $\text{AnonCreds} = (\text{AnonCreds.Init}, \text{AnonCreds.Issue}, \text{AnonCreds.Verify})$ construction in Figure 18, Figure 19 and Figure 20. We assume that all the involved parties retrieve $(\text{crs}_{\text{ISIS}}, \text{crs}_{\text{Com}})$ by querying $H_{\text{crs}}(0)$.

Key generation. The key generation algorithm runs $(a_1, \mathbf{B}) \leftarrow \text{NTRU.TrapGen}(q, d)$ and samples $\mathbf{a}_2 \leftarrow \mathcal{R}_q^m$. It further generates vectors $\mathbf{c}_0 \leftarrow \mathcal{R}_q^{\ell_m}$ and $\mathbf{c}_1 \leftarrow \mathcal{R}_q^{\ell_r}$. Then, we define the public key $\text{ipk} = (a_1, \mathbf{a}_2, \mathbf{c}_0, \mathbf{c}_1)$ and secret key $\text{isk} = \mathbf{B}$.

Credential issuing. Next, we give an intuition on the issuing protocol. The holder \mathcal{H} is given $\text{ipk}, \text{attrs} = (\mathbf{a}_1, \dots, \mathbf{a}_l)$, and the set of indices $\text{id}x \subseteq [l]$ for which the attributes in attrs are public. Then, it computes the polynomial vector

$$\mathbf{m} = \text{Coeffs}^{-1}(H_M(\mathbf{a}_1), \dots, H_M(\mathbf{a}_l)) \in \mathcal{R}_q^{\ell_m}. \quad (35)$$

<p>AnonCreds.Init Input: number of attributes l Output: (ipk, isk) 1: $(a_1, \mathbf{B}) \leftarrow \text{NTRU.TrapGen}(q, d)$ 2: $\mathbf{a}_2 \leftarrow \mathcal{R}_q^m$ 3: $(\mathbf{c}_0, \mathbf{c}_1) \leftarrow \mathcal{R}_q^{\ell_m} \times \mathcal{R}_q^{\ell_r}$ 4: $\text{ipk} \leftarrow (a_1, \mathbf{a}_2, \mathbf{c}_0, \mathbf{c}_1)$ 5: $\text{isk} \leftarrow \mathbf{B}$ 6: return (ipk, isk)</p>

Fig. 18: AnonCreds.Init algorithm.

Further, to commit to \mathbf{m} , \mathcal{H} samples randomness $\mathbf{r} \leftarrow S_\psi^{\ell_r}$ and computes the commitment to \mathbf{m} as

$$u := \mathbf{c}_0^T \mathbf{m} + \mathbf{c}_1^T \mathbf{r} \in \mathcal{R}_q.$$

The next step is to prove knowledge of a commitment opening of u . Since we want to use the proof system $\Pi_{\text{NIZK}}^{\text{ISIS}}$, we need to translate this relation from \mathcal{R}_q to \mathbb{Z}_q . Since $\text{Coeffs}(\mathbf{m})_{\text{idx}}$ (resp. $\text{Coeffs}(\mathbf{m})_{\overline{\text{idx}}}$) are public (resp. hidden), we transform into a relation of the form (23) by defining:

$$\begin{aligned} P &:= [\text{rot}(\mathbf{c}_0^T)_{\overline{\text{idx}}} \mid \text{rot}(\mathbf{c}_1^T)], \\ \vec{s} &:= (\text{Coeffs}(\mathbf{m})_{\overline{\text{idx}}}, \text{Coeffs}(\mathbf{r})), \\ \vec{u} &:= \text{Coeffs}(u) - \text{rot}(\mathbf{c}_0^T)_{\text{idx}} \text{Coeffs}(\mathbf{m})_{\text{idx}}. \end{aligned}$$

Then, using the properties of multiplication matrices we deduce that:

$$\begin{aligned} \text{Coeffs}(u) &= \text{rot}(\mathbf{c}_0^T) \text{Coeffs}(\mathbf{m}) + \text{rot}(\mathbf{c}_1^T) \text{Coeffs}(\mathbf{r}) \\ &= \text{rot}(\mathbf{c}_0^T)_{\text{idx}} \text{Coeffs}(\mathbf{m})_{\text{idx}} + \text{rot}(\mathbf{c}_0^T)_{\overline{\text{idx}}} \text{Coeffs}(\mathbf{m})_{\overline{\text{idx}}} + \text{rot}(\mathbf{c}_1^T) \text{Coeffs}(\mathbf{r}) \end{aligned}$$

and thus:

$$\begin{aligned} P\vec{s} &= \text{rot}(\mathbf{c}_0^T)_{\overline{\text{idx}}} \text{Coeffs}(\mathbf{m})_{\overline{\text{idx}}} + \text{rot}(\mathbf{c}_1^T) \text{Coeffs}(\mathbf{r}) \\ &= \text{Coeffs}(u) - \text{rot}(\mathbf{c}_0^T)_{\text{idx}} \text{Coeffs}(\mathbf{m})_{\text{idx}} \\ &= \vec{u}. \end{aligned}$$

Finally, since in $\Pi_{\text{NIZK}}^{\text{Com}}$ we need to prove relations modulo \hat{q}_1 , we prove instead:

$$\left(\frac{\hat{q}_1}{q} \cdot P \right) \vec{s} = \frac{\hat{q}_1}{q} \cdot \vec{u} \pmod{\hat{q}_1}.$$

Hence, the holder sends the commitment $u \in \mathcal{R}_q$ and proof π .

The issuer \mathcal{I} , given secret key isk , attributes attrs' , the set of indices idx and the message (u, π) from \mathcal{H} , first verifies the proof π , where both P and \vec{u} can be manually computed by \mathcal{I} . Further, it generates a uniformly random $x \leftarrow [N]$ and samples $\mathbf{s} \leftarrow \text{GSampler}(a_1, \mathbf{a}_2, \mathbf{B}, \mathbf{s}, f(x) + u)$, which by definition satisfies

$$[a_1 \mid \mathbf{a}_2^T \mid 1] \mathbf{s} = f(x) + u.$$

Then, \mathcal{I} outputs (x, \mathbf{s}) to \mathcal{H} . The holder then checks whether \mathbf{s} is short and satisfies the equation above. If so, then the credential is a triple $\text{cred} = (\mathbf{s}, \mathbf{r}, x)$.




Fig. 19: AnonCreds.Issue protocol.

Verification. We describe the protocol between the holder $\mathcal{H}(\text{ipk}, \text{attrs}, \text{cred}, \text{idx})$ and the verifier \mathcal{V} . Informally, the holder wants to prove knowledge of a short vector $\mathbf{s} \in \mathcal{R}_q^{m+2}$, binary polynomial $f(x) \in \mathbf{B}$, a short vector $\mathbf{r} \in \mathcal{R}_q^{\ell_r}$, and private parts of \mathbf{m} defined as in (35) (dictated by indices idx of public attributes) which satisfy:

$$[a_1 | \mathbf{a}_2^T | 1] \mathbf{s} = f(x) + \mathbf{c}_0^T \mathbf{m} + \mathbf{c}_1^T \mathbf{r}.$$

Now, we need to transform these relations to the ones of the form (12) in order to use $\Pi_{\text{NIZK}}^{\text{ISIS}}$. From the properties of multiplication matrices we get

$$\begin{aligned} & \text{rot}([a_1 | \mathbf{a}_2^T | 1]) \text{Coeffs}(\mathbf{s}) \\ &= \text{Coeffs}([a_1 | \mathbf{a}_2^T | 1] \mathbf{s}) \\ &= \text{Coeffs}(f(x)) + \text{rot}(\mathbf{c}_0^T) \text{Coeffs}(\mathbf{m}) + \text{rot}(\mathbf{c}_1^T) \text{Coeffs}(\mathbf{r}) \\ &= \text{Coeffs}(f(x)) + \text{rot}(\mathbf{c}_0^T)_{\text{idx}} \text{Coeffs}(\mathbf{m})_{\text{idx}} + \text{rot}(\mathbf{c}_0^T)_{\overline{\text{idx}}} \text{Coeffs}(\mathbf{m})_{\overline{\text{idx}}} + \text{rot}(\mathbf{c}_1^T) \text{Coeffs}(\mathbf{r}). \end{aligned}$$

Hence, let us define

$$\begin{aligned} P &:= \text{rot}([a_1 | \mathbf{a}_2^T | 1]), \quad C := [\text{rot}(\mathbf{c}_0^T)_{\text{idx}} \mid \text{rot}(\mathbf{c}_0^T)_{\overline{\text{idx}}} \mid \text{rot}(\mathbf{c}_1^T)], \quad \vec{m} := \text{Coeffs}(\mathbf{m})_{\text{idx}}, \\ \vec{s} &:= \text{Coeffs}(\mathbf{s}), \quad \vec{r} := (\text{Coeffs}(\mathbf{m})_{\overline{\text{idx}}}, \text{Coeffs}(\mathbf{r})), \quad \vec{u} := \text{enc}(x). \end{aligned}$$

One can check that

$$P \vec{s} = B \vec{u} + C \begin{bmatrix} \vec{m} \\ \vec{r} \end{bmatrix}$$




Fig. 20: AnonCreds.Verify protocol.

and also $\|\vec{s}\| \leq \mathfrak{s} \sqrt{(m+2)d}$ with high probability, $\|\vec{r}\| \leq \psi \sqrt{h \cdot |\text{idx}| + \ell_r d}$ and $\vec{u} \in \{0, 1\}^t$. Then, to obtain a relation of the form (12), we lift the equation from \mathbb{Z}_q to $\mathbb{Z}_{\hat{q}_2}$ as before, i.e. we prove instead:

$$\left(\frac{\hat{q}_2}{q} \cdot P \right) \vec{s} = \left(\frac{\hat{q}_2}{q} \cdot B \right) \vec{u} + \left(\frac{\hat{q}_2}{q} \cdot C \right) \begin{bmatrix} \vec{m}' \\ \vec{r}' \end{bmatrix} \pmod{\hat{q}_2}.$$

Then, we end up with exactly the relation (12). Thus, the holder outputs the proof π from $\Pi_{\text{NIZK}}^{\text{ISIS}}$ to the verifier \mathcal{V} w.r.t. the statement

$$(\hat{q}_2/q \cdot P, \hat{q}_2/q \cdot C, \vec{m}', \hat{q}_2/q \cdot B, \text{Bounds}, \text{idx}),$$

where $\text{Bounds} = \left(\mathfrak{s} \sqrt{(m+2)d}, \psi \sqrt{h \cdot |\text{idx}| + \ell_r d} \right)$.

Finally, \mathcal{V} given the public key ipk , set of attributes attrs' , set of indices idx and a proof π , manually re-computes the statement $(\hat{q}_2/q \cdot P, \hat{q}_2/q \cdot C, \vec{m}', \hat{q}_2/q \cdot B, \text{Bounds}, \text{idx})$ and checks whether the proof π is valid.

8.2 Security Analysis

Lemma 8.1 (Correctness). *Let $\delta, d = O(\lambda)$, $\epsilon = 2^{-\delta}/(4d)$ and $\mathfrak{s} \geq 1.17\sqrt{q} \cdot \eta'_\epsilon(\mathbb{Z})$. If both $\Pi_{\text{NIZK}}^{\text{ISIS}}$ and $\Pi_{\text{NIZK}}^{\text{Com}}$ are correct, then AnonCreds is correct.*

Proof. We start by remarking that any honest execution of `AnonCreds.Init` will succeed since it only entails random sampling and an invocation of `NTRU.TrapGen` which, by Lemma 2.6, always succeeds.

We next show that given a successful execution of `AnonCreds.Init`, any honest execution of `AnonCreds.Issue` succeeds with overwhelming probability (Equation 10): the first round $\mathcal{H}_1(\text{ipk}, \text{attrs}, \text{idx})$ performed by the holder always succeeds since it only involves random sampling, algebraic operations and generation of a proof. Not to mention the fact that for the vector \vec{s} generated honestly by \mathcal{H}_1 , we must have

$$\|\vec{s}\| = \|(\text{Coeffs}(\mathbf{m})_{\overline{\text{idx}}}, \text{Coeffs}(\mathbf{r}))\| \leq \psi \sqrt{h \cdot |\overline{\text{idx}}| + \ell_r d}.$$

Further, the round $\mathcal{I}(\text{isk}, \text{attrs}', \text{idx}, \rho_1)$ executed by the issuer might fail if

$$\text{Verify}_{\text{Com}}^{\text{HCom}} \left(\text{crs}_{\text{Com}}, (\hat{q}_1/q \cdot P, \hat{q}_1/q \cdot \vec{u}, \psi \sqrt{h \cdot |\overline{\text{idx}}| + \ell_r d}), \pi \right) = 0,$$

but by the correctness of $\Pi_{\text{NIZK}}^{\text{Com}}$ (c.f. Lemma 6.1) this happens with negligible probability. Here, we used the fact that if both \mathcal{H} and \mathcal{I} are given attributes `attrs` and `attrs'`, which agree on indices in `idx`, then we must have that \vec{u} computed by \mathcal{H} , and \vec{u} computed by \mathcal{I} , are identical. Finally, the last round $\mathcal{H}_2(\text{ipk}, \rho_2, \text{st}_{\mathcal{H}})$ performed by the holder may fail if $\|\mathbf{s}\| > \mathfrak{s} \sqrt{(m+2)d}$. Recall that $\mathbf{s} \in \mathcal{R}_q^{m+2}$ is generated such that

$$[a_1 | \mathbf{a}_2^T | 1] \mathbf{s} = f(x) + u.$$

Note that \mathbf{s}_2 is randomly sampled from a distribution statistically close to a discrete Gaussian distribution by Lemma 2.7. Consequently, by Lemma 2.3, \mathbf{s} satisfies the condition $\|\mathbf{s}\| \leq \mathfrak{s} \sqrt{(m+2)d}$, with an overwhelming probability. Also, since $\mathbf{r} \leftarrow S_{\psi}^{\ell_r}$, we get $\|\mathbf{r}\| \leq \psi \sqrt{\ell_r d}$.

Finally we show that given a successful execution of `AnonCreds.Init` and of `AnonCreds.Issue`, any honest execution of `AnonCreds.Verify` succeeds with overwhelming probability (Equation 11): the first round $\mathcal{H}(\text{ipk}, \text{attrs}, \text{cred}, \text{idx})$ executed by the holder always succeeds since it only involves random sampling, algebraic operations and generation of a proof. Similarly as before, the vector \vec{r} computed honestly by \mathcal{H} satisfies

$$\|\vec{r}\| = \|(\text{Coeffs}(\mathbf{m})_{\overline{\text{idx}}}, \text{Coeffs}(\mathbf{r}))\| \leq \psi \sqrt{h \cdot |\overline{\text{idx}}| + \ell_r d}.$$

Furthermore, the round $\mathcal{V}(\text{ipk}, \text{attrs}', \text{idx}, \pi)$ executed by the verifier might fail if $\text{Verify}_{\text{ISIS}}^{\text{HISIS}}(\text{crs}_{\text{ISIS}}, (\hat{q}_2/q \cdot P, \hat{q}_2/q \cdot C, \vec{m}', \hat{q}_2/q \cdot B, \text{Bounds}, \text{idx}), \pi) = 0$ but by the correctness of $\Pi_{\text{NIZK}}^{\text{ISIS}}$ (c.f. Lemma 5.5) this happens with negligible probability. Similarly as before, we used the observation that if both \mathcal{H} and \mathcal{V} are given attributes `attrs` and `attrs'`, which agree on indices in `idx`, then we must have \vec{m} (computed by \mathcal{H}) is equal to \vec{m}' (computed by \mathcal{V}). \square

Theorem 8.2 (Anonymity). *Suppose $\Pi_{\text{NIZK}}^{\text{Com}}$ and $\Pi_{\text{NIZK}}^{\text{ISIS}}$ are zero-knowledge and $\text{MLWE}_{1, \ell_r, \psi}$ assumption is hard. Then, `AnonCreds` is anonymous.*

Proof. Let \mathcal{A} be a PPT adversary against the anonymity game. We prove the statement by introducing a sequence of games where we denote ϵ_i the probability that \mathcal{A} outputs wins the i -th game.

Game₁: This is the standard anonymity game. By definition, the probability that \mathcal{A} outputs $\mathbf{b}' = \mathbf{b}$ is ϵ_1 .

Game₂: In this game, instead of running $\text{Prove}_{\text{ISIS}}^{\text{HISIS}}$, the challenger simulates the holder's behaviour in `AnonCreds.Verify` _{\mathcal{H}, \mathcal{V}} using the zero-knowledge simulator ($\text{Sim}_0, \text{Sim}_1$) of $\Pi_{\text{NIZK}}^{\text{ISIS}}$. Namely, when \mathcal{A} makes a random oracle query to H_{ISIS} , the challenger calls Sim_0 . Furthermore, in the `AnonCreds.Verify` _{\mathcal{H}, \mathcal{V}} protocol, the challenger now outputs a simulated proof $\pi_{\mathbf{b}} \leftarrow \text{Sim}_1(\text{crs}_{\text{ISIS}}, (\hat{q}_2/q \cdot P, \hat{q}_2/q \cdot C, \vec{m}', \hat{q}_2/q))$ to the adversary. Note that Sim_1 is only called once during the game.

One can naturally construct a PPT adversary \mathcal{A}' which wins the zero-knowledge game of $\Pi_{\text{NIZK}}^{\text{ISIS}}$ with probability $|\epsilon_2 - \epsilon_1|$. Concretely, the adversary first programs the output of $\text{H}_{\text{crs}}(0)$ to use the crs_{ISIS} provided by the zero-knowledge game. Next, it internally runs \mathcal{A} and simulates the challenger's behaviour given access to either the oracles $(\text{H}_{\text{ISIS}}, \text{Prove})$ or $(\text{Sim}_0, \text{Sim}_1)$. Note that by the verification checks made by the challenger, \mathcal{A} queries only valid statements. Thus, by assumption we obtain:

$$|\epsilon_2 - \epsilon_1| \leq \text{Adv}_{\Pi_{\text{NIZK}}^{\text{ISIS}}}^{\text{ZK}}(\mathcal{A}') = \text{negl}(\lambda).$$

Game₃: In this game, instead of running $\text{Prove}_{\text{Com}}^{\text{H}_{\text{Com}}}$ in AnonCreds.Issue , the challenger simulates the proof π_b of a validity of a commitment u_b using the zero-knowledge simulator $(\text{Sim}_0, \text{Sim}_1)$ of $\Pi_{\text{NIZK}}^{\text{Com}}$. Namely, when \mathcal{A} makes random oracle query to H_{Com} , the challenger calls Sim_0 . Further, the challenger computes simulated proofs $\pi_b := \text{Sim}_1(\text{crs}_{\text{Com}}, (P, \bar{u}))$. Finally, the challenger outputs (u_b, π_b) as the first message pair to \mathcal{A} .

Similarly as before, one can construct a PPT adversary \mathcal{A}' which wins the zero-knowledge game of $\Pi_{\text{NIZK}}^{\text{Com}}$ with probability $|\epsilon_3 - \epsilon_2|$. Hence, by assumption on $\Pi_{\text{NIZK}}^{\text{Com}}$ we get

$$|\epsilon_3 - \epsilon_2| \leq \text{Adv}_{\Pi_{\text{NIZK}}^{\text{Com}}}^{\text{ZK}}(\mathcal{A}') = \text{negl}(\lambda).$$

Game₄: Here, the challenger simulates the commitment u . Namely, \mathcal{C} picks the commitment u' uniformly at random from \mathcal{R}_q and sets $u := u' + \mathbf{c}_0^T \mathbf{m}$. Then, there exists a natural PPT adversary $\mathcal{A}^{\text{MLWE}}$ which can distinguish $\mathbf{c}_1^T \mathbf{r}$ from a uniformly random polynomial u with probability at least $|\epsilon_4 - \epsilon_3|$, i.e. solves Module-LWE. Hence,

$$|\epsilon_4 - \epsilon_3| \leq \text{Adv}_{1, \ell_r, \psi}^{\text{MLWE}}(\mathcal{A}^{\text{MLWE}}) = \text{negl}(\lambda).$$

Game₅: Now, the challenger directly samples u uniformly at random from \mathcal{R}_q . Clearly, we get $\epsilon_5 = \epsilon_4$. The key observation here is that in **Game₅**, the behaviour of the challenger is independent of the bit \mathbf{b} . Hence, we conclude that $\epsilon_5 = 1/2$ and the statement holds by the hybrid argument. \square

Theorem 8.3 (One-More Unforgeability). *Let $h = O(\lambda)$. Suppose $\Pi_{\text{NIZK}}^{\text{Com}}$ is multi-proof extractable and the proof system modulus \hat{q}_2 for $\Pi_{\text{NIZK}}^{\text{ISIS}}$ satisfies*

$$\hat{q}_2 > \max \left(\mathfrak{s} \sqrt{(m+2)d}, \psi \sqrt{(\ell_m + \ell_r)d}, 16m_1 \hat{d} \mathcal{B}_3, \frac{2}{\omega_{\min}(\lambda)^2} \mathcal{B}_3^2 \right)$$

where m_1, m_2 are the parameters defined in Section 5. Then, AnonCreds is one-more unforgeable under the $\text{Int-NTRU-ISIS}_{\mathbf{f}}^{\text{PP}}$ and $\text{MSIS}_{n, m_1 + m_2, \bar{\mathcal{B}}}$ assumptions where $\text{pp} := (q, d, m, \ell_m, \ell_r, N, \mathfrak{s}, \mathfrak{s} \sqrt{(m+2)d}, \psi \sqrt{(\ell_m + \ell_r)d})$ and $\bar{\mathcal{B}} = 4\nu \sqrt{\mathcal{B}_1^2 + \mathcal{B}_2^2}$.

Proof. Suppose there is a PPT adversary \mathcal{A} which wins the one-more unforgeability game with non-negligible probability ϵ . Suppose \mathcal{A} makes at most Q_{Issue} type queries and $Q_{\text{H}_{\text{ISIS}}}, Q_{\text{H}_{\text{Com}}}, Q_{\text{H}_{\text{M}}}$ queries to the random oracles $\text{H}_{\text{ISIS}}, \text{H}_{\text{Com}},$ and H_{M} respectively. Without loss of generality, assume that \mathcal{A} never repeats a random oracle query.

Hybrid games. We prove the statement by introducing a sequence of games. In the following, let E_i be the event that \mathcal{A} wins in **Game_i** and denote \mathcal{C}_i as the challenger in **Game_i**.

Game₀: This is the standard one-more unforgeability game. Hence, by definition we have $\Pr[\text{E}_0] = \epsilon$.

Game₁: At the beginning of the game, \mathcal{C}_1 samples $\vec{h}_j \leftarrow \{-\psi, -\psi + 1, \dots, \psi\}^h$ for all $j \in [Q_{\text{H}_{\text{M}}}]$. When the adversary queries \mathbf{a}_j' as the j -th random oracle query to H_{M} , then the challenger just outputs h_j . However, if some collision occurs, i.e. $\vec{h}_i = \vec{h}_j$ for distinct indices i, j , then \mathcal{C}_1 aborts. By the union bound, we have

$$\Pr[\text{E}_1] \geq \Pr[\text{E}_0] - \frac{|Q_{\text{H}_{\text{M}}}|^2}{(2\psi + 1)^h} = \Pr[\text{E}_0] - \text{negl}(\lambda).$$

Game₂: In this game, the challenger programs the simulated $\text{c}\tilde{\text{r}}\text{s}$ used for multi-proof extractability. Namely, recall that in **Game₁** we had $\text{H}_{\text{crs}}(0) = (\text{crs}_{\text{SIS}}, \text{crs}_{\text{Com}})$. Here, the challenger \mathcal{C}_2 honestly generates $\text{crs}_{\text{NIZK}}^{\text{S}}$, but it also runs the CRS simulator \mathcal{S}_{crs} and obtains $(\text{c}\tilde{\text{r}}\text{s}, \text{td}) \leftarrow \mathcal{S}_{\text{crs}}(1^\lambda)$. Then, it programs $\text{H}_{\text{crs}}(0) := (\text{crs}_{\text{SIS}}, \text{c}\tilde{\text{r}}\text{s})$ and keeps td .

It is easy to see that **Game₃** and **Game₂** are indistinguishable by the CRS indistinguishability property of the multi-proof extractability, i.e. one can construct a PPT adversary \mathcal{A}^{crs} such that

$$\Pr[\text{E}_2] \geq \Pr[\text{E}_1] - \text{Adv}_{\Pi_{\text{NIZK}}^{\text{Com}}}^{\text{crs}}(\mathcal{A}^{\text{crs}}) = \Pr[\text{E}_1] - \text{negl}(\lambda).$$

Game₃: In this game, the challenger \mathcal{C}_3 uses the multi-proof extractability of $\Pi_{\text{NIZK}}^{\text{Com}}$ to extract witnesses for all proofs π sent by \mathcal{A} as a part of issue queries. Namely, when \mathcal{A} submits the j -th issue query $(\text{attrs}'_j, \text{idx}_j, u_j, \pi_j)$, where $j \in [Q_{\text{Issue}}]$, the challenger additionally runs

$$w_j \leftarrow \text{Multi-Extract} \left(1^\lambda, Q_{\text{HCom}}, Q_{\text{Issue}}, 1/\Pr[\text{E}_2], \text{td}, (\hat{q}_1/q \cdot P_j, \hat{q}_1/q \cdot \vec{u}_j), \pi_j \right).$$

Let $\text{Abort}_{\text{extract}}$ be the event that $w_j \notin R^{\text{Com}}(\hat{q}_1/q \cdot P_j, \hat{q}_1/q \cdot \vec{u}_j)$ for some $j \in [Q_{\text{Issue}}]$, where P_j, \vec{u}_j are computed as in the honest execution of the issuer. If $\text{Abort}_{\text{extract}}$ occurs then \mathcal{C}_4 aborts the game and overwrites the adversary's forgery to be \perp . Otherwise, it proceeds as \mathcal{C}_3 . Observe that \mathcal{C}_4 does not make any use of the extracted witnesses w_j . In the following, let $\mathbf{m}_j \in \mathcal{R}_q^{\ell_m}$ and $\mathbf{r}_j \in \mathcal{R}_q^{\ell_r}$ be the vectors which satisfy:

$$w_j = \left(\text{Coeffs}(\mathbf{m}_j)_{\overline{\text{idx}}_j}, \text{Coeffs}(\mathbf{r}_j) \right) \quad \text{and} \quad \text{Coeffs}(\mathbf{m}_j)_{\text{idx}_j} = \text{Coeffs}(\mathbf{m}'_j)_{\text{idx}_j}$$

where $\mathbf{m}'_j = \text{Coeffs}^{-1}(\text{H}_M(\mathbf{a}_{j,1}'), \dots, \text{H}_M(\mathbf{a}_{j,l}'))$. In particular, we have

$$u_j = \mathbf{c}_0^T \mathbf{m}_j + \mathbf{c}_1^T \mathbf{r}_j \quad \text{for } j = 1, 2, \dots, Q_{\text{Issue}}.$$

The challenger stores the pairs $(\mathbf{m}_j, \mathbf{r}_j)_{j \in [Q_{\text{Issue}}]}$. Intuitively, these pairs are the extracted message/randomness committed and sent from the holder to the issuer. From now on, we will informally say that $(\mathbf{m}_j, \mathbf{r}_j)$ is the j -th issuing query to \mathcal{C}_3 , rather than (u_j, π_j) . By definition we have

$$\|(\mathbf{m}_j, \mathbf{r}_j)\| = \left\| \left(\text{Coeffs}(\mathbf{m}_j)_{\overline{\text{idx}}_j}, \text{Coeffs}(\mathbf{r}_j), \text{Coeffs}(\mathbf{m}'_j)_{\text{idx}_j} \right) \right\| \leq \psi \sqrt{(\ell_m + \ell_r)d}.$$

Arguing identically as in [67, Lemma 3.6] and assuming that $\Pr[\text{E}_2]$ is non-negligible, the runtime of \mathcal{C}_3 is still $\text{poly}(\lambda)$ and also

$$\Pr[\text{E}_3] \geq \frac{1}{2} \Pr[\text{E}_2] - \text{negl}(\lambda).$$

Introducing a wrapper. Now, we provide a wrapper algorithm \mathcal{P}^* for \mathcal{A} . Namely, \mathcal{P}^* is given the crs_{SIS} for the proof system $\Pi_{\text{NIZK}}^{\text{SIS}}$, and runs \mathcal{A} as a challenger for **Game₃** (where crs_{SIS} is programmed as a part of $\text{H}_{\text{crs}}(0) := (\text{crs}_{\text{SIS}}, \text{c}\tilde{\text{r}}\text{s})$). It all simulates all random oracle queries, apart from the one to H_{SIS} , using lazy sampling.

Suppose that at the end \mathcal{A} outputs a set $\mathcal{L} := \{\langle \text{attrs}_i^*, \text{idx}_i^* \rangle\}$. For each $i \in [Q_{\text{Issue}} + 1]$, write $\text{attrs}_i^* = (\mathbf{a}_{i,1}^*, \dots, \mathbf{a}_{i,l}^*)$ and define

$$\mathbf{m}_i^* := \text{Coeffs}^{-1}(\text{H}_M(\mathbf{a}_{i,1}^*), \dots, \text{H}_M(\mathbf{a}_{i,l}^*)) \in \mathcal{R}_q^{\ell_m}. \quad (36)$$

We let the algorithm \mathcal{P}^* find an index t which satisfies

$$\text{Coeffs}(\mathbf{m}_t^*)_{\text{idx}_t^*} \notin \left\{ \text{Coeffs}(\mathbf{m}_1)_{\text{idx}_1^*}, \dots, \text{Coeffs}(\mathbf{m}_{Q_{\text{Issue}}})_{\text{idx}_{Q_{\text{Issue}}}^*} \right\}. \quad (37)$$

If it cannot find one, then it outputs $(\perp, \perp, \perp, 0)$. Otherwise, the algorithm outputs a pair $(\mathcal{I}, y, v, \mathcal{Q})$ defined as follows. First, let

$$\mathbf{x} := (\hat{q}_2/q \cdot P, \hat{q}_2/q \cdot C, \vec{m}, \hat{q}_2/q \cdot B, \text{Bounds}, \text{aux} := \text{idx}_t^*) \quad (38)$$

be the statement for relation (12) where

$$P := \text{rot}([a_1 | \mathbf{a}_2^T | 1]), \quad C := [\text{rot}(\mathbf{c}_0^T)_{\text{idx}_t^*} \mid \text{rot}(\mathbf{c}_0^T)_{\overline{\text{idx}_t^*}} \mid \text{rot}(\mathbf{c}_1^T)], \quad \vec{m} := \text{Coeffs}(\mathbf{m})_{\text{idx}_t^*}.$$

Also, let $\pi_t := (a_1, a_2, a_3, a_4, a_5)$ be the proof sent by the adversary \mathcal{A} in $\text{AnonCreds.Verify}_{\mathcal{A}, \mathcal{P}^*}(\langle \cdot, \cdot, \cdot \rangle, (\text{ipk}, \text{attrs}_t^*, \text{idx}_t^*))$. Define $\mathcal{I} := (I_1, I_2, I_3, I_4, I_5)$ as

$$I_1 := (1, \text{crs}_{\text{ISIS}}, \mathbf{x}, a_1), I_2 := (2, \text{crs}_{\text{ISIS}}, \mathbf{x}, a_1, a_2), \dots, I_5 := (5, \text{crs}_{\text{ISIS}}, \mathbf{x}, a_1, a_2, a_3, a_4, a_5).$$

Then, it obtains the challenges by querying

$$\begin{aligned} (\mathfrak{R}_0, \mathfrak{R}_1) &:= \text{H}_{\text{ISIS}}(I_1), \\ (\gamma_{i,j})_{i \in [\tau], j \in [256+d+3]} &:= \text{H}_{\text{ISIS}}(I_2), \\ (\mu_1, \dots, \mu_\tau) &:= \text{H}_{\text{ISIS}}(I_3), \\ c &:= \text{H}_{\text{ISIS}}(I_4) \end{aligned}$$

and defines

$$\text{tr} := (\text{crs}_{\text{ISIS}}, \mathbf{x}, a_1, (\mathfrak{R}_0, \mathfrak{R}_1), a_2, (\gamma_{i,j}), a_3, (\mu_i), a_4, c, a_5), \quad \text{and} \quad v := V(\text{crs}_{\text{ISIS}}, \mathbf{x}, \text{tr})$$

where $V(\text{crs}_{\text{ISIS}}, \mathbf{x}, \text{tr}) = 1$ if tr is the accepting transcript for statement \mathbf{x} with public parameters crs_{ISIS} , and 0 otherwise. This implies that \mathcal{P}^* makes at most $Q_{\text{H}_{\text{ISIS}}} + 4$ queries to H_{ISIS} . Finally, $\mathcal{Q} := (\mathbf{m}_j, \mathbf{r}_j)_{j \in [Q_{\text{Issue}}]}$ and \mathcal{P}^* outputs

$$(\mathcal{I}, \text{tr}, v, \mathcal{Q}).$$

Also, given such a tuple from \mathcal{P}^* , where $v = 1$, we will denote:

$$\mathcal{I}(\vec{m}) := \vec{m} \quad \text{and} \quad \mathcal{I}(\text{aux}) := \text{aux} = \text{idx}_t^*.$$

Below, we provide a few basic properties of the wrapper \mathcal{P}^* to keep in mind for the remainder of the proof.

Lemma 8.4. $\Pr[(\mathcal{I}, \text{tr}, v, \mathcal{Q}) \leftarrow \mathcal{P}^*(\text{crs}_{\text{ISIS}}) \wedge v = 1] \geq \Pr[\text{E}_3]$.

Proof. Suppose \mathcal{A} wins Game_3 . Namely, the set \mathcal{L} is of size $Q_{\text{Issue}} + 1$ such that (by definition of the one-more unforgeability game):

- for all $i \in [Q_{\text{Issue}} + 1]$, $\langle \perp, 1 \rangle \leftarrow \text{AnonCreds.Verify}_{\mathcal{A}, \mathcal{P}^*}(\langle \cdot, \cdot, \cdot \rangle, (\text{ipk}, \text{attrs}_i^*, \text{idx}_i^*))$,
- for each distinct pairs $\langle \text{attrs}_i^*, \text{idx}_i^* \rangle \in \mathcal{L}$ and $\langle \text{attrs}_j^*, \text{idx}_j^* \rangle \in \mathcal{L}$, there exists at least one index i^* in both idx_i^* and idx_j^* such that $\text{attrs}_{i,i^*} \neq \text{attrs}_{j,i^*}$.

We claim there exists an index $t \in [Q_{\text{Issue}} + 1]$ such that (37). Suppose it is not the case. Then, by the pigeonhole principle there are two distinct indices $s, t \in [Q_{\text{Issue}} + 1]$ and $j \in [Q_{\text{Issue}}]$ such that:

$$\text{Coeffs}(\mathbf{m}_s^*)_{\text{idx}_s^*} = \text{Coeffs}(\mathbf{m}_j)_{\text{idx}_s^*} \quad \text{and} \quad \text{Coeffs}(\mathbf{m}_t^*)_{\text{idx}_t^*} = \text{Coeffs}(\mathbf{m}_j)_{\text{idx}_t^*}.$$

It implies that for all $i \in \text{idx}_s^* \cap \text{idx}_t^*$, $\text{H}_M(\mathbf{a}_{s,i}^*) = \text{H}_M(\mathbf{a}_{t,i}^*)$. Since we excluded the case of getting a hash collision in Game_1 , we have $\mathbf{a}_{s,i}^* = \mathbf{a}_{t,i}^*$ which leads to the contradiction with the second winning condition of one-more unforgeability. Thus, \mathcal{P}^* outputs $v = 1$. \square

Lemma 8.5. *The following holds:*

$$\Pr \left[(\mathcal{I}, \text{tr}, 1, \mathcal{Q}) \leftarrow \mathcal{P}^*(\text{crs}_{\text{ISIS}}) \wedge \mathcal{I}(\vec{m}) \notin \{ \text{Coeffs}(\mathbf{m})_{\mathcal{I}(\text{aux})} : \exists \mathbf{r} \in \mathcal{R}_q^{\ell_r} \text{ s.t. } (\mathbf{m}, \mathbf{r}) \in \mathcal{Q} \} \right] = 1.$$

Proof. The statement can be shown in a similar manner as for Lemma 8.4. \square

Next, we state crucial forking property of \mathcal{P}^* .

Black-box access to: $\mathcal{P}^*(\text{crs}_{\text{ISIS}})$

1. Run $\mathcal{P}^*(\text{crs}_{\text{ISIS}})$ with randomness ρ as follows: relay the $Q_{\text{HISIS}} + 4$ random oracle queries to the random oracle and record all query-response pairs. Obtain $(\mathcal{I}, \text{tr}, v, \mathcal{Q})$. Set $i = I_4$ and let c_i be the response to query i .
2. If $v = 0$, abort and return $v = 0$. Else, $(\vec{m}, \text{aux}) := (\mathcal{I}(\vec{m}), \mathcal{I}(\text{aux}))$.
3. Repeat:
 - (a) sample $c'_i \in \mathcal{C} \setminus \{c_i\}$ without replacement,
 - (b) run $\mathcal{P}^*(\text{crs}_{\text{ISIS}})$ with randomness ρ as follows to obtain $(\mathcal{I}', \text{tr}', v', \mathcal{Q}')$: answer the query to i with c'_i while answering all the other queries consistently if the query was performed by \mathcal{P}^* in the previous run, and with fresh random values otherwise,
 - (c) **If the input for some j -th issuing query to \mathcal{P}^* is of the form $(\mathbf{m}'_j, \mathbf{r}'_j)$ where $\vec{m} = \text{Coeffs}(\mathbf{m}'_j)_{\text{aux}}$, then stop running \mathcal{P}^* and directly go back to Step 3(a),**
until an additional challenge c'_i with $v' = 1$ and $I'_K = I_K$ has been found, or until all challenges $c'_i \in \mathcal{C}$ have been tried.
4. If the latter case occurs, output $v = 0$.
5. Otherwise, return $((\mathcal{I}, \text{tr}, v, \mathcal{Q}), (\mathcal{I}', \text{tr}', v', \mathcal{Q}'))$.

Fig. 21: Basic extractors $\mathcal{E}(\text{crs}_{\text{ISIS}})$ (defined with black text) and $\mathcal{E}'(\text{crs}_{\text{ISIS}})$ (defined with black and blue text). Here, $K \in \{1, 2, 3, 4\}$ is a fixed constant.

Lemma 8.6. *Consider the algorithms \mathcal{E} and \mathcal{E}' in Figure 21. Then, $\Pr[0 \leftarrow \mathcal{E}(\text{crs}_{\text{ISIS}})] = \Pr[0 \leftarrow \mathcal{E}'(\text{crs}_{\text{ISIS}})]$.*

Proof. Clearly, if \mathcal{E} outputs $v = 0$ then so does \mathcal{E}' . Hence, suppose that when running \mathcal{E}' , in particular Step 3, for some j -th issuing query to \mathcal{P}^* , the input $(\mathbf{m}'_j, \mathbf{r}'_j)$ satisfies $\vec{m} = \text{Coeffs}(\mathbf{m}'_j)_{\text{aux}}$. In this case, \mathcal{E}' stops the execution of \mathcal{P}^* and goes back to Step 3(a), while \mathcal{E} continues running \mathcal{P}^* . Assume \mathcal{P}^* returns $(\mathcal{I}', \text{tr}', v', \mathcal{Q}')$ and $I'_4 = I_4$. By our assumption on the j -th query,

$$\vec{m} \in \{\text{Coeffs}(\mathbf{m}'_j)_{\text{aux}} : \exists \mathbf{r}' \in \mathcal{R}_q^{\ell_r} \text{ such that } (\mathbf{m}', \mathbf{r}') \in \mathcal{Q}'\}. \quad (39)$$

Thus, by Lemma 8.5 this implies $v' = 0$. Therefore, \mathcal{E} would still go back to Step 3(a), identically as \mathcal{E}' . This concludes the proof. \square

Even though we will not use this lemma as a black-box, its core idea will be applied throughout the extraction procedure. In particular, the key observation is that, by construction, \mathcal{E}' after Step 2 never answers an issuing query for any $(\mathbf{m}'_j, \mathbf{r}'_j)$ such that $\vec{m} = \text{Coeffs}(\mathbf{m}'_j)_{\text{aux}}$. This will be crucial when reducing to Int-NTRU-ISIS_f . Obviously, the lemma can be applied similarly if we reprogram the random oracle for I_1, I_2 or I_3 .

Defining the extractor. Since we will now analyse the protocol from Figures 9 and 10, let us denote for simplicity:

$$\begin{aligned} \hat{q} &:= \hat{q}_2, & P &:= \frac{\hat{q}}{q} \cdot P, & C &:= \frac{\hat{q}}{q} \cdot C, & B &:= \frac{\hat{q}}{q} \cdot B \\ \mathcal{B}_s &:= \mathfrak{s}\sqrt{(m+2)d}, & \mathcal{B}_r &:= \psi\sqrt{\mathfrak{h} \cdot |\text{idx}| + \ell_r d}. \end{aligned}$$

If we want to apply the proof system $\Pi_{\text{NIZK}}^{\text{ISIS}}$, we define the variables used in Section 5 as follows ¹⁹:

$$\mathfrak{n} := d, \quad \mathfrak{m} := (m+2)d + \hat{d}, \quad \ell_m := |\text{idx}| \cdot \mathfrak{h}, \quad \ell_r := (l - |\text{idx}|) \cdot 2\lambda + \ell_r d + \hat{d}.$$

¹⁹ Note that we increase the values of \mathfrak{m} and ℓ_r by \hat{d} since instead of proving, e.g. $\|\vec{s}\| \leq \mathcal{B}_s$, we prove $\|\vec{s}\| = \mathcal{B}_s$ using the transformation described in Section 5.2.

Further, \hat{q} is product of two primes \hat{p}_1, \hat{p}_2 as described in Section 5.1.

Intuitively, we will use the extraction algorithm from [57, Appendix B] to extract from the prover \mathcal{P}^* . Namely, suppose in the first run $(\mathcal{I}, \text{tr}, 1, \mathcal{Q}) \leftarrow \mathcal{P}^*(\text{crs}_{\text{ISIS}})$ and let $(\vec{m}, \text{aux}) := (\mathcal{I}(\vec{m}), \mathcal{I}(\text{aux}))$. Then, the extractor wants to obtain short $\mathbf{m}^*, \mathbf{r}^*$, a short vector \mathbf{s}^* and $f(x^*) \in \mathbf{B}$ such that

$$[a_1 \mathbf{a}_2^T \ 1] \mathbf{s}^* = f(x^*) + \mathbf{c}_0^T \mathbf{m}^* + \mathbf{c}_1^T \mathbf{r}^* \quad \text{and} \quad \vec{m} = \text{Coeffs}(\mathbf{m}^*)_{\text{aux}}.$$

The extractor will be constructed similarly as \mathcal{E}' in Lemma 8.6, i.e. when the extractor has to answer any issuing query for input $(\mathbf{m}, \mathbf{r})^{20}$, it stops running the current execution of \mathcal{P}^* if $\vec{m} = \text{Coeffs}(\mathbf{m})_{\text{aux}}$. This way, we make sure that when our ideal extractor outputs $(\mathbf{m}^*, \mathbf{r}^*, \mathbf{s}^*, f(x^*))$ satisfying the conditions above, it never answered any issuing query on input of the form (\mathbf{m}^*, \cdot) since $\vec{m} = \text{Coeffs}(\mathbf{m}^*)_{\text{aux}}$. This gives us a natural reduction to Int-NTRU-ISIS_f .

We start by defining the extractor $\mathcal{E}_4(\text{crs}_{\text{ISIS}})$, which is a $(Q_{\text{HISIS}} + 4)$ -query random oracle algorithm in Fig. 22. Informally, \mathcal{E}_4 finds the witness for a single quadratic equation with automorphism described in (18) by constructing a $(1, 1, 1, 3)$ -tree of transcripts using the extraction from [6].

Black-box access to: $\mathcal{P}^*(\text{crs}_{\text{ISIS}})$

1. Run $\mathcal{P}^*(\text{crs}_{\text{ISIS}})$ with randomness ρ as follows: relay the $Q_{\text{HISIS}} + 4$ random oracle queries to the random oracles and record all query-response pairs. Obtain $(\mathcal{I}, \text{tr}, v, \mathcal{Q})$. Set $i = I_4$ and let c_i be the response to query i .
2. If $v = 0$, abort and return $v = 0$. Else, $(\vec{m}, \text{aux}) := (\mathcal{I}(\vec{m}), \mathcal{I}(\text{aux}))$.
3. Repeat:
 - (a) sample $c'_i \in \mathcal{C} \setminus \{c_i\}$ without replacement,
 - (b) run $\mathcal{P}^*(\text{crs}_{\text{ISIS}})$ with randomness ρ as follows to obtain $(\mathcal{I}', \text{tr}', v', \mathcal{Q}')$: answer the query to i with c'_i while answering all the other queries consistently if the query was performed by \mathcal{P}^* in the previous run, and with fresh random value otherwise,
 - (c) **If the input for some j -th issuing query to \mathcal{P}^* is of the form $(\mathbf{m}'_j, \mathbf{r}'_j)$ where $\vec{m} = \text{Coeffs}(\mathbf{m}'_j)_{\text{aux}}$, then stop running \mathcal{P}^* and directly go back to Step 3(a),**
 until either 2 additional challenges c'_i with $v' = 1$ and $I'_4 = I_4$ have been found or until all challenges $c'_i \in \mathcal{C}$ have been tried.
4. If the latter case occurs, output $v = 0$.
5. For $i = 0, 1, 2$, $\text{tr}_i := (\text{crs}_{\text{ISIS}}, x, a_1, (\mathfrak{R}_0, \mathfrak{R}_1), a_2, (\gamma_{i,j}), a_3, (\mu_i), a_4, c^{(i)}, a_5^{(i)})$ be the extracted transcripts.
6. Parse $a_5^{(i)} := (\mathbf{z}_1^{(i)}, \mathbf{z}_2^{(i)})$ and a_1, a_2, a_3, a_4 as in the protocol description in Figure 10,
7. Define $\bar{c} := c^{(1)} - c^{(0)}$, $\bar{\mathbf{s}}_i = \frac{\mathbf{z}_i^{(1)} - \mathbf{z}_i^{(0)}}{c^{(1)} - c^{(0)}}$ for $i = 1, 2$ and $\bar{\mathbf{m}} = (\bar{\mathbf{y}}_3, \bar{\mathbf{g}}) := \mathbf{t}_B - \mathbf{B}\bar{\mathbf{s}}_2$.
8. If $\mathbf{z}_i^{(1)} - c^{(1)}\bar{\mathbf{s}}_i \neq \mathbf{z}_i^{(2)} - c^{(2)}\bar{\mathbf{s}}_i$ for some $i = 1, 2$:
 - return \mathcal{I} , $\left[\begin{array}{l} \bar{c} \cdot \left(\left(\mathbf{z}_1^{(2)} - \mathbf{z}_1^{(1)} \right) - \left(c^{(2)} - c^{(1)} \right) \bar{\mathbf{s}}_1 \right) \\ \bar{c} \cdot \left(\left(\mathbf{z}_2^{(2)} - \mathbf{z}_2^{(1)} \right) - \left(c^{(2)} - c^{(1)} \right) \bar{\mathbf{s}}_2 \right) \end{array} \right]$ as the MSIS solution to $[\mathbf{A}_1 \mid \mathbf{A}_2]$ and $v = 1$.
9. Otherwise, return \mathcal{I} , $(\bar{\mathbf{s}}_1, \bar{\mathbf{s}}_2, \bar{\mathbf{m}}, \bar{c})$ and $v = 1$.

Fig. 22: Subextractor $\mathcal{E}_4(\text{crs}_{\text{ISIS}})$ as a $(Q_{\text{HISIS}} + 4)$ -query random oracle algorithm.

²⁰ Recall that the extractor, as an issuer, is only given its commitment u and proof π , but we use the multi-proof extractability property to extract the secret \mathbf{m}, \mathbf{r} .

Lemma 8.7. For any crs_{SIS} , the extractor $\mathcal{E}_4(\text{crs}_{\text{SIS}})$ makes an expected number of at most $3 + 2Q_{\text{H}_{\text{SIS}}}$ queries to $\mathcal{P}^*(\text{crs}_{\text{SIS}})$. Furthermore, for uniformly random crs_{SIS} , \mathcal{E}_4 outputs $v = 1$ with probability at least

$$\Pr[\mathbf{E}_3] - (Q_{\text{H}_{\text{SIS}}} + 1) \cdot \frac{2}{|\mathcal{C}|}.$$

Next, let $\mathcal{I} = (I_1, I_2, I_3, I_4, I_5)$ be the index vector obtained by \mathcal{E}_4 in the first step and denote

$$(\mathfrak{R}_0, \mathfrak{R}_1) := \mathbf{H}_1(I_1), \quad (\gamma_{i,j}) := \mathbf{H}_2(I_2), \quad (\mu_i) := \mathbf{H}_3(I_3).$$

as the corresponding (recorded) random oracle responses. Further, parse

$$I_4 := (4, \text{crs}_{\text{SIS}}, \mathbf{x}, a_1, (\mathfrak{R}_0, \mathfrak{R}_1), a_2, (\gamma_{i,j}), a_3, (\mu_i), a_4).$$

Then, conditioned on $v = 1$, the extractor either returns $(\bar{\mathbf{s}}_1 := (\bar{\mathbf{s}}, \bar{\mathbf{r}}, \bar{\mathbf{u}}), \bar{\mathbf{s}}_2, \bar{\mathbf{m}}, \bar{\mathbf{c}}) \in \mathcal{R}_q^{m_1} \times \mathcal{R}_q^{m_2} \times \mathcal{R}_q^{256/\hat{d}+\tau} \times \bar{\mathcal{C}}$ which satisfies the following relations:

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_y \\ \mathbf{t}_g \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \bar{\mathbf{s}}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B}_y \\ \mathbf{B}_g \end{bmatrix} \bar{\mathbf{s}}_2 + \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{y}}_3 \\ \bar{\mathbf{g}} \end{bmatrix}, \quad \|\bar{\mathbf{c}}\bar{\mathbf{s}}_1\| \leq 2\mathcal{B}_1, \quad \|\bar{\mathbf{c}}\bar{\mathbf{s}}_2\| \leq 2\mathcal{B}_2$$

and²¹

$$\begin{aligned} 0 = \sum_{i=1}^{\tau} \mu_i & \left(\sum_{j=1}^{256} \gamma_{i,j} \cdot \left(\sigma(\mathbf{t}_j)^T \bar{\mathbf{s}}_1 + \sigma(\mathbf{e}_j)^T \bar{\mathbf{y}}_3 - z_j \right) \right. \\ & + \sum_{j=1}^n \gamma_{i,256+j} \cdot \left(\sigma(\mathbf{p}_j)^T \bar{\mathbf{s}} - \sigma(\beta_j)^T \bar{\mathbf{u}} - m_{C,j} - \sigma(\mathbf{c}_{r,j})^T \bar{\mathbf{r}} \right) \\ & + \gamma_{i,256+n+1} \cdot \left(\sigma(\bar{\mathbf{s}})^T \bar{\mathbf{s}} - \mathcal{B}_s \right) + \gamma_{i,256+n+2} \cdot \left(\sigma(\bar{\mathbf{r}})^T \bar{\mathbf{r}} - \mathcal{B}_r \right) \\ & \left. + \gamma_{i,256+n+3} \cdot \sigma(\bar{\mathbf{u}} - \mathbf{x})^T \bar{\mathbf{u}} + g_i - h_i \right), \end{aligned} \tag{40}$$

or a $\text{MSIS}_{n, m_1+m_2, \bar{\mathcal{B}}}$ solution for the matrix $[\mathbf{A}_1 \mid \mathbf{A}_2]$ where $\bar{\mathcal{B}} := 4\nu\sqrt{\mathcal{B}_1^2 + \mathcal{B}_2^2}$.

Proof. We first focus on the probability that \mathcal{E}_4 outputs $v = 0$. The first observation is that by Lemma 8.6, we can simply remove Step 3(c). Then, the statement holds by proving almost identically as in [57, Lemma B.8]. \square

Next, we define the subextractor \mathcal{E}_3 which is informally responsible for proving knowledge of a solution of multiple quadratic equations (41). We present the subextractor in Fig. 23. Here, \mathcal{E}_3 uses the early abort feature of \mathcal{E}_4 , as shown in [6]. That is, \mathcal{E}_4 computes the index vector \mathcal{I} by running \mathcal{P}^* as the first step. This allows the executions in the repeat loop of \mathcal{E}_3 to abort right after a single run of \mathcal{P}^* if $I'_3 \neq I_3$. This allows to have the runtime of the extractor to be linear in the number of random oracle queries $Q_{\text{H}_{\text{SIS}}}$.

The main intuition is that by running \mathcal{E}_4 once, \mathcal{E}_3 obtains a candidate witness \mathbf{w} (or a Module-SIS solution, but we ignore the latter case). Since the ABDLOP commitment $(\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g)$ to the witness was sent before getting (μ_i) and the commitment itself is binding, this means that the extracted witness must be *independent* of the challenges μ_1, \dots, μ_τ . Therefore, if for some index $i \in [\tau]$, (41) does not hold, then only with probability at most $\hat{p}_1^{-\hat{d}/2}$ Equation 40, which is the relation that \mathbf{w} satisfies by construction of \mathcal{E}_4 , is true. Formally, the proof simply combines Lemma 8.6 with [57, Lemma B.9] so we refer to [57] for more details.

²¹ See Equation 18 for more intuition on the following equation.

Black-box access to: $\mathcal{E}_4(\text{crs}_{\text{SIS}})$

1. Run $\mathcal{E}_4(\text{crs}_{\text{SIS}})$ with randomness $\rho \leftarrow R$ as follows: relay the $Q_{\text{H}_{\text{SIS}}} + 4$ random oracle queries to the random oracles and record all query-response pairs. Obtain (\mathcal{I}, y, v) . Set $i = I_3$ and let (μ_1, \dots, μ_τ) be the response to query i .
2. If $v = 0$, abort and return $v = 0$. Else, $(\vec{m}, \text{aux}) := (\mathcal{I}(\vec{m}), \mathcal{I}(\text{aux}))$.
3. Repeat:
 - (a) sample $(\mu'_1, \dots, \mu'_\tau) \leftarrow \hat{\mathcal{R}}_q^\tau$,
 - (b) run $\mathcal{E}_4(\text{crs}_{\text{SIS}})$ with randomness ρ as follows to obtain (\mathcal{I}', y', v') , aborting right after (or during) the initial run of $\mathcal{P}^*(\text{crs}_{\text{SIS}})$ if one of the following conditions holds:
 - $I'_3 \neq I_3$,
 - the input for some j -th issuing query to \mathcal{P}^* is of the form $(\mathbf{m}'_j, \mathbf{r}'_j)$ where $\vec{m} = \text{Coeffs}(\mathbf{m}'_j)_{\text{aux}}$, answer the query to i with $(\mu'_1, \dots, \mu'_\tau)$ while answering all the other queries consistently if the query was performed by \mathcal{E}_4 in the previous run, and with fresh random value otherwise,
 until a challenge $(\mu'_1, \dots, \mu'_\tau)$ with $v' = 1$ and $I'_3 = I_3$ has been found.
4. If y (resp. y') is a MSIS solution for the matrix $[\mathbf{A}_1 \mid \mathbf{A}_2]$, return \mathcal{I} , y (resp. y') and $v = 1$.
5. Parse $y = (\bar{\mathbf{s}}_1, \bar{\mathbf{s}}_2, \bar{\mathbf{m}} = (\bar{\mathbf{y}}_3, \bar{\mathbf{g}}), \bar{c})$ and $y' = (\bar{\mathbf{s}}'_1, \bar{\mathbf{s}}'_2, \bar{\mathbf{m}}' = (\bar{\mathbf{y}}'_3, \bar{\mathbf{g}}'), \bar{c}')$.
6. If $\bar{\mathbf{s}}_1 \neq \bar{\mathbf{s}}'_1$ or $\bar{\mathbf{s}}_2 \neq \bar{\mathbf{s}}'_2$, return \mathcal{I} , $\begin{bmatrix} \bar{c}\bar{c}'(\bar{\mathbf{s}}_1 - \bar{\mathbf{s}}'_1) \\ \bar{c}\bar{c}'(\bar{\mathbf{s}}_2 - \bar{\mathbf{s}}'_2) \end{bmatrix}$ as the MSIS solution to $[\mathbf{A}_1 \mid \mathbf{A}_2]$ and $v = 1$.
7. Parse $\bar{\mathbf{s}}_1 := (\bar{\mathbf{s}}, \bar{\mathbf{r}}, \bar{\mathbf{u}})$, $(\mathfrak{R}_0, \mathfrak{R}_1) := \text{H}_1(I_1), (\gamma_{i,j}) := \text{H}_2(I_2)$ and $I_3 := (3, \text{crs}_{\text{SIS}}, \mathbf{x}, \mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w}, \bar{z}_3, h_1, \dots, h_\tau)$.
8. If for $i = 1, 2, \dots, \tau$ we have

$$\begin{aligned}
 h_i &= g_i + \sum_{j=1}^{256} \gamma_{i,j} \cdot \left(\sigma(\mathbf{r}_j)^T \bar{\mathbf{s}}_1 + \sigma(\mathbf{e}_j)^T \bar{\mathbf{y}}_3 - z_{3,j} \right) \\
 &\quad + \sum_{j=1}^n \gamma_{i,256+j} \cdot \left(\sigma(\mathbf{p}_j)^T \bar{\mathbf{s}} - \sigma(\beta_j)^T \bar{\mathbf{u}} - m_{C,j} - \sigma(\mathbf{c}_{r,j})^T \bar{\mathbf{r}} \right) \\
 &\quad + \gamma_{i,256+n+1} \cdot \left(\sigma(\bar{\mathbf{s}})^T \bar{\mathbf{s}} - \mathcal{B}_s \right) + \gamma_{i,256+n+2} \cdot \left(\sigma(\bar{\mathbf{r}})^T \bar{\mathbf{r}} - \mathcal{B}_r \right) \\
 &\quad + \gamma_{i,256+n+3} \cdot \sigma(\bar{\mathbf{u}} - \mathbf{x})^T \bar{\mathbf{u}} \in \hat{\mathcal{R}}_q.
 \end{aligned}$$

return \mathcal{I} , $(\bar{\mathbf{s}}_1, \bar{\mathbf{s}}_2, \bar{\mathbf{m}}, \bar{c})$ and $v = 1$.

9. Otherwise, return $v = 0$.

Fig. 23: Subextractor $\mathcal{E}_3(\text{crs}_{\text{SIS}})$ as a $(Q_{\text{H}_{\text{SIS}}} + 4)$ -query random oracle algorithm.

Lemma 8.8. For any crs_{ISIS} , the extractor \mathcal{E}_3 makes an expected number of at most $6 + 5Q_{\text{H}_{\text{ISIS}}}$ queries to $\mathcal{P}^*(\text{crs}_{\text{ISIS}})$. Also, for uniformly random crs_{ISIS} , $\mathcal{E}_3(\text{crs}_{\text{ISIS}})$ outputs $v = 1$ with probability at least

$$\Pr[\mathbf{E}_3] - (Q_{\text{H}_{\text{ISIS}}} + 1) \cdot \left(\frac{2}{|\mathcal{C}|} + \hat{p}_1^{-\hat{d}/2} \right).$$

Next, let $\mathcal{I} = (I_1, I_2, I_3, I_4, I_5)$ be the index vector obtained by \mathcal{E}_3 in the first step and denote

$$(\mathfrak{R}_0, \mathfrak{R}_1) := \mathbf{H}_1(I_1), (\gamma_{i,j}) := \mathbf{H}_2(I_2), I_3 := (3, \text{crs}_{\text{ISIS}}, \mathbf{x}, \mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w}, \vec{z}_3, h_1, \dots, h_\tau).$$

Then, conditioned on $v = 1$, the extractor either returns $(\bar{\mathbf{s}}_1 := (\bar{\mathbf{s}}, \bar{\mathbf{r}}, \bar{\mathbf{u}}), \bar{\mathbf{s}}_2, \bar{\mathbf{m}}, \bar{c}) \in \mathcal{R}_q^{m_1} \times \mathcal{R}_q^{m_2} \times \mathcal{R}_q^{256/\hat{d}+\tau} \times \bar{\mathcal{C}}$ which satisfies the following relations:

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_y \\ \mathbf{t}_g \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \bar{\mathbf{s}}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B}_y \\ \mathbf{B}_g \end{bmatrix} \bar{\mathbf{s}}_2 + \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{y}}_3 \\ \bar{\mathbf{g}} \end{bmatrix}, \quad \|\bar{c}\bar{\mathbf{s}}_1\| \leq 2\mathcal{B}_1, \quad \|\bar{c}\bar{\mathbf{s}}_2\| \leq 2\mathcal{B}_2$$

and for $i = 1, 2, \dots, \tau$ we have²²

$$\begin{aligned} h_i &= g_i + \sum_{j=1}^{256} \gamma_{i,j} \cdot \left(\sigma(\mathbf{r}_j)^T \bar{\mathbf{s}}_1 + \sigma(\mathbf{e}_j)^T \bar{\mathbf{y}}_3 - z_{3,j} \right) \\ &\quad + \sum_{j=1}^n \gamma_{i,256+j} \cdot \left(\sigma(\mathbf{p}_j)^T \bar{\mathbf{s}} - \sigma(\beta_j)^T \bar{\mathbf{u}} - m_{C,j} - \sigma(\mathbf{c}_{r,j})^T \bar{\mathbf{r}} \right) \\ &\quad + \gamma_{i,256+n+1} \cdot \left(\sigma(\bar{\mathbf{s}})^T \bar{\mathbf{s}} - \mathcal{B}_s \right) + \gamma_{i,256+n+2} \cdot \left(\sigma(\bar{\mathbf{r}})^T \bar{\mathbf{r}} - \mathcal{B}_r \right) \\ &\quad + \gamma_{i,256+n+3} \cdot \sigma(\bar{\mathbf{u}} - \mathbf{x})^T \bar{\mathbf{u}} \in \hat{\mathcal{R}}_q, \end{aligned} \tag{41}$$

or a $\text{MSIS}_{n, m_1+m_2, \bar{\mathcal{B}}}$ solution for the matrix $[\mathbf{A}_1 \mid \mathbf{A}_2]$ where $\bar{\mathcal{B}} := 4\nu\sqrt{\mathcal{B}_1^2 + \mathcal{B}_2^2}$.

For the next step, we define the subextractor \mathcal{E}_2 which is informally responsible for extracting secret vectors $\bar{\mathbf{s}}, \bar{\mathbf{r}}, \bar{\mathbf{u}}$ that satisfy (42). The rough intuition for \mathcal{E}_2 is as follows. Suppose \mathcal{E}_2 runs \mathcal{E}_3 which outputs a witness \mathbf{w} which satisfies (41) (or a Module-SIS solution but then we are done). As before, since the ABDLOP commitment $(\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g)$ to the witness was sent before getting $(\gamma_{i,j})$ and the commitment itself is binding, this means that the extracted witness must be *independent* of the challenges $(\gamma_{i,j})$. Suppose the relevant part of \mathbf{w} does not satisfy (24). This implies that with probability at most $\hat{p}_1^{-\tau}$ all τ equations in (41) hold. Formally, the analysis for \mathcal{E}_2 is almost identical to [57, Lemma B.10], combined with Lemma 8.6, and thus we refer to the aforementioned result for more details.

Lemma 8.9. For any crs_{ISIS} , the extractor \mathcal{E}_2 makes an expected number of at most $12 + 11Q_{\text{H}_{\text{ISIS}}}$ queries to $\mathcal{P}^*(\text{crs}_{\text{ISIS}})$. Also, for uniformly random crs_{ISIS} , \mathcal{E}_2 outputs $v = 1$ with probability at least

$$\Pr[\mathbf{E}_3] - (Q_{\text{H}_{\text{ISIS}}} + 1) \cdot \left(\frac{2}{|\mathcal{C}|} + \hat{p}_1^{-\hat{d}/2} + \hat{p}_1^{-\tau} \right).$$

Next, let $\mathcal{I} = (I_1, I_2, I_3, I_4, I_5)$ be the index vector obtained by \mathcal{E}_2 in the first step and denote

$$(\mathfrak{R}_0, \mathfrak{R}_1) := \mathbf{H}_1(I_1), \quad I_2 := (2, \text{crs}_{\text{ISIS}}, \mathbf{x}, \mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w}, \vec{z}_3).$$

Then, conditioned on $v = 1$, the extractor either returns $(\bar{\mathbf{s}}_1 := (\bar{\mathbf{s}}, \bar{\mathbf{r}}, \bar{\mathbf{u}}), \bar{\mathbf{s}}_2, \bar{\mathbf{m}}, \bar{c}) \in \mathcal{R}_q^{m_1} \times \mathcal{R}_q^{m_2} \times \mathcal{R}_q^{256/\hat{d}+\tau} \times \bar{\mathcal{C}}$ which satisfies the following relations:

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_y \\ \mathbf{t}_g \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \bar{\mathbf{s}}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B}_y \\ \mathbf{B}_g \end{bmatrix} \bar{\mathbf{s}}_2 + \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{y}}_3 \\ \bar{\mathbf{g}} \end{bmatrix}, \quad \|\bar{c}\bar{\mathbf{s}}_1\| \leq 2\mathcal{B}_1, \quad \|\bar{c}\bar{\mathbf{s}}_2\| \leq 2\mathcal{B}_2$$

²² See the description of Equation 17 for more details.

Black-box access to: $\mathcal{E}_3(\text{crs}_{\text{ISIS}})$

1. Run $\mathcal{E}_3(\text{crs}_{\text{ISIS}})$ with randomness $\rho \leftarrow R$ as follows: relay the $Q_{\text{H}_{\text{ISIS}}} + 4$ random oracle queries to the random oracles and record all query-response pairs. Obtain (\mathcal{I}, y, v) . Set $i = I_2$ and let $(\gamma_{i,j})$ be the response to query i .
2. If $v = 0$, abort and return $v = 0$. Else, $(\vec{m}, \text{aux}) := (\mathcal{I}(\vec{m}), \mathcal{I}(\text{aux}))$.
3. Repeat:
 - (a) sample $(\gamma'_{i,j}) \leftarrow \mathbb{Z}_{\hat{q}}^{\tau \times (256+d+3)}$,
 - (b) run $\mathcal{E}_3(\text{crs}_{\text{ISIS}})$ with randomness ρ as follows to obtain (\mathcal{I}', y', v') , aborting right after (or during) the initial run of $\mathcal{P}^*(\text{crs}_{\text{ISIS}})$ if one of the following conditions holds:
 - $I'_2 \neq I_2$,
 - the input for some j -th issuing query to \mathcal{P}^* is of the form $(\mathbf{m}'_j, \mathbf{r}'_j)$ where $\vec{m} = \text{Coeffs}(\mathbf{m}'_j)_{\text{aux}}$, answer the query to i with $(\gamma'_{i,j})$ while answering all the other queries consistently if the query was performed by \mathcal{E}_4 in the previous run, and with fresh random value otherwise,
4. If y (resp. y') is a MSIS solution for the matrix $[\mathbf{A}_1 \mid \mathbf{A}_2]$, return \mathcal{I}, y (resp. y') and $v = 1$.
5. Parse $y = (\bar{\mathbf{s}}_1, \bar{\mathbf{s}}_2, \bar{\mathbf{m}} = (\bar{\mathbf{y}}_3, \bar{\mathbf{g}}), \bar{c})$ and $y' = (\bar{\mathbf{s}}'_1, \bar{\mathbf{s}}'_2, \bar{\mathbf{m}}' = (\bar{\mathbf{y}}'_3, \bar{\mathbf{g}}'), \bar{c}')$.
6. If $\bar{\mathbf{s}}_1 \neq \bar{\mathbf{s}}'_1$ or $\bar{\mathbf{s}}_2 \neq \bar{\mathbf{s}}'_2$, return $\mathcal{I}, \begin{bmatrix} \bar{c}\bar{c}'(\bar{\mathbf{s}}_1 - \bar{\mathbf{s}}'_1) \\ \bar{c}\bar{c}'(\bar{\mathbf{s}}_2 - \bar{\mathbf{s}}'_2) \end{bmatrix}$ as the MSIS solution to $[\mathbf{A}_1 \mid \mathbf{A}_2]$ and $v = 1$.
7. Parse $(\mathfrak{R}_0, \mathfrak{R}_1) := \text{H}_1(I_1)$ and $I_2 := (2, \text{crs}_{\text{ISIS}}, x, \mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w}, \vec{z}_3)$.
8. Parse $\bar{\mathbf{s}}_1 := (\bar{\mathbf{s}}, \bar{\mathbf{r}}, \bar{\mathbf{u}})$ and define $\bar{\mathbf{s}}'_1 := \text{Coeffs}(\bar{\mathbf{s}}_1)$, $\bar{\mathbf{s}} := \text{Coeffs}(\bar{\mathbf{s}})$, $\bar{\mathbf{r}} := \text{Coeffs}(\bar{\mathbf{r}})$, $\bar{\mathbf{u}} := \text{Coeffs}(\bar{\mathbf{u}})$
9. If all the following equations hold over $\mathbb{Z}_{\hat{q}}$:

$$\vec{z}_3 = \bar{y}_3 + (\mathfrak{R}_0 - \mathfrak{R}_1)\bar{\mathbf{s}}_1,$$

$$P\bar{\mathbf{s}} = B\bar{\mathbf{u}} + C \begin{bmatrix} \vec{m} \\ \bar{\mathbf{r}} \end{bmatrix},$$

$$\langle \bar{\mathbf{s}}, \bar{\mathbf{s}} \rangle = \mathcal{B}_s^2,$$

$$\langle \bar{\mathbf{r}}, \bar{\mathbf{r}} \rangle = \mathcal{B}_r^2,$$

$$\langle \bar{\mathbf{u}}, \bar{\mathbf{u}} - \vec{1} \rangle = 0,$$

return $\mathcal{I}, (\bar{\mathbf{s}}_1, \bar{\mathbf{s}}_2, \bar{\mathbf{m}}, \bar{c})$ and $v = 1$.

10. Otherwise, return $v = 0$.

Fig. 24: Subextractor $\mathcal{E}_2(\text{crs}_{\text{ISIS}})$ as a $(Q_{\text{H}_{\text{ISIS}}} + 4)$ -query random oracle algorithm.

and

$$\begin{cases} \vec{z}_3 = \vec{y}_3 + (\mathfrak{R}_0 - \mathfrak{R}_1) \vec{s}_1 \\ P\vec{s} = B\vec{u} + C \begin{bmatrix} \vec{m} \\ \vec{r} \end{bmatrix} \\ \langle \vec{s}, \vec{s} \rangle = \mathcal{B}_s^2 \\ \langle \vec{r}, \vec{r} \rangle = \mathcal{B}_r^2 \\ \langle \vec{u}, \vec{u} - \vec{1} \rangle = 0 \end{cases} \quad (42)$$

where

$$\vec{s}_1 := \text{Coeffs}(\bar{s}_1), \quad \vec{s} := \text{Coeffs}(\bar{s}), \quad \vec{r} := \text{Coeffs}(\bar{r}), \quad \vec{u} := \text{Coeffs}(\bar{u}),$$

or a MSIS $_{n, m_1+m_2, \bar{B}}$ solution for the matrix $[\mathbf{A}_1 \mid \mathbf{A}_2]$ where $\bar{B} := 4\nu\sqrt{\mathcal{B}_1^2 + \mathcal{B}_2^2}$.

Finally, we define the extractor \mathcal{E}_1 in Figure 25 which is informally responsible for extracting from the approximate range proof (c.f. Section 5.1). Intuitively, suppose \mathcal{E}_1 runs the subextractor \mathcal{E}_2 which outputs a witness $(\vec{s}_1 = (\vec{s}, \vec{r}, \vec{u}), \vec{y}_3)$ that satisfies (42) (or a Module-SIS solution), and in particular

$$\vec{z}_3 = \vec{y}_3 + \mathfrak{R} \vec{s}_1 \quad \text{and} \quad \|\vec{z}\| \leq \mathcal{B}_3$$

where $\mathfrak{R} := \mathfrak{R}_0 - \mathfrak{R}_1$. Identically as in the previous case, since the ABDLOP commitment $(\mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g)$ to the witness was sent before getting $(\mathfrak{R}_0, \mathfrak{R}_1)$ and the commitment is binding, the extracted witness must be independent of the challenges $(\mathfrak{R}_0, \mathfrak{R}_1)$. Now, suppose that $\|\vec{s}_1\| \geq \frac{\sqrt{2}}{\omega_{\min}(\lambda)} \cdot \mathcal{B}_3$. Then, by Lemma 5.4 only with negligible probability at most $2^{-\lambda}$ we can have

$$\|\vec{y}_3 + \mathfrak{R} \vec{s}_1 \bmod \hat{q}\| \leq \mathcal{B}_3.$$

In order to use the lemma, we have the requirement $\hat{q} > 16m_1\hat{d}\mathcal{B}_3$. Thus, with overwhelming probability we must have $\|\vec{s}_1\| \leq \frac{\sqrt{2}}{\omega_{\min}(\lambda)} \cdot \mathcal{B}_3$. If \mathcal{B}_3 is relatively small compared to the proof system modulus \hat{q} , then by combining with (40) we deduce the norm bounds on \vec{s}, \vec{r} and that \vec{u} has binary coefficients. For example, if $\mathcal{B}_3 < \frac{\omega_{\min}(\lambda)}{\sqrt{2}} \cdot \sqrt{\hat{q}}$ then

$$-\hat{q} < -\mathcal{B}_s^2 \leq \|\vec{s}\|^2 - \mathcal{B}_s^2 \leq \|\vec{s}_1\|^2 - \mathcal{B}_s^2 \leq \frac{2}{\omega_{\min}(\lambda)^2} \cdot \mathcal{B}_3^2 < \hat{q}$$

so we deduce that $\|\vec{s}\|^2 = \mathcal{B}_s^2$ over integers (and similarly for \vec{r} and \vec{u}).

Formally, the analysis for \mathcal{E}_1 is almost identical to [57, Lemma B.11], combined with Lemma 8.6, so we skip the proof.

Lemma 8.10. *For any crs_{SIS} , the extractor $\mathcal{E}_1(\text{crs}_{\text{SIS}})$ makes an expected number of at most $24 + 23Q_{\text{H}_{\text{SIS}}}$ queries to $\mathcal{P}^*(\text{crs}_{\text{SIS}})$. Also, for uniformly random crs_{SIS} , \mathcal{E}_1 outputs $v = 1$ with probability at least*

$$\delta := \Pr[\text{E}_3] - (Q_{\text{H}_{\text{SIS}}} + 1) \cdot \left(\frac{2}{|\mathcal{C}|} + \hat{p}_1^{-\hat{d}/2} + \hat{p}_1^{-\tau} + 2^{-\lambda} \right).$$

Next, let $\mathcal{I} = (I_1, I_2, I_3, I_4, I_5)$ be the index vector obtained by \mathcal{E}_2 in the first step and denote $I_1 := (1, \text{crs}_{\text{SIS}}, x, \mathbf{t}_A, \mathbf{t}_y, \mathbf{t}_g, \mathbf{w})$. Then, conditioned on $v = 1$, the extractor either returns $(\bar{s}_1 := (\bar{s}, \bar{r}, \bar{u}), \bar{s}_2, \bar{\mathbf{m}}, \bar{c}) \in \mathcal{R}_q^{m_1} \times \mathcal{R}_q^{m_2} \times \mathcal{R}_q^{256/\hat{d}+\tau} \times \bar{\mathcal{C}}$ which satisfies the following relations:

$$\begin{bmatrix} \mathbf{t}_A \\ \mathbf{t}_y \\ \mathbf{t}_g \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \bar{s}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B}_y \\ \mathbf{B}_g \end{bmatrix} \bar{s}_2 + \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{y}}_3 \\ \bar{\mathbf{g}} \end{bmatrix}, \quad \|\bar{c}\bar{s}_1\| \leq 2\mathcal{B}_1, \quad \|\bar{c}\bar{s}_2\| \leq 2\mathcal{B}_2$$

Black-box access to: $\mathcal{E}_2(\text{crs}_{\text{ISIS}})$

1. Run $\mathcal{E}_2(\text{crs}_{\text{ISIS}})$ with randomness $\rho \leftarrow R$ as follows: relay the $Q_{\text{H}_{\text{ISIS}}} + 4$ random oracle queries to the random oracles and record all query-response pairs. Obtain (\mathcal{I}, y, v) . Set $i = I_1$ and let $(\mathfrak{R}_0, \mathfrak{R}_1)$ be the response to query i .
2. If $v = 0$, abort and return $v = 0$. Else, $(\vec{m}, \text{aux}) := (\mathcal{I}(\vec{m}), \mathcal{I}(\text{aux}))$.
3. Repeat:
 - (a) sample $(\mathfrak{R}'_0, \mathfrak{R}'_1) \leftarrow \{0, 1\}^{256 \times m_1 \bar{d}} \times \{0, 1\}^{256 \times m_1 \bar{d}}$,
 - (b) run $\mathcal{E}_2(\text{crs}_{\text{ISIS}})$ with randomness ρ as follows to obtain (\mathcal{I}', y', v') , aborting right after (or during) the initial run of \mathcal{P}^* if one of the following conditions holds:
 - $I'_1 \neq I_1$,
 - the input for some j -th issuing query to \mathcal{P}^* is of the form $(\mathbf{m}'_j, \mathbf{r}'_j)$ where $\vec{m} = \text{Coeffs}(\mathbf{m}'_j)_{\text{aux}}$, answer the query to i with $(\mathfrak{R}'_0, \mathfrak{R}'_1)$ while answering all the other queries consistently if the query was performed by \mathcal{E}_4 in the previous run, and with fresh random value otherwise,
4. If y (resp. y') is a MSIS solution for the matrix $[\mathbf{A}_1 \mid \mathbf{A}_2]$, return \mathcal{I} , y (resp. y') and $v = 1$.
5. Parse $y = (\bar{\mathbf{s}}_1, \bar{\mathbf{s}}_2, \bar{\mathbf{m}} = (\bar{\mathbf{y}}_3, \bar{\mathbf{g}}), \bar{c})$ and $y' = (\bar{\mathbf{s}}'_1, \bar{\mathbf{s}}'_2, \bar{\mathbf{m}}' = (\bar{\mathbf{y}}'_3, \bar{\mathbf{g}}'), \bar{c}')$.
6. If $\bar{\mathbf{s}}_1 \neq \bar{\mathbf{s}}'_1$ or $\bar{\mathbf{s}}_2 \neq \bar{\mathbf{s}}'_2$, return \mathcal{I} , $\begin{bmatrix} \bar{c}\bar{c}'(\bar{\mathbf{s}}_1 - \bar{\mathbf{s}}'_1) \\ \bar{c}\bar{c}'(\bar{\mathbf{s}}_2 - \bar{\mathbf{s}}'_2) \end{bmatrix}$ as the MSIS solution to $[\mathbf{A}_1 \mid \mathbf{A}_2]$ and $v = 1$.
7. Define $\vec{s}_1 := \text{Coeffs}(\bar{\mathbf{s}}_1)$
8. If $\|\vec{s}_1\| \leq \frac{\sqrt{2}}{\omega_{\min}(\lambda)} \cdot \mathcal{B}_3$, return \mathcal{I} , $(\bar{\mathbf{s}}_1, \bar{\mathbf{s}}_2, \bar{\mathbf{m}}, \bar{c})$ and $v = 1$.
9. Otherwise, return $v = 0$.

Fig. 25: Subextractor $\mathcal{E}_1(\text{crs}_{\text{ISIS}})$ as a $(Q_{\text{H}_{\text{ISIS}}} + 4)$ -query random oracle algorithm.

and

$$\begin{cases} \|\vec{s}_1\| \geq \frac{\sqrt{2}}{\omega_{\min}(\lambda)} \cdot \mathcal{B}_3 \\ P\vec{s} = B\vec{u} + C \begin{bmatrix} \vec{m} \\ \vec{r} \end{bmatrix} \\ \langle \vec{s}, \vec{s} \rangle = \mathcal{B}_s^2 \\ \langle \vec{r}, \vec{r} \rangle = \mathcal{B}_r^2 \\ \langle \vec{u}, \vec{u} - \vec{1} \rangle = 0 \end{cases} \quad (43)$$

where

$$\vec{s}_1 := \text{Coeffs}(\bar{\mathbf{s}}_1), \quad \vec{s} := \text{Coeffs}(\bar{\mathbf{s}}), \quad \vec{r} := \text{Coeffs}(\bar{\mathbf{r}}), \quad \vec{u} := \text{Coeffs}(\bar{\mathbf{u}}),$$

or a MSIS $_{n, m_1 + m_2, \bar{\mathbf{B}}}$ solution for the matrix $[\mathbf{A}_1 \mid \mathbf{A}_2]$ where $\bar{\mathbf{B}} := 4\nu\sqrt{\mathcal{B}_1^2 + \mathcal{B}_2^2}$.

For the next observation, we define the following events parameterised by crs_{ISIS} :

$$\begin{aligned} \text{Accept}_{\text{crs}_{\text{ISIS}}} &:= [(\mathcal{I}, \text{tr}, v, \mathcal{Q}) \leftarrow \mathcal{P}^*(\text{crs}_{\text{ISIS}}) \wedge v = 1] \\ \text{NoBadQuery}_{\text{crs}_{\text{ISIS}}} &:= \left[\begin{array}{l} (\mathcal{I}, \text{tr}, v, \mathcal{Q}) \leftarrow \mathcal{P}^*(\text{crs}_{\text{ISIS}}) \wedge v = 1 \\ \wedge \mathcal{E}_1(\text{crs}_{\text{ISIS}}) \text{ never responds to any issuing query} \\ \text{from } \mathcal{A} \text{ on input } (\mathbf{m}', \cdot) \text{ s.t. } \mathcal{I}(\vec{m}) = \text{Coeffs}(\mathbf{m}')_{\mathcal{I}(\text{aux})} \end{array} \right] \end{aligned}$$

Lemma 8.11. *For any crs_{ISIS} , we have*

$$\Pr [\text{Accept}_{\text{crs}_{\text{ISIS}}} \wedge \text{NoBadQuery}_{\text{crs}_{\text{ISIS}}}] = \Pr [\text{Accept}_{\text{crs}_{\text{ISIS}}}] .$$

Proof. It follows in the similar vein as the proof of Lemma 8.6. Indeed, if $(\mathcal{I}, \text{tr}, 1, \mathcal{Q}) \leftarrow \mathcal{P}^*(\text{crs}_{\text{ISIS}})$ and $(\vec{m}, \text{aux}) := (\mathcal{I}(\vec{m}), \mathcal{I}(\text{aux}))$ then by the conditions described with blue text for the extractors from Figures 22 to 25, the extractor \mathcal{E}_1 never answers to any issuing query on input of the form (\mathbf{m}', \cdot) which satisfies $\vec{m} = \text{Coeffs}(\mathbf{m}')_{\text{aux}}$. We call such a query a *bad query*.

To illustrate our reasoning more precisely, we show that $\mathcal{E}_3(\text{crs}_{\text{ISIS}})$ never answers a bad query. The case for \mathcal{E}_2 and \mathcal{E}_1 follows similarly. In Step 1 of \mathcal{E}_3 , it runs $\mathcal{E}_4(\text{crs}_{\text{ISIS}})$. Clearly, \mathcal{E}_4 never responds to a bad query by construction. Now, if $\text{Accept}_{\text{crs}_{\text{ISIS}}}$ holds then Step 2 of \mathcal{E}_3 passes.

Further, we consider the loop in Step 3 of \mathcal{E}_3 . That is, the algorithm samples uniformly random challenges (μ_i) and runs $\mathcal{E}_4(\text{crs}_{\text{ISIS}})$ with the same random coins, but with reprogrammed random oracle on input I_3 . So now, we look at \mathcal{E}_4 . It runs again \mathcal{P}^* (with the reprogrammed random oracle) but by the blue condition in Step 3(b) of \mathcal{E}_3 , if \mathcal{P}^* gets a bad query, we abort the current run of \mathcal{E}_4 and go back to Step 3(a) of \mathcal{E}_3 . Otherwise, \mathcal{P}^* outputs a tuple $(\mathcal{I}', \text{tr}', v', \mathcal{Q}')$. By the condition in Step 3(b) of \mathcal{E}_3 , if $I'_3 \neq I_3$ then we go back to Step 3(a) of \mathcal{E}_3 . We also abort the current execution of \mathcal{E}_4 if $v' = 0$ by Step 2 of \mathcal{E}_4 .

Suppose we have not aborted the execution of \mathcal{E}_4 yet, so we are currently in the loop Step 3 of \mathcal{E}_4 . Since $I'_3 = I_3$, this implies that $\mathcal{I}'(\vec{m}) = \vec{m}$ and $\mathcal{I}'(\text{aux}) = \text{aux}$. Therefore, Step 3(c) of \mathcal{E}_4 makes sure that \mathcal{P}^* does not answer any bad query. This reasoning applies to each iteration of the loop in Step 3 of \mathcal{E}_3 . Hence, we conclude that \mathcal{E}_3 never responds to a bad query assuming $\text{Accept}_{\text{crs}_{\text{ISIS}}}$ holds. \square

Now, we are ready to define our strict polynomial-time reduction from Game_3 to Int-NTRU-ISIS_f and MSIS. The reduction is given the instances of the aforementioned problems and hardwires them in Game_3 . Namely, the Int-NTRU-ISIS_f challenges $(a_1, \mathbf{a}_2, \mathbf{c}_0, \mathbf{c}_1)$ will be put inside the public key ipk , while the MSIS challenge matrix $[\mathbf{A}_1 \ \mathbf{A}_2]$ will be programmed in the crs_{ISIS} . Then, the reduction runs $\mathcal{E}_1(\text{crs}_{\text{ISIS}})$ but it aborts after $2(24 + 23Q_{\text{H}_{\text{ISIS}}})/\delta$ queries to \mathcal{P}^* . Further, when there is any issuing query on input (\mathbf{m}, \mathbf{r}) that \mathcal{E}_1 must answer, the reduction makes a preimage query (\mathbf{m}, \mathbf{r}) to the Int-NTRU-ISIS_f oracle $\mathcal{O}_{\text{pre}}(\mathbf{m}, \mathbf{r})$. This implies that \mathcal{B} makes at most $O\left(\frac{Q_{\text{Issue}}Q_{\text{H}_{\text{ISIS}}}}{\delta}\right)$ preimage queries and by Lemma 8.11, the reduction never makes a bad query. Since δ is non-negligible, the reduction runs in strict polynomial time. Finally, by Lemma 2.14, the probability that the reduction outputs a valid Int-NTRU-ISIS_f or MSIS solution is at least $\delta/2$. \square

8.3 Concrete Instantiation

In this section, we propose concrete parameters to instantiate the anonymous credentials from the construction above where we aim for 128-bit (classical) security. Note that one can directly obtain a blind signature by simply setting the number of attributes l to be one. Then, the anonymity and one-more unforgeability map to the blindness and one-more unforgeability properties of the blind signature respectively.

We split the instantiation procedure into three parts: (i) the core construction, (ii) the proof of knowledge $\Pi_{\text{NIZK}}^{\text{ISIS}}$, (iii) multi-proof extractable $\Pi_{\text{NIZK}}^{\text{Com}}$. As in many prior works, we measure the hardness of MSIS and MLWE by the root Hermite factor (we refer to [37, Section 3.2.4] for more explanation on the methodology and the references therein) where e.g. $\delta \approx 1.0045$ corresponds to 128 bits of security. We bound the number of (random oracle and signing) queries by 2^{64} . Further, we assume hardness of the lattice problems with respect to expected polynomial time adversaries the same as for strict polynomial time algorithms (which is the case for currently the most efficient lattice attacks). In particular, this means that we do not need to apply Lemma 2.14 to the knowledge extractors which would additionally result with worse parameters.

Core construction. We propose the parameters in Table 3. Namely, we set $(d, q) = (4096, \approx 2^{34})$ where q is a prime congruent to 5 modulo 8. We provide a construction for $l = 8$ attributes. One then puts each attribute inside the hash function H_M which outputs vectors of length $h = 512$ with coefficients between -2 and 2 . For correctness, we select $\delta = 412$ (it will have a bigger meaning in setting parameters for one-more unforgeability). As for anonymity, we require $\text{MLWE}_{1, \ell_r, \psi}$ to be hard. With the parameters proposed, the root Hermite factor is 1.00115 which corresponds to around more than 600 bits of security. We describe correctness and zero-knowledge of the proof systems separately below.

In the remainder, we focus on one-more unforgeability (c.f. Theorem 8.3) which is one of the most challenging tasks. To begin with, note that the reduction for one-more unforgeability runs in time $\text{poly}(\lambda) \cdot 1/\varepsilon$,

parameter	value
d	4096
q	$17179861781 \approx 2^{34}$
N	2^{640}
t	640
h	512
l	8
ℓ_r	2
ℓ_m	1
m	3
ψ	2
s	$\approx 2^{27}$
δ	284
κ	56
α	12
M	3
T_{\max}	678

Table 3: Parameters for the main construction. More explanation behind these variables can be found in Table 2 and Theorem 7.3.

where ε is the advantage of the adversary. Since we set $\varepsilon = 2^{-128}$, we now need to consider 256-bit security for one-more unforgeability. To this end, let us first focus on all the distinguishing advantages between the security hybrids. With our parameters, we have

$$\frac{|Q_{H_M}|^2}{(2\psi + 1)^h} < 2^{-1100}$$

and thus we ignore this part. If we assume 257-bit security for CRS indistinguishability (Game_2) and 259-bit security for straight-line extractability (Game_3), then we aim for around 259-bit security for Game_3 . Further, we will set the “knowledge error” term in Lemma 8.10 to be around 2^{-260} . Thus, we end up with 260-bit security for either the Int-NTRU-ISIS_f or the MSIS problems. The latter one will be analysed together with $\Pi_{\text{NIZK}}^{\text{ISIS}}$ so we focus on Int-NTRU-ISIS_f. To this end, we look at Theorem 7.3 where the bound on the number of queries is $Q := Q_{\text{Issue}} \cdot Q_{H_{\text{ISIS}}} = 2^{128}$.

We instantiate Theorem 7.3 with parameters from Table 3. We analyse each subtracted term in the inequality. First, with our parameters the $\text{MLWE}_{1,m+1,\psi}$ problem has the root Hermite factor 1.0005 which corresponds to more than 1000 bits of security. Thus, we loosely bound:

$$\frac{\ell_m + \ell_r}{6Q} \text{Adv}_{1,m+1,\kappa}^{\text{MLWE}}(\mathcal{B}) < 2^{-1000}.$$

For the next term, we had set $T_{\max} = 678$, so that $(1 - 1/M)^{T_{\max}} < 2^{-396}$. Thus, the term $2^\lambda/6$ in Theorem 7.3 can be written alternatively as $(1 - 1/M)^{T_{\max}}/6 < 2^{-398}$. Further, we can bound the rest of the terms with our parameters:

$$\begin{aligned} \frac{T_{\max}^2 Q}{12N} &\approx 2^{-493}, & \frac{\ell_m}{3} \cdot \left(\frac{q}{(2\kappa + 1)^{m+2}} \right)^d &\approx 2^{-414}, \\ \left(Q - \frac{2}{3} \right) T_{\max} \left(\frac{\varepsilon}{2M} + \frac{2\varepsilon}{M} + 2^{-\delta+1} \right) &\approx 2^{-394}. \end{aligned}$$

Hence, the advantage of the reduction to the standard NTRU-ISIS_f is at least

$$2^{-260}/6Q - (2^{-1000} + 2^{-398} + 2^{-493} + 2^{-414} + 2^{-394}) \approx 2^{-390}.$$

parameter	explanation	value
\hat{d}	ring dimension of $\hat{\mathcal{R}}$	128
\hat{p}_1	the smallest prime divisor of \hat{q}_2	17179861781
\hat{q}_2	proof system modulus divisible by q	$\approx 2^{88}$
ξ	infinity norm on the challenges in \mathcal{C}	17
ν	parameter introduced in (13)	350
τ	number of repetitions for soundness	10
n	length of \mathbf{t}_A	30
m_2	length of the randomness vector \mathbf{s}_2	68

Table 4: Parameters for the proof system $\Pi_{\text{NIZK}}^{\text{ISIS}}$.

parameter	explanation	value
\hat{d}	ring dimension of $\hat{\mathcal{R}}$	128
\hat{p}_1	the smallest prime divisor of \hat{q}_1	16253
\hat{p}_2	second prime divisor of \hat{q}_1	17179861781
\hat{q}_1	proof system modulus divisible by q	$\approx 2^{48}$
n	length of \mathbf{t}_A	18
m_2	length of the randomness vector \mathbf{s}_2	51
\bar{p}	prime modulus used for LHC	68527215
n_1	width of the matrix $\hat{\mathbf{A}}_1$ from LHC	28
n_2	width of the matrix $\hat{\mathbf{A}}_2$ from LHC	28
η_1	infinity norm of the encryption randomness $\hat{\mathbf{e}}_{1,j}$ in LHC	1
η_2	infinity norm of the encryption randomness $\hat{\mathbf{e}}_{1,j}$ in LHC	1

Table 5: Parameters for the multi-proof extractable $\Pi_{\text{NIZK}}^{\text{Com}}$.

We heuristically assume that NTRU-ISIS_f is as hard as the standard $\text{MSIS}_{1,m+2,\mathcal{B}}$ (or rather NTRU-SIS [69]). Recall that the MSIS bound is

$$\mathcal{B} = \mathfrak{s}\sqrt{(m+2) \cdot d} + \mathcal{B}_m \psi d \sqrt{(\ell_m + \ell_r)(m+2)} \approx 2^{33.83} < q.$$

As for the computational hardness, the root Hermite factor is 1.001425 which corresponds to more than 490 bits of security.

Parameters for $\Pi_{\text{NIZK}}^{\text{ISIS}}$. As described above, we pick parameters so that the underlying proof system satisfies correctness, zero-knowledge (with 130-bit security) and knowledge soundness (with 260-bit security). To this end, we adapt the SAGE script provided in [57, Section 6.1] to compute the proof size produced by $\Pi_{\text{NIZK}}^{\text{ISIS}}$. We further optimise the proof by applying small tricks from [57], such as the Dilithium-G compression [32] and reducing the number of garbage terms h_i by a factor of two, as well as the recent bimodal Gaussian technique described in [56, Section 3]. This allows us to moderately reduce the standard deviations related to rejection sampling.

We set the ring dimension \hat{d} to be 128. Further, if we bound the number of oracle queries by 2^{64} , then the knowledge error term (excluding the random oracle query bound) in Lemma 8.10 must be smaller than $\approx 2^{-324}$. We set $\hat{p}_1 = q \approx 2^{88}$, so $\tau = 10$. As for the challenge space \mathcal{C} , in order to accommodate the 324 bits of security, we experimentally pick $(\xi, \nu) = (17, 350)$. This gives us a challenge space of size at least 2^{327} . As for the approximate range proofs, we computed $\omega_{\min}(192) = 1$ and $\omega_{\max}(192) = \sqrt{530}$ using the cumulative distribution function of chi-squared distribution with 256 degrees of freedom, as discussed in [40]. Finally, the underlying computational assumption, i.e. MSIS , must hold with around 324-bit security. We satisfy this condition by picking $\hat{q} \approx 2^{88}$ and $n = 30$. The reason for such a large modulus, which we believe is the main bottleneck, is the requirement on \hat{q} described in Theorem 8.3.

Multi-proof extractable $\Pi_{\text{NIZK}}^{\text{Com}}$. We propose the parameters in Table 5. Recall that we need to satisfy correctness, CRS indistinguishability (with 257-bit security), zero-knowledge (with 130-bit security) and straightline extractability (with 259-bit security). Similarly as above, we apply the parameter selection strategy from [57, 56]. We additionally need to take into account the extractable linear homomorphic commitment (LHC) which is required for Katsumata transform [46].

We first set parameters so that the underlying proof system [57] works²³. Then, we pick parameters related to the MLWE-based LHC exactly as in [46, Lemma 3.13]. First, we need large enough prime \bar{p} such that we can encrypt both \mathbf{z}_1 and \mathbf{z}_2 . Hence, we set \hat{p} such that

$$\max \left(\mathfrak{s}_1 \sqrt{2m_1 \hat{d}}, \mathfrak{s}_2 \sqrt{2m_2 \hat{d}} \right) \leq (\bar{p} - 1)/4.$$

Then, we need to ensure that LHC provides no decryption errors. Consider the application where it is used, i.e. Theorem 6.3. To this end, we require for $i \in \{1, 2\}$ (using notation from there):

$$\begin{aligned} & \|(\bar{\mathbf{c}}_{i,2} + \bar{\mathbf{w}}_{i,2}) - \bar{\mathbf{D}}_{i,1}(\bar{\mathbf{c}}_{i,1} + \bar{\mathbf{w}}_{i,1})\|_\infty \\ &= \|\bar{p} \cdot (\bar{\mathbf{D}}_{i,2} \bar{\mathbf{z}}_{i,1} - \bar{\mathbf{D}}_{i,1} \bar{\mathbf{z}}_2 + \bar{\mathbf{z}}_3) + \mathbf{z}_i\| \\ &\leq \bar{p} \cdot \left(\sqrt{n_i d} \eta_i \cdot \bar{\mathfrak{s}}_i \sqrt{2n_i d} + \sqrt{m_i d} \eta_i \cdot \bar{\mathfrak{s}}_i \sqrt{2m_i d} + \bar{\mathfrak{s}}_i \sqrt{2m_i d} \right) \\ &\leq \sqrt{2\bar{p}} \left(n_i d \eta_i + m_i d \eta_i + \sqrt{m_i d} \right) \bar{\mathfrak{s}}_i < q/2 \end{aligned} \quad (44)$$

and the lower-bounds on the standard deviations $\bar{\mathfrak{s}}_i$ are dictated by the correctness and zero-knowledge conditions in Section 6.2. Recall that we need the proof system modulus large enough to prove the norm bounds, i.e. it has to be $(\psi \sqrt{(\ell_m + \ell_r)d})^2 \cdot O(\lambda)$. However, this condition is much milder than the one in (44). What are still unknown from the inequality above are the values of η_i and n_i . These need to be chosen such that $\text{MLWE}_{m_i, m_i + n_i, \eta_i}$ is hard in order to provide both zero-knowledge and CRS indistinguishability. The latter one becomes the bottleneck since we aim for 257-bit security. We found that it is possible to instantiate the construction, and in particular satisfy all the aforementioned conditions, with the proof system modulus $q \approx 2^{48}$. Next, we pick $\eta_1 = \eta_2 = 1$ which results in setting $n_1 = n_2 = 28$ for the Module-LWE problem.

In order to keep the rejection rate small, we additionally apply the bimodal Gaussian technique from [56]. Note that, similarly as the randomness vector \mathbf{s}_2 for the ABDLOP commitment, the randomness vectors $\bar{\mathbf{e}}_{i,j}$ related to encrypting \mathbf{s}_i are freshly created *every time* the prover algorithm is called²⁴. Hence, leaking one bit of information on $\bar{\mathbf{e}}_{i,j}$ should not significantly decrease security of the protocol. Based on the observation above, we can apply bimodal Gaussian rejection as described in [56, Section 3.1]. with this modification we can still prove zero-knowledge but under a new assumption, called Extended-MLWE. Hence, by combining this observation with [56], we can simply set $\alpha_1, \alpha_2, \alpha_3, \bar{\alpha}, \bar{\alpha}_2$ all to be one. Thus, the repetition rate of the $\Pi_{\text{NIZK}}^{\text{Com}}$ is only 12.

As for straight-line extractability, we require the knowledge error term (without the random oracle query bound) in (34) to be at most 2^{-324} . To this end, we select parameters similarly as for $\Pi_{\text{NIZK}}^{\text{ISIS}}$, e.g. the challenge space or the bounds for approximate range proofs.

Conclusion. Finally, we are ready to present various signature sizes in Table 6 (see $\text{NTRU-ISIS}_f.\text{AnonCreds}$). The public key ipk size is a single polynomial in \mathcal{R}_q which is a public output of NTRU.TrapGen . All other parts of ipk can be generated from a seed. Hence, we obtain the public key of size $d \log q / 2^{13} \approx 4096 \cdot 34 / 2^{13} = 17\text{KB}$.

As for the communication complexity, the proof size of π is around 398KB, about 80% of which is linked to LHC. We highlight that one can further reduce the proof overhead, which comes with extractable linear homomorphic commitments, by applying the NTRU-type construction from [46, Section 3.5]. Then, as estimated in the aforementioned paper, it would reduce the overhead by around a factor of two, which

²³ We note that we additionally apply the Dilithium compression techniques [57, Appendix A] which were not explicitly described in Section 6.

²⁴ In other words, Com in Fig. 11 is a one-time commitment [60].

scheme	signature	communication complexity	public key size
NTRU-ISIS _f .AnonCreds	243KB	473KB	17KB
Int-NTRU-ISIS _f .AnonCreds	62KB	107KB	3.5KB

Table 6: Signature, communication and public key sizes for the concrete instantiations of our anonymous credentials.

parameter	value
d	1024
\hat{d}	128
q	$33641 \approx 2^{15}$
\hat{q}_1	$\approx 2^{41}$
\hat{q}_2	$\approx 2^{50}$
N	2^{256}
t	256
h	128
l	8
ℓ_r	2
ℓ_m	1
ψ	2
s	658
\mathcal{B}	$2^{14.86}$

Table 7: Parameters for our alternative variant Int-NTRU-ISIS_f.AnonCreds.

gives roughly less than 150KB. However, security of the protocol would rely on so-called *decisional small matrix ratio* assumption, which does not yet seem to be a well-studied assumption. Since our construction itself already relies on new assumptions (such as NTRU-ISIS_f or Extended-MLWE [56]), we stick with the more standard but less efficient instantiation of LHC. By counting the sizes of (s, x) sent by the signer, which is around 75KB, the overall communication size is ≈ 473 KB. Finally, the signature size, which is the proof produced from Π_{NIZK} , is around 243KB.

Reducing to Int-NTRU-ISIS_f only. There are a few artefacts of the security proof which make the parameters, and consequently the signature size, unappealing. The main bottleneck lies in the non-tight reduction from Int-NTRU-ISIS_f to NTRU-ISIS_f in Theorem 7.3. Let us start with the parameter m which is the length of \mathbf{a}_2 . Namely, it allows us to do a game hop where we switch from uniformly random matrix $[\mathbf{c}_0^T | \mathbf{c}_1^T]$ to the one computed as $[\mathbf{c}_0^T | \mathbf{c}_1^T] = [a_1 | \mathbf{a}_2^T | 1] \mathbf{D}$ for a short random matrix \mathbf{D} . Note that if $m = 0$ then this reasoning would not go through. In other words, we cannot reduce distinguishing $[\mathbf{c}_0^T | \mathbf{c}_1^T] = [a_1 | 1] \mathbf{D}$ from a uniformly random one to Module-LWE. The reason is that in the first case we know the trapdoor for a_1 but not in the latter one. However, it is completely non-trivial whether the scheme becomes insecure when $m = 0$. This “artificial” problem has already appeared in the literature in the context of lattice-based group signatures [68, 61]. Another issue, which requires *long enough* m , is the witness indistinguishability argument of Theorem 7.3. Namely, we want to make sure that with an overwhelming probability, for a uniformly random vector \mathbf{d}_0 with coefficients in $[-\kappa, \kappa]$, there exists another one \mathbf{d}_1 from the same set such that $[a_1 | \mathbf{a}_2^T | 1] \mathbf{d}_0 = [a_1 | \mathbf{a}_2^T | 1] \mathbf{d}_1$. This requires us to pick $m > \frac{\log q}{\log(2\kappa+1)} - 2$ which has significant impact on the parameters. Not to mention the fact that one has to deal with the security loss Q from Theorem 7.3.

In order to investigate the limits of our construction, we propose an alternative scheme (which we call Int-NTRU-ISIS_f.AnonCreds in Table 6) where we ignore the issues pointed out above. In short, we heuristically

assume that the reduced Int-NTRU-ISIS_f is as hard as the Module-SIS problem for the matrix $[A \ C]$ and the norm bound $\mathcal{B} := \sqrt{\mathcal{B}_s^2 + \mathcal{B}_m^2}$. In particular, we can set $m = 0$.

We set parameters following the guidelines from the previous instantiation, and we summarise our selection in Table 7. For convenience, we aim for only 124-bit security, which allows us to pick smaller ring dimension $d = 1024$. In this case, we need to ensure 254-bit security for the Int-NTRU-ISIS_f, i.e. the MSIS problem. Then, we need to make sure the “knowledge error” term (c.f. Lemmas 6.5 and 8.10) is at most 2^{-318} to accommodate for the factor of $Q_{\text{HISIS}} = 2^{64}$, which comes from the Fiat-Shamir transform.

Expected polynomial-time Multi-Extract. Unfortunately, even in our milder variant we still need to account for the security loss depending on the adversary’s advantage ε , since the reduction sub-algorithm, i.e. Multi-Extract runs in time $\text{poly}(\lambda)/\varepsilon$. Because we already assume hardness of Module-SIS against expected PPT adversaries, it makes sense to make Multi-Extract expected polynomial time as well. Namely, in Figure 15, we allow the main loop to run until it either outputs the correct witness w or the weak opening, or it runs out of challenges to pick. The key observation is that the Multi-Extract algorithm is *only* run by the reduction if the initial proof is valid. Hence, by the heavy rows argument [28], or its more fine-grained version in [6, Lemma 4], we deduce that the expected runtime of Multi-Extract is $\text{poly}(\lambda)$ and does not depend on ε anymore. In particular, we can reduce the bit security by 128.

For fair comparison with certain related works, e.g. [3, 13], we ignore the Fiat-Shamir transform loss Q_{HISIS} which was also not explicitly considered in the aforementioned constructions. Thus, the signature is a proof size for $\Pi_{\text{NIZK}}^{\text{ISIS}}$ with ≈ 128 -bit security. In Table 1 we provide the signature sizes for both cases: either reducing to NTRU-ISIS_f or Int-NTRU-ISIS_f using the same methodology as described throughout this section.

Acknowledgments

We would like to thank the reviewers for the useful feedback. This work is supported by the EU H2020 ERC Project 101002845 PLAZA.

References

- [1] Shweta Agrawal, Dan Boneh, and Xavier Boyen. “Lattice Basis Delegation in Fixed Dimension and Shorter-Ciphertext Hierarchical IBE”. In: *CRYPTO*. 2010, pp. 98–115.
- [2] Shweta Agrawal, Damien Stehlé, and Anshu Yadav. “Towards Practical and Round-Optimal Lattice-Based Threshold and Blind Signatures”. In: *IACR Cryptol. ePrint Arch.* (2021), p. 381.
- [3] Shweta Agrawal et al. *Practical, Round-Optimal Lattice-Based Blind Signatures*. 2022.
- [4] Miklós Ajtai. “Generating Hard Instances of Lattice Problems (Extended Abstract)”. In: *STOC*. 1996, pp. 99–108.
- [5] Martin R. Albrecht et al. “Lattice-Based SNARKs: Publicly Verifiable, Preprocessing, and Recursively Composable”. In: *CRYPTO*. Lecture Notes in Computer Science. 2022.
- [6] Thomas Attema, Serge Fehr, and Michael Klooß. “Fiat-Shamir Transformation of Multi-round Interactive Proofs”. In: *TCC (1)*. Vol. 13747. Lecture Notes in Computer Science. Springer, 2022, pp. 113–142.
- [7] Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. “Practical Product Proofs for Lattice Commitments”. In: *CRYPTO (2)*. Vol. 12171. Lecture Notes in Computer Science. Springer, 2020, pp. 470–499.
- [8] Man Ho Au, Willy Susilo, and Yi Mu. “Constant-Size Dynamic k-TAA”. In: *IACR Cryptol. ePrint Arch.* (2008), p. 136. URL: <http://eprint.iacr.org/2008/136>.
- [9] Wojciech Banaszczyk. “New bounds in some transference theorems in the geometry of numbers”. In: *Mathematische Annalen* 296.1 (1993), pp. 625–635. ISSN: 1432-1807. DOI: 10.1007/BF01445125. URL: <https://doi.org/10.1007/BF01445125>.
- [10] Carsten Baum et al. “More Efficient Commitments from Structured Lattice Assumptions”. In: *SCN*. 2018, pp. 368–385.
- [11] Mihir Bellare, Daniele Micciancio, and Bogdan Warinschi. “Foundations of Group Signatures: Formal Definitions, Simplified Requirements, and a Construction Based on General Assumptions”. In: *EUROCRYPT*. 2003, pp. 614–629.

- [12] Ward Beullens and Gregor Seiler. “LaBRADOR: Compact Proofs for R1CS from Module-SIS”. In: *IACR Cryptol. ePrint Arch.* (2022), p. 1341.
- [13] Ward Beullens et al. *Lattice-Based Blind Signatures: Short, Efficient, and Round-Optimal*. Cryptology ePrint Archive, Paper 2023/077. <https://eprint.iacr.org/2023/077>. 2023. URL: <https://eprint.iacr.org/2023/077>.
- [14] Dan Boneh, Xavier Boyen, and Hovav Shacham. “Short Group Signatures”. In: *IACR Cryptol. ePrint Arch.* (2004), p. 174. URL: <http://eprint.iacr.org/2004/174>.
- [15] Jonathan Bootle, Vadim Lyubashevsky, and Gregor Seiler. “Algebraic Techniques for Short(er) Exact Lattice-Based Zero-Knowledge Proofs”. In: *CRYPTO (1)*. Vol. 11692. Lecture Notes in Computer Science. Springer, 2019, pp. 176–202.
- [16] Jonathan Bootle et al. “LWE Without Modular Reduction and Improved Side-Channel Attacks Against BLISS”. In: *ASIACRYPT (1)*. Vol. 11272. Lecture Notes in Computer Science. Springer, 2018, pp. 494–524.
- [17] Joppe W. Bos et al. “CRYSTALS - Kyber: A CCA-Secure Module-Lattice-Based KEM”. In: *2018 IEEE European Symposium on Security and Privacy, EuroS&P*. 2018, pp. 353–367.
- [18] Cecilia Boschini, Akira Takahashi, and Mehdi Tibouchi. *MuSig-L: Lattice-Based Multi-Signature With Single-Round Online Phase*. Cryptology ePrint Archive, Paper 2022/1036. <https://eprint.iacr.org/2022/1036>. 2022. URL: <https://eprint.iacr.org/2022/1036>.
- [19] Katharina Boudgoust et al. “Towards Classical Hardness of Module-LWE: The Linear Rank Case”. In: *ASIACRYPT (2)*. Vol. 12492. Lecture Notes in Computer Science. Springer, 2020, pp. 289–317.
- [20] Jan Camenisch, Manu Drijvers, and Anja Lehmann. “Anonymous Attestation Using the Strong Diffie Hellman Assumption Revisited”. In: *TRUST*. Vol. 9824. Lecture Notes in Computer Science. Springer, 2016, pp. 1–20.
- [21] Jan Camenisch, Markulf Kohlweiss, and Claudio Soriente. “An Accumulator Based on Bilinear Maps and Efficient Revocation for Anonymous Credentials”. In: *Public Key Cryptography - PKC 2009, 12th International Conference on Practice and Theory in Public Key Cryptography, Irvine, CA, USA, March 18-20, 2009. Proceedings*. Ed. by Stanislaw Jarecki and Gene Tsudik. Vol. 5443. Lecture Notes in Computer Science. Springer, 2009, pp. 481–500. DOI: 10.1007/978-3-642-00468-1_27. URL: https://doi.org/10.1007/978-3-642-00468-1_27.
- [22] Jan Camenisch, Markulf Kohlweiss, and Claudio Soriente. “Solving Revocation with Efficient Update of Anonymous Credentials”. In: *Security and Cryptography for Networks, 7th International Conference, SCN 2010, Amalfi, Italy, September 13-15, 2010. Proceedings*. Ed. by Juan A. Garay and Roberto De Prisco. Vol. 6280. Lecture Notes in Computer Science. Springer, 2010, pp. 454–471. DOI: 10.1007/978-3-642-15317-4_28. URL: https://doi.org/10.1007/978-3-642-15317-4_28.
- [23] Jan Camenisch and Anna Lysyanskaya. “A Signature Scheme with Efficient Protocols”. In: *Security in Communication Networks, Third International Conference, SCN 2002, Amalfi, Italy, September 11-13, 2002. Revised Papers*. Ed. by Stelvio Cimato, Clemente Galdi, and Giuseppe Persiano. Vol. 2576. Lecture Notes in Computer Science. Springer, 2002, pp. 268–289. DOI: 10.1007/3-540-36413-7_20. URL: https://doi.org/10.1007/3-540-36413-7_20.
- [24] Jan Camenisch and Anna Lysyanskaya. “An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation”. In: *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*. Ed. by Birgit Pfitzmann. Vol. 2045. Lecture Notes in Computer Science. Springer, 2001, pp. 93–118. DOI: 10.1007/3-540-44987-6_7. URL: https://doi.org/10.1007/3-540-44987-6_7.
- [25] Jan Camenisch and Anna Lysyanskaya. “Dynamic Accumulators and Application to Efficient Revocation of Anonymous Credentials”. In: *Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 2002, Proceedings*. Ed. by Moti Yung. Vol. 2442. Lecture Notes in Computer Science. Springer, 2002, pp. 61–76. DOI: 10.1007/3-540-45708-9_5. URL: https://doi.org/10.1007/3-540-45708-9_5.
- [26] David Chaum. “Security Without Identification: Transaction Systems to Make Big Brother Obsolete”. In: *Commun. ACM* 28.10 (1985), pp. 1030–1044. DOI: 10.1145/4372.4373. URL: <https://doi.org/10.1145/4372.4373>.
- [27] Lidong Chen. “Access with Pseudonyms”. In: *Cryptography: Policy and Algorithms, International Conference, Brisbane, Queensland, Australia, July 3-5, 1995, Proceedings*. Ed. by Ed Dawson and Jovan Dj. Golic. Vol. 1029. Lecture Notes in Computer Science. Springer, 1995, pp. 232–243. DOI: 10.1007/BFb0032362. URL: <https://doi.org/10.1007/BFb0032362>.
- [28] Ivan Damgård. *On Σ -Protocols*. Lecture on Cryptologic Protocol Theory; Faculty of Science, University of Aarhus. 2010.

- [29] Ivan Damgård. “Payment Systems and Credential Mechanisms with Provable Security Against Abuse by Individuals”. In: *Advances in Cryptology - CRYPTO '88, 8th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1988, Proceedings*. Ed. by Shafi Goldwasser. Vol. 403. Lecture Notes in Computer Science. Springer, 1988, pp. 328–335. DOI: 10.1007/0-387-34799-2_26. URL: https://doi.org/10.1007/0-387-34799-2_26.
- [30] Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. “Efficient Identity-Based Encryption over NTRU Lattices”. In: *ASIACRYPT*. 2014, pp. 22–41.
- [31] Léo Ducas and Phong Q. Nguyen. “Learning a Zonotope and More: Cryptanalysis of NTRUSign Countermeasures”. In: *ASIACRYPT*. 2012, pp. 433–450.
- [32] Leo Ducas et al. *CRYSTALS – Dilithium: Digital Signatures from Module Lattices*. Cryptology ePrint Archive, Report 2017/633. <https://ia.cr/2017/633>. “The Dilithium-G” scheme can be found in the June 2017 version of this report. The direct url is <https://eprint.iacr.org/eprint-bin/getfile.pl?entry=2017/633&version=20170627:201152&file=633.pdf>. 2017.
- [33] Léo Ducas et al. “CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme”. In: *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2018.1 (2018), pp. 238–268.
- [34] Muhammed F. Esgin, Ngoc Khanh Nguyen, and Gregor Seiler. “Practical Exact Proofs from Lattices: New Techniques to Exploit Fully-Splitting Rings”. In: *ASIACRYPT (2)*. 2020, pp. 259–288.
- [35] Muhammed F. Esgin et al. “Lattice-Based Zero-Knowledge Proofs: New Techniques for Shorter and Faster Constructions and Applications”. In: *CRYPTO (1)*. Springer, 2019, pp. 115–146.
- [36] Muhammed F. Esgin et al. “MatRiCT: Efficient, Scalable and Post-Quantum Blockchain Confidential Transactions Protocol”. In: *CCS*. ACM, 2019, pp. 567–584.
- [37] Muhammed Fethullah Esgin. “Practice-Oriented Techniques in Lattice-Based Cryptography”. In: (2020). DOI: 10.26180/5eb8f525b3562. URL: https://bridges.monash.edu/articles/thesis/Practice-Oriented_Techniques_in_Lattice-Based_Cryptography/12279728.
- [38] Amos Fiat and Adi Shamir. “How to Prove Yourself: Practical Solutions to Identification and Signature Problems”. In: *CRYPTO*. 1986, pp. 186–194.
- [39] Marc Fischlin. “Round-Optimal Composable Blind Signatures in the Common Reference String Model”. In: *CRYPTO*. Vol. 4117. Lecture Notes in Computer Science. Springer, 2006, pp. 60–77.
- [40] Craig Gentry, Shai Halevi, and Vadim Lyubashevsky. “Practical Non-interactive Publicly Verifiable Secret Sharing with Thousands of Parties”. In: *IACR Cryptol. ePrint Arch.* (2021), p. 1397.
- [41] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. “Trapdoors for hard lattices and new cryptographic constructions”. In: *STOC*. 2008, pp. 197–206.
- [42] Gottfried Herold and Alexander May. “LP Solutions of Vectorial Integer Subset Sums - Cryptanalysis of Galbraith’s Binary Matrix LWE”. In: *Public Key Cryptography (1)*. Vol. 10174. Lecture Notes in Computer Science. Springer, 2017, pp. 3–15.
- [43] Hyperledger Foundation. *Hyperledger Aries*. <https://www.hyperledger.org/use/aries>. [Online; accessed 06-October-2022].
- [44] Hyperledger Foundation. *Hyperledger Indy*. <https://www.hyperledger.org/use/hyperledger-indy>. [Online; accessed 06-October-2022].
- [45] Corentin Jeudy, Adeline Roux-Langlois, and Olivier Sanders. *Lattice Signature with Efficient Protocols, Application to Anonymous Credentials*. Cryptology ePrint Archive, Paper 2022/509. <https://eprint.iacr.org/2022/509>. 2022. URL: <https://eprint.iacr.org/2022/509>.
- [46] Shuichi Katsumata. “A New Simple Technique to Bootstrap Various Lattice Zero-Knowledge Proofs to QROM Secure NIZKs”. In: *CRYPTO (2)*. Vol. 12826. Lecture Notes in Computer Science. Springer, 2021, pp. 580–610.
- [47] Qiqi Lai et al. *Lattice-based Commit-Transferrable Signatures and Applications to Anonymous Credentials*. Cryptology ePrint Archive, Paper 2023/766. <https://eprint.iacr.org/2023/766>. 2023. URL: <https://eprint.iacr.org/2023/766>.
- [48] Adeline Langlois and Damien Stehlé. “Worst-case to average-case reductions for module lattices”. In: *Designs, Codes and Cryptography* 75.3 (2015), pp. 565–599.
- [49] Benoît Libert et al. “Practical “Signatures with Efficient Protocols” from Simple Assumptions”. In: *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security, AsiaCCS 2016, Xi’an, China, May 30 - June 3, 2016*. Ed. by Xiaofeng Chen, XiaoFeng Wang, and Xinyi Huang. ACM, 2016, pp. 511–522. DOI: 10.1145/2897845.2897898. URL: <https://doi.org/10.1145/2897845.2897898>.
- [50] Tobias Looker et al. *The BBS Signature Scheme*. <https://www.ietf.org/id/draft-looker-cfrg-bbs-signatures-01.html>. [Online; accessed 06-October-2022]. 2022.

- [51] Anna Lysyanskaya et al. “Pseudonym Systems”. In: *Selected Areas in Cryptography, 6th Annual International Workshop, SAC'99, Kingston, Ontario, Canada, August 9-10, 1999, Proceedings*. Ed. by Howard M. Heys and Carlisle M. Adams. Vol. 1758. Lecture Notes in Computer Science. Springer, 1999, pp. 184–199. DOI: 10.1007/3-540-46513-8\14. URL: <https://doi.org/10.1007/3-540-46513-8\14>.
- [52] Vadim Lyubashevsky. “Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures”. In: *ASIACRYPT*. 2009, pp. 598–616.
- [53] Vadim Lyubashevsky. “Lattice Signatures Without Trapdoors”. In: *EUROCRYPT*. 2012, pp. 738–755.
- [54] Vadim Lyubashevsky and Daniele Micciancio. *Asymptotically Efficient Lattice-Based Digital Signatures*. Cryptology ePrint Archive, Paper 2013/746. <https://eprint.iacr.org/2013/746>. 2013. URL: <https://eprint.iacr.org/2013/746>.
- [55] Vadim Lyubashevsky and Daniele Micciancio. “Asymptotically Efficient Lattice-Based Digital Signatures”. In: *J. Cryptol.* 31.3 (2018). Preliminary version appeared in TCC 2008, pp. 774–797. URL: <https://eprint.iacr.org/2013/746>.
- [56] Vadim Lyubashevsky and Ngoc Khanh Nguyen. *BLOOM: Bimodal Lattice One-Out-of-Many Proofs and Applications*. Cryptology ePrint Archive, Paper 2022/1307. <https://eprint.iacr.org/2022/1307>. 2022. URL: <https://eprint.iacr.org/2022/1307>.
- [57] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. “Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General”. In: *IACR Cryptol. ePrint Arch.* (2022). Appears in Crypto 2022, p. 284.
- [58] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. “Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General”. In: *CRYPTO (2)*. Vol. 13508. Lecture Notes in Computer Science. Springer, 2022, pp. 71–101.
- [59] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. “SMILE: Set Membership from Ideal Lattices with Applications to Ring Signatures and Confidential Transactions”. In: *CRYPTO (2)*. Vol. 12826. Lecture Notes in Computer Science. Springer, 2021, pp. 611–640.
- [60] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. “Shorter Lattice-Based Zero-Knowledge Proofs via One-Time Commitments”. In: *Public Key Cryptography (1)*. Springer, 2021, pp. 215–241.
- [61] Vadim Lyubashevsky et al. “Shorter Lattice-Based Group Signatures via “Almost Free” Encryption and Other Optimizations”. In: *ASIACRYPT (4)*. Springer, 2021, pp. 218–248.
- [62] MATTR. *MATTR*. <https://github.com/mattrglobal>. [Online; accessed 06-October-2022].
- [63] MATTR. *bbs-signatures*. <https://github.com/mattrglobal/bbs-signatures>. [Online; accessed 06-October-2022].
- [64] Daniele Micciancio and Chris Peikert. “Trapdoors for Lattices: Simpler, Tighter, Faster, Smaller”. In: *EUROCRYPT*. 2012, pp. 700–718.
- [65] Daniele Micciancio and Oded Regev. “Worst-Case to Average-Case Reductions Based on Gaussian Measures”. In: *SIAM J. Comput.* 37.1 (2007), pp. 267–302.
- [66] NFCW. *Digital identity market revenues to reach US\$53bn in 2026*. <https://www.nfcw.com/2022/01/31/375825/digital-identity-market-revenues-to-reach-us53bn-in-2026/>. [Online; accessed 06-October-2022]. 2022.
- [67] Rafael del Pino and Shuichi Katsumata. *A New Framework For More Efficient Round-Optimal Lattice-Based (Partially) Blind Signature via Trapdoor Sampling*. Cryptology ePrint Archive, Paper 2022/834. <https://eprint.iacr.org/2022/834>. 2022. URL: <https://eprint.iacr.org/2022/834>.
- [68] Rafaël del Pino, Vadim Lyubashevsky, and Gregor Seiler. “Lattice-Based Group Signatures and Zero-Knowledge Proofs of Automorphism Stability”. In: *ACM Conference on Computer and Communications Security*. ACM, 2018, pp. 574–591.
- [69] Thomas Prest et al. *FALCON*. Tech. rep. <https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions>. National Institute of Standards and Technology, 2017.
- [70] Oded Regev. “On lattices, learning with errors, random linear codes, and cryptography”. In: *J. ACM* 56.6 (2009).
- [71] The AnonCreds Specification Working Group. *The AnonCreds Specification*. <https://github.com/AnonCreds-WG/anoncreds-spec>. [Online; accessed 06-October-2022]. 2022.
- [72] Veramo. *Veramo core development*. <https://github.com/uport-project>. [Online; accessed 06-October-2022].
- [73] Hoeteck Wee and David J. Wu. *Succinct Vector, Polynomial, and Functional Commitments from Lattices*. Cryptology ePrint Archive, Paper 2022/1515. <https://eprint.iacr.org/2022/1515>. 2022. URL: <https://eprint.iacr.org/2022/1515>.

- [74] Rupeng Yang et al. “Efficient Lattice-Based Zero-Knowledge Arguments with Standard Soundness: Construction and Applications”. In: *CRYPTO (1)*. Springer, 2019, pp. 147–175.