





MyOPE: Malicious securitY for Oblivious Polynomial Evaluation

Malika Izabachène¹ , Anca Nitulescu² ,
Paola de Perthuis^{1,3} , and David Pointcheval³ 

¹ Cosmian, Paris, France

² Protocol Labs, Paris, France

³ DIENS, École normale supérieure, CNRS, Inria, PSL University, Paris, France

Abstract. Oblivious Polynomial Evaluation (OPE) schemes are interactive protocols between a sender with a private polynomial and a receiver with a private evaluation point where the receiver learns the evaluation of the polynomial in their point and no additional information. In this work, we introduce MyOPE, a “short-sighted” non-interactive polynomial evaluation scheme with a poly-logarithmic communication complexity in the presence of malicious senders. In addition to strong privacy guarantees, MyOPE enforces honest sender behavior and consistency by adding verifiability to the calculations.

The main building block for this new verifiable OPE is an inner product argument (IPA) over rings that guarantees an inner product relation holds between committed vectors. Our IPA works for vectors with elements from generic rings of polynomials and has constant-size proofs that consist in one commitment only while the verification, once the validity of the vector-commitments has been checked, consists in one quadratic equation only.

We further demonstrate the applications of our IPA for verifiable OPE using Fully Homomorphic Encryption (FHE) over rings of polynomials: we prove the correctness of an inner product between the vector of powers of the evaluation point and the vector of polynomial coefficients, along with other inner-products necessary in this application’s proof.

MyOPE builds on generic secure encoding techniques for succinct commitments, that allow real-world FHE parameters and Residue Number System (RNS) optimizations, suitable for high-degree polynomials.

1 Introduction

1.1 Oblivious Polynomial Evaluation

Oblivious Polynomial Evaluation (OPE) is a protocol that allows two parties, the sender and the receiver, to evaluate a polynomial $f(X)$ of fixed public degree N secretly chosen by the sender in a point m known only by the receiver. The receiver obtains the value $f(m)$ without learning anything else about the polynomial f and without giving the sender any information about the point m .

OPE is an important building block for various 2-party computation (2-PC) schemes that generally require multiple executions of an OPE protocol for the same polynomial and different evaluation points, such as for Private Set Intersection (PSI), data mining [LP00], privacy-preserving keyword search [JL09], set membership (related to PSI) or and RSA key generation [Gil99], to mention a few. Our building blocks for OPE can also be used to attain Symmetric Private Information Retrieval (SPIR). However, the standard definition of receiver privacy does not preclude the sender from cheating by using a polynomial of higher degree than expected, changing the polynomial between multiple executions, or sending polynomial evaluation results in a wrong order, thus potentially easily leaking private information if the protocol includes a step to return the intersection result to the sender in the PSI application. Therefore, extending the security to malicious senders is essential in practical contexts.

With generic 2PC techniques in the malicious setting, some efficiency overheads incur, with the cut-and-choose technique or an expensive preprocessing to generate correlated randomness, that needs to be regularly repeated, and for the techniques using FHE ciphertexts to get better asymptotic communication and less interaction, this level of security had not yet been guaranteed.

1.2 Inner-Product Arguments over Rings

An inner product argument ensures the correctness for an inner product evaluation between two committed vectors. Following the observation that a polynomial evaluation can be written as an inner product between the vector of its coefficients and the vector of the consecutive powers of the evaluation point, we can use such protocols to enforce honest behavior in OPE schemes.

Inner-product arguments [BCC⁺16, BBB⁺18, WTs⁺18, LMR19, BCMS20, DRZ20, BMM⁺21] are core components of many other primitives, including zero-knowledge proofs and polynomial/vector commitment schemes. While all of these IPAs follow the same strategy using folding techniques, they only achieve logarithmic-size proofs without privacy, and they support only inner product of vectors with elements from a field. While the verification time is linear in most of the mentioned works, the recent results by [DRZ20, BMM⁺21] achieve logarithmic-time verification using a trusted setup.

We design a new ring Inner-Product Argument (ring-IPA) that allows the verification of inner product evaluations for vectors with elements from a ring. Our scheme relies on a trusted setup in order to achieve succinct proofs and verification. We also offer a new commitment scheme for vectors with elements from generic rings of polynomials, compatible with FHE ciphertexts, as an improvement to previous such schemes that only work for vectors over fields or groups of elliptic curve elements.

Our ring-IPA construction improves on proof sizes and verification times achieving constant-size proofs and constant-time verification, independent of the size of the vectors. Moreover, the ring-IPA can be used for vectors over rings of polynomials, compatible with the rings used by various FHE schemes, providing proofs for the evaluation of inner products over ciphertexts, which can then also be seen as ciphertexts of inner products, from their homomorphic properties.

1.3 Related Work

Despite its broad applicability, the study of the OPE functionality includes few practical and secure protocols, initiated in [NP99] and further continued in works like [CL01, ZB05, HL09].

While [NP99] proposed a first construction for OPE, it relies on a newly introduced intractability assumption: the noisy polynomial interpolation. Naor and Pinkas conjectured that it could be reduced to a more widely studied assumption, the polynomial reconstruction problem. Nevertheless, as shown in [BN00], this conjecture seems not to hold in general.

OPE Schemes with Active Security. Among recent OPE schemes, to the best of our knowledge, some of the best schemes with security against malicious adversaries are given by [HL09, Haz18]. However, [HL09] has at least 17 rounds of interaction and the parties send each other $\mathcal{O}(\lambda N)$ Paillier encryptions, where λ is the security parameter and N the degree of the polynomial, and their claimed efficiency holds only for sufficiently low degree polynomials. [Haz18] shows an OPE scheme for polynomial evaluation in the exponent of a DLog group using algebraic Pseudo-Random Functions (PRF). They focus on improving the computational efficiency of [HL09] by reducing the number of modular exponentiations, and removing the trusted setup requirement, while preserving the same number of rounds of interaction and communication complexity as in [HL09], and apply their scheme to private set membership 2PC. [PRTY20] gives malicious security for PSI with symmetric set sizes, but the communication is linear in the set sizes. There are also efficient schemes like [BC22], but they don't have the sublinear communication complexity and reduced interactivity we get with FHE methods.

Verifiable Computation (VC). Introduced by [GGP10], VC schemes are cryptographic systems that enable checking the integrity of results from delegated computations. More recent works [FNP20, BCFK21] have improved the efficiency of VC schemes to work for computations over encrypted data. These schemes, however, require proving the entire FHE circuit evaluation which is very

expensive. Moreover, they neither allow using practical parameters for the FHE scheme, nor speedups through classical optimizations such as Residue Number System (RNS).

Then, [GNS21] also gives constructions for general algebraic circuits over ring elements, but using a general approach which is not optimized for our inner-product and OPE application, which we can instantiate with a small cleartext modulus $t = 3$ (when theirs are greater to have big enough ideal subsets), and for which we select parameters compatible with OPE requirements like the computation privacy, when this other work is focused on non-private algebraic circuit calculations, and would thus not hide the polynomial.

1.4 Our Contribution

Our first contribution is a generic framework based on any secure encoding that allows building an inner product argument (IPA) over vectors with elements coming from wider spaces, not only from fields as defined in prior works. Depending on the instantiation of the underlying secure encoding scheme, we are able to obtain IPA schemes for vectors of ciphertexts from Fully Homomorphic Encryption (FHE) schemes, such as the Fan-Vercauteren scheme [FV12], that can provide privacy under the Ring-LWE (RLWE) assumption. Other FHE schemes relying on the RLWE assumption could also be used.

Equipped with our new constant-size and constant-time inner product argument, we next apply our techniques to enhance the security of OPE schemes to malicious senders, by enforcing an honest behavior when evaluating the polynomial. We focus on OPE schemes with minimal communication requirements and without an offline pre-processing phase, based on Fully Homomorphic Encryption (FHE) on an encrypted point.

More precisely, we introduce MyOPE, a scheme for verifiable oblivious evaluation of polynomials of high degrees N that achieves $\mathcal{O}((\log(N))^2)$ communication cost, for a constant security level (when considering the privacy, correctness, and soundness properties). This sublinear communication improves on the state-of-the-art, with just a $\log N$ factor with respect to schemes without an active security.

As a straightforward use case, we illustrate our verifiable OPE application to Private Set Intersection (PSI) in the unbalanced-set setting: basically, the sender, who owns the larger set $\mathcal{X} = \{x_1, \dots, x_N\}$, defines the polynomial $f = \prod_{i=1}^N (X - x_i)$, and the receiver asks for evaluations $f(y_j)$ for each element in $\mathcal{Y} = \{y_1, \dots, y_K\}$ to detect the common roots. Our communication complexity then becomes $\mathcal{O}(K \cdot (\log(N))^2)$, with just a $\log N$ factor compared to the most efficient PSI algorithm [CLR17] without verifiability, where K is the size of the small set and N is the size of the large set.

Our scheme guarantees the client's privacy with post-quantum FHE ciphertexts, under the Ring-LWE assumption, while the soundness of the proof can rely on a variety of secure encoding schemes, from pairings (under the Power Knowledge of Exponent Assumption) to any linear-only encryption scheme such as the Paillier scheme [Pai99] (under the integer factoring) and the Castagnos-Laguillaumie scheme [CL15] (under some class group problems), but also based on the post-quantum Learning With Errors (LWE) problem [GMNO18], thus making the entire scheme post-quantum secure.

1.5 Technical Overview

Inner Product Argument. Our first tool is of independent interest: it allows to prove inner-product evaluation of two vectors \vec{u} and \vec{v} , with respect to their commitments U and \bar{V} respectively. Our vector commitments will be based on commitments keys that encode powers of a secret point s , using any secure encoding scheme, $[1], [s], \dots, [s^{n-1}]$, defined and published once for all. We define $U = [\sum_i u_i s^i] = \sum_i u_i [s^i]$ while \bar{V} will be in the reverse order: $\bar{V} = \sum_i v_i [s^{n-i}]$. The notation $[.]$ is an informal representation of a secure encoding, that leads to a computationally binding commitment. The hiding property is achieved with an additional secret component

and Schnorr-like proofs. It will additionally have bilinear properties, which lead to $U \times \bar{V} = \sum_{i,j} (u_i v_j) [s^{n+i-j}] = \langle \vec{u}, \vec{v} \rangle [s^n] + \sum_{i \neq j} (u_i v_j) [s^{n+i-j}]$. By showing $U \times \bar{V} - \alpha [s^n]$ has no term in $[s^n]$, one proves that $\langle \vec{u}, \vec{v} \rangle = \alpha$. Whereas the analysis will be performed on polynomials evaluated in a secret point s , the Schwartz-Zippel lemma [Sch80, Zip79] will lift the relations on the polynomials, as non-zero polynomials are unlikely evaluated to 0. But this assumes good properties for the algebraic structures, which are not always satisfied. A favorable situation is considered first, in fields of large characteristic, then more complex structures are discussed in appendix C, in rings. Globally, the soundness of the proofs relies on the secure encodings (which can require the Power Knowledge of Exponent Assumption when pairings are used, or the linear-only property of any encryption scheme) and the Schwartz-Zippel lemma (that requires no computational assumption, but appropriate alg

Verifiable Polynomial Evaluation. Now, from a polynomial $f = \sum f_i X^i$, which can be encoded as a vector $\vec{v} = (f_i)_i$, and thus committed as \bar{V} (in reverse order), and an element m , which can be encoded as a vector $\vec{u} = (m^i)_i$, and thus committed as U , $f(m) = \langle \vec{u}, \vec{v} \rangle$, correctness can be proven as above with respect to the binding commitments U and \bar{V} . Again, efficiency will depend on the actual algebraic structures. For the reader's convenience, indices for vectors start at 0, to consider the constant monomial in polynomials.

Receiver Privacy: Fully Homomorphic Encryption. When both f and m are public, which easily allows getting confidence in U and \bar{V} , the above approach is convincing. When receiver privacy is expected, with a private point m , the receiver could send encryptions M_i of the m^i under an additively homomorphic encryption scheme: if W is a commitment of the vector $\vec{w} = (M_i)_i$, using the linear property of the encryption scheme, $\langle \vec{w}, \vec{v} \rangle = \sum f_i M_i$ is the encryption of $f(m) = \sum f_i m^i$, which can be proven with respect to W and \bar{V} , as above. Once convinced, the receiver can decrypt it to get $f(m)$. Their privacy is guaranteed by the semantic security of the encryption scheme. But sending all the M_i 's implies a huge communication cost. One can then use a Fully Homomorphic Encryption (FHE) to let the sender generate the vector: the receiver provides M as an encryption (under their own key) of m , and FHE allows the sender to compute M_i . The additive homomorphism allows continuing as above.

But why should the receiver trust the sender to have correctly computed the M_i 's as the encryptions of the successive powers of m and the commitment W correctly? Let us assume each M_i , in \vec{w} committed in W , encrypts the plaintext m_i . One can use another verifiable inner product to check $m_i = m^i$: from a random common public element n , chosen after the publication of W , and the vector $\vec{z} = (n^i)_i$ publicly committed into \bar{Z} (in reverse order), the receiver can verifiably compute the inner product $\langle \vec{w}, \vec{z} \rangle$ with respect to W and \bar{Z} . Since it is proven correct, it decrypts to $\sum m_i n^i$, while one would like it to be $\sum m^i n^i$. In case of equality, this means that polynomials $\sum m_i X^i$ and $\sum m^i X^i$ evaluate the same way on the random point n . By applying the Schwartz-Zippel lemma in the plaintext-space, this means that $m_i = m^i$ with overwhelming probability. This concludes the security analysis.

To draw the above conclusion, we assumed the Schwartz-Zippel lemma could be applied in both the plaintext-space (for getting $m_i = m^i$) and the ciphertext-space (for verifying the two inner product evaluations between $(M_i)_i$ and $(f_i)_i$, and between $(M_i)_i$ and $(n^i)_i$). Furthermore, for effective application of FHE, additional constraints might be added to the ciphertext-space, such as Residue Number System (RNS) representation. All these questions will be addressed below, dealing with arbitrary rings.

Sender Privacy: Noise-Flooding and Hiding Commitments. OPE also expects sender privacy, with no leakage about the polynomial f . However, the commitment U might leak some information, unless it guarantees the additional *hiding* property. Furthermore, the homomorphic evaluation $f(m)$ in the ciphertexts may leak more than just the result, and possibly the evaluation steps,

as the final noise in $\langle \vec{w}, \vec{v} \rangle$ leaks them. To avoid such a leakage, the sender can add extra super-polynomial noise to $\langle \vec{w}, \vec{v} \rangle$. This is the so-called *noise-flooding* technique [Gen09]. One will of course have to prove this does not impact the decrypted result, requiring a small enough norm for the added noise. Which results in another inner-product proof, as the L_2 -norm is an inner product square root.

Efficiency for two Secure Encoding Instantiations. Any Secure Encoding scheme can be used in our construction: depending on the instantiation choice different assumptions will be used, and if the Secure Encoding scheme uses a modulus which is not a multiple of q , then the size will have a $\mathcal{O}(\log(N))$ complexity. We compare two example instantiations in table 1.

Secure Encoding	Pairings	Paillier Encryption
Assumptions	Power Knowledge of Exponent	Hardness of Integer Factoring + Linear-Only Encryption
Field/Ring Modulus q	prime	composite (for RNS)
Size complexity	$\mathcal{O}(1)$	$\mathcal{O}(\log(N))$

Fig. 1. Comparison of pairing and Paillier encryption instantiations of the secure encoding scheme. The size complexity (of both the individual encodings and the total proof) is given in N , the size of the inner-product vectors.

Efficiency in Practice. With the Fan-Vercauteren FHE scheme [FV12], from the plaintext ring $\mathcal{R}_t = \mathbb{Z}_t[X]/r(X)$, where $r(X) = X^n + 1$, into the ciphertext ring $\mathcal{R}_q = \mathbb{Z}_q[X]/r(X)$, the core parameters are the integers n, q , and t , to guarantee the semantic security under the Ring-LWE assumption. For the PSI application, one needs to encode elements from \mathcal{X} and \mathcal{Y} into \mathcal{R}_t . Since $n > 128$, $t = 3$ will be big enough. Each ciphertext is $2n \log_2 q$ bits long.

With a MyOPE instantiation for polynomials of degree $N = 2^{30}$, we set $n = 2^{14}$ which leads to q over 610 bits to obtain an appropriate FHE semantic security (according to the LWE estimator [APS15]) and decryption correctness, including with noise-flooding. To exploit RNS optimizations [BEHZ16], as proposed in SEAL⁴, q can be the product of 11 primes on less than 60 bits each.

Then, the size of the FHE ciphertext to be sent is of about 3 MBytes. The result of the sender and proof of their honest behavior consists of about 5 MBytes, for a prime q and soundness of 2^{-128} . If RNS is used, with a composite q , the Schwartz-Zippel guarantees are lower, and our commitments will have to be repeated several times for the soundness: the result and proof then consists of some 170MBytes, from our analysis in appendix F.

2 Preliminaries

2.1 MyOPE: Verifiable OPE

We first formalize the notion of *verifiable oblivious polynomial evaluation*. A MyOPE = (OPE.Setup, OPE.KeyGen, OPE.QueryGen, OPE.Compute, OPE.Verify, OPE.Decode) scheme for polynomial evaluation consists of the following algorithms:

OPE.Setup(1^λ) \rightarrow (PK, SK): Given the security parameter λ , output a pair of keys independent on polynomials to compute. The public key PK will be provided as input to all the subsequent algorithms.

OPE.KeyGen(PK, f) \rightarrow pk_f : Given the polynomial f , output a public key pk_f .

OPE.QueryGen(PK, pk_f, x) \rightarrow σ_x : Given the public key pk_f and the input x , encode the evaluation point x into σ_x , and output it.

OPE.Compute(PK, f, σ_x) \rightarrow σ_y : Given the polynomial f and the encoded input, output an encoded value σ_y of the result y .

⁴ <https://github.com/microsoft/SEAL>

$\text{Exp}_{II, \mathcal{A}}^{\text{SND}}(\lambda):$ $(\text{PK}, \text{SK}) \leftarrow \text{Setup}(1^\lambda); (f, x, \text{st}) \leftarrow \mathcal{A}_1(\text{PK})$ $\text{pk}_f \leftarrow \text{KeyGen}(\text{PK}, f); \sigma_x \leftarrow \text{QueryGen}(\text{PK}, \text{pk}_f, x); \sigma_y \leftarrow \mathcal{A}_2(\text{PK}, \text{pk}_f, \sigma_x, \text{st});$ if $\text{Verify}(\text{PK}, [\text{SK}], \text{pk}_f, \sigma_x, \sigma_y) = 1$ and $\text{Decode}(\text{SK}, \sigma_y) \neq f(x)$, then return 1 else return 0.
$\text{Exp}_{II, \mathcal{A}}^{\text{R-Privacy}}(\lambda):$ $b \xleftarrow{\$} \{0, 1\}; (\text{PK}, \text{SK}) \leftarrow \text{Setup}(1^\lambda); (f, x_0, x_1, \text{st}) \leftarrow \mathcal{A}_1(\text{PK})$ $\text{pk}_f \leftarrow \text{KeyGen}(\text{PK}, f); \sigma_b \leftarrow \text{QueryGen}(\text{PK}, \text{pk}_f, x_b); b' \leftarrow \mathcal{A}_2(\text{PK}, \sigma_b, \text{st})$ if $f(x_0) \neq f(x_1)$, then return \perp else return $(b = b')$
$\text{Exp}_{II, \mathcal{A}}^{\text{S-Privacy}}(\lambda):$ $b \xleftarrow{\$} \{0, 1\}; (\text{PK}, \text{SK}) \leftarrow \text{Setup}(1^\lambda); (f_0, f_1, \text{st}) \leftarrow \mathcal{A}_1(\text{PK})$ $\text{pk}_f \leftarrow \text{KeyGen}(\text{PK}, f_b); b' \leftarrow \mathcal{A}_2^{\text{CO}(\cdot)}(\text{PK}, \text{SK}, \text{pk}_f, \text{st})$ if f_0 or f_1 invalid functions, if some σ_x asked to CO decodes to an x such that $f_0(x) \neq f_1(x)$, then return \perp else return $(b = b')$
$\text{CO}(\sigma_x):$ return $\sigma_y \leftarrow \text{Compute}(\text{PK}, f_b, \sigma_x)$

Fig. 2. Security Games

$\text{OPE.Verify}(\text{PK}, [\text{SK}], \text{pk}_f, \sigma_x, \sigma_y) \rightarrow \text{acc}$: Given the secret key SK , in case of designated-verifier scheme, the public key pk_f for polynomial f , and the encoding σ_x of the evaluation point, accept (with $\text{acc} = 1$) or reject (with $\text{acc} = 0$) an output encoding σ_y .

$\text{OPE.Decode}(\text{SK}, \sigma_y) \rightarrow y$: Given the secret key SK for polynomial f , and an output encoding σ_y , output the result y .

For concrete use, between a sender with input polynomial f and a receiver with evaluation point x :

- The **Setup** algorithm is first run by the receiver (in case of designated-verifier) or a trusted party.
- The **Sender** executes the **KeyGen** algorithm with their input polynomial f .
- The **Receiver** runs **QueryGen** on their input x .
- The **Sender** runs **Compute** algorithm to obtain an encoding of the result σ_y .
- The **Receiver** can verify and decode the result with algorithms **Verify** and **Decode**.

Note that the **QueryGen** and **Decode** correspond to the encryption and decryption algorithms of a (fully) homomorphic encryption scheme.

Correctness. The *correctness* of a MyOPE scheme requires that if one runs **Compute** on an honestly generated query encoding of x , after honest **Setup** and **KeyGen** executions for f , then the output must verify and its decoding should be $y = f(x)$.

Soundness. The verifiability of a MyOPE guarantees the receiver of correct computation, even in front of a malicious sender, once **Setup** and **KeyGen** have been run honestly. This is done by means of a proof, the encoded value of the result σ_y should contain a proof of correctness of y .

Definition 1 (Soundness (SND)). Let II be an instance of our VC protocol and $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ a two-stage adversary. Protocol II is SND-secure if the advantage $\text{Adv}_{II, \mathcal{A}}^{\text{SND}}(\lambda) = \Pr[1 \leftarrow \text{Exp}_{II, \mathcal{A}}^{\text{SND}}(1^\lambda)]$ is negligible.

However, one may also expect some privacy properties which are now defined.

Receiver Privacy. This notion ensures that the input x of the receiver remains hidden during the protocol execution, for an honestly generated pk_f .

Definition 2 (Receiver Privacy (R-Privacy)). Let II be an instance of our MyOPE protocol and $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ a two-stage adversary, where \mathcal{A}_2 has adaptive access to the **Compute-oracle**

on legitimate queries only. Protocol Π is R-Privacy-secure if the advantage $\text{Adv}_{\Pi, \mathcal{A}}^{\text{R-Privacy}}(\lambda) = \Pr[1 \leftarrow \text{Exp}_{\Pi, \mathcal{A}}^{\text{R-Privacy}}(1^\lambda)]$ is negligible.

Sender Privacy. This notion ensures that the polynomial f of the sender remains hidden during the protocol execution, for adaptive legitimate requests by the receiver. We indeed exclude Compute-queries that trivially help to distinguish between two functions.

Definition 3 (Sender Privacy (S-Privacy)). Let Π be an instance of our MyOPE protocol and $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ a two-stage adversary, where \mathcal{A}_2 has adaptive access to the Compute-oracle on legitimate queries only. Protocol Π is S-Privacy-secure if the advantage $\text{Adv}_{\Pi, \mathcal{A}}^{\text{S-Privacy}}(\lambda) = \Pr[1 \leftarrow \text{Exp}_{\Pi, \mathcal{A}}^{\text{S-Privacy}}(1^\lambda)]$ is negligible.

2.2 Building Blocks

Let us recall the generic definitions of verifiable computation to illustrate our inner-product argument and verifiable oblivious polynomial evaluation, with fully homomorphic encryption. First we will need compact binding encodings with verifiable commitments.

Verifiable Commitments. A first tool is verifiable commitments $\text{Com} = (\text{Setup}, \text{Commit}, \text{Verify}, \text{Open})$ for elements in a space \mathcal{X} :

Com.Setup $(1^\lambda) \rightarrow (\text{ck}, [\text{vk}])$: Given the security parameter, output the commitment key ck and possibly a secret verification key vk in case of designated-verifier proof.

Com.Commit $(\text{ck}, x) \rightarrow (c_x, w)$: Given the commitment key and an element $x \in \mathcal{X}$, output a commitment c_x and an opening value w .

Com.Verify $(\text{ck}, [\text{vk}], c) \rightarrow \text{acc}$: Given the commitment key, optionally the verification key, and a commitment c , accept (with $\text{acc} = 1$) or reject (with $\text{acc} = 0$) a commitment c .

Com.Open $(\text{ck}, x, c, w) \rightarrow \text{acc}$: Given the commitment key, an element x , a commitment c , and the opening value w , accept (with $\text{acc} = 1$) or reject (with $\text{acc} = 0$) the commitment c for x .

The *correctness* property means that the **Verify** and **Open** algorithms accept when commitments have been honestly generated on inputs in the appropriate space \mathcal{X} . On the other hand, the *binding* property means that no adversary can make **Verify** accept on committed elements that open outside the appropriate space \mathcal{X} nor make **Open** accept on two different values. Additionally, one may expect the *hiding* property which means that c_x does not reveal any information about x (at least computationally). We stress that the target space \mathcal{X} is verified: an acceptable commitment necessarily encodes an element in \mathcal{X} .

rIPA: Inner Product Argument for Rings. We are first interested in the verifiable computation of inner products between committed vectors. As already noted, when there is no privacy issue, the commitment algorithm can simply be the identity function. But for efficiency reasons, we expect a more compact binding verifiable commitment to encode inputs. Then, from $\text{pk}_x \leftarrow c_x$ with a commitment c_x of \vec{x} (from **Com.Commit** (ck, x)) and $\sigma_y \leftarrow (\vec{y}, c_y)$ with a commitment c_y of \vec{y} , the sender generates $\sigma_z = (z, \pi)$, where π is a proof of $z = \langle \vec{x}, \vec{y} \rangle$, when c_x and c_y are valid commitments of appropriate vectors \vec{x} and \vec{y} (in the correct vector spaces). One may additionally expect receiver and/or sender privacy, with private \vec{y} and/or \vec{x} .

Verifiable OPE: Oblivious Polynomial Evaluation. This is another case of VC, with a polynomial $f(X) \in \mathbb{R}[X^n]$ as function, and $\mathbf{m} \in \mathbb{R}$ as evaluation point. Only $f(\mathbf{m})$ is learnt by the receiver, and no other information leaks. It also ensures the computations were executed as pledged in the protocol by providing verifiability.

FHE: Fully Homomorphic Encryption. We will achieve privacy-preserving VCs using Fully Homomorphic Encryption (FHE). FHE has been introduced in [Gen09]. This is particular case of classical public-public encryption, with a KeyGen algorithm that generates a key-pair $(\mathbf{pk}, \mathbf{sk})$ as well as encryption and decryption algorithms Enc and Dec, but with an additional Eval algorithm to operate on ciphertexts to build a ciphertext of $f((x_i)_i)$ from the ciphertexts of the x_i 's. Since the initial construction, major improvements have been made, with now practical and efficient solutions. In this work we will use the Fan-Vercauteren (FV) FHE scheme [FV12] for our analyses, with $\mathcal{R}_t = \mathbb{Z}_t[X]/r(X)$ as the plaintext message space and $\mathcal{R}_q \times \mathcal{R}_q$ as the ciphertext space, where $\mathcal{R}_q = \mathbb{Z}_q[X]/r(X)$, with $r(X) = X^n + 1$ for some well-chosen integer n (usually a power of 2). Semantic security relies on the Ring-LWE assumption. To enable some optimizations as the RNS representation used in SEAL, we will allow q with small prime factors on 60 bits. Detail about this FHE scheme is given in appendix A. We use the notation $[a]_q = a \bmod q$ for $a \in \mathbb{Z}$, and, for a ring element \mathbf{a} , $[\mathbf{a}]_q$ will represent the ring element with $[\cdot]_q$ applied to all its coefficients. Essentially, the public key is then $\mathbf{pk} = (\mathbf{p}, \mathbf{p}') = ([-(\mathbf{a} \cdot \mathbf{s} + \mathbf{e})]_q, \mathbf{a}) \in \mathcal{R}_q^2$, for random polynomials $\mathbf{a} \xleftarrow{\$} \mathcal{R}_q$, $\mathbf{s}, \mathbf{e} \leftarrow \chi$ (for χ a discrete centered Gaussian distribution in \mathcal{R}_q), while the secret key is $\mathbf{sk} = \mathbf{s}$. To encrypt a message $\mathbf{m} \in \mathcal{R}_t$, one computes $(\mathbf{c}, \mathbf{c}') = ([\mathbf{p} \cdot \mathbf{u} + \mathbf{e}_1 + \Delta \cdot \mathbf{m}]_q, [\mathbf{p}' \cdot \mathbf{u} + \mathbf{e}_2]_q)$ with small noises $\mathbf{u}, \mathbf{e}_1, \mathbf{e}_2 \leftarrow \chi$, where $\Delta = \lfloor q/t \rfloor$. One can see that with $\mathbf{d} = [\mathbf{c} + \mathbf{c}' \cdot \mathbf{s}]_q = [\Delta \cdot \mathbf{m} - \mathbf{e} \cdot \mathbf{u} + \mathbf{e}_2 \cdot \mathbf{s} + \mathbf{e}_1]_q = \Delta \cdot \mathbf{m} + \mathbf{v}$, with $\|\mathbf{v}\|$, the resulting noise vector, having a small enough norm for correct decryption, $\mathbf{m}' = \llbracket \mathbf{d}/\Delta \rrbracket_t = \llbracket \mathbf{m} + \lfloor \mathbf{v}/\Delta \rfloor \rrbracket_t$ leads to \mathbf{m} , if the error term \mathbf{v} is small enough (with an infinity norm $\|\mathbf{v}\|_\infty$ less than $\Delta/2$). We do not detail Eval, but for our analysis to hold, we will need the following property, in the particular case where $L = 2N$ (N being the maximal degree of our OPE polynomial), which can be guaranteed with an appropriate parameter choice, for d one less than the minimal circuit depth enabling bootstrapping:

Definition 4 ((L, d)- \mathcal{R}_t -Linear-Homomorphism). For any $\mathbf{a}_i, \mathbf{m}_i \in \mathcal{R}_t$, and some ciphertexts $(\mathbf{c}_i, \mathbf{c}'_i) \in \mathcal{R}_q^2$ of \mathbf{m}_i obtained from a circuit of multiplicative depth at most d ,

$$\text{Dec} \left(\mathbf{sk}, \sum_{i=1}^L \mathbf{a}_i \cdot (\mathbf{c}_i, \mathbf{c}'_i) \right) = \sum_{i=1}^L \mathbf{a}_i \cdot \mathbf{m}_i \in \mathcal{R}_t$$

2.3 Secure Encoding Schemes

Our main ingredient will be secure, or linear-only, encodings introduced by [BCI⁺13, GGPR13] for succinct non-interactive arguments of knowledge (SNARKs). We provide a general definition over rings. An encoding scheme over a ring \mathbb{R} consists in a tuple of algorithms:

- $(\mathbf{pk}, \mathbf{sk}, \mathbf{vk}) \leftarrow \text{Gen}(1^\lambda)$, a key generation algorithm that takes as input a security parameter and outputs public information \mathbf{pk} , a secret key \mathbf{sk} , and a verification key \mathbf{vk} , that can be either public or private;
- $E \leftarrow \text{E}_{\mathbf{sk}}(a)$, a (probabilistic) encoding algorithm mapping a ring element $a \in \mathbb{R}$ in the encoding space \mathcal{E} , using the secret key \mathbf{sk} .

It should then satisfy a few properties:

- L -Linearly homomorphic, with an algorithm $\text{Eval}_{\mathbf{pk}}$ that homomorphically combine encodings into the encoding of the same linear combination of the inputs;
- L -Quadratic root verification, with an algorithm $\text{QCheck}_{\mathbf{vk}}$ that can check a quadratic relation between the encoded elements, just from the encodings;
- Image verification, with an algorithm $\text{Verify}_{\mathbf{vk}}$ that check the validity of the encoding. This certifies the membership of the encoded element in the appropriate space.

According to the verification key that can be either public or private, the verification processes will be either public or private. The encoding is linearly-homomorphic, but for a *secure* encoding, one expects no one to be able to derive new valid encodings except from linear combinations,

hence the *linear-only* property: any new *valid* encoding E of some $a \in \mathbb{R}$ will necessarily satisfy $a = \sum_i c_i a_i$, for extractable elements $c_i \in \mathbb{R}$. Intuitively, when an encoding E passes the verification test, one can extract the linear combination of the given initial encodings.

The above properties will be enough for a binding commitment, but additional blinding factors will be required for hiding commitments, together with zero-knowledge proofs to keep the above verifications possible, without leaking more information.

In appendix B, we provide a more formal definition, with two illustrations. First, encodings over \mathbb{Z}_q , with q a prime large enough, in a pairing-friendly setting $\text{pk} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g, \mathbf{g}, e)$, using the Knowledge of Exponent Assumption [Dam92]. If we denote $G = e(g, \mathbf{g})$ and $\text{vk} = \mathbf{g}^\alpha$ for the secret key $\alpha \xleftarrow{\$} \mathbb{Z}_q$, the encoding function can be defined as $\text{E}_{\text{sk}}(a) = (g^a, g^{\alpha \cdot a}, \mathbf{g}^a) \in \mathbb{G}_1^2 \times \mathbb{G}_2$. Image verification can be publicly done with $e(g^a, \mathbf{g}) = e(g, \mathbf{g}^a)$ and $e(g^a, \text{vk}) = e(g^{\alpha \cdot a}, \mathbf{g})$. It is clearly L -linearly-homomorphic for any L . The bilinear map e allows public quadratic root verification, on the elements g^a and \mathbf{g}^a . To hide the content of an encoding, one just needs the encoding $E' = (g, g^\alpha, \mathbf{g})$ of 1, multiplied by a private random factor in $\mathcal{M} = \mathbb{Z}_q$. Classical Schnorr-proof can then be applied. Such encodings just consist of three group elements (2 in \mathbb{G}_1 and 1 in \mathbb{G}_2).

We then discuss the situation where q is the product of smaller primes. Then, the hardness of discrete logarithm does not hold anymore, but one can use linear-only encryption schemes, which limit to linear combinations only. We develop more the case of the Paillier encryption scheme [Pai99] with large RSA modulus \mathcal{N} . Such secure encodings just consist of two Paillier ciphertexts in $\mathbb{Z}_{\mathcal{N}^2}$ each. We could also use a secure encoding scheme based on the Learning With Errors (LWE) problem to make the whole scheme post-quantum secure.

When the receiver has published secure encodings of successive powers $\text{E}_{\text{sk}}(1), \text{E}_{\text{sk}}(s), \dots, \text{E}_{\text{sk}}(s^{n-1})$, for a random secret point s , the sender can only generate valid commitments of polynomials f of degree at most $n-1$, as $\text{E}_{\text{sk}}(f(s))$, thanks to the linear property of the encoding and the image verification. Quadratic root verification allows the sender to verify any quadratic relations between polynomials committed by the sender. Note that the initial secure encodings of successive powers can either be generated by a trusted third party, when using the above pairing-based secure encodings that are publicly verifiable, or by the receiver when using a linear-only encryption scheme.

In the body of the paper, for the sake of clarity, hiding commitments and zero-knowledge proofs will be ignored, as their computational and communication impacts are minimal, since they only deal with few scalars.

3 Verifiable Commitments

A major contribution of this paper is the construction of commitments of multivariate polynomials over rings so that succinct proofs can later be described. This is in the same vein as in [FNP20], but the latter is only defined for secure encodings based in pairings, whereas we describe here the construction from any secure encodings. With a linear-only encryption scheme, they will not be publicly verifiable anymore, but this will be useful to build compact commitments of polynomials over \mathbb{Z}_q in 2PC protocols, whatever the integer q (large prime or composite).

We provide here the intuition of our approach for polynomials in $\mathcal{R}_1 = \mathbb{Z}_q[X^{n-1}]$ (polynomials of degree at most $n-1$), while more polynomial spaces will be used in the global protocol. We stress again that commitments are specific to a space \mathcal{X} and when valid they ensure the committed element actually lies in \mathcal{X} . The verifier first generates secure encodings $E_i \leftarrow \text{E}_{\text{sk}}(s^i)$, for $i \in \llbracket 0; n-1 \rrbracket$ and a random secret element $s \xleftarrow{\$} \mathbb{Z}_q^*$. Thanks to the linear-only extractability, when a player generates a valid encoding E , being only given (E_0, \dots, E_{n-1}) , one can extract (c_i) such that E is an encoding of $c_0 + c_1 s + \dots + c_{n-1} s^{n-1}$ in \mathbb{Z}_q , and thus of the polynomial $\mathbf{c} = \sum_i c_i X^i$ in \mathcal{R}_1 . The encoding E is thus a commitment of $\mathbf{c} \in \mathcal{R}_1$: the list of initial encodings

E_i specifies a basis of the exact space \mathcal{X} we target. Here, \mathcal{R}_1 is spanned by $(1, X, \dots, X^{n-1})$ in \mathbb{Z}_q .

In addition, thanks to the quadratic verification on the encodings, if we have four polynomials u, v, m and r such that $m = u \cdot v \bmod r$, which means that $m = u \cdot v + r \cdot q$ for some polynomial q , where all the polynomials are of degree at most $n - 1$, we can check such a product: from valid commitments U and V of u and v , R and Q of r and q , respectively, and M of the polynomial m , all of degree at most $n - 1$, as they are all simple encodings, $\text{QCheck}_{\text{vk}}(X_1 X_2 + X_3 X_4 - X_5, U, V, R, Q, M) = \text{true}$ implies that $m(s) = u(s) \cdot v(s) + r(s) \cdot q(s)$.

Under the Schwartz-Zippel lemma, if q is a large prime, the probability to have this equality in a random point $s \in \mathbb{Z}_q$ whereas $m \neq u \cdot v + r \cdot q$ in $\mathbb{Z}_q[X]$ is bounded by $2n/q$, as the total degree of the relation is at most $2n$. Hence, the probability over s to have a false positive is bounded by $2n/q$. This is negligible in the large prime case but if we want to use RNS optimizations when computing modulo q , we need to take q a product of primes, and will hence need more repetitions with probability bounds stated by the Schwartz-Zippel lemma. Detail about this case is given in appendix C, along with a complete description of binding and hiding polynomial commitments, for univariate and bivariate polynomials, with multiple evaluation points when necessary.

4 Inner Product Arguments

4.1 Description of our ring-IPA

Our main tool is verifiable computation of inner products, from commitments on vectors, in various structures. To this aim, we convert vectors in polynomials to commit them as explained above. We stress that for the moment, we do not consider privacy, nor FHE ciphertexts, but just vectors in clear.

To start off, let us consider vectors in a field \mathbb{Z}_q (with a prime q). We will extend our method to the case where q is a product of primes, and \mathbb{Z}_q a ring, in the appendix, adding necessary repetitions. To commit such vectors, we will consider them as coefficients of a polynomial, and then commit the corresponding polynomials, as above. Let us consider $\mathbf{A} = (a_0, \dots, a_N)$ and $\mathbf{B} = (b_0, \dots, b_N)$ in \mathbb{Z}_q^{N+1} (equivalent to $\mathcal{R}_3 = \mathbb{Z}_q[Y^N]$, as defined in appendix C), two vectors whose inner-product is equal to $c = \langle \mathbf{A}, \mathbf{B} \rangle$ in \mathbb{Z}_q . As explained in Section 1.4, the commitments \bar{A} of \mathbf{A} and B of \mathbf{B} with secure encodings are $\bar{A} = \bar{a}(s)$ and $B = \mathbf{b}(s)$ for the polynomials $\bar{a}(Y) = \sum_{j=0}^N a_j Y^{N-j}$ and $\mathbf{b}(Y) = \sum_{j=0}^N b_j Y^j$ in \mathcal{R}_3 . Note that coefficients of \mathbf{A} are set into \bar{a} in a reversed order:

$$\bar{a}(Y) \cdot \mathbf{b}(Y) = \sum_{i,j=0}^N a_i b_j \cdot Y^{N+j-i} = \sum_{j=0}^N a_j b_j Y^N + \sum_{0 \leq i \neq j \leq N} a_i b_j Y^{N+j-i}.$$

Let us define the polynomial $\mathbf{d}(Y) = \bar{a}(Y) \cdot \mathbf{b}(Y) - cY^N$ of degree at most $2N$. If c is correct, \mathbf{d} is in the subspace $\mathcal{R}_4 = \mathbb{Z}_q[Y^{2N \setminus N}]$ (the polynomials of degree at most $2N$, without monomial of degree N). By publishing a commitment D of \mathbf{d} , that is verifiably in \mathcal{R}_4 , one can verify the above quadratic relation, using \bar{A} , B , c , and D , and get convinced of the inner product value c . The proof of correct computation of $c = \langle \mathbf{A}, \mathbf{B} \rangle$ with respect to the given commitments \bar{A} and B just consists of $\pi = \{D\}$ (1 commitment only), and the verification consists in checking the validity of the commitments and one quadratic relation.

Inner-Product Arguments Algorithms. More generally, one can define rIPA scheme on vectors $\mathbf{A}, \mathbf{B} \in \mathcal{X}^{N+1}$ and the result $c = \langle \mathbf{A}, \mathbf{B} \rangle$ in a space \mathcal{X} for which \mathcal{X}^{N+1} has either vectorial space (when \mathcal{X} is a field) or module (when \mathcal{X} is a ring) structure:

rIPA.Setup(1^λ) generates the verifiable commitment keys for the acceptable bases for \mathcal{R}_3 and \mathcal{R}_4 in PK to allow verifying commitments in these spaces. According to the encoding, verification will need SK or not;

$\text{rIPA.KeyGen}(\text{PK}, \mathbf{A})$, from a vector \mathbf{A} , outputs $\bar{\mathbf{A}}$, a commitment of \mathbf{A} (in the reverse order) using the commitment scheme Com ;
 $\text{rIPA.QueryGen}(\text{PK}, \mathbf{B})$, from a vector \mathbf{B} , outputs (\mathbf{B}, B) , where B is a commitment of \mathbf{B} ;
 $\text{rIPA.Compute}(\text{PK}, \mathbf{A}, (\mathbf{B}, B))$, from the two vectors \mathbf{A} and \mathbf{B} , outputs $c = \langle \mathbf{A}, \mathbf{B} \rangle$ and $\pi = \{D\}$ (where D is a commitment of \mathbf{d} , as defined above);
 $\text{rIPA.Verify}(\text{PK}, [\text{SK}], \bar{\mathbf{A}}, B, (c, \pi))$, with $\pi = \{D\}$, checks the relation $\mathbf{d}(Y) = \mathbf{a}(Y) \cdot \mathbf{b}(Y) - c \cdot Y^N$ from $\bar{\mathbf{A}}, B, D$ and c .

Since there is no privacy in this protocol, rIPA.Compute directly outputs the result $c = \langle \mathbf{A}, \mathbf{B} \rangle$: there is no need of private rIPA.Decode .

Polynomial Evaluation. It can be turned into a polynomial evaluation $y = P(x)$, with one vector \mathbf{A} containing the coefficients of P and the other vector \mathbf{B} built from the powers of x , and the expected inner product being y .

Infinity Norm Evaluation. It also provides a setting to compute the L_2 -norm $\|\mathbf{e}\|_2$, of $\mathbf{e} \in \mathbb{Z}_q[X^{n-1}]$, as $\|\mathbf{e}\|_2^2 = \langle \mathbf{E}, \mathbf{E} \rangle$, for the vector \mathbf{E} of the polynomial's coefficients. This leads to an approximation of the infinity norm with $\|\mathbf{e}\|_\infty \leq \|\mathbf{e}\|_2 \leq \sqrt{n} \cdot \|\mathbf{e}\|_\infty \leq \sqrt{n} \cdot \|\mathbf{e}\|_2$. One just needs $E = \mathbf{e}(s)$ and $\bar{E} = \bar{\mathbf{e}}(s)$, where

$$\mathbf{e}(X) = \sum_0^{n-1} e_i X^i \quad \bar{\mathbf{e}}(X) = \sum_0^{n-1} e_i X^{n-1-i} = X^{n-1} \cdot \mathbf{e}(1/X)$$

which can be verified with the existence, for a random challenge $\beta \xleftarrow{\$} \mathbb{Z}_q^*$, of polynomials \mathbf{e}' and $\bar{\mathbf{e}}'$ that satisfy, with $e = \mathbf{e}(1/\beta)$,

$$\begin{aligned} \mathbf{e}(X) - e &= \mathbf{e}(X) - \mathbf{e}(1/\beta) = (X - 1/\beta) \cdot \mathbf{e}'(X) \\ \bar{\mathbf{e}}(X) - \beta^{n-1} \cdot e &= \bar{\mathbf{e}}(X) - \beta^{n-1} \cdot \mathbf{e}(1/\beta) = \bar{\mathbf{e}}(X) - \bar{\mathbf{e}}(\beta) = (X - \beta) \cdot \mathbf{e}''(X). \end{aligned}$$

Indeed, as \mathbf{e} and $\bar{\mathbf{e}}$ have been committed in E and \bar{E} before the random choice of β , the first equation guarantees that $e = \mathbf{e}(1/\beta)$ while the second equation guarantees $\bar{\mathbf{e}}(\beta) = \beta^{n-1}e = \beta^{n-1}\mathbf{e}(1/\beta)$. The Schwartz-Zippel lemma ensures that the polynomials \mathbf{e} and $\bar{\mathbf{e}}$, of degree $n-1$, satisfy with high probability $\bar{\mathbf{e}}(X) = X^{n-1}\mathbf{e}(1/X)$: $\bar{\mathbf{e}}$ is indeed \mathbf{e} with order of coefficients reversed. From the commitment E of \mathbf{e} and the result $\|\mathbf{e}\|_2$ to be proven (or a commitment of it), the proof consists of the commitments $E' = \mathbf{e}'(s)$ and $E'' = \mathbf{e}''(s)$ (to verify the two above equations), plus the inner-product proof (with the commitment D as above) with the additional commitment \bar{E} and the scalar e . The validity of the proof requires the verification of the validity of the commitments and three quadratic relations (the two above, and the one for the inner product). For strong privacy, one can first commit the scalar $\|\mathbf{e}\|_2^2$, prove the correct computation of this hidden value with the above approach, and then perform a zero-knowledge range proof for the committed value, to show it is of appropriate size.

4.2 Inner Product Arguments with Privacy

If we now consider vectors $\mathbf{A} = (\mathbf{a}_0, \dots, \mathbf{a}_N)$ and $\mathbf{B} = (\mathbf{b}_0, \dots, \mathbf{b}_N)$ in \mathcal{R}_t^{N+1} , where $\mathcal{R}_t = \mathbb{Z}_t[X]/r(X)$, they can be committed with bivariate polynomials in $\mathbb{Z}_t[X, Y]$, using secure encodings with monomials $s^i s'^j$: $\bar{\mathbf{A}} = \bar{\mathbf{a}}(s, s')$ and $B = \mathbf{b}(s, s')$, where

$$\bar{\mathbf{a}}(X, Y) = \sum_{i=0}^{n-1} \sum_{j=0}^N a_{j,i} X^i Y^{N-j} \quad \mathbf{b}(X, Y) = \sum_{i=0}^{n-1} \sum_{j=0}^N b_{j,i} X^i Y^j$$

One can get $\mathbf{p} = \langle \mathbf{A}, \mathbf{B} \rangle \in \mathcal{R}_t$. If one wants to keep vector \mathbf{B} private, the latter can be encrypted with the FV FHE scheme, in $(\mathbf{c}_i, \mathbf{c}'_i) \in \mathcal{R}_q \times \mathcal{R}_q$, for $i = 0, \dots, N$. Thanks to the linear-homomorphism of the FHE, $\text{Dec}(\langle \mathbf{A}, \mathbf{C} \rangle, \langle \mathbf{A}, \mathbf{C}' \rangle) = \langle \mathbf{A}, \mathbf{B} \rangle$, where $\mathbf{C} = (\mathbf{c}_0, \dots, \mathbf{c}_N)$ and $\mathbf{C}' =$

(c'_0, \dots, c'_N) are in \mathcal{R}_q^{N+1} (equivalent to $\mathcal{R}_2 = \mathbb{Z}_q[X^{n-1}, Y^N]$, as denoted in appendix C). One now needs the verifiability of $\langle A, C \rangle$ and $\langle A, C' \rangle$ in \mathcal{R}_q : we consider $A = (a_0, \dots, a_N) \in \mathcal{R}_t^{N+1}$ and $C = (c_0, \dots, c_N) \in \mathcal{R}_q^{N+1}$, and we want to compute $d = \langle A, C \rangle \in \mathcal{R}_q$ and prove it. One can similarly operate to compute and prove $d' = \langle A, C' \rangle$.

We set both polynomials in $\mathcal{R}_2 = \mathbb{Z}_q[X^{n-1}, Y^N]$,

$$\bar{a}(X, Y) = \sum_{i=0}^{n-1} \sum_{j=0}^N a_{j,i} X^i Y^{N-j} \quad c(X, Y) = \sum_{i=0}^{n-1} \sum_{j=0}^N c_{j,i} X^i Y^j.$$

They are committed as $\bar{A} = \bar{a}(s, s')$ and $C = c(s, s')$. The result of the inner product $d \in \mathcal{R}_q$ is committed into D . However, in $\mathbb{Z}_q[X]$, the result of the inner product is equal to $d + q \cdot r$, where r is the public quotient polynomial in rings \mathcal{R}_t and \mathcal{R}_q , and q is the quotient, committed into Q . We want to prove $\langle A, C \rangle = d + q \cdot r$ in $\mathbb{Z}_q[X]$.

Then, for a random scalar $\sigma \in \mathbb{Z}_q$, one has the following relations, with $\bar{a}'(Y) = \bar{a}(\sigma, Y)$ and $c'(Y) = c(\sigma, Y)$, committed into \bar{A}', C' ,

$$\bar{a}(X, Y) - \bar{a}'(Y) = (X - \sigma) \cdot \bar{a}''(X, Y) \quad c(X, Y) - c'(Y) = (X - \sigma) \cdot c''(X, Y)$$

for some polynomials \bar{a}'' and c'' that can be computed from \bar{a}, \bar{a}', c, c' and committed into \bar{A}'', C'' , as well as the polynomial $X - \sigma$, so the receiver can check the above quadratic relations. Then, we also have

$$\begin{aligned} \bar{a}'(Y) \cdot c'(Y) &= \sum_{j=0}^N a'_j \cdot c'_j \cdot Y^N + \sum_{0 \leq i \neq j \leq N} a'_i \cdot c'_j \cdot Y^{N+j-i} \\ \text{and } \sum_{j=0}^N a'_j \cdot c'_j &= \sum_{j=0}^N a_j(\sigma) \cdot c_j(\sigma) = d(\sigma) + q(\sigma) \cdot r(\sigma) \end{aligned}$$

Setting $\delta = d(\sigma)$, $\phi = q(\sigma)$ and $\rho = r(\sigma)$, the values can be sent and checked with respect to d, Q , and r , as $q(X) - \phi = (X - \sigma) \cdot q'(X)$. If we additionally set $e(Y) = \bar{a}'(Y) \cdot c'(Y) - (\delta + \phi \cdot \rho) \cdot Y^N$, committed in E , by proving it relies in $\mathcal{R}_4 = \mathbb{Z}_q[Y^{2N \setminus N}]$, this proves the result d of the inner product in \mathcal{R}_q .

The proof of correct computation of $d = \langle A, C \rangle$ in \mathcal{R}_q , with respect to the given commitments \bar{A} and C just consists of $\pi = \{Q, Q', \bar{A}', \bar{A}'', C', C'', E, \phi\}$ (7 commitments and a scalar), for publicly generated or computed σ, δ and ρ , and the verification consists in checking the validity of the commitments and four quadratic relations. The same is needed for $d' = \langle A, C' \rangle$. By then decrypting (d, d') one should get $p = \langle A, B \rangle \in \mathcal{R}_t$.

Inner-Product Arguments with Privacy. More formally, we can define an inner products argument with privacy for one of the vectors rIPAwP. Given vectors A, B in the ring \mathcal{R}_t^{N+1} we prove the correctness of their inner product result $p \in \mathcal{R}_t$, while keeping B private:

- rIPAwP.Setup(1^λ) generates the parameters for the FV FHE with plaintext space \mathcal{R}_t and ciphertext space \mathcal{R}_q . The secure encodings on the acceptable bases for $\mathcal{R}_1, \mathcal{R}_2$ and \mathcal{R}_4 with the FHE encryption key are put in PK to allow encryption and evaluations on ciphertexts, as well as the generation of commitments in these spaces. The verification key of the secure encodings (if needed) and the FHE decryption key are put in SK;
- rIPAwP.KeyGen(PK, A), from a vector A, outputs \bar{A} , a commitment of A (in the reverse order);
- rIPAwP.QueryGen(PK, B), from a vector B, outputs $((C, C'), (C, C'))$, where in $C = (c_i)_i, C' = (c'_i)_i$ with (c_i, c'_i) the ciphertext of b_i , for $i = 0, \dots, N$, and then C, C' the commitment of C and C' respectively;
- rIPAwP.Compute(PK, A, (C, C', C, C')), from the two pairs of vectors (A, C) and (A, C'), outputs $d = \langle A, C \rangle$ and $d' = \langle A, C' \rangle$ and π , for proving the correct inner-product evaluations;
- rIPAwP.Verify(PK, [SK], $\bar{A}, (C, C'), (d, d', \pi)$), checks the proof π from the initial commitments $\bar{A}, (C, C')$ and the additional ones in π ;

$\text{rPAwP.Decode}(\text{SK}, (d, d'), \pi)$, from the FHE decryption key, decrypts the pair (d, d') to get $p = \langle A, B \rangle$.

In this case, rPAwP.Compute outputs an encryption of the expected result, hence the need of the private rPAwP.Decode . We used again the Schwartz-Zippel lemma to translate equalities between evaluated polynomials into equalities between polynomials, based on the unpredictability of σ . But according to q (large prime or product of smaller primes), one may reduce the bad cases by using multiple σ_κ 's.

4.3 Verifiability of the Committed Ciphertext

As already explained in the overall description of our protocol in Section 1.4, before verifying the correct inner products $d = \langle A, C \rangle$ and $d' = \langle A, C' \rangle$ and decrypt the pair (d, d') to get $\langle A, B \rangle$, one may want to be sure that each (c_i, c'_i) , ciphertext that would decrypt to m_i , is actually a correct encryption of m^i in \mathcal{R}_t . This means that B should be the vector (m^0, \dots, m^N) in \mathcal{R}_t^{N+1} . Indeed, the sender receives an encryption of m (and possibly some additional information) and generates the vectors C and C' thanks to the linearity of the FHE scheme. Why would they be honest?

To verify that, we use the above inner-product proof between each vector of ciphertexts $C = (c_0, \dots, c_N)$ or $C' = (c'_0, \dots, c'_N)$ and a vector of powers $N = (n^0, n^1, \dots, n^N)$ derived from a public random $n \in \mathcal{R}_t$ drawn by the verifier (or generated from a hash). Neither of these vectors need to be kept private as they are generated from information both parties have.

Let $u = \langle N, C \rangle$ and $u' = \langle N, C' \rangle$ be the results of the inner products in \mathcal{R}_q , proven as above (with 12 commitments and 2 scalars, and 8 quadratic relations to check). From the linear-homomorphism of the FHE, with appropriate parameters, $\text{Dec}(u, u') = \sum_{j=0}^N n^j \cdot m_j$. The verifier checks this decryption is $\sum_{j=0}^N n^j \cdot m^j$, with appropriate error (bounded as expected). This leads to $\sum_{j=0}^N n^j \cdot m_j = \sum_{j=0}^N n^j \cdot m^j$ in \mathcal{R}_t . Note that \mathcal{R}_t is unfortunately not a field, but possibly a ring that is the product of large fields only: $r = X^{2^k} + 1$ is not irreducible in any $\mathbb{Z}_t[X]$ for a prime t , but for well-chosen prime, all the factors of $X^n + 1$ may have large degrees in $\mathbb{Z}_t[X]$: according to [BGM93], with $t + 1 = 2^\alpha(2\tau + 1)$, for any integer τ , $\alpha \geq 2$, and $n = 2^k$, then all the factors of $X^{2^k} + 1$ have degree $2^{k+1-\alpha}$. If one chooses $t = 3 \pmod 8$, $\alpha = 2$, and in particular $t = 3$ (with $\tau = 0$), there are just two irreducible factors of degree $2^{k-1} = n/2$ in $\mathbb{Z}_t[X]$: the above polynomial thus has all zero coefficients by the Schwartz-Zippel lemma, excepted with probability $2N/t^{n/2}$, as the polynomial is of degree N and n is randomly chosen among $t^{n/2}$ possible values in each of the two fields. Hence, $m_i = m^i$ in \mathcal{R}_t , excepted with probability bounded by $2N/t^{n/2}$, which is clearly negligible. Note that one cannot use $t = 2$ as $r(X) = X^n + 1$ is divisible by $X + 1$ in $\mathbb{Z}_2[X]$.

By checking the noise in (u, u') , as expected with reasonable margin, as one knows the expected plaintext, one gets an upper-bound on individual errors. Even if this might be larger than initially expected, one can guarantee appropriate noise in the (c_i, c'_i) 's to satisfy the linear-homomorphism, as we will take additional margin to take care of the noise-flooding. If the sender tries to cheat with larger noise in the (c_i, c'_i) 's, they may reduce the privacy impact of the noise-flooding. But soundness remains guaranteed.

5 Verifiable OPE with Privacy

We now have the tools to allow the receiver/verifier with their private input message m to learn in a verifiable way the inner product of the vector $M = (m^j)_j$ with a private vector $F = (f_j)_j$, for indices j in $\llbracket 0; N \rrbracket$, both committed by the sender/prover.

5.1 Complete Protocol

More details and more applications are provided in appendix D, but we sketch here a full verifiable OPE protocol, where we assume all the global parameters set, and the sender's polynomial $F = (f_j)_j$ committed in a hiding way in \bar{F} . Once the receiver has encrypted the input $\mathbf{m} \in \mathcal{R}_t$ under their own FHE key and sent $\text{Enc}(\mathbf{m}) = (c, c') \in \mathcal{R}_q^2$ to the sender, the latter

- computes the $(u_j, u'_j) = \text{Enc}(\mathbf{m}^j)$, for $j \in \llbracket 0; N \rrbracket$, from (c, c') thanks to the homomorphic properties of the encryption scheme; generates the vectors $\mathbf{U} = (u_j)$ and $\mathbf{U}' = (u'_j)$ as well as their commitments U and U' ; and provides a proof of valid computation of the inner products $\mathbf{b} = \langle \mathbf{N}, \mathbf{U} \rangle$ and $\mathbf{b}' = \langle \mathbf{N}, \mathbf{U}' \rangle$, for a common vector $\mathbf{N} = (\mathbf{n}^j)_j$ for a random $\mathbf{n} \xleftarrow{\$} \mathcal{R}_t$ (chosen with a random hash function on the previous information), with respect to the commitments U, U' : ignoring scalars, the proof consists of 12 commitments (to be sent and checked), the ciphertext $(\mathbf{b}, \mathbf{b}')$, and 6 quadratic relations to be verified;
- generates a zero-ciphertext $(\mathbf{z}, \mathbf{z}')$ with a proof of small norm of the error, and provides a proof of valid computation of the noisy inner products $\mathbf{d} + \mathbf{z}$ and $\mathbf{d}' + \mathbf{z}'$, where $\mathbf{d} = \langle \mathbf{F}, \mathbf{U} \rangle$ and $\mathbf{d}' = \langle \mathbf{F}, \mathbf{U}' \rangle$, with respect to the commitments \bar{F}, U, U' : the proof consists of 9 new commitments (and also the original commitment of the polynomial) (to be sent and checked) and 5 quadratic relations to be verified, along with the result $(\mathbf{d}, \mathbf{d}')$, once the noise-flooding has been proven. The latter consists of 15 commitments (to be sent and checked) and 9 quadratic relations to be verified;

By first verifying $\text{Dec}(\mathbf{b}, \mathbf{b}') = \sum (\mathbf{n}\mathbf{m})^j$ and the appropriate noise, the receiver gets convinced (U, U') commits to correct encryptions of the powers \mathbf{m}^j . Then, with the verification of the inner products and the small noise, one gets the guarantee that $\text{Dec}(\mathbf{d} + \mathbf{z}, \mathbf{d}' + \mathbf{z}') = \langle \mathbf{F}, \mathbf{M} \rangle = F(\mathbf{m})$. As there are common commitments in the successive phases and one actually considers the noise components of $(\mathbf{z}^*, \mathbf{z}'^*)$ instead of these directly in practical applications, the global proof consists of 33 commitments and the verification checks them plus 20 quadratic relations (ignoring scalars and zero-knowledge proofs on scalars), as shown in appendix F (figures 5 and 6). This is thus independent of the degree of the polynomials.

5.2 Security Remarks

The soundness of this protocol is guaranteed by the proofs of valid computations of inner products, first to ensure the content of the commitments U and U' (with 2 inner products), and then to convince of the correct computation of the ciphertext $(\mathbf{d} + \mathbf{z}, \mathbf{d}' + \mathbf{z}')$ (with 2 inner products). The small additional noise $(\mathbf{z}, \mathbf{z}')$ is also proven by inner products to bound the infinity norms of the 3 polynomials $\mathbf{u}, \mathbf{e}_1, \mathbf{e}_2$ involved in $(\mathbf{z}, \mathbf{z}') = (\mathbf{p} \cdot \mathbf{u} + \mathbf{e}_1 \bmod q, \mathbf{p}' \cdot \mathbf{u} + \mathbf{e}_2 \bmod q)$, and proven with linear relations:

Theorem 5. *Our MyOPE scheme is SND-secure against malicious adversaries (see Definition 1), under the security of the secure encoding (and namely the quadratic root verification and the image verification properties).*

The complete proof is provided in appendix D, together with the analysis of the privacy properties, as stated in definitions 2 and 3. First, the receiver's privacy is ensured by the semantic security of the FHE encryption of \mathbf{m} that protects its input. This is a computational security, under the Ring-LWE assumption. Second, the sender's privacy is guaranteed by the hiding commitment \bar{F} and the noise $(\mathbf{z}, \mathbf{z}')$ that hides the evaluated circuit. They both provide statistical privacy to the sender.

For a non-interactive proof, the random elements chosen by the receiver can be generated by a hash function, using the Fiat-Shamir paradigm. Then, the security holds in the random oracle model.

5.3 FHE Security Analysis

Semantic Security. [FV12] gives us a condition in the relationship between parameters q , σ , n that will grant us the security relying on the difficulty of the RLWE problem. For a fixed root-Hermite factor δ , which we take such that $\log_2(\delta) = 1.8/(\lambda + 110)$, where λ is the security parameter, that we will take equal to $\lambda = 128$ in our applications, and if ε is the advantage of the distinguishing attack in [LP11], which we take equal to $\varepsilon = 2^{-64}$ for a corresponding $\lambda = 128$ in applications, and $\alpha = \sqrt{\ln(1/\varepsilon)/\pi}$ ($\varepsilon = 2^{-64}$ leads to $\alpha \approx 3.758$), we then have the condition:

$$\alpha \cdot \frac{q}{\sigma} < 2^{2\sqrt{n \log_2(q) \log_2(\delta)}}. \quad (1)$$

So that a parameter n can be derived from a (q, σ) pair and reciprocally. In particular, taking $\sigma = \alpha q = \sqrt{n}$ meets the above condition, thus granting the asymptotic security. For the practical security, in order to grant privacy with specific parameter sets, we will use [APS15]'s estimator, with results shown in figure 3.

Correctness. We now study the parameters needed for the correctness of the computations on FV ciphertexts.

Correctness for Basic FV. From [FV12], assuming χ is B -bounded, we find that the decryption of ciphertexts obtained from a d -depth circuit of somewhat homomorphic operations will be correct if d verifies:

$$4\beta(\varepsilon)\delta_{\mathcal{R}}^d \cdot (\delta_{\mathcal{R}} + 1.25)^{d+1} \cdot t^{d-1} < \frac{q}{\sigma}$$

where $\beta(\varepsilon)$ is drawn from the security parameter for [LP11]'s distinguishing attack ε , taken equal to $\varepsilon = 2^{-64}$ when $\lambda = 128$ which yields $\beta(\varepsilon) \approx 9.2$. For a cyclotomic polynomial r in \mathcal{R} , the above becomes:

$$4\beta(\varepsilon)n^d \cdot (n + 1.25)^{d+1} \cdot t^{d-1} < \frac{q}{\sigma} \quad (2)$$

Which gives an upper bound on d . Then, if we want full FHE capabilities from bootstrapping, [FV12], we will need the minimum allowed circuit depth d_{\min} to verify the above equation, where $d_{\min} = d_{\text{bs}} + 1$ (where d_{bs} is the depth of the bootstrapping operation) is given in [FV12]'s third theorem to have full FHE bootstrapping capacities with the condition: $d_{\min} \geq \text{BitSize}(\lceil \nu \cdot t \rceil) + \text{HammingWeight}(t) + 2$, with $\nu = \gamma \cdot (H(r) \cdot h + 1)$ with $2 < \gamma < 3$, $H(r) = 1$ for r a cyclotomic polynomial, and h the hamming weight of the FV scheme's secret key \mathbf{s} , for which we can take $h = 63$ according to [FV12]. Taking some margin on h , that we can consider as high as $h = 169$ for better security, with $t = 3$ in our scheme, and using a cyclotomic polynomial ring, we can thus take $d_{\min} = 12$. Replacing relation (1) in equation (2), with a security parameter $\lambda = 128$, $\beta(\varepsilon) \approx 9.2$, $\alpha \approx 3.8$, $t = 3$, we find the relation:

$$4\alpha \cdot \beta(\varepsilon) \cdot \delta_{\mathcal{R}}^{d_{\min}} (\delta_{\mathcal{R}} + 1.25)^{d_{\min}+1} t^{d_{\min}-1} < 2^{2\sqrt{n \log_2(q) \log_2(\delta)}}$$

is verified with:

$$25 + 25 \log_2(n) < 0.1743 \cdot \sqrt{n \log_2(q)}$$

So if we choose $n = 2^{14}$, then FHE capabilities will be granted with q on 283 bits or more.

Correctness with $2N$ -Linearity. Then, as we also want to be able to perform $L = 2N$ additions (N being the public degree of the sender's polynomial) on FV ciphertexts without additional bootstrapping in our protocol, we need to check that the error growth they give on bootstrapped

ciphertexts still allows decryption. In a nutshell, we will need the $(L, d_{\min} - 1)$ - \mathcal{R}_t linear homomorphism property: $\text{Dec}\left(\sum_{i=1}^L \mathbf{a}_i \cdot (\mathbf{c}_i, \mathbf{c}'_i)\right) = \sum_{i=1}^L \mathbf{a}_i \cdot \mathbf{m}_i$ for any $\mathbf{a}_i \in \mathcal{R}_t$ and ciphertexts $(\mathbf{c}_i, \mathbf{c}'_i) \in \mathcal{R}_q^2$ generated with a circuit of multiplicative depth $d_{\min} - 1$ (applying the bootstrapping operation), encrypting $\mathbf{m}_i \in \mathcal{R}_t$.

In order to decrypt the linear combination, we compute:

$$\begin{aligned} \left[\sum_{i=1}^L \mathbf{a}_i \cdot \mathbf{c}_i + s \cdot \sum_{i=1}^L \mathbf{a}_i \cdot \mathbf{c}'_i \right]_q &= \left[\sum_{i=1}^L \mathbf{a}_i \left(\mathbf{c}_i + s \cdot \mathbf{c}'_i \right) \right]_q \\ &= \left[\sum_{i=1}^L \mathbf{a}_i (\Delta \cdot \mathbf{m}_i + \mathbf{v}_i) \right]_q = \left[\Delta \cdot \sum_{i=1}^L \mathbf{a}_i \cdot \mathbf{m}_i + \sum_{i=1}^L \mathbf{a}_i \cdot \mathbf{v}_i \right]_q \end{aligned}$$

where each \mathbf{v}_i is the noise contained on the bootstrapped ciphertext $(\mathbf{c}_i, \mathbf{c}'_i)$, of infinite norm bounded by $2\beta(\epsilon)\sigma \cdot \delta_{\mathcal{R}}^{d_{\min}-1} \cdot (\delta_{\mathcal{R}} + 1.25)^{d_{\min}} \cdot t^{d_{\min}-3}$.

The decryption will be correct if:

$$\left\| \sum_{i=1}^L \mathbf{a}_i \cdot \mathbf{v}_i \right\|_{\infty} \leq \sum_{i=1}^L n \|\mathbf{a}_i\|_{\infty} \cdot \|\mathbf{v}_i\|_{\infty} \leq nt \sum_{i=1}^L \|\mathbf{v}_i\|_{\infty} \leq ntL \|\mathbf{v}\|_{\infty} \leq \Delta/2,$$

for a noise bounded by $\|\mathbf{v}\|_{\infty} \leq 2\beta(\epsilon)\sigma \cdot \delta_{\mathcal{R}}^{d_{\min}-1} \cdot (\delta_{\mathcal{R}} + 1.25)^{d_{\min}} \cdot t^{d_{\min}-3}$.

So the decryption works if:

$$4nL\beta(\epsilon)\sigma \cdot \delta_{\mathcal{R}}^{d_{\min}-1} \cdot (\delta_{\mathcal{R}} + 1.25)^{d_{\min}} \cdot t^{d_{\min}-1} \leq q$$

With a cyclotomic polynomial ring \mathcal{R} , and approximating $\beta(\epsilon)$ with 9.2 this becomes:

$$36.8 \times L\sigma \cdot n^{d_{\min}} \cdot (n + 1.25)^{d_{\min}} \cdot t^{d_{\min}-1} \leq q \quad (3)$$

With $\sigma \leq \sqrt{n}$, $n \geq 2^9$, $t = 3$ and d_{\min} from the above calculations, we get the previous inequality with the following condition, with $L = 2N$:

$$23.8 + \log_2(N) + 24.5 \log_2(n) \leq \log_2(q)$$

As an example, if $N = 2^{40}$ and $n = 2^{14}$, then taking q on 407 bits or more will grant this condition.

Correctness with Noise-Flooding. A linear combination of ciphertexts can leak the coefficients, from the evolution of the final noise, which can be recovered by the owner of the decryption key. To avoid this leakage, one can add super-polynomial noise to the result, this is the so-called *noise-flooding* technique: the sender will generate encryption of 0, *i.e.* polynomials $(\mathbf{z}^*, \mathbf{z}'^*)$ of the form

$$(\mathbf{z}^*, \mathbf{z}'^*) = (\mathbf{p} \cdot \mathbf{u} + \mathbf{e}_1 \bmod q, \mathbf{p}' \cdot \mathbf{u} + \mathbf{e}_2 \bmod q),$$

with coefficients of $\mathbf{u}, \mathbf{e}_1, \mathbf{e}_2$ follow the appropriate distribution for their own privacy: according to a Gaussian distribution on \mathbb{Z} with standard deviation 2^λ larger than the error in the result, that we bounded by $B' = 2Ln\beta(\epsilon)\delta_{\mathcal{R}}^{d_{\text{bs}}} \cdot (\delta_{\mathcal{R}} + 1.25)^{d_{\text{bs}}+1} \cdot t^{d_{\text{bs}}-1} \cdot \sigma$. Using the verifiable inner-product for provable L_2 -norm, one can prove that $\|\mathbf{u}\|_2, \|\mathbf{e}_1\|_2, \|\mathbf{e}_2\|_2$ are lower than $2^\lambda B'$, which guarantees the infinity norms are also lower than that.

Asymptotic Parameters. One now needs $q \geq 2^\lambda \times 2tB'$: With the above B' , this means that

$$4Ln2^\lambda\beta(\epsilon)\delta_{\mathcal{R}}^{d_{\text{bs}}} \cdot (\delta_{\mathcal{R}} + 1.25)^{d_{\text{bs}}+1} \cdot t^{d_{\text{bs}}} \cdot \sigma \leq q \quad (4)$$

guarantees the (L, d_{bs}) - \mathcal{R}_t linear homomorphism property and noise-flooding, and combined with the required hardness of the RLWE problem in equation 1, we now require:

$$4\alpha Ln2^\lambda\beta(\epsilon)\delta_{\mathcal{R}}^{d_{\text{bs}}} \cdot (\delta_{\mathcal{R}} + 1.25)^{d_{\text{bs}}+1} \cdot t^{d_{\text{bs}}} < 2^2\sqrt{n\log_2(q)\log_2(\delta)}$$

with $L = 2N$ (N being the public degree of the sender's polynomial), a cyclotomic polynomial ring, $\lambda = 128$, $t = 3$, $n \geq 2^9$, $d_{\text{bs}} = 11$ (from the above $d_{\text{min}} = 12$ calculation in section 5.3), $\alpha \approx 3.8$, $\beta(\epsilon) \approx 9.2$, the above equation is verified if:

$$151.7 + 28\log_2(n) + \log_2(N) < 0.1743 \cdot \sqrt{n\log_2(q)}$$

So for instance for $n = 2^{14}$ and $N = 2^{40}$, we are sure that q on 685 bits or more will grant full FHE functionalities, and we also get the following FHE ciphertext size complexity: $n\log_2(q) = \mathcal{O}(\log(N)^2)$.

5.4 Succinctness

As explained above, the proof is succinct, with a constant number of elements, and independent of the degree of the polynomial, but the size of the commitments may depend on q for some choices of Secure Encodings. In section 5.3, we studied asymptotic bounds for q and n , to get correctness, and get $\log q = \mathcal{O}(\log N)$, and $n = \mathcal{O}(\log N)$. Then the size of a ciphertext is in $\mathcal{O}(n\log q) = \mathcal{O}((\log N)^2)$ bits and this gives the receiver's communication complexity. Then, the sender essentially sends back the result (1 ciphertext = $\mathcal{O}((\log N)^2)$) and the proof which consists in a constant number of commitments in $\mathcal{O}(\log q) = \mathcal{O}(\log N)$. So, globally, the communication complexity is in $\mathcal{O}((\log N)^2)$.

We estimated practical sizes using security bounds on the privacy of FHE given by the LWE estimator [APS15]. For N between 2^{20} and 2^{40} , a prime q should be on 600 to 620 bits, which would lead to 3MBytes for the FHE ciphertext to be sent, and about 5MBytes for the result and its proof.

In the appendix, we give more details with a composite q , which allows RNS optimizations for FHE. Then, the size of the proof increases because of the repetitions of the commitments, as the Schwartz-Zippel lemma provides a smaller soundness, but it remains in a 90 to 520MBytes range for N less than 2^{40} , from our analysis in appendix F using our correctness formula, the [APS15] estimator, and the Schwartz-Zippel lemma for the proof soundness. In table 3, we

N	n	q	Receiver Communications		Sender Communications		
			FHE Security	Total Size	ν_c	ν_e	Total Size
20	14	600	115	3 MB	33	20	5 MB
25	14	605	114	3 MB	33	20	5 MB
30	14	610	113	3 MB	33	20	5 MB
35	14	615	111	3 MB	33	20	5 MB
40	14	620	110	3 MB	33	20	5 MB

Fig. 3. Parameters and security bounds for the FHE ciphertexts and the proof, with a prime ciphertext modulus q and the plaintext modulus $t = 3$. The proof size column encompasses all the elements communicated by the sender, including the result. This is for a 2^{-128} -soundness. $|x|$ is the bit-length of x . As q is a large prime, one can use an encoding scheme based on pairings: \mathbb{G}_2 elements are encoded on 880 bits for a 128-bit security, as in [Gui21]'s recommendations.

give parameters and resulting sizes for our OPE construction. We achieve the correctness from

inequality (4)’s exact requirements (without the following approximations), and the privacy provided by the FHE security is calculated using [APS15]’s estimator. Soundness requirements given by the Schwartz-Zippel lemma then provide the number of required repetitions on which the number of commitments ν_c and the number of checked equations ν_e depend.

References

- APS15. Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.
- BBB⁺18. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018.
- BC22. Dung Bui and Geoffroy Couteau. Private set intersection from pseudorandom correlation generators. Cryptology ePrint Archive, Report 2022/334, 2022. <https://ia.cr/2022/334>.
- BCC⁺16. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 327–357. Springer, Heidelberg, May 2016.
- BCFK21. Alexandre Bois, Ignacio Cascudo, Dario Fiore, and Dongwoo Kim. Flexible and efficient verifiable computation on encrypted data. In Juan Garay, editor, *PKC 2021, Part II*, volume 12711 of *LNCS*, pages 528–558. Springer, Heidelberg, May 2021.
- BCI⁺13. Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. Succinct non-interactive arguments via linear interactive proofs. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 315–333. Springer, Heidelberg, March 2013.
- BCMS20. Benedikt Bünz, Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. Recursive proof composition from accumulation schemes. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 1–18. Springer, Heidelberg, November 2020.
- BEHZ16. Jean-Claude Bajard, Julien Eynard, M. Anwar Hasan, and Vincent Zucca. A full RNS variant of FV like somewhat homomorphic encryption schemes. In Roberto Avanzi and Howard M. Heys, editors, *SAC 2016*, volume 10532 of *LNCS*, pages 423–442. Springer, Heidelberg, August 2016.
- BGM93. Ian F. Blake, Shuhong Gao, and Ronald C. Mullin. Explicit factorization of $x^{2^k} + 1$ over \mathbb{F}_p with prime $p = 3 \pmod{4}$. *Appl. Algebra Eng. Commun. Comput.*, 4:89–94, 1993.
- BMM⁺21. Benedikt Bünz, Mary Maller, Pratyush Mishra, Nirvan Tyagi, and Psi Vesely. Proofs for inner pairing products and applications. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 65–97. Springer, Heidelberg, December 2021.
- BN00. Daniel Bleichenbacher and Phong Q. Nguyen. Noisy polynomial interpolation and noisy Chinese remaindering. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 53–69. Springer, Heidelberg, May 2000.
- CL01. Yan-Cheng Chang and Chi-Jen Lu. Oblivious polynomial evaluation and oblivious neural learning. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 369–384. Springer, Heidelberg, December 2001.
- CL15. Guilhem Castagnos and Fabien Laguillaumie. Linearly homomorphic encryption from DDH. In Kaisa Nyberg, editor, *CT-RSA 2015*, volume 9048 of *LNCS*, pages 487–505. Springer, Heidelberg, April 2015.
- CLR17. Hao Chen, Kim Laine, and Peter Rindal. Fast private set intersection from homomorphic encryption. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1243–1255. ACM Press, October / November 2017.
- Dam92. Ivan Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In Joan Feigenbaum, editor, *CRYPTO ’91*, volume 576 of *LNCS*, pages 445–456. Springer, Heidelberg, August 1992.
- DRZ20. Vanesa Daza, Carla Ràfols, and Alexandros Zacharakis. Updateable inner product argument with logarithmic verifier and applications. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 527–557. Springer, Heidelberg, May 2020.
- FNPO4. Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. Efficient private matching and set intersection. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 1–19. Springer, Heidelberg, May 2004.
- FNP20. Dario Fiore, Anca Nitulescu, and David Pointcheval. Boosting verifiable computation on encrypted data. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part II*, volume 12111 of *LNCS*, pages 124–154. Springer, Heidelberg, May 2020.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO ’86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.

- FV12. Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144, 2012. <https://eprint.iacr.org/2012/144>.
- Gen09. Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
- GGP10. Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 465–482. Springer, Heidelberg, August 2010.
- GGPR13. Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.
- Gil99. Niv Gilboa. Two party RSA key generation. In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 116–129. Springer, Heidelberg, August 1999.
- GMNO18. Rosario Gennaro, Michele Minelli, Anca Nitulescu, and Michele Orrù. Lattice-based zk-SNARKs from square span programs. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 556–573. ACM Press, October 2018.
- GNS21. Chaya Ganesh, Anca Nitulescu, and Eduardo Soria-Vazquez. Rinocchio: SNARKs for ring arithmetic. Cryptology ePrint Archive, Report 2021/322, 2021. <https://eprint.iacr.org/2021/322>.
- Gui21. Aurore Guillevic. Pairing friendly curves, February 2021. <https://members.loria.fr/AGuillevic/pairing-friendly-curves/>.
- Haz18. Carmit Hazay. Oblivious polynomial evaluation and secure set-intersection from algebraic PRFs. *Journal of Cryptology*, 31(2):537–586, April 2018.
- HL09. Carmit Hazay and Yehuda Lindell. Efficient oblivious polynomial evaluation with simulation-based security. Cryptology ePrint Archive, Report 2009/459, 2009. <https://eprint.iacr.org/2009/459>.
- JL09. Stanislaw Jarecki and Xiaomin Liu. Efficient oblivious pseudorandom function with applications to adaptive OT and secure computation of set intersection. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 577–594. Springer, Heidelberg, March 2009.
- LMR19. Russell W. F. Lai, Giulio Malavolta, and Viktoria Ronge. Succinct arguments for bilinear group arithmetic: Practical structure-preserving cryptography. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2057–2074. ACM Press, November 2019.
- LP00. Yehuda Lindell and Benny Pinkas. Privacy preserving data mining. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 36–54. Springer, Heidelberg, August 2000.
- LP11. Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In Aggelos Kiayias, editor, *CT-RSA 2011*, volume 6558 of *LNCS*, pages 319–339. Springer, Heidelberg, February 2011.
- NP99. Moni Naor and Benny Pinkas. Oblivious transfer and polynomial evaluation. In *31st ACM STOC*, pages 245–254. ACM Press, May 1999.
- Pai99. Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In Jacques Stern, editor, *EUROCRYPT'99*, volume 1592 of *LNCS*, pages 223–238. Springer, Heidelberg, May 1999.
- PRTY20. Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. PSI from PaXoS: Fast, malicious private set intersection. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 739–767. Springer, Heidelberg, May 2020.
- Sch80. Jack T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of ACM*, 27(4):701–717, 1980.
- WTs⁺18. Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. In *2018 IEEE Symposium on Security and Privacy*, pages 926–943. IEEE Computer Society Press, May 2018.
- ZB05. Huafei Zhu and Feng Bao. Augmented oblivious polynomial evaluation protocol and its applications. In Sabrina De Capitani di Vimercati, Paul F. Syverson, and Dieter Gollmann, editors, *ESORICS 2005*, volume 3679 of *LNCS*, pages 222–230. Springer, Heidelberg, September 2005.
- Zip79. Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation, EUROSAM '79*, volume 72 of *LNCS*, pages 216–226. Springer, 1979.

Appendix

A Fully Homomorphic Encryption: FV Scheme

As explained in the preliminaries, we consider the Fan-Vercauteren (FV) Fully Homomorphic Encryption (FHE) scheme [FV12].

A.1 Description

Notations. Let \mathcal{R} be the ring $\mathbb{Z}[X]/r(X)$, where $r(X) = X^n + 1$ for $n = 2^k$. Given a polynomial $\mathbf{p} \in \mathcal{R}$, we denote $\|\mathbf{p}\|_\infty$ the infinity norm, *i.e.* the max of its coefficients. We also define the polynomial multiplication expansion factor $\delta_{\mathcal{R}}$ as

$$\delta_{\mathcal{R}} = \max_{\mathbf{c}, \mathbf{d} \in \mathcal{R}} \{\|\mathbf{c} \cdot \mathbf{d}\|_\infty / (\|\mathbf{c}\|_\infty \cdot \|\mathbf{d}\|_\infty)\}.$$

By taking the cyclotomic polynomial $r(X) = X^n + 1$, the worst-case bound for this expansion factor is $\delta_{\mathcal{R}} = n$.

The Fan-Vercauteren FHE. In FV scheme, the plaintext space is $\mathcal{R}_t = \mathbb{Z}_t[X]/r(X)$ while the ciphertext space is $\mathcal{R}_q \times \mathcal{R}_q$, where $\mathcal{R}_q = \mathbb{Z}_q[X]/r(X)$. FV encryption scheme is thus described by parameters $\Gamma = (q, t, n, \sigma)$, with $0 < \sigma < 1$, the noise parameter. They all will depend on several constraints, namely the expected multiplicative depth for the correctness and the hardness of the Ring-LWE problem for the semantic security.

For efficiency reasons, one may want to use FHE with RNS [BEHZ16]. We then assume $q = \prod_{i=1}^{\ell} p_i$ for distinct prime factors $p_1 < \dots < p_\ell$, assumed larger than p (that will be assumed a lower bound on the p_i 's all along this paper). We additionally take $t = 3$ in our applications (though bigger primes could also be used, but still with the constraint that $t = 3 \pmod{8}$ as explained in Section 4.3). We denote $\Delta = \lfloor q/t \rfloor$ and $\chi = D_{\mathbb{Z}, \sigma}^n$ the centered discrete Gaussian distribution over \mathbb{Z} with standard deviation σ for each coordinate. A distribution is B -bounded if it lies on a $[-B; B]$ support, and we consider χ to be statistically indistinguishable from a B -bounded distribution for B large enough, in practice taken to be $B = 10\sigma$.

The FV encryption scheme [FV12] consists in the following algorithms, with the notation $[a]_q = a \pmod{q}$ for $a \in \mathbb{Z}$, and, for a ring element \mathbf{a} , $[\mathbf{a}]_q$ representing the ring element with $[\cdot]_q$ applied to all its coefficients:

KeyGen($1^\lambda, \Gamma$) \rightarrow (**sk**, **pk**): On input the security parameter λ and the efficiency parameters set Γ , sample $\mathbf{a} \leftarrow \mathcal{R}_q$, $\mathbf{s} \leftarrow \mathcal{R}_2$ and set $(\mathbf{p}, \mathbf{p}') = ([-(\mathbf{a} \cdot \mathbf{s} + \mathbf{e})]_q, \mathbf{a}) \in \mathcal{R}_q^2$. **sk** = \mathbf{s} is the secret key, where the coefficients of \mathbf{e} are taken from χ . The public key **pk** contains $(\mathbf{p}, \mathbf{p}')$ with a relinearization key **rlk**.

Enc(**pk**, \mathbf{m}) \rightarrow (\mathbf{c}, \mathbf{c}'): Using $(\mathbf{p}, \mathbf{p}')$ from **pk** and a message $\mathbf{m} \in \mathcal{R}_t$, compute the ciphertext $(\mathbf{c}, \mathbf{c}') = ([\mathbf{p} \cdot \mathbf{u} + \mathbf{e}_1 + \Delta \cdot \mathbf{m}]_q, [\mathbf{p}' \cdot \mathbf{u} + \mathbf{e}_2]_q)$, with coefficients of $\mathbf{e}_1, \mathbf{e}_2$ also taken from χ , and $\mathbf{u} \leftarrow \mathcal{R}_2$.

Dec(**sk**, $(\mathbf{c}, \mathbf{c}')$) \rightarrow \mathbf{m} : Given **sk** = \mathbf{s} , compute

$$\mathbf{d} = [\mathbf{c} + \mathbf{c}' \cdot \mathbf{s}]_q = [\Delta \cdot \mathbf{m} - \mathbf{e} \cdot \mathbf{u} + \mathbf{e}_2 \cdot \mathbf{s} + \mathbf{e}_1]_q = \Delta \cdot \mathbf{m} + \mathbf{v}$$

with $\|\mathbf{v}\| \leq 2 \cdot \delta_{\mathcal{R}} \cdot B^2 + B$ considering χ is B -bounded. Then, compute:

$$\mathbf{m}' = \llbracket \mathbf{d} / \Delta \rrbracket_t = \llbracket (\Delta \cdot \mathbf{m} + \mathbf{v}) / \Delta \rrbracket_t = \llbracket \mathbf{m} + \lfloor \mathbf{v} / \Delta \rfloor \rrbracket_t$$

where $\mathbf{v} = -\mathbf{e} \cdot \mathbf{u} + \mathbf{e}_2 \cdot \mathbf{s} + \mathbf{e}_1$ is the error term: $\mathbf{m}' = \mathbf{m}$ if $\|\mathbf{v}\|_\infty \leq \Delta/2$.

Eval(**pk**, $f, (c_i, c'_i)_{i=1, \dots, \ell}$) \rightarrow (c_f, c'_f) : Given **pk**, an arithmetic circuit for a function f with bounded multiplicative depth and ℓ ciphertexts $(c_i, c'_i)_{i=1, \dots, \ell}$ output the ciphertext (c_f, c'_f) .

The addition of two ciphertexts is a ciphertext of the sum of the plaintexts. Multiplicative operations have also been shown possible with additional information to relinearize the ciphertext after a product using rlk (the relinearization key included in pk). For Eval to be correct, $\text{Dec}(\text{sk}, (c_f, c'_f))$ should return $f(m_1, \dots, m_\ell)$ where $\text{Dec}(\text{sk}, (c_i, c'_i)) = m_i$ for $i = 1, \dots, \ell$, with overwhelming probability.

B Secure Encodings

We here give a bit more detail about secure encodings sketched in Section 2.3. Some examples are described in the Appendix B.2.

B.1 Definitions

Definition 6 (Encoding Scheme). *An encoding scheme over a ring \mathbb{R} consists in a tuple of algorithms (Gen, E) .*

- $(\text{pk}, \text{sk}, \text{vk}) \leftarrow \text{Gen}(1^\lambda)$, a key generation algorithm that takes as input a security parameter and outputs public information pk , a secret key sk , and a verification key vk , that can be either public or private;
- $E \leftarrow \text{E}_{\text{sk}}(a)$, a (probabilistic) encoding algorithm mapping a ring element $a \in \mathbb{R}$ in the encoding space \mathcal{E} , using the secret key sk .

Properties. An encoding scheme should satisfy the following properties, with efficient and correct algorithms:

- L -Linearly homomorphic: An algorithm $\text{Eval}_{\text{pk}}(E_1, \dots, E_L; c_1, \dots, c_L)$, on input public information pk , encodings $E_1 = \text{E}_{\text{sk}}(a_1), \dots, E_L = \text{E}_{\text{sk}}(a_L)$, and coefficients $c_1, \dots, c_L \in \mathbb{R}$, outputs an encoding of $\sum_{i=1}^L c_i \cdot a_i$;
- L -Quadratic root verification: An algorithm $\text{QCheck}_{\text{vk}}(Q, E_1, \dots, E_L)$, on input the verification key vk , a quadratic polynomial $Q \in \mathbb{R}[X_1, \dots, X_L]$ and encodings $E_1 = \text{E}_{\text{sk}}(a_1), \dots, E_L = \text{E}_{\text{sk}}(a_L)$, checks whether or not the relation $Q(a_1, \dots, a_L) = 0$ is satisfied in \mathbb{R} ;
- Image verification: An algorithm $\text{Verify}_{\text{vk}}(E)$, on input the verification key vk and an element E , verifies E is an actual encoding of some element in \mathbb{R} : this algorithm not only verifies that $E \in \mathcal{E}$, but also that there exists an element $a \in \mathbb{R}$ such that E can be an encoding of a .

According to the verification key that can be either public or private, the verification processes will be either public or private.

Secure Encodings. We now formally define the soundness properties for the above verification algorithms, in terms of knowledge-soundness:

Definition 7 (Linear-Only Extractability). *An encoding scheme (Gen, E) over \mathbb{R} is extractable if for any PPT adversary \mathcal{A} , there exists a PPT extractor $\text{Ext}_{\mathcal{A}}$ such that the following probability is negligible in the security parameter:*

$$\Pr \left[\text{QCheck}_{\text{vk}} \left(X - \sum_{i=1}^n c_i X_i, E, E_1, \dots, E_n \right) = \text{false} \mid \text{Verify}_{\text{vk}}(E) = \text{true} \right]$$

on the probability space $(\text{pk}, \text{sk}, \text{vk}) \leftarrow \text{Gen}(1^\lambda)$, $a_1, \dots, a_n \xleftarrow{\$} \mathbb{R}$, $E_i \leftarrow \text{E}_{\text{sk}}(a_i)$, for $i = 1, \dots, n$, and $(E; c_1, \dots, c_n) \leftarrow (\mathcal{A} \parallel \text{Ext}_{\mathcal{A}})(\text{pk}, (E_i)_i)$.

While the encoding is linearly-homomorphic, the above extractability property requires that it be impossible to derive new valid encodings excepted by linear combinations: any new *valid* encoding E of some $a \in \mathbb{R}$ will necessarily satisfy $a = \sum_i c_i a_i$, for extractable elements $c_i \in \mathbb{R}$. Intuitively, when an encoding E passes the verification test, one can extract the linear combination of the given initial encodings.

Zero-Knowledge Proofs. The above properties will be enough for a binding commitment, but additional blinding factors in an appropriate masking set \mathcal{M} will be required for hiding commitments, which will depend on the ring \mathbb{R} (see below for the particular case of $\mathbb{R} = \mathbb{Z}_q$). Then, Zero-Knowledge proofs will be needed for Quadratic root verifications with private linear combinations: $\text{ZKLQCheck}_{\text{vk}}(Q, E_1, \dots, E_L; E'_1, \dots, E'_\mu)$, on input a quadratic polynomial $Q \in \mathbb{R}[X_1, \dots, X_L]$ and encodings $E_1 = \mathbf{E}_{\text{sk}}(a_1), \dots, E_L = \mathbf{E}_{\text{sk}}(a_L), E'_1 = \mathbf{E}_{\text{sk}}(b_1), \dots, E'_\mu = \mathbf{E}_{\text{sk}}(b_\mu)$. The verification key vk on the verifier side and the private coefficients c_1, \dots, c_μ on the prover side prove that the relation $Q(a_1, \dots, a_L) = \sum c_i \cdot b_i$ is satisfied. The c_i 's will be the blinding factors in \mathcal{M} .

B.2 Examples

In the following we illustrate the definition of secure encodings with two distinct constructions. the first provides public verifiability whereas the second will be for designated-verifier only.

Encodings over $\mathbb{R} = \mathbb{Z}_q$ from Pairings. The Knowledge of Exponent Assumption, introduced by Damgård [Dam92] states that given g and g^α in a group \mathbb{G} , it is hard to create c and \hat{c} in that group \mathbb{G} so that $\hat{c} = c^\alpha$, without knowing a such that $c = g^a$. The only way to compute \hat{c} being to generate $(g^\alpha)^a$.

In a pairing-friendly setting $\text{pk} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, q, g, \mathbf{g}, e)$, one can check the appropriate relation between c and \hat{c} , with \mathbf{g}^α . We can thus consider such a pairing-friendly setting with a large prime q , we denote $G = e(g, \mathbf{g})$. The verification key is $\text{vk} = \mathbf{g}^\alpha$ for the secret key $\alpha \xleftarrow{\$} \mathbb{Z}_q$: the encoding function \mathbf{E}_{sk} is $\mathbf{E}_{\text{sk}}(a) = (g^a, g^{\alpha \cdot a}, \mathbf{g}^a) \in \mathbb{G}_1^2 \times \mathbb{G}_2$. It is clearly L -linearly-homomorphic for any L . The bilinear map e allows public quadratic root verification, on the elements g^a and \mathbf{g}^a : for example, $Q = X_1 \cdot X_2 - X_3$ on encodings of a_1, a_2 and a_3 , it can be done with $e(g^{a_1}, \mathbf{g}^{a_2}) \cdot e(g^{-a_3}, \mathbf{g}) = e(g, \mathbf{g})^{Q(a_1, a_2, a_3)}$. This must be done after image verification of any individual encoding with $e(g^a, \mathbf{g}) = e(g, \mathbf{g}^a)$ and $e(g^a, \text{vk}) = e(g^{\alpha \cdot a}, \mathbf{g})$. They are both public, as vk is public.

To hide the content of an encoding, one just needs the encoding $E' = (g, g^\alpha, \mathbf{g})$ of 1, multiplied by a private random factor in $\mathcal{M} = \mathbb{Z}_q$. To prove the existence of μ such private coefficients c_i such that $Q(a_1, \dots, a_L) = \sum c_i b_i \pmod q$ is satisfied, one has to prove the knowledge of $(c_i)_i$ such that

$$V = e(g, \mathbf{g})^{Q(a_1, \dots, a_L)} = \prod e(g^{b_i}, \mathbf{g})^{c_i}$$

which can be done with a classical Schnorr-proof, from the encodings $\mathbf{E}_{\text{sk}}(b_i)$, as $V = e(g, \mathbf{g})^{Q(a_1, \dots, a_L)}$ can be computed thanks to the pairing. Using the Fiat-Shamir paradigm [FS87], this proof can be non-interactive: with random exponents $k_i \xleftarrow{\$} \mathbb{Z}_q$, the prover sets $D = \prod e(g^{b_i}, \mathbf{g})^{k_i}$, gets a random challenge $e = \mathcal{H}(V, \{\mathbf{E}_{\text{sk}}(b_i)\}, D) \in \llbracket 0; 2^\lambda \rrbracket$, with $2^\lambda < q$, and generates the proof $\Pi = (\{s_i = k_i - ec_i \pmod q\}, e) \in \mathbb{Z}_q^\mu \times \llbracket 0; 2^\lambda \rrbracket$. The verifier can compute

$$D' = \prod e(g^{b_i}, \mathbf{g})^{s_i} \times V^e$$

and check whether $e \stackrel{?}{=} \mathcal{H}(V, \{\mathbf{E}_{\text{sk}}(b_i)\}, D')$. Indeed, if the statement is true

$$\begin{aligned} \prod e(g^{b_i}, \mathbf{g})^{s_i} \times V^e &= \prod e(g^{b_i}, \mathbf{g})^{s_i} \times \left(\prod e(g^{b_i}, \mathbf{g})^{c_i} \right)^e \\ &= \prod e(g^{b_i}, \mathbf{g})^{s_i + ec_i} = \prod e(g^{b_i}, \mathbf{g})^{k_i} = D. \end{aligned}$$

The linear-only extractability relies on the above Knowledge of Exponent Assumption, that requires the hardness of the discrete logarithm in the bilinear generic group model, which additionally requires q to be large enough (at least 2λ bits, for a security parameter λ). An encoding is a tuple of elements in $\mathcal{E} = \mathbb{G}_1^2 \times \mathbb{G}_2$, and a zero-knowledge proof of μ scalars contains

μ elements from \mathbb{Z}_q and the challenge in $[[0; 2^\lambda]]$. In case of multiple proofs, one can use a unique global challenge e , and the same (k_i, s_i) can be used for the same private scalars c_i . Hence, globally, the size is thus μ' elements from \mathbb{Z}_q where μ' is the total number of private scalars, independently of the number of equations, plus one challenge. The requirement of a large prime will be prohibitive when used with Fully Homomorphic Encryption, as RNS optimizations do not apply.

Encodings over $R = \mathbb{Z}_q$ from a Linear-Only Encryption Scheme. Given a linear-only encryption scheme (Enc, Dec) from \mathbb{Z}_q onto \mathcal{C} , for any integer q , one chooses a random private element $\alpha \xleftarrow{\$} \mathbb{Z}_q^*$, to be the secret key of the encoding. Then, pk is the public key of the encryption scheme, while vk is α together with the secret decryption key of the encryption scheme. It is thus private. Then, the encoding function E_{sk} is $E_{\text{sk}}(a) = (\text{Enc}(a), \text{Enc}(\alpha \cdot a)) \in \mathcal{C} \times \mathcal{C}$. It is clearly linearly-homomorphic from an additively-homomorphic encryption scheme. The decryption algorithm allows any root verification using the decryption key in vk , while the image verification tests whether the two decryption values verify the secret ratio α . The linear-only extractability depends on the specific encryption scheme: but either the encryption scheme is fully/somewhat homomorphic, or this property is satisfied. An encoding is a pair of elements in $\mathcal{E} = \mathcal{C} \times \mathcal{C}$. Such schemes can be obtained from class groups, Paillier encryption or TFHE, among other examples.

Encodings from TFHE The Torus Fully Homomorphic Encryption (TFHE) scheme can be used without enabling the multiplication of ciphertexts, thus turning the scheme in a linear-Only Encryption scheme. One can use the *Concrete* library for the implementation. Using this scheme, which relies on the LWE (Learning With Errors) problem, makes our construction entirely post-quantum. To use these encodings over $R = \mathbb{Z}_q$, one uses q as the cleartext modulus of the TFHE scheme, and then uses a larger ciphertext modulus.

B.3 Encodings from Paillier Encryption

More concretely, one can use Paillier encryption scheme [Pai99] with a large RSA modulus \mathcal{N} , the encryption of $x \in \mathbb{Z}_{\mathcal{N}}$ is $E = (1 + \mathcal{N})^x \cdot r^{\mathcal{N}} \bmod \mathcal{N}^2$, for $r \xleftarrow{\$} \mathbb{Z}_{\mathcal{N}}^*$. The ciphertext space is thus $\mathcal{C} = \mathbb{Z}_{\mathcal{N}^2}^*$. For the decryption, one needs the value $\lambda = \lambda(\mathcal{N})$, where λ is Carmichael's function, which is equivalent of the knowledge of the factorisation of \mathcal{N} . As $\lambda(\mathcal{N}^2) = \mathcal{N}\lambda$, $E^\lambda = (1 + \mathcal{N})^{x\lambda} = 1 + x\lambda \cdot \mathcal{N} \bmod \mathcal{N}^2$. If λ is invertible modulo \mathcal{N} , one can extract x modulo \mathcal{N} .

For the encoding, the public key pk is thus the modulus \mathcal{N} , the secret key sk is a random element $\alpha \xleftarrow{\$} \mathbb{Z}_{\mathcal{N}}^*$, and the verification key vk is the decryption key $\lambda(\mathcal{N})$ and the secret value α :

- $E_{\text{sk}}(a) = ((1 + \mathcal{N})^a \cdot r_0^{\mathcal{N}} \bmod \mathcal{N}^2, (1 + \mathcal{N})^{\alpha \cdot a} \cdot r_1^{\mathcal{N}} \bmod \mathcal{N}^2)$, for $r_0, r_1 \xleftarrow{\$} \mathbb{Z}_{\mathcal{N}}^*$;
- $\text{Dec}_{\text{vk}}(C_0)$ decrypts x_0 from C_0 as $\frac{C_0^\lambda \bmod \mathcal{N}^2 - 1}{\mathcal{N}} \times \lambda^{-1} \bmod \mathcal{N}$.
- $\text{Verify}_{\text{vk}, \text{sk}}(C_0, C_1)$ first decrypts both ciphertexts using $\lambda(\mathcal{N})$ to get $x_0, x_1 \bmod \mathcal{N}$, and checks whether $x_1 = \alpha \cdot x_0 \bmod \mathcal{N}$.

Decryption and verification can be optimized using the CRT as in [Pai99]. An encoding is thus $4 \log \mathcal{N}$ bit-long.

We want an encoding on \mathbb{Z}_q , one can thus take $\mathcal{N} > q$ to encode elements $x \in [[0; q - 1]]$. Decoding first decrypts both ciphertexts modulo \mathcal{N} , with the elements in $[-\mathcal{N}/2; \mathcal{N}/2]$, checks the relation with α modulo \mathcal{N} , and reduces it again modulo q in $[[0; q - 1]]$ to extract the encoded value in \mathbb{Z}_q . For further verifications (quadratic checks), one just considers the decryption of the first ciphertext in $[-\mathcal{N}/2; \mathcal{N}/2]$, and relations among the plaintexts.

From L ciphertexts $E_i = (1 + \mathcal{N})^{x_i} \cdot r_i^{\mathcal{N}} \bmod \mathcal{N}^2$, for $r_i \xleftarrow{\$} \mathbb{Z}_{\mathcal{N}}^*$, one can compute the linear combination with coefficients smaller than q , $\prod E_i^{c_i} = (1 + \mathcal{N})^{\sum c_i x_i} \cdot (\prod r_i^{c_i})^{\mathcal{N}}$, which is an

encryption of $\sum c_i x_i \bmod \mathcal{N}$. It decodes to $\sum c_i x_i$ if $|\sum c_i x_i| < \mathcal{N}/2$: we thus have to take $\mathcal{N} > 2L \cdot q^2$ if the encodings are fresh encodings that encrypt elements in $\llbracket 0; q-1 \rrbracket$. For a quadratic check, using vk , the verifier can decrypt all the encodings modulo \mathcal{N} and reduce them modulo q to check the relation modulo q . There is no more constraint on \mathcal{N} .

To hide the content of an encoding, even with respect to the owner of the secret key, one uses a random encoding of a random private mask in $\llbracket 0; 2^\lambda L q^2 \rrbracket$, to act as a statistically hiding one-time pad, furthermore randomized with \mathcal{N} -th powers to remove any information in the initial random coins. In this case, one needs $\mathcal{N} > 2L \cdot 2^\lambda q^2$ for correct decryption modulo \mathcal{N} , without wrapping around modulo \mathcal{N} before the reduction modulo q . Hence $\mathcal{M} = \llbracket 0; 2^\lambda L q^2 \rrbracket$.

About the zero-knowledge verification $\text{ZKLQCheck}(Q, (E_i)_i; (E'_i)_i)$, to prove the existence of μ coefficients $c_i \in \mathcal{M}$ such that:

$$Q(a_1, \dots, a_L) = \sum c_i b_i \bmod q,$$

on the encodings $E_1 = \text{E}_{\text{sk}}(a_1), \dots, E_L = \text{E}_{\text{sk}}(a_L), E'_1 = \text{E}_{\text{sk}}(b_1), \dots, E'_\mu = \text{E}_{\text{sk}}(b_\mu)$, one has to prove the knowledge of (c_i) in

$$V = \text{Eval}(\{E'_i\}, \{c_i\}), \text{ for } i \in \llbracket 1; \mu \rrbracket,$$

where the verifier knows, after decryption of the encodings $(E_i)_i$ and quadratic computations modulo q ,

$$V' = (1 + \mathcal{N})^{Q(a_1, \dots, a_L) \bmod q} \bmod \mathcal{N}^2.$$

One wants to prove that V and V' decrypt to the same value modulo q : $V = V' \times (1 + \mathcal{N})^{qr_1} \cdot r_0^\mathcal{N} \bmod \mathcal{N}^2$, but for unknown values r_0, r_1 . We stress that we only consider the first ciphertext in the encodings. One thus proves the knowledge of $c_i \in \mathcal{M}$ for $i \in \llbracket 1; \mu \rrbracket$ such that $V = \prod E_i^{c_i} = V' \times (1 + \mathcal{N})^{qr_1} \cdot r_0^\mathcal{N} \bmod \mathcal{N}^2$. The c_i 's are the masks in $\llbracket 0; 2^\lambda L q^2 \rrbracket$, but one can use their q -reduction. Then, as the encodings E'_i can encrypt elements in $\llbracket -Lq^2; Lq^2 \rrbracket, |r_1 q| \leq \mu L q^3$.

For the latter zero-knowledge proof, the prover chooses random $k_i \xleftarrow{\$} \mathbb{Z}_q$, with additional noise $\nu' \xleftarrow{\$} \llbracket 0; 2^\lambda L q^2 \rrbracket$ and $\nu \xleftarrow{\$} \mathbb{Z}_\mathcal{N}^*$, to hide any extra information beyond the modulo q relation, and sets $D = \prod E_i^{k_i} \cdot (1 + \mathcal{N})^{q\nu'} \nu^\mathcal{N} \bmod \mathcal{N}^2$, gets a random challenge e (possibly derived from $(Q, (E_i)_i, (E'_i)_i, D)$ in $\llbracket 0; 2^{\lambda'} - 1 \rrbracket$ and outputs the proof $\Pi = (D, (s_i = k_i - ec_i \bmod q)_i) \in \mathbb{Z}_{\mathcal{N}^2} \times \mathbb{Z}_q^\mu$.

For the moment, we use a different space $\llbracket 0; 2^{\lambda'} - 1 \rrbracket$ for the challenge, with λ' possibly smaller than λ or $-\log \varepsilon_s$, in which case $-\log \varepsilon_s / \lambda'$ parallel repetitions should be performed for correct soundness. One can check

$$e \leftarrow \mathcal{H}(Q, (E_i)_i, (E'_i)_i, D) \quad D' \leftarrow \prod E_i^{s_i} \cdot (V')^e \bmod \mathcal{N}^2 \quad \text{Dec}(D/D') \stackrel{?}{=} 0 \bmod q$$

Indeed,

$$\begin{aligned} D' &= \prod E_i^{s_i} \times (V')^e = \prod E_i^{s_i} \times V^e \cdot (1 + \mathcal{N})^{-eqr_1} \cdot r_0^{-e\mathcal{N}} \\ &= \prod E_i^{k_i - ec_i} \times \prod E_i^{ec_i} \times (1 + \mathcal{N})^{-eqr_1} \cdot r_0^{-e\mathcal{N}} \\ &= \prod E_i^{k_i} \cdot (1 + \mathcal{N})^{-eqr_1} \cdot r_0^{-e\mathcal{N}} = D \cdot (1 + \mathcal{N})^{-q(er_1 + \nu')} \cdot (r_0^e \nu)^{-\mathcal{N}} \bmod \mathcal{N}^2. \end{aligned}$$

One must make sure that $2eqr_1 \leq \mathcal{N}$: with $\mathcal{N} \geq 2\mu L \cdot pq^3$ (where p is the smallest prime factor of q and $2^{\lambda'} < p$), there is no reduction modulo \mathcal{N} before the reduction modulo q . The zero-knowledge proof Π of μ scalars thus contains $2 \log \mathcal{N} + \mu \times \log q$ bits.

In case of multiple equations involving the same secret c_i , the same challenge is used, and the same (k_i, s_i) , reducing the bit-size to $2 \log \mathcal{N} \times \#\text{Equations} + \log q \times \#\text{Secrets}$.

Proofs for a Hiding Commitment in \mathcal{R}_2 . Let us illustrate on the proof of validity of a hiding commitment in $\mathcal{R}_2 = \mathbb{Z}_q[X^{n-1}, Y^N]$, as presented in Section C.3. $\text{Commit}^*(u, \mathcal{R}_2)$, outputs the tuple $C^* = (E_u^* = (E_k^*, E_k^{(2*)})_k, \Pi_u^*)$, where for all indices $k \in \llbracket 1; K \rrbracket$ and $m \in \llbracket 1; M \rrbracket$, with $\rho_k, \rho'_k \xleftarrow{\$} \mathcal{M}$ and $\sigma_k, \sigma'_k \in \mathbb{Z}_{\mathcal{N}}^*$:

$$E_k^* \leftarrow \prod_{j,i} E_{k,j,i}^{u_{j,i}} \times E_{k,0,0}^{\rho_k} \times \sigma_k^{\mathcal{N}} \bmod \mathcal{N}^2$$

$$E_k^{(2*)} \leftarrow \prod_{j,i} E_{k,j,i}^{(2) u_{j,i}} \times E_{k,0,0}^{\rho'_k} \times \sigma'_k{}^{\mathcal{N}} \bmod \mathcal{N}^2$$

with $\text{ZKLQCheck}(X_1 - X_2 \cdot X_3, E_k^{(2*)}, E_k^*, E_{k,0,0}^{(2)}; E_{k,0,0}, E_{k,0,0}^{(2)}) = \text{true}$ for which the verifier can compute the plaintexts in $E_k^{(2*)}$ and E_k^* , and then the quadratic relation a_k , that should be $a_k = \rho'_k \times 1 - \rho_k \times r_k^{(2)} \bmod q$:

$$V_k' = (1 + \mathcal{N})^{a_k} = (1 + \mathcal{N})^{\rho'_k \times 1 - \rho_k \times r_k^{(2)}} \bmod q \bmod \mathcal{N}^2$$

and the encodings in the proof Π_u^* :

$$V_{k,m}^* \leftarrow (1 + \mathcal{N})^{v_m(s_k, t_k) + \rho_{k,m} \bmod q} \cdot \sigma_{k,m}^{\mathcal{N}} \bmod \mathcal{N}^2$$

$$W_{k,m}^* \leftarrow (1 + \mathcal{N})^{w_m(s_k) + \rho'_{k,m} \bmod q} \cdot \sigma'_{k,m}{}^{\mathcal{N}} \bmod \mathcal{N}^2$$

for random $\rho_{k,m}, \rho'_{k,m} \xleftarrow{\$} \mathcal{M}$, chosen by the prover for their privacy, and some unknown $\sigma_{k,m}, \sigma'_{k,m} \in \mathbb{Z}_{\mathcal{N}}^*$, using random $x_m, y_m \xleftarrow{\$} \mathbb{Z}_q$ given by a hash function, and

$$\text{ZKLQCheck}(X_1 - (X_2 - y_m) \cdot X_3 - (X_4 - x_m) \cdot X_5,$$

$$E_k^*, E_{k,1,0}, V_{k,m}^*, E_{k,0,1}, W_{k,m}^*; E_{k,0,0}, E_{k,0,0}, E_{k,1,0}^{(m)}, E_{k,0,1}^{(m)}) = \text{true}$$

where

$$E_{k,1,0}^{(m)} = E_{k,1,0} \cdot E_{k,0,0}^{-y_m} \bmod \mathcal{N}^2 \quad E_{k,0,1}^{(m)} = E_{k,0,1} \cdot E_{k,0,0}^{-x_m} \bmod \mathcal{N}^2$$

Again, the verifier can compute the plaintexts in the input encodings, and the plaintext $b_{k,m}$ to be proven, that should be $b_{k,m} = \rho_k + u(x_m, y_m) - \rho_{k,m}(t_k - y_m) - \rho'_{k,m}(s_k - x_m) \bmod q$. They build $V_{k,m}'$:

$$V_{k,m}' = (1 + \mathcal{N})^{b_{k,m}} = (1 + \mathcal{N})^{\rho_k + u(x_m, y_m) - \rho_{k,m}(t_k - y_m) - \rho'_{k,m}(s_k - x_m) \bmod q} \bmod \mathcal{N}^2$$

As the same ρ_k and the same $u_m = u(x_m, y_m)$ are used many times, the prover first randomly chooses $T_k, T_k', v_m, T_{k,m}, T_{k,m}' \xleftarrow{\$} \llbracket 0; q-1 \rrbracket$, $\nu'_k, \nu'_{k,m} \xleftarrow{\$} \llbracket 0; 2^\lambda Lq^2 - 1 \rrbracket$, $\nu_k, \nu_{k,m} \xleftarrow{\$} \mathbb{Z}_{\mathcal{N}}^*$ and sets

$$D_k = E_{k,0,0}^{T_k'} \cdot (E_{k,0,0}^{(2)})^{-T_k} (1 + \mathcal{N})^{q\nu'_k \nu_k^{\mathcal{N}}} \bmod \mathcal{N}^2$$

$$D_{k,m} = E_{k,0,0}^{T_k} \times E_{k,0,0}^{v_m} \times E_{k,1,0}^{(m) - T_{k,m}} \times E_{k,0,1}^{(m) - T_{k,m}'} (1 + \mathcal{N})^{q\nu'_{k,m} \nu_{k,m}^{\mathcal{N}}} \bmod \mathcal{N}^2$$

The huge range for $\nu'_k, \nu'_{k,m} \xleftarrow{\$} \llbracket 0; 2^\lambda Lq^2 - 1 \rrbracket$ is to hide the random values even if then encodings come from L -linear combinations (which is not the case in this specific proof, but will be for relation (11), with v_κ).

A challenge $e \in \llbracket 0; 2^\lambda - 1 \rrbracket$ is provided by the verifier or drawn from a hash function evaluated on the statement to be proven and all the $(D_k)_k$ and $(D_{k,m})_{k,m}$, and the proof eventually consists of $\Pi = ((D_k)_k, (D_{k,m})_{k,m}, (S_k, S'_k)_k, (w_m)_m, (S_{k,m}, S'_{k,m})_{k,m})$, where

$$S_k = T_k - e\rho_k \bmod q \quad S'_k = T_k' - e\rho'_k \bmod q$$

$$w_m = v_m - eu_m \bmod q$$

$$S_{k,m} = T_{k,m} - e\rho_{k,m} \bmod q \quad S'_{k,m} = T_{k,m}' - e\rho'_{k,m} \bmod q$$

It thus contains $K(M+1)$ ciphertexts (the number of equations), of $2 \log \mathcal{N}$ bits, and $2K(M+1) + M$ scalars in $\llbracket 0; q-1 \rrbracket$ (the number of private masks), of less than $\log q$ bits.

The verifier can first compute the challenge e , and

$$\begin{aligned} D'_k &= E_{k,0,0}^{S'_k} \cdot (E_{k,0,0}^{(2)})^{-S_k} \times (V'_k)^e \bmod \mathcal{N}^2 \\ D'_{k,m} &= E_{k,0,0}^{S_k} \times E_{k,0,0}^{w_m} \times E_{k,1,0}^{(m)} \cdot (E_{k,1,0}^{(m)})^{-S_{k,m}} \times E_{k,0,1}^{(m)} \cdot (E_{k,0,1}^{(m)})^{-S'_{k,m}} \times (V'_{k,m})^e \bmod \mathcal{N}^2 \end{aligned}$$

They then decrypt all the D_k/D'_k and $D_{k,m}/D'_{k,m}$, that should be 0 modulo q , as there is no reduction modulo \mathcal{N} before the reduction modulo q , thanks to the constraint on $\mathcal{N} > 2\mu L \cdot 2^\lambda q^3$.

Soundness. After a rewinding, with a different challenge $\tilde{e} \neq e$, such that D'_k and \tilde{D}'_k decrypt to the same value modulo q , and $D'_{k,m}$ and $\tilde{D}'_{k,m}$ decrypt to the same value modulo q , where

$$\begin{aligned} D'_k &= E_{k,0,0}^{S'_k} \cdot (E_{k,0,0}^{(2)})^{-S_k} \times (V'_k)^e \bmod \mathcal{N}^2 \\ \tilde{D}'_k &= E_{k,0,0}^{\tilde{S}'_k} \cdot (E_{k,0,0}^{(2)})^{-\tilde{S}_k} \times (V'_k)^{\tilde{e}} \bmod \mathcal{N}^2 \\ D'_{k,m} &= E_{k,0,0}^{S_k} \times E_{k,0,0}^{w_m} \times E_{k,1,0}^{(m)} \cdot (E_{k,1,0}^{(m)})^{-S_{k,m}} \times E_{k,0,1}^{(m)} \cdot (E_{k,0,1}^{(m)})^{-S'_{k,m}} \times (V'_{k,m})^e \bmod \mathcal{N}^2 \\ \tilde{D}'_{k,m} &= E_{k,0,0}^{\tilde{S}_k} \times E_{k,0,0}^{\tilde{w}_m} \times E_{k,1,0}^{(m)} \cdot (E_{k,1,0}^{(m)})^{-\tilde{S}_{k,m}} \times E_{k,0,1}^{(m)} \cdot (E_{k,0,1}^{(m)})^{-\tilde{S}'_{k,m}} \times (V'_{k,m})^{\tilde{e}} \bmod \mathcal{N}^2 \end{aligned}$$

we have both

$$\begin{aligned} A_k &= E_{k,0,0}^{S'_k - \tilde{S}'_k} \cdot (E_{k,0,0}^{(2)})^{\tilde{S}_k - S_k} \times (V'_k)^{e - \tilde{e}} \bmod \mathcal{N}^2 \\ A_{k,m} &= E_{k,0,0}^{S_k - \tilde{S}_k} \times E_{k,0,0}^{w_m - \tilde{w}_m} \times E_{k,1,0}^{(m)} \cdot (E_{k,1,0}^{(m)})^{\tilde{S}_{k,m} - S_{k,m}} \times E_{k,0,1}^{(m)} \cdot (E_{k,0,1}^{(m)})^{\tilde{S}'_{k,m} - S'_{k,m}} \times (V'_{k,m})^{e - \tilde{e}} \bmod \mathcal{N}^2 \end{aligned}$$

decrypt to 0 modulo q , while the small size of the answers and the evaluation of V'_k and $V'_{k,m}$ by the verifier on small scalars guarantees no wrap-up modulo \mathcal{N} : $q\alpha_k, q\alpha_{k,m} < 2\mu L q^3 + 2^{\lambda'} q < \mathcal{N}$, even with encodings generated from L -linear combinations in

$$A_k = (1 + \mathcal{N})^{q\alpha_k} \beta_k^{\mathcal{N}} \bmod \mathcal{N}^2 \quad A_{k,m} = (1 + \mathcal{N})^{q\alpha_{k,m}} \beta_{k,m}^{\mathcal{N}} \bmod \mathcal{N}^2$$

If we assume $\tilde{e} - e$ invertible modulo q , and set $\varepsilon = (\tilde{e} - e)^{-1} \bmod q$, we can compute

$$\begin{aligned} \rho'_k &\leftarrow (S'_k - \tilde{S}'_k) \cdot \varepsilon \bmod q & \rho_k &\leftarrow (S_k - \tilde{S}_k) \cdot \varepsilon \bmod q \\ u_m &\leftarrow (w_m - \tilde{w}_m) \cdot \varepsilon \bmod q \\ \rho'_{k,m} &\leftarrow (S'_{k,m} - \tilde{S}'_{k,m}) \cdot \varepsilon \bmod q \end{aligned}$$

all smaller than q

$$\begin{aligned} E_{k,0,0}^{\rho'_k} \cdot (E_{k,0,0}^{(2)})^{-\rho_k} &= V'_k \cdot (1 + \mathcal{N})^{q\alpha'_k} \gamma_k^{\mathcal{N}} \bmod \mathcal{N}^2 \\ E_{k,0,0}^{\rho_k} \times E_{k,0,0}^{u_m} \times E_{k,1,0}^{(m)} \cdot (E_{k,1,0}^{(m)})^{-\rho_{k,m}} \times E_{k,0,1}^{(m)} \cdot (E_{k,0,1}^{(m)})^{-\rho'_{k,m}} &= V'_{k,m} \cdot (1 + \mathcal{N})^{q\alpha'_{k,m}} \gamma_{k,m}^{\mathcal{N}} \bmod \mathcal{N}^2 \end{aligned}$$

Eventually, if $\mathcal{N} > 2\mu L q^3$, all the exponents to $(1 + \mathcal{N})$ remain smaller than \mathcal{N} , which proves the soundness, until $\tilde{e} - e$ is invertible: one can take $2^{\lambda'}$ smaller than the smallest prime factor of q , so that any non-trivial difference will always be invertible, and iterate several times in parallel with multiple challenges e , to increase soundness. For \mathcal{N} , it is safe to take it larger than $2\mu L 2^\lambda q^3$. As μ will always be smaller than 4, in each individual equations, we can take $\mathcal{N} > L 2^{\lambda+3} q^3$.

Zero-Knowledge. Thanks to the random masks in \mathcal{M} , the plaintexts are statistically hidden. The proofs contain tuples $((D_k)_k, (D_{k,m})_{k,m}, (S_k, S'_k)_k, (w_m)_m, (S_{k,m}, S'_{k,m})_{k,m})$, where the ciphertexts statistically hide their representations modulo q , and the scalars are uniformly distributed in $\llbracket 0; q-1 \rrbracket$. Hence, a statistical zero-knowledge proof that guarantees the hiding property of the commitments.

C Commitments

We here develop more on the sketch provided in Section 3. In particular, when q is large and prime, as already shown, under the Schwartz-Zippel lemma, the probability to have equalities between polynomials evaluated in a random point $s \in \mathbb{Z}_q$ whereas equalities do not hold between the polynomials of total degree at most $2n$ is less than $2n/q$, which is negligible. Then, the soundness is straightforward. When q is the product of ℓ distinct prime factors greater than p , this is less than $2n\ell/p$. This probability is unfortunately non-negligible for polynomial prime factors. In particular, for SEAL implementation of FV FHE, the prime factors are at most 60-bit long.

According to the expected soundness parameter ε , to reduce the probability of false positive cases, the natural solution is to iterate K times, with multiple evaluation points s_k , for $k \in \llbracket 1; K \rrbracket$, so that $(2n\ell/p)^K \leq \varepsilon$. But then, we have to make sure the same polynomials are committed in each point. This is discussed below.

C.1 Binding Commitments

Because of the linear-only combinations, one can limit encodings in sub-spaces. As we can also do quadratic verifications, we will be able to check products of two polynomials. From an encoding scheme (Gen, E) over a ring \mathbb{R} , we can define a compact binding commitment scheme over multivariate polynomials. More precisely, such commitment schemes will be defined by 4 algorithms:

- $\text{Setup}(1^\lambda, \mathcal{R}, (\mathcal{R}_\pi)_\pi)$ generates the public key pk and the verification key vk , according to the polynomial space \mathcal{R} , and the authorized subspaces $(\mathcal{R}_\pi)_\pi$;
- $\text{Commit}_{\text{pk}}(u, \mathcal{R}_\pi)$, for a polynomial $u \in \mathcal{R}_\pi \subseteq \mathcal{R}$, outputs a commitment C of u ;
- $\text{Validity}_{\text{vk}}(C, \mathcal{R}_\pi)$, for a commitment C , verifies whether this is a valid commitment for an (unknown) polynomial u in the appropriate subspace \mathcal{R}_π ;
- $\text{Quadratic}_{\text{vk}}(Q, C_1, \dots, C_\ell)$, for valid commitments C_i of (unknown) polynomials u_i and a quadratic polynomial Q in ℓ variables, verifies whether $Q(u_1, \dots, u_\ell) = 0$ in \mathcal{R} .

For the above definition to make sense, for any adversary \mathcal{A} , there exists an extractor $\text{Ext}_{\mathcal{A}}$ such that for any valid commitment C generated by \mathcal{A} , $\text{Ext}_{\mathcal{A}}$ outputs the committed polynomial $u \in \mathcal{R}_\pi$.

While \mathcal{R} will usually be a global ring of polynomials, such as $\mathbb{R}[X]$ or $\mathbb{R}[X, Y]$, the sub-spaces \mathcal{R}_π will only be the module generated by a limited basis, $\mathbb{R}[X^{n-1}]$, $\mathbb{R}[Y^N]$, or $\mathbb{R}[X^{n-1}, Y^N]$, where the explicit exponents limit the degrees.

A Warm-Up with the Ring of Polynomials $\mathcal{R} = \mathbb{R}[X]$. Before dealing with multivariate polynomials, we start with univariate polynomials, to illustrate some requirements and some issues:

- $\text{Setup}(1^\lambda, \mathcal{R} = \mathbb{R}[X], (\mathcal{R}_1 = \mathbb{R}[X^{n-1}]))$ first runs $(\text{pk}', \text{sk}', \text{vk}') \leftarrow \text{Gen}(1^\lambda)$, chooses a random element $s \xleftarrow{\$} \mathbb{R}^*$ and, for $i \in \llbracket 0; n-1 \rrbracket$, sets $E_i \leftarrow \text{E}_{\text{sk}'}(s^i)$. Then, the public key of the commitment scheme is $\text{pk} = (\text{pk}', \{E_i\}_i)$, while the verification key is $\text{vk} = \text{vk}'$;
- $\text{Commit}_{\text{pk}}(u, \mathcal{R}_1)$, for a polynomial $u = \sum_{i=0}^{n-1} u_i X^i \in \mathcal{R}_1 \subset \mathcal{R}$ of degree at most $n-1$, outputs $E = \text{Eval}(\{E_i\}_i, \{u_i\}_i) = \text{E}(\sum_i u_i s^i) = \text{E}(u(s))$;

- $\text{Validity}_{\text{vk}}(E)$ is exactly the `Verify` of the encoding scheme, as it outputs whether this is a valid encoding or not, and thus a valid commitment or not.

Thanks to the linear-only extractability, when a player generates a valid encoding (or commitment) E , being only given (E_0, \dots, E_{n-1}) , one can extract (c_i) such that E is an encoding of $c_0 + c_1s + \dots + c_{n-1}s^{n-1}$ in \mathbb{R} , and thus of the polynomial $\mathbf{c} = \sum_i c_i X^i$ in \mathcal{R} . Our above commitment scheme on polynomials is thus extractable under the linear-only extractability of the secure encoding scheme.

In addition, thanks to the quadratic verification on the encodings, if we have four polynomials $\mathbf{u}, \mathbf{v}, \mathbf{m}$ and \mathbf{r} such that $\mathbf{m} = \mathbf{u} \cdot \mathbf{v} \bmod \mathbf{r}$, which means that $\mathbf{m} = \mathbf{u} \cdot \mathbf{v} + \mathbf{r} \cdot \mathbf{q}$, where all the polynomials are of degree at most $n-1$, we can check such a product: from valid commitments U and V of \mathbf{u} and \mathbf{v} , R and Q of \mathbf{r} and \mathbf{q} , respectively, and M of the polynomial \mathbf{m} , all of degree at most $n-1$, as they are all simple encodings, $\text{QCheck}(X_1 X_2 + X_3 X_4 - X_5, U, V, R, Q, M) = \text{true}$ implies that $\mathbf{m}(s) = \mathbf{u}(s) \cdot \mathbf{v}(s) + \mathbf{r}(s) \cdot \mathbf{q}(s)$. According to the ring \mathbb{R} , this might imply the expected relation, or not. If \mathbb{R} is a large enough field, this is true.

However, in the following, we will be interested in the particular case of $\mathbb{R} = \mathbb{Z}_q$, with $q = p_1 \cdot \dots \cdot p_\ell$ a product of ℓ prime integers $p_1 < \dots < p_\ell$, greater than p . Having $\mathbf{m}(s) = \mathbf{u}(s) \cdot \mathbf{v}(s) + \mathbf{r}(s) \cdot \mathbf{q}(s) \bmod q$ while $\mathbf{m} \neq \mathbf{u} \cdot \mathbf{v} + \mathbf{r} \cdot \mathbf{q}$ in $\mathbb{Z}_q[X]$ means there is an index $j \in \llbracket 1; \ell \rrbracket$ such that $\mathbf{m}(s) = \mathbf{u}(s) \cdot \mathbf{v}(s) + \mathbf{r}(s) \cdot \mathbf{q}(s) \bmod p_j$ while $\mathbf{m} \neq \mathbf{u} \cdot \mathbf{v} + \mathbf{r} \cdot \mathbf{q}$ in $\mathbb{Z}_{p_j}[X]$. Under the Schwartz-Zippel lemma [Sch80, Zip79], this probability is bounded by $2n/p_j$ for each j , as the total degree of the relation is at most $2n$. Hence, the probability over s to have a false positive is bounded by $2n\ell/p$. This probability is unfortunately non-negligible for polynomial prime factors.

C.2 Commitments on Bivariate Polynomials

We now build commitments on bivariate polynomials over \mathbb{Z}_q for a composite q , providing a way to deal with scalars and univariate polynomials, when they are evaluated in one-fixed coordinate or a fixed point. In addition, a bivariate polynomial is a way to encode multiple univariate polynomials: given N polynomials $\mathbf{u}_j = \sum_i u_{j,i} X^i \in \mathbb{Z}_q[X]$, of degree $n-1$, we can consider the polynomial in $\mathbb{Z}_q[X, Y]$: $\mathbf{u}(X, Y) = \sum_j Y^j \mathbf{u}_j(X) = \sum_{j,i} u_{j,i} X^i Y^j$.

As explained above, in order to reduce the probability of errors with the Schwartz-Zippel lemma, we will use encodings in K multiple points. We will additionally prove they all encode the same polynomial.

$\text{Setup}(1^\lambda, \mathcal{R} = \mathbb{Z}_q[X, Y], \mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N])$ first runs $(\text{pk}', \text{sk}', \text{vk}') \leftarrow \text{Gen}(1^\lambda)$, chooses K tuples of random elements $s_k, t_k \xleftarrow{\$} \mathbb{Z}_q^*$ and, for $k \in \llbracket 1; K \rrbracket$, $i \in \llbracket 0; n-1 \rrbracket$, and $j \in \llbracket 0; N \rrbracket$, sets $E_{k,j,i} \leftarrow \text{E}_{\text{sk}'}(s_k^i \cdot t_k^j)$. To explicitly limit to some degrees, such as $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$ or even to univariate polynomials. One also chooses K tuples of random elements $r_k^{(3)} \xleftarrow{\$} \mathbb{Z}_q^*$. Then, for $k \in \llbracket 1; K \rrbracket$, $i \in \llbracket 0; n-1 \rrbracket$, and $j \in \llbracket 0; N \rrbracket$, one sets $E_{k,j,i}^{(3)} \leftarrow \text{E}_{\text{sk}'}(r_k^{(3)} \cdot s_k^i \cdot t_k^j)$ for bivariate polynomials in $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$. The public key of the commitment scheme is composed of the above encodings $\text{pk} = (\text{pk}', \{E_{k,j,i}, E_{k,j,i}^{(3)}\}_{k,j,i})$, and the verification key is $\text{vk} = \text{vk}'$. For our application, other sub-spaces will be added, but in this section, we focus on $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$. More spaces are detailed in Appendix C, with other random values $r_k^{(l)}$.

$\text{Commit}(\mathbf{u}, \mathbb{Z}_q[X^{n-1}, Y^N])$, for a polynomial $\mathbf{u} = \sum_{j=0}^N \sum_{i=0}^{n-1} u_{j,i} X^i Y^j$ in $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$, outputs $C = (E_{\mathbf{u}} = (E_k, E_k^{(2)})_k, \Pi_{\mathbf{u}})$, for $k \in \llbracket 1; K \rrbracket$, with $\Pi_{\mathbf{u}}$ detailed later, where

$$\begin{aligned} E_k &\leftarrow \text{Eval}(\{E_{k,j,i}\}_{j,i}, \{u_{j,i}\}_{j,i}) = \text{E}(\mathbf{u}(s_k, t_k)) \\ E_k^{(3)} &\leftarrow \text{Eval}(\{E_{k,j,i}^{(3)}\}_{j,i}, \{u_{j,i}\}_{j,i}) = \text{E}(r_k^{(3)} \cdot \mathbf{u}(s_k, t_k)) \\ &\text{such that } \text{QCheck}(X_1 - X_2 \cdot X_3, E_k^{(3)}, E_k, E_{k,0,0}^{(3)}) = \text{true} \end{aligned}$$

The idea behind the twin encodings $E_u = (E_k, E_k^{(3)})_k$ is that the first element E_k is compatible between all the polynomials in $\mathcal{R} = \mathbb{Z}_q[X, Y]$, independently of the constraints on the allowed monomials, while the second element restrains the polynomial space: the limited basis in $\{E_{k,j,i}^{(3)}\}_{j,i}$ and the relation with $E_{k,0,0}^{(3)}$ limits to $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$. Note there is still a non-negligible probability of false positives when accepting a twin encoding, as remarked above, as the quadratic check might accept the relation on the specific points without the relation holding on the polynomials. We will say an encoding E_k is *valid* when the relation really holds on the extracted polynomials from the twin encodings $(E_k, E_k^{(3)})$. Intuitively, for each index k , the probability of false positive (accepting an invalid twin encoding) is bounded by $\ell\mathcal{D}/p$, where \mathcal{D} will be the maximal total degree of the polynomials that one will be able to generate from the elements in the basis provided as input, as $E_{k,0,0}^{(3)}$ encodes a polynomial of degree 0. An additional proof Π_u (either provided from an interactive protocol or built in a non-interactive way) is thus appended to the commitment to make the validity check sound, with error-probability ε_c for each commitment: a huge fraction of the encodings are valid and encode the same polynomial.

In this section, we target the soundness of the commitments, whereas our ultimate goal will be a soundness bound ε_s for our global proof. To achieve this soundness bound, we will first require all the commitments to be correct, excepted an error probability $\varepsilon_s/3$, the relations between the commitments to also hold excepted with error probability $\varepsilon_s/3$, and the bounds on the noise-flooding to be small enough with error probability $\varepsilon_s/3$. Hence, we use a specific soundness parameter $\varepsilon_c \leq \varepsilon_s/3\nu_c$ for commitments, where ν_c will be the total number of commitments.

Validity(C) first verifies the twin encodings, which checks the appropriate sub-spaces for each E_k : the quadratic check with the limited bases in $\{E_{k,j,i}^{(3)}\}_{k,j,i}$ guarantees the limited list of monomials, but with an error probability bounded by $2\ell\mathcal{D}/t$. Note that in this section, $\mathcal{D} = N + n - 1$, because of the limited \mathcal{R}_3 , but monomials with higher degrees will be required later. With K large enough, we can expect a large number of valid encodings E_k : let us assume more than $K/5$ encodings are not valid, this will remain undetected with probability at most $(\ell\mathcal{D}/p)^{K/5}$. If we constrain p with $p \geq 2^S \times 2\ell\mathcal{D}$ (where S will be seen as a security margin, all along this analysis), the error probability is less than $2^{-(S+1)K/5}$. Hence, the number of valid encodings E_k is greater than $4K/5$ excepted with probability upper-bounded by $2^{-(S+1)K/5}$.

But this is not enough to amplify the above Schwartz-Zippel lemma: one also has to make sure that all the valid E_k encode the same polynomial u to exploit the iterations in the quadratic check between encoded polynomials:

$$\begin{aligned} u_k(X, Y) - u(X', Y') &= u(X, Y) - u(X', Y') \\ &= u(X, Y) - u(X, Y') + u(X, Y') - u(X', Y') \\ &= (Y - Y') \cdot v(X, Y, Y') + (X - X') \cdot w(X, X', Y'). \end{aligned}$$

This can be checked with random $x_m, y_m \xleftarrow{\$} \mathbb{Z}_q$, sent by the verifier:

$$u(X, Y) - u(x_m, y_m) = (Y - y_m) \cdot v_m(X, Y) + (X - x_m) \cdot w_m(X),$$

from the proof

$$\Pi_u = (u_m \leftarrow u(x_m, y_m), (V_{k,m} \leftarrow \mathbb{E}(v_m(s_k, t_k)), W_{k,m} \leftarrow \mathbb{E}(w_m(s_k))))_k)_m$$

that can be checked as

$$\begin{aligned} \text{QCheck}(X_1 - u_m - (X_2 - y_m) \cdot X_3 - (X_4 - x_m) \cdot X_5, \\ E_k, E_{k,1,0}, V_{k,m}, E_{k,0,1}, W_{k,m}) = \text{true} \end{aligned}$$

For each $k \in \llbracket 1; K \rrbracket$, if the above relations do not hold at the polynomial level (which means for the polynomials $v_{k,m}$, $w_{k,m}$, encoded in $V_{k,m}$, and $W_{k,m}$ respectively, possibly in any subspace)

for more than $1/3$ of the $m \in \llbracket 1; M \rrbracket$, they will remain undetected with probability at most $(\ell(\mathcal{D} + 1)/p)^{M/3} \leq 2^{-(S+1)M/3}$. Otherwise

$$u_k(X, Y) - u_m = (Y - y_m) \cdot v_{k,m} + (X - x_m) \cdot w_{k,m}$$

for more than $2M/3$ indices m . Then, for any $k' \neq k$, that correspond to valid encodings E_k and $E_{k'}$, at least $M/3$ common values (x_m, y_m) satisfy both relations in k and k' , hence $u_k(x_m, y_m) = u_{k'}(x_m, y_m) = u_m$. As a consequence, as the polynomials u_k and $u_{k'}$ were committed before seeing (x_m, y_m) , there are $M/3$ random points in which the two polynomials are equal. Then $u_k = u_{k'}$, excepted with probability upper-bounded by $(\ell\mathcal{D}/p)^{M/3} \leq 2^{-2M}$. A false acceptance for some pair (k, k') of consecutive valid encodings is upper-bounded by $2K \cdot 2^{-(S+1)M/3}$. Globally, the probability of having a false positive among the valid encodings E_k is bounded by $2K \cdot 2^{-(S+1)M/3}$.

We thus complete the commitments in \mathcal{R}_3 : in addition to the twin encodings $E_u = (E_k, E_k^{(3)})_k$, after having received (or seen) $(x_m, y_m)_m$, one completes the commitment with the proofs Π_u , for $k \in \llbracket 1; K \rrbracket$, $m \in \llbracket 1; M \rrbracket$

– in $\text{Commit}(u, \mathbb{Z}_q[X^{n-1}, Y^N])$, $\Pi_u = (u_m, (V_{k,m}, W_{k,m})_k)_m$;

Then, we can state the following security result, which can be extended to other subspaces:

Theorem 8 (Knowledge-Soundness of Commitments in \mathcal{R}_2). *For any commitment $C = (E = (E_k, E_k^{(3)})_k, \Pi = (u_m, (V_{k,m}, W_{k,m})_k)_m)$, if C successfully passes all the validity checks and quadratic root checks, on randomly chosen $(x_m, y_m) \xleftarrow{\$} \mathbb{Z}_q^2$, there exists a polynomial $u \in \mathcal{R}_3$ such that at least $4K/5$ of the twin encodings actually encode u (which can be extracted from the extractability of valid encodings), excepted with probability less than $2^{-(S+1)K/5} + 2K \cdot 2^{-(S+1)M/3}$, where $q = \prod_{i=1}^{\ell} p_i$, with all $p_i \geq 2^S \times 2\ell\mathcal{D}$, and \mathcal{D} is the maximal degree of the polynomials in the subspace.*

Quadratic Root Checks. According to the above analysis, if we assume $p \geq 2^S \times 2\ell\mathcal{D}$, for all the accepted commitments, we know that we have at least $4K/5$ valid encodings of the same polynomials in all twin encodings:

- if the number of invalid encodings is greater than $K/5$, they will remain undetected with probability less than $2^{-(S+1)K/5}$.
- all these valid encodings contain the same polynomial u , excepted with probability less than $2K \cdot 2^{-(S+1)M/3}$.

Therefore, when all the verifications succeed for the commitment, there are at least $4K/5$ valid encodings on the same polynomial u , excepted with small error probability. Hence, when 4 commitments are involved in a quadratic check, at least $K/5$ common indices k correspond to valid encodings on the same polynomials in the 4 commitments. Those polynomials then satisfy the relation excepted with probability less than $(2\ell\mathcal{D}/p)^{K/5} \leq 2^{-S \cdot K/5}$.

More Parameters. Commitments in $\mathbb{Z}_q[X^{n-1}]$ thus consist of $2K + KM = K(M + 2)$ encodings, and M scalars in \mathbb{Z}_q while commitments in $\mathbb{Z}_q[X^{n-1}, Y^N]$ consist of $2K + 2KM = 2K(M + 1)$ encodings, and M scalars in \mathbb{Z}_q .

The above analysis was for commitments that appear in quadratic checks involving $c = 4$ commitments, hence one required $4K/5$ valid encodings. When other values of c in $\{2, 3\}$, it is enough to have $cK/(c + 1)$ valid encodings in the c commitments to have $K/(c + 1)$ common indices, and then the soundness of the equation is $2^{-S \cdot K/(c+1)}$.

Corollary 9. *When all the tests pass in a quadratic check between at most c prover-generated commitments, that are all valid, the relation is really satisfied on the committed polynomials (which can be extracted from the extractability of valid encodings), excepted with probability less than $2^{-S \cdot K/(c+1)}$.*

Bounds on K and M are detailed in Appendix F.

C.3 Hiding Commitments

These commitments with encodings are strongly binding, because of the knowledge-soundness, but are not hiding, as sometimes expected from commitments. Because of the quadratic verification, if the verifier hesitates between two polynomials in a commitment C , they can commit them and do a simple linear verification, as they know \mathbf{vk} .

To statistically hide the content, one has to add random blinding elements from the appropriate masking set \mathcal{M} to every encodings, to get hiding encodings. We illustrate this here with $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$. We later detail the case of $\mathcal{R}_4 = \mathbb{Z}_q[Y^N]$.

$\text{Commit}^*(\mathbf{u}, \mathbb{Z}_q[X^{n-1}, Y^N])$, the hiding commitment for a bivariate polynomial $\mathbf{u}(X, Y) = \sum_{j=0}^N \sum_{i=0}^{n-1} u_{j,i} X^i Y^j \in \mathcal{R}_2 = \mathbb{Z}_q[X^{n-1}, Y^N]$, outputs the tuple $C^* = (E_{\mathbf{u}}^* = (E_k^*, E_k^{(3*)}, \pi_k)_k, \Pi_{\mathbf{u}}^*)$, where for all indices $k \in \llbracket 1; K \rrbracket$ and $m \in \llbracket 1; M \rrbracket$, with $\rho_k, \rho'_k \xleftarrow{\$} \mathcal{M}$:

$$\begin{aligned} E_k^* &\leftarrow \text{Eval}(\{E_{k,j,i}\}_{j,i}, \{E_{k,0,0}\}, \{u_{j,i}\}_{j,i}, \{\rho_k\}) = \mathbb{E}(\mathbf{u}(s_k, t_k) + \rho_k) \\ E_k^{(3*)} &\leftarrow \text{Eval}(\{E_{k,j,i}^{(2)}\}_{j,i}, \{E_{k,0,0}\}, \{u_{j,i}\}_{j,i}, \{\rho'_k\}) = \mathbb{E}(r_k^{(3)} \cdot \mathbf{u}(s_k, t_k) + \rho'_k) \\ \text{with } \pi_k &= \{\text{ZKLQCheck}(X_1 - X_2 \cdot X_3, E_k^{(3*)}, E_k^*, E_{k,0,0}^{(3)}; E_{k,0,0}, E_{k,0,0}^{(3)}) = \text{true}\} \end{aligned}$$

as the quadratic relation is equal to $\rho'_k \times 1 - \rho_k \times r_k^{(3)}$, with private scalars ρ_k and ρ'_k , and a proof, using random $x_m, y_m \xleftarrow{\$} \mathbb{Z}_q$ invertible modulo \mathcal{N}^2 sent by the verifier, of

$$\begin{aligned} \mathbf{u}(X, Y) - \mathbf{u}(x_m, y_m) &= (Y - y_m) \cdot \mathbf{v}_m(X, Y) + (X - x_m) \cdot \mathbf{w}_m(X) \\ &= (Y - y_m) \cdot \mathbf{v}_m + (X - x_m) \cdot \mathbf{w}_m \end{aligned}$$

which can be verified with

$$\Pi_{\mathbf{u}}^* = (V_{k,m}^* \leftarrow \mathbb{E}(\mathbf{v}_m(s_k, t_k) + \rho_{k,m}), W_{k,m}^* \leftarrow \mathbb{E}(\mathbf{w}_m(s_k) + \rho'_{k,m}), \pi_{k,m})_{k,m}$$

for random $\rho_{k,m}, \rho'_{k,m} \xleftarrow{\$} \mathcal{M}$, chosen by the prover for their privacy, and

$$\begin{aligned} \pi_{k,m} &= \{\text{ZKLQCheck}(X_1 - (X_2 - y_m) \cdot X_3 - (X_4 - x_m) \cdot X_5, \\ &\quad E_k^*, E_{k,1,0}, V_{k,m}^*, E_{k,0,1}, W_{k,m}^*; E_{k,0,0}, E_{k,0,0}, E_{k,1,0}^{(m)}, E_{k,0,1}^{(m)}) = \text{true}\} \end{aligned}$$

where anyone can compute:

$$\begin{aligned} E_{k,1,0}^{(m)} &= \text{Eval}(\{E_{k,1,0}, E_{k,0,0}\}, \{1, -y_m\}) = \mathbb{E}(t_k - y_m) \\ E_{k,0,1}^{(m)} &= \text{Eval}(\{E_{k,0,1}, E_{k,0,0}\}, \{1, -x_m\}) = \mathbb{E}(s_k - x_m) \end{aligned}$$

as the quadratic relation is equal to $\rho_k + \mathbf{u}(x_m, y_m) - \rho_{k,m}(t_k - y_m) - \rho'_{k,m}(s_k - x_m)$. One thus proves their knowledge of the 4 private scalars, $\rho_k \in \mathcal{M}$ (the same as above), $u_m = \mathbf{u}(x_m, y_m) \in \mathbb{R}$ (the same for all the k 's), and $\rho_{k,m}, \rho'_{k,m} \in \mathcal{M}$.

Let us assume that a reasonable fraction of the twin encodings $(E_k^*, E_k^{(3*)})$ are valid for a polynomial \mathbf{u}_k^* in $\mathbb{Z}_q[X^{n-1}, Y^N]$. For each k , the quadratic check guarantees that

$$\begin{aligned} \mathbf{u}_k^*(X, Y) &= \mathbf{u}_k(X, Y) + \rho_k = (Y - y_m) \cdot \mathbf{v}_{k,m} + (X - x_m) \cdot \mathbf{w}_{k,m} \\ &\quad + \rho_k + u_m - \rho_{k,m} \cdot (Y - y_m) - \rho'_{k,m} \cdot (X - x_m) \end{aligned}$$

excepted with the same error probability as in Theorem 8. On the other hand, one can state:

Theorem 10 (Hiding Property of Commitments in \mathcal{R}_3). *For any commitment $C^* = (E^* = (E_k^*, E_k^{(3*)}, \pi_k)_k, \Pi^* = (V_{k,m}^*, W_{k,m}^*, \pi_{k,m})_{k,m})$, thanks to the random masks, and the zero-knowledge proofs, all the encodings are statistically indistinguishable from random encodings.*

Note that for a hiding commitment, we have zero-knowledge proofs on $K(M + 1)$ equations involving globally $2K(M + 1) + M$ private scalars (but at most 4 in each equation). In Appendix B.3, we detail the zero-knowledge proofs when the encoding relies on the Paillier's encryption scheme: soundness is $2^{-\lambda'}$, for $\lambda' < \log_2 p$, then one needs to iterate $-\log_2(\varepsilon_s/3\nu_c)/\lambda'$ times. Each proof consists of $K(M + 1)$ Paillier ciphertexts and $2K(M + 1) + M$ scalars in \mathbb{Z}_q , iterated $\log(3\nu_c/\varepsilon_s)/\log p$ times. Soundness also implies the RSA modulus needs to be larger than $L \cdot 2^{\lambda+3}q^3$.

C.4 Quadratic Root Check

The check $\text{Quadratic}(Q, C_1, \dots, C_\ell)$, on non-hiding commitments, for a quadratic polynomial Q , that verifies whether the committed polynomials P_i 's in the C_i 's satisfy the relation $Q(P_1, \dots, P_\ell) = 0$, is performed as a QCheck on the encodings for each index k . Since a huge fraction of the indices k encode the same polynomials, for the prover-generated commitments, iterations amplify the soundness.

When some hiding commitments are involved, one additionally has to prove the appropriate blinding factors. More concretely, let us consider the binding commitment D of \mathbf{d} , and the hiding commitments C^* and H^* of \mathbf{c}^* and \mathbf{h}^* respectively, with blinding factors $\rho_k, \sigma_k \in \mathcal{M}$ respectively. The relation $\mathbf{c}^* = \mathbf{d} \times \mathbf{h}^*$ can be checked as:

$$\begin{aligned} \text{Quadratic}(X_1 - X_2 \cdot X_3, C^*, D, H^*) \\ = \text{ZKLQCheck}(X_1 - X_2 \cdot X_3, C_k^*, D_k, H_k^*; E_{k,0,0}, D_k) \text{ for all } k \end{aligned}$$

as the quadratic relation is equal to $\rho_k \cdot 1 - \sigma_k \cdot \mathbf{d}$, where ρ_k and σ_k are the same private values as the ones used in the validity verification of the hiding commitments C^* and H^* .

We stress that we only allow quadratic verifications where at most one polynomial is committed in a hiding way in each quadratic product. In the end, we know that the quadratic equations among polynomials are satisfied in at least $K/5$ random points. They are thus satisfied by the polynomials excepted with probability less than $2^{-S \cdot K/5}$, which is chosen to be much less than $\varepsilon_s/3\nu_e$, where ν_e is the global number of equations to be checked.

C.5 Commitments in Additional Subspaces

As already explained, in our application, we are considering $q = p_1 \cdot \dots \cdot p_\ell$, a composite modulus q , with ℓ prime factors $p_1 < \dots < p_\ell$ larger than p . We will work in subspaces of $\mathcal{R} = \mathbb{Z}_q[X^{2n-2}, Y^{2N}]$. We will consider the subspaces $\mathcal{R}_1 = \mathbb{Z}_q[X^{n-1}]$, $\mathcal{R}_2 = \mathbb{Z}_q[X^{(2n-2) \setminus (n-1)}]$, $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$, $\mathcal{R}_4 = \mathbb{Z}_q[Y^N]$, and $\mathbb{Z}_q[Y^{2N}]$ with no term in Y^N , denoted $\mathcal{R}_5 = \mathbb{Z}_q[Y^{2N \setminus N}]$, hence $\mathcal{D} = 2N$, so we will assume $p \geq 2^S \times 4N\ell$. Details on the complete construction of the commitments with system setup and detailed contents and checks are given in the following section.

In our proof of inner product, we will use polynomials in multiple subspaces of $\mathcal{R} = \mathbb{Z}_q[X, Y]$. In addition to $\mathcal{R}_2 = \mathbb{Z}_q[X^{n-1}, Y^N]$, we will use $\mathcal{R}_1 = \mathbb{Z}_q[X^{n-1}]$, $\mathcal{R}_2 = \mathbb{Z}_q[X^{(2n-2) \setminus (n-1)}]$, $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$, $\mathcal{R}_4 = \mathbb{Z}_q[Y^N]$, and $\mathbb{Z}_q[Y^{2N}]$ with no term in Y^N , denoted $\mathcal{R}_5 = \mathbb{Z}_q[Y^{2N \setminus N}]$. For this, we draw additional random elements $r_k^{(1)}, r_k^{(2)}, r_k^{(4)}, r_k^{(5)} \xleftarrow{\$} \mathbb{Z}_q^*$ for $k \in \llbracket 1; K \rrbracket$, and set:

$$\begin{aligned} E_{k,i}^{(1)} &\leftarrow \mathbb{E}(r_k^{(1)} \cdot s_k^i) & i \in \llbracket 0; n-1 \rrbracket \\ E_{k,i}^{(2)} &\leftarrow \mathbb{E}(r_k^{(2)} \cdot s_k^i) & i \in \llbracket 0; 2n-2 \rrbracket \setminus \{n-1\} \\ E_{k,j}^{(4)} &\leftarrow \mathbb{E}(r_k^{(4)} \cdot t_k^j) & j \in \llbracket 0; N \rrbracket \\ E_{k,j}^{(5)} &\leftarrow \mathbb{E}(r_k^{(5)} \cdot t_k^j) & j \in \llbracket 0; 2N \rrbracket \setminus \{N\} \end{aligned}$$

and we extend

$$E_{k,j,i} \leftarrow \mathbf{E}(s_k^i \cdot t_k^j) \quad i \in \llbracket 0; 2n - 2 \rrbracket, j \in \llbracket 0; 2N \rrbracket$$

Note that \mathcal{D} becomes $\max\{N + n - 1, 2N, 2n - 2\} = 2N$, for $N \geq n$. We then commit the polynomials $\text{Commit}(u, \mathbb{Z}_q[Y^N])$ or $\text{Commit}(u, \mathbb{Z}_q[Y^{2N \setminus N}])$ with appropriate twin encodings, that can be verified as above: for each pair $k \neq k'$ of valid encodings:

$$u_k(Y) - u_{k'}(Y') = u(Y) - u(Y') = (Y - Y') \cdot v(Y, Y')$$

and so for a random $y_m \xleftarrow{\$} \mathbb{Z}_q$ chosen by the prover:

$$u(Y) - u(y_m) = (Y - y_m) \cdot v(Y, y_m) = (Y - y_m) \cdot v_m,$$

from

$$\Pi_u = (u_m \leftarrow u(y_m), (V_{k,m} \leftarrow \mathbf{E}(v_m(t_k)))_k)_m$$

as

$$\text{QCheck}(X_1 - u_m - (X_2 - y_m) \cdot X_3, E_k, E_{k,1,0}, V_{k,m}) = \text{true}$$

and the rest of the proof follows as in Section C.2, with thus $\Pi_u = (u_m, (V_{k,m})_k)_m$, for $k \in \llbracket 1; K \rrbracket$, $m \in \llbracket 1; M \rrbracket$.

C.6 Complete Construction of the Commitment

In the following, we are considering $q = p_1 \cdot \dots \cdot p_\ell$, a composite modulus q , with ℓ distinct prime factors $p_1 < \dots < p_\ell$ larger than p . We will work in subspaces of $\mathcal{R} = \mathbb{Z}_q[X^{2n-2}, Y^{2N}]$. We will consider the subspaces $\mathcal{R}_1 = \mathbb{Z}_q[X^{n-1}]$, $\mathcal{R}_2 = \mathbb{Z}_q[X^{(2n-2) \setminus (n-1)}]$, $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$, $\mathcal{R}_4 = \mathbb{Z}_q[Y^N]$, and $\mathcal{R}_5 = \mathbb{Z}_q[Y^{2N \setminus N}]$, hence $\mathcal{D} = 2N$, so we will assume $p \geq 128N\ell$.

Setup of the System: $\text{Setup}(1^\lambda, \mathcal{R}, (\mathcal{R}_i)_i)$ first runs $(\mathbf{pk}', \mathbf{vk}') \leftarrow \text{Gen}(1^\lambda)$, chooses K tuples of random elements $s_k, t_k \xleftarrow{\$} \mathbb{Z}_q^*$, as well as K tuples of random elements $r_k^{(1)}, r_k^{(2)}, r_k^{(3)}, r_k^{(4)}, r_k^{(5)} \xleftarrow{\$} \mathbb{Z}_q^*$, to limit the combinations of the bases. Then, for $k \in \llbracket 1; K \rrbracket$, one sets

$$\begin{aligned} \mathcal{R} &= \mathbb{Z}_q[X^{n-1}, Y^N] & E_{k,j,i} &\leftarrow \mathbf{E}(s_k^i \cdot t_k^j) & (i, j) &\in (\llbracket 0; n - 1 \rrbracket \times \llbracket 0; N \rrbracket) \\ &+ \mathbb{Z}_q[X^{2n-2}] + \mathbb{Z}_q[Y^{2N}] & & & &\cup (\llbracket n; 2n - 2 \rrbracket \times \{0\}) \\ & & & & &\cup (\{0\} \times \llbracket N + 1; 2N \rrbracket) \\ \mathcal{R}_1 &= \mathbb{Z}_q[X^{n-1}] & E_{k,i}^{(1)} &\leftarrow \mathbf{E}(r_k^{(1)} \cdot s_k^i) & i &\in \llbracket 0; n - 1 \rrbracket \\ \mathcal{R}_2 &= \mathbb{Z}_q[X^{(2n-2) \setminus (n-1)}] & E_{k,i}^{(2)} &\leftarrow \mathbf{E}(r_k^{(2)} \cdot s_k^i) & i &\in \llbracket 0; 2n - 2 \rrbracket \setminus \{n - 1\} \\ \mathcal{R}_3 &= \mathbb{Z}_q[X^{n-1}, Y^N] & E_{k,j,i}^{(3)} &\leftarrow \mathbf{E}(r_k^{(3)} \cdot s_k^i \cdot t_k^j) & i &\in \llbracket 0; n - 1 \rrbracket, j \in \llbracket 0; N \rrbracket \\ \mathcal{R}_4 &= \mathbb{Z}_q[Y^N] & E_{k,j}^{(4)} &\leftarrow \mathbf{E}(r_k^{(4)} \cdot t_k^j) & j &\in \llbracket 0; N \rrbracket \\ \mathcal{R}_5 &= \mathbb{Z}_q[Y^{2N \setminus N}] & E_{k,j}^{(5)} &\leftarrow \mathbf{E}(r_k^{(5)} \cdot t_k^j) & j &\in \llbracket 0; 2N \rrbracket \setminus \{N\} \end{aligned}$$

Then, the public key of the commitment scheme is set to \mathbf{pk}' with all these encodings, while the verification key is \mathbf{vk}' . For \mathcal{R} , we can limit encoded elements to $\mathbb{Z}_q[X^{n-1}, Y^N] + \mathbb{Z}_q[X^{2n-2}] + \mathbb{Z}_q[Y^{2N}]$, with $(n + 1) \times N + 2n - 1$ encodings in the public key, sent once for many proofs.

Commitment Generation: there are two binding commitment algorithms, with or without the hiding property.

Non-Hiding Commitment. We have defined commitments on specific subspaces: $\mathcal{R}_1 = \mathbb{Z}_q[X^{n-1}]$, $\mathcal{R}_2 = \mathbb{Z}_q[X^{(2n-2)\setminus(n-1)}]$, $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$, $\mathcal{R}_4 = \mathbb{Z}_q[Y^N]$, and $\mathcal{R}_5 = \mathbb{Z}_q[Y^{2N\setminus N}]$:

$\text{Commit}(u, \mathbb{Z}_q[X^{n-1}])$: it outputs $C = (E_u = (E_k, E_k^{(1)})_k, \Pi_u)$, for $k \in \llbracket 1; K \rrbracket$, where

$$E_k \leftarrow \mathbf{E}(u(s_k)) \quad E_k^{(1)} \leftarrow \mathbf{E}(r_k^{(1)} \cdot u(s_k)),$$

and for all $m \in \llbracket 1; M \rrbracket$: $x_m \xleftarrow{\$} \mathbb{Z}_q$ chosen by the verifier (or from a hash function for a non-interactive proof), and then

$$\Pi_u = (u_m \leftarrow u(x_m), (W_{k,m} \leftarrow \mathbf{E}(w_m(s_k)))_k)_m$$

where w_m is such that $u(X) - u_m = (X - x_m) \cdot w_m(X)$: in total, these are $2K$ encodings, plus M scalars and KM encodings for the proof.

$\text{Commit}(u, \mathbb{Z}_q[X^{(2n-2)\setminus(n-1)}])$: as the previous case but replacing $r^{(1)}$ with $r^{(2)}$ and analogously (1) indices with (2) indices.

$\text{Commit}(u, \mathbb{Z}_q[X^{n-1}, Y^N])$: it outputs $C = (E_u = (E_k, E_k^{(3)})_k, \Pi_u)$, for $k \in \llbracket 1; K \rrbracket$, where

$$E_k \leftarrow \mathbf{E}(u(s_k, t_k)) \quad E_k^{(3)} \leftarrow \mathbf{E}(r_k^{(3)} \cdot u(s_k, t_k)),$$

and for all $m \in \llbracket 1; M \rrbracket$: $(x_m, y_m) \xleftarrow{\$} \mathbb{Z}_q^2$ chosen by the verifier (or from a hash function for a non-interactive proof), and then

$$\Pi_u = (u_m \leftarrow u(x_m, y_m), (V_{k,m} \leftarrow \mathbf{E}(v_m(s_k, t_k)), W_{k,m} \leftarrow \mathbf{E}(w_m(s_k)))_k)_m$$

where v_m and w_m are such that

$$u(X, Y) - u_m = (Y - y_m) \cdot v_m(X, Y) + (X - x_m) \cdot w_m(X).$$

In total, these are $2K$ encodings, plus M scalars and $2KM$ encodings for the proof.

$\text{Commit}(u, \mathbb{Z}_q[Y^N])$: it outputs $C = (E_u = (E_k, E_k^{(4)})_k, \Pi_u)$, for $k \in \llbracket 1; K \rrbracket$, where:

$$E_k \leftarrow \mathbf{E}(u(t_k)) \quad E_k^{(4)} \leftarrow \mathbf{E}(r_k^{(4)} \cdot u(t_k)),$$

and for all $m \in \llbracket 1; M \rrbracket$: $y_m \xleftarrow{\$} \mathbb{Z}_q$ chosen by the verifier (or from a hash function for a non-interactive proof), and then

$$\Pi_u = (u_m \leftarrow u(y_m), (V_{k,m} \leftarrow \mathbf{E}(v_m(t_k)))_k)_m$$

where v_m is such that $u(Y) - u_m = (Y - y_m) \cdot v_m(Y)$. In total, these are $2K$ encodings, plus M scalars and KM encodings for the proof.

$\text{Commit}(u, \mathbb{Z}_q[Y^{2N\setminus N}])$: as the previous case but replacing $r^{(4)}$ with $r^{(5)}$ and analogously (4) indices with (5) indices.

Hiding Commitment. We only focus on $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ and \mathcal{R}_4 :

$\text{Commit}^*(u, \mathbb{Z}_q[X^{n-1}, Y^N])$: outputs $C^* = (E_u^* = (E_k^*, E_k^{(3*)}, \pi_k)_k, \Pi_u^*)$, for $k \in \llbracket 1; K \rrbracket$, where, with $\rho_k, \rho'_k \xleftarrow{\$} \mathcal{M}$, where \mathcal{M} is the appropriate masking set:

$$\begin{aligned} E_k^* &\leftarrow \text{Eval}(\{E_{k,j,i}\}_{j,i}, \{E_{k,0,0}\}, \{u_{j,i}\}_{j,i}, \{\rho_k\}) = \mathbf{E}(u(s_k, t_k) + \rho_k) \\ E_k^{(3*)} &\leftarrow \text{Eval}(\{E_{k,j,i}^{(3)}\}_{j,i}, \{E_{k,0,0}\}, \{u_{j,i}\}_{j,i}, \{\rho'_k\}) = \mathbf{E}(r_k^{(3)} \cdot u(s_k, t_k) + \rho'_k) \\ \text{with } \pi_k &= \{\text{ZKLQCheck}(X_1 - X_2 \cdot X_3, E_k^{(3*)}, E_k^*, E_{k,0,0}^{(3)}; E_{k,0,0}, E_{k,0,0}^{(3)}) = \text{true}\} \end{aligned}$$

and for all $m \in \llbracket 1; M \rrbracket$: $(x_m, y_m) \xleftarrow{\$} \mathbb{Z}_q^2$ with x_m and y_m invertible modulo \mathcal{N}^2 chosen by the verifier (or from a hash function for a non-interactive proof), and then for random $\rho_{k,m}, \rho'_{k,m} \xleftarrow{\$} \mathbb{Z}_q^*$ chosen by the prover for their privacy:

$$\Pi_u^* = (V_{k,m}^* \leftarrow \mathbb{E}(v_m(s_k, t_k) + \rho_{k,m}), W_{k,m}^* \leftarrow \mathbb{E}(w_m(s_k) + \rho'_{k,m}))_{k,m}$$

where v_m and w_m are such that

$$u(X, Y) - u(x_m, y_m) = (Y - y_m) \cdot v_m + (X - x_m) \cdot w_m$$

with

$$\pi_{k,m} = \{\text{ZKLQCheck}(X_1 - (X_2 - y_m) \cdot X_3 - (X_4 - x_m) \cdot X_5, \\ E_k^*, E_{k,1,0}, V_{k,m}^*, E_{k,0,1}, W_{k,m}^*; E_{k,0,0}, E_{k,0,0}, E_{k,1,0}^{(m)}, E_{k,0,1}^{(m)}) = \text{true}\}$$

KM additional zero-knowledge proofs of knowledge of 4 scalars $\rho_{k,m}, \rho'_{k,m}, \rho_k$, and $u_m = u(x_m, y_m)$. In total, this is $2K$ encodings, plus KM encodings for the proof, and KM zero-knowledge proofs of 4 scalars.

$\text{Commit}^*(u, \mathbb{Z}_q[Y^N])$: outputs $C^* = (E_u^* = (E_k^*, E_k^{(4*)}, \pi_k)_k, \Pi_u^*)$, for $k \in \llbracket 1; K \rrbracket$, where, with $\rho_k, \rho'_k \xleftarrow{\$} \mathcal{M}$:

$$E_k \leftarrow \mathbb{E}(u(t_k) + \rho_k) \qquad E_k^{(4*)} \leftarrow \mathbb{E}(r_k^{(4)} \cdot u(t_k) + \rho'_k)$$

with $\pi_k = \{\text{ZKLQCheck}(X_1 - X_2 \cdot X_3, E_k^{(4*)}, E_k^*, E_{k,0,0}^{(4)}; E_{k,0,0}, E_{k,0,0}^{(4)}) = \text{true}\}$, and for all $m \in \llbracket 1; M \rrbracket$: $y_m \xleftarrow{\$} \mathbb{Z}_q$ invertible modulo \mathcal{N}^2 chosen by the verifier (or from a hash function for a non-interactive proof), and then for random $\rho_{k,m} \xleftarrow{\$} \mathbb{Z}_q^*$ chosen by the prover for their privacy:

$$\Pi_u^* = (V_{k,m}^* \leftarrow \mathbb{E}(v_m(t_k) + \rho_{k,m}))_{k,m}$$

for v_m such that $u(Y) - u(y_m) = (Y - y_m) \cdot v_m(Y)$ with KM additional zero-knowledge proofs of knowledge of 3 scalars $\rho_{k,m}, \rho_k$, and $u_m = u(y_m)$. In total, this is $2K$ encodings, plus KM encodings for the proof, and KM zero-knowledge proofs of 3 scalars.

$\text{Commit}^*(u, \mathbb{Z}_q[X^{(n-1)}])$ and $\text{Commit}^*(u, \mathbb{Z}_q[X^{(2n-2) \setminus (n-1)}])$ are analogous to the previous hiding commitment description, but in the other variable.

Validity Check: it also depends on hiding or non-hiding commitments and on the space \mathcal{R}_π .

Non-Hiding Commitment. $\text{Validity}(C, \mathcal{R}_\pi)$ first checks the twin encodings, for $k \in \llbracket 1; K \rrbracket$:

$$\begin{array}{ll} \mathcal{R}_1, \mathcal{R}_2 & \text{QCheck}(X_1 - X_2 \cdot X_3, E_k^{(1/2)}, E_k, E_{k,0}^{(1/2)}) = \text{true} \\ \mathcal{R}_3 & \text{QCheck}(X_1 - X_2 \cdot X_3, E_k^{(3)}, E_k, E_{k,0,0}^{(3)}) = \text{true} \\ \mathcal{R}_4, \mathcal{R}_5 & \text{QCheck}(X_1 - X_2 \cdot X_3, E_k^{(4/5)}, E_k, E_{k,0}^{(4/5)}) = \text{true} \end{array}$$

and then, for $k \in \llbracket 1; K \rrbracket$ and $m \in \llbracket 1; M \rrbracket$, either

$$\begin{array}{ll} \mathcal{R}_1, \mathcal{R}_2 & \text{QCheck}(X_1 - u_m - (X_2 - x_m) \cdot X_3, E_k, E_{k,0,1}, W_{k,m}) = \text{true} \\ \mathcal{R}_3 & \text{QCheck}(X_1 - u_m - (X_2 - y_m) \cdot X_3 - (X_4 - x_m) \cdot X_5, \\ & E_k, E_{k,1,0}, V_{k,m}, E_{k,0,1}, W_{k,m}) = \text{true} \\ \mathcal{R}_4, \mathcal{R}_5 & \text{QCheck}(X_1 - u_m - (X_2 - y_m) \cdot X_3, E_k, E_{k,1,0}, V_{k,m}) = \text{true} \end{array}$$

Hiding Commitment. $\text{Validity}^*(C, \mathcal{R}_\pi)$ first checks the twin encodings, for $k \in \llbracket 1; K \rrbracket$:

$$\begin{aligned} \mathcal{R}_1 & \quad \text{ZKLQCheck}(X_1 - X_2 \cdot X_3, E_k^{(1*)}, E_k^*, E_{k,0,0}^{(1)}; E_{k,0,0}, E^{(2)}) = \text{true} \\ \mathcal{R}_2 & \quad \text{ZKLQCheck}(X_1 - X_2 \cdot X_3, E_k^{(2*)}, E_k^*, E_{k,0,0}^{(2)}; E_{k,0,0}, E^{(2)}) = \text{true} \\ \mathcal{R}_3 & \quad \text{ZKLQCheck}(X_1 - X_2 \cdot X_3, E_k^{(3*)}, E_k^*, E_{k,0,0}^{(3)}; E_{k,0,0}, E^{(3)}) = \text{true} \\ \mathcal{R}_4 & \quad \text{ZKLQCheck}(X_1 - X_2 \cdot X_3, E_k^{(4*)}, E_k^*, E_{k,0,0}^{(4)}; E_{k,0,0}, E^{(4)}) = \text{true} \end{aligned}$$

and then, for $k \in \llbracket 1; K \rrbracket$ and $m \in \llbracket 1; M \rrbracket$, the zero-knowledge proofs

$$\begin{aligned} \mathcal{R}_1, \mathcal{R}_2 & \quad \text{ZKLQCheck}(X_1 - (X_2 - x_m) \cdot X_3, E_k^*, E_{k,0,1}, V_{k,m}^*; E_{k,0,0}, E_{k,0,0}, E_{k,0,1}^{(m)}) \\ & \quad = \text{true} \\ \mathcal{R}_3 & \quad \text{ZKLQCheck}(X_1 - (X_2 - y_m) \cdot X_3 - (X_4 - x_m) \cdot X_5, E_k^*, E_{k,1,0}, V_{k,m}^*, \\ & \quad E_{k,0,1}, W_{k,m}^*; E_{k,0,0}, E_{k,0,0}, E_{k,1,0}^{(m)}, E_{k,0,1}^{(m)}) = \text{true} \\ \mathcal{R}_4 & \quad \text{ZKLQCheck}(X_1 - (X_2 - y_m) \cdot X_3, E_k^*, E_{k,1,0}, V_{k,m}^*; E_{k,0,0}, E_{k,0,0}, E_{k,1,0}^{(m)}) \\ & \quad = \text{true} \end{aligned}$$

D More Detail on the Protocol

Here is more detail on some steps of our protocol, when explained succinctly in Section 5, with a more general description than just OPE and vector of powers.

We indeed consider the receiver/verifier with their private input message \mathbf{m} to learn in a verifiable way the inner product of the vector $\Phi(\mathbf{m}) = (\phi(\mathbf{m}, j))_{j \in \llbracket 0; N \rrbracket}$ with a private vector $\mathbf{F} = (f_j)_{j \in \llbracket 0; N \rrbracket}$, committed by the sender/prover, where ϕ is a function known by them both, depending on the application. If $\phi(\mathbf{m}, j) = \mathbf{m}^j$, the inner product corresponds to the Oblivious Polynomial Evaluation (OPE) of the polynomial \mathbf{F} with coefficients $(f_j)_j$ on the message \mathbf{m} ; if $\phi(\mathbf{m}, j) = \delta_{\mu, j}$, with δ the Kronecker symbol and $\mu \in \llbracket 0; N \rrbracket$ such that \mathbf{m} is the representation of μ , it provides the μ -th coefficient of \mathbf{f} , which would coincide with a Symmetric Private Information Retrieval (SPIR) application (see Appendix E). We now show how the sender can provide this evaluation with fully homomorphic encryption in a provable way.

More precisely, we consider the above FV scheme that encrypts messages from $\mathcal{R}_t = \mathbb{Z}_t[X]/r(X)$, where $r = X^n + 1$, into $\mathcal{R}_q = \mathbb{Z}_q[X]/r(X)$.

The receiver encrypts the input $\mathbf{m} \in \mathcal{R}_t$ under their own key, and sends $\text{Enc}(\mathbf{m}) = (c, c') \in \mathcal{R}_q^2$, in order to get back the homomorphic inner product of $\Phi(\mathbf{m}) = (\phi(\mathbf{m}, j))_{j \in \llbracket 0; N \rrbracket} \in \mathcal{R}_t^{N+1}$ with $\mathbf{F} = (f_j)_{j \in \llbracket 0; N \rrbracket} \in \mathcal{R}_t^{N+1}$, committed by the sender, with a proof of correct value $(d, d') = \text{Enc}(\langle \Phi(\mathbf{m}), \mathbf{F} \rangle) \in \mathcal{R}_q^2$, that thereafter decrypts to $\langle \Phi(\mathbf{m}), \mathbf{F} \rangle \in \mathcal{R}_t$, using the receiver's decryption key.

From (c, c') , the sender is able to compute $(u_j, u'_j) = \text{Enc}(\phi(\mathbf{m}, j))$, for $j \in \llbracket 0; N \rrbracket$, in any way they want, using the homomorphic properties of the encryption scheme:

$$(u_j, u'_j) = \text{Enc}(\phi(\mathbf{m}, j)) = \left(\sum_{i=0}^n u_{j,i} \cdot X^i, \sum_{i=0}^n u'_{j,i} \cdot X^i \right) \in \mathcal{R}_q^2.$$

Thereafter, the goal is to evaluate the inner products:

$$\begin{aligned} (d, d') &= \text{Enc}(\langle \Phi(m), F \rangle) = \text{Enc} \left(\sum_{j=0}^N f_j \cdot \phi(m, j) \right) = \sum_{j=0}^N f_j \cdot \text{Enc}(\phi(m, j)) \\ &= \sum_{j=0}^N f_j \cdot (u_j, u'_j) = \left(\sum_{j=0}^N f_j \cdot u_j, \sum_{j=0}^N f_j \cdot u'_j \right). \end{aligned}$$

As already explained, the sender can add noise to hide F in the evaluation and prove the correct evaluation of $\tilde{d} = d + z^*$ and $\tilde{d}' = d' + z'^*$ so that (\tilde{d}, \tilde{d}') decrypts to $\langle \Phi(m), F \rangle$. The pair (z^*, z'^*) will be committed in a hidden way using three noise polynomials in \mathcal{R}_1 from the FV encryption scheme for encodings of zero, with a proof of correct noise, in the appropriate range, to ensure the correct decryption. As such we will take $(z^*, z'^*) = (\mathbf{p} \cdot \mathbf{u} + \mathbf{e}_1, \mathbf{p}' \cdot \mathbf{u} + \mathbf{e}_2)$, knowing that the quadratic relation checks will be made modulo q , where \mathbf{p} and \mathbf{p}' are public polynomials from the FV scheme setup, and commit the noise polynomials $\mathbf{u}, \mathbf{e}_1, \mathbf{e}_2 \in \mathcal{R}_1$ into U^*, E_1^* and E_2^* respectively.

This verifiable evaluation will be done in two steps: first, the sender will prove the correct evaluation of all the (u_j, u'_j) , while committed in a very compact way. Then, a proof of the inner product is provided, as described in Section 4. We will later explain how everything remains succinct.

We also claim this protocol to be SND-secure, proving theorem 5, with steps provided in the following sections.

D.1 Commitment of F

One can first note that ring elements in \mathcal{R}_t and \mathcal{R}_q are polynomials of degree at most $n-1$, and can be encoded into $\mathcal{R}_1 = \mathbb{Z}_q[X^{n-1}]$. The sender's vector $F \in \mathcal{R}_t^{N+1}$ can be committed using the polynomial $f = \sum_{i=0}^N f_j \cdot Y^j$ in $\mathcal{R}_t[Y]$, where $f_j = \sum_{i=0}^{n-1} f_{j,i} \cdot X^i$, can be committed in a hidden way in $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$ as $f^* = \sum_{j=0}^N \sum_{i=0}^{n-1} f_{j,i} \cdot X^i Y^{N-j}$ (with reverse order coefficients) into F^* .

The sender can also compute and publish the noisy inner products \tilde{d} and \tilde{d}' in \mathcal{R}_q , the expected ciphertext of the result, for some private $\mathbf{q}^*, \mathbf{q}'^*, z^*, z'^* \in \mathcal{R}_q$:

$$\tilde{d} = \sum_{j=0}^N f_j \cdot u_j + z^* - \mathbf{q}^* \cdot r \quad \tilde{d}' = \sum_{j=0}^N f_j \cdot u'_j + z'^* - \mathbf{q}'^* \cdot r.$$

With $(z^*, z'^*) = ([\mathbf{p} \cdot \mathbf{u} + \mathbf{e}_1]_q, [\mathbf{p}' \cdot \mathbf{u} + \mathbf{e}_2]_q)$, the prover-owned polynomials $\mathbf{u}^*, \mathbf{e}_1^*, \mathbf{e}_2^*, \mathbf{q}^*, \mathbf{q}'^*$ are committed in a hidden way in $U^*, E_1^*, E_2^*, Q^*, Q'^*$, from $\mathcal{R}_1 = \mathbb{Z}_q[X^{n-1}]$. The inner products must be proven – which is the goal of the following second part – after we have proven the correct computation of the (u_j, u'_j) 's, as both together will prove correctness of the ciphertext (\tilde{d}, \tilde{d}') .

D.2 Validity of the (u_j, u'_j) 's

From the ciphertexts (u_j, u'_j) that the prover computed on their own, they define the polynomials:

$$u(X, Y) = \sum_{j=0}^N u_j(X) \cdot Y^j = \sum_{j=0}^N \sum_{i=0}^{n-1} u_{j,i} X^i Y^j \quad u'(X, Y) = \sum_{j=0}^N u'_j(X) \cdot Y^j$$

and commit them from the subspace $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$ into U and U' respectively. Our first step is to prove that the polynomials committed in U and U' by the sender indeed satisfy,

for any exponents $j \in \llbracket 0; N \rrbracket$, $(u_j, u'_j) = \text{Enc}(\phi(m, j))$, or more precisely that the decryptions $\text{Dec}(u_j, u'_j)$ that we denote $\varphi_{m,j}$ are indeed $\phi(m, j) \in \mathcal{R}_t$, for m initially encrypted by the verifier in (c, c') .

The verifier first chooses and sends a random polynomial $\mathbf{n} \xleftarrow{\$} \mathcal{R}_t$ (or it can be generated by a hash function on the previous data to remove interaction). Both parties can compute, $\mathbf{n}_j = \mathbf{n}^j = \sum_{i=0}^{n-1} n_{j,i} \cdot X^i$ in \mathcal{R}_t , for all $j \in \llbracket 0; N \rrbracket$, and commit them from $\mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$ as $\mathbf{t} = \sum_{j=0}^N \sum_{i=0}^{n-1} n_{j,i} \cdot X^i Y^{N-j}$ (with reverse order coefficients) into T . Then, the prover computes and sends $(\mathbf{b}, \mathbf{b}') = \sum_{j=0}^N \mathbf{n}_j \cdot (u_j, u'_j)$ in \mathcal{R}_q^2 .

Since this is symmetric, we now focus on the first component (without $'$), and a similar analysis will have to be done on the second component (with $'$): the prover also computes and commits \mathbf{l} so that $\sum_{j=0}^N \mathbf{n}_j \cdot u_j = \mathbf{b} + \mathbf{l} \cdot r$, into the commitment L from $\mathcal{R}_1 = \mathbb{Z}_q[X^{n-1}]$.

The verifier chooses and sends a sequence of random scalars $\beta_\kappa \xleftarrow{\$} \mathbb{Z}_q$, for $\kappa \in \llbracket 1; \Lambda \rrbracket$. They will define $v_{\kappa,j} = u_j(\beta_\kappa) \in \mathbb{Z}_q$ for $j \in \llbracket 0; N \rrbracket$. We thus have the polynomial:

$$\mathbf{v}_\kappa(Y) = \sum_{j=0}^N v_{\kappa,j} \cdot Y^j = \sum_{j=0}^N u_j(\beta_\kappa) \cdot Y^j = \sum_{j=0}^N \sum_{i=0}^{n-1} u_{j,i} \beta_\kappa^i \cdot Y^j = \mathbf{u}(\beta_\kappa, Y)$$

that can be committed as V_κ from $\mathcal{R}_4 = \mathbb{Z}_q[Y^N]$, and proven correct with the quadratic check

$$\mathbf{u}(X, Y) - \mathbf{v}_\kappa(Y) = \mathbf{u}(X, Y) - \mathbf{u}(\beta_\kappa, Y) = (X - \beta_\kappa) \cdot \mathbf{w}_\kappa(X, Y) \quad (5)$$

on the commitment W_κ of $\mathbf{w}_\kappa \in \mathcal{R}_3 = \mathbb{Z}_q[X^{n-1}, Y^N]$ and the public commitment C_κ of $X - \beta_\kappa$, and thus using only 3 prover-generated commitments (U , V_κ and W_κ , as C_κ is public). With $\mathbf{t}_\kappa = \mathbf{t}(\beta_\kappa, Y) = \sum_{j=0}^N \mathbf{n}_j(\beta_\kappa) \cdot Y^j$ publicly computed and committed in T_κ , we can note that

$$\begin{aligned} \mathbf{v}_\kappa(Y) \cdot \mathbf{t}_\kappa(Y) &= \sum_{0 \leq i, j \leq N} v_{\kappa,i} \cdot \mathbf{n}_j(\beta_\kappa) \cdot Y^{N+i-j} \\ &= \sum_{0 \leq i \leq N} v_{\kappa,i} \cdot \mathbf{n}_i(\beta_\kappa) \cdot Y^N + \sum_{0 \leq i \neq j \leq N} v_{\kappa,i} \cdot \mathbf{n}_j(\beta_\kappa) \cdot Y^{N+i-j} \end{aligned}$$

and

$$\sum_{0 \leq i \leq N} v_{\kappa,i} \cdot \mathbf{n}_i(\beta_\kappa) = \sum_{0 \leq i \leq N} u_i(\beta_\kappa) \cdot \mathbf{n}_i(\beta_\kappa) = b_\kappa + l_\kappa \cdot r_\kappa$$

where $b_\kappa = \mathbf{b}(\beta_\kappa)$, and $r_\kappa = \mathbf{r}(\beta_\kappa)$ can be publicly computed in \mathbb{Z}_q , and $l_\kappa = \mathbf{l}_\kappa(\beta_\kappa)$ will have to be checked with respect to the commitment L with the following relation:

$$\mathbf{l}(X) - l_\kappa = (X - \beta_\kappa) \cdot \mathbf{l}_\kappa(X), \quad (6)$$

also committing \mathbf{l}_κ from $\mathcal{R}_1 = \mathbb{Z}_q[X^{n-1}]$ into L_κ .

We can now check that:

$$\mathbf{v}_\kappa(Y) \cdot \mathbf{t}_\kappa(Y) - (b_\kappa + l_\kappa \cdot r_\kappa) \cdot Y^N = \sum_{0 \leq i \neq j \leq N} v_{\kappa,i} \cdot \mathbf{n}_j(\beta_\kappa) \cdot Y^{N+i-j} = \mathbf{y}_\kappa(Y) \quad (7)$$

with a commitment Y_κ of \mathbf{y}_κ in $\mathcal{R}_5 = \mathbb{Z}_q[Y^{2N \setminus N}]$ and C_N a public commitment of Y^N , which makes only 2 prover-generated commitments (V_κ and Y_κ).

We stress that because of the repetitions in each commitment (up to K , according to the kind of relations they are involved in, and the number of prover-generated commitments), we know (see Corollary 9) that when all the tests pass, the above equations (5) and (7) are all satisfied with error probability less than $2\Lambda \cdot 2^{-(S+1)K/4}$.

They altogether prove that, for each $\kappa \in \llbracket 1; \Lambda \rrbracket$, $b_\kappa + l_\kappa \cdot r_\kappa$ is the coefficient of Y^N in $\mathbf{v}_\kappa(Y) \cdot \mathbf{t}_\kappa(Y)$, and so $\sum_j u_j(\beta_\kappa) \cdot \mathbf{n}_j(\beta_\kappa) = \mathbf{b}(\beta_\kappa) + \mathbf{l}(\beta_\kappa) \cdot \mathbf{r}(\beta_\kappa) \pmod q$ for the random β_κ , on

polynomials committed beforehand, of degree n . Hence, $\sum_j u_j \cdot n_j = \mathbf{b} + \mathbf{l} \cdot \mathbf{r}$, excepted with error probability bounded by $2\ell n/p \leq n/2^{S+1}N$, from the Schwartz-Zippel lemma.

For Λ large enough (to be set later), after all these steps, on both \mathbf{b} and \mathbf{b}' , one gets the proof that, in \mathcal{R}_q ,

$$\mathbf{b} = \sum_{j=0}^N n_j \cdot u_j \quad \mathbf{b}' = \sum_{j=0}^N n_j \cdot u'_j : \quad (\mathbf{b}, \mathbf{b}') = \sum_{j=0}^N n_j \cdot (u_j, u'_j).$$

These relations hold for both \mathbf{b} and \mathbf{b}' excepted with error probability bounded by $2 \times (2\Lambda \cdot 2^{-(S+1)K/4} + (n/2^{S+1}N)^\Lambda)$.

D.3 Validity of $(\tilde{\mathbf{d}}, \tilde{\mathbf{d}}')$

As above, we focus on $\tilde{\mathbf{d}}$ with respect to all $(u_j)_j$ and $\mathbf{z}^*, \mathbf{q}^*$. The same should be done for $\tilde{\mathbf{d}}'$ with respect to all $(u'_j)_j$ and $\mathbf{z}'^*, \mathbf{q}'^*$, but for the same $(f_j)_j$ and \mathbf{z}, \mathbf{r} :

$$\tilde{\mathbf{d}} = \sum_{j=0}^N f_j \cdot u_j + \mathbf{z}^* - \mathbf{q}^* \cdot \mathbf{r}.$$

This following analysis is quite similar to the previous one, considering $\mathbf{f}^*, (\mathbf{z}^*, \mathbf{q}^*), \tilde{\mathbf{d}}$ instead of $\mathbf{t}, \mathbf{p}, \mathbf{b}$, and with hidden intermediate values, as $(f_j)_j$ is private to the prover, contrarily to \mathbf{n} which was publicly chosen by the verifier. One first proves the quadratic relation on the hidden commitment of $\mathbf{g}_\kappa^*(Y) = \mathbf{f}^*(\beta_\kappa, Y) = \sum_{j=0}^N g_{\kappa,j} \cdot Y^{N-j}$ in $\mathcal{R}_4 = \mathbb{Z}_q[Y^N]$, where $g_{\kappa,j} = f_j(\beta_\kappa)$, and the hiding commitment of $\mathbf{h}_\kappa^*(X, Y)$ in $\mathcal{R}_2 = \mathbb{Z}_q[X^n, Y^N]$, so that

$$\mathbf{f}^*(X, Y) - \mathbf{g}_\kappa^*(Y) = \mathbf{f}^*(X, Y) - \mathbf{f}^*(\beta_\kappa, Y) = (X - \beta_\kappa) \cdot \mathbf{h}_\kappa^*(X, Y) \quad (8)$$

with only 3 prover-generated commitments. And similarly, one proves in a hidden way:

$$\mathbf{q}^*(X) - q_\kappa^* = \mathbf{q}^*(X) - \mathbf{q}^*(\beta_\kappa) = (X - \beta_\kappa) \cdot \mathbf{q}_\kappa^*(X) \quad (9)$$

$$\mathbf{z}^*(X) - z_\kappa^* = \mathbf{z}^*(X) - \mathbf{z}^*(\beta_\kappa) = (X - \beta_\kappa) \cdot \mathbf{z}_\kappa^*(X) \quad (10)$$

with the private $q_\kappa^* = \mathbf{q}_\kappa^*(\beta_\kappa)$ and $z_\kappa^* = \mathbf{z}_\kappa^*(\beta_\kappa)$ in \mathbb{Z}_q that will be also used below. There are only 2 prover-generated commitments in each relation. We have the following relations:

$$\begin{aligned} v_\kappa(Y) \cdot \mathbf{g}_\kappa^*(Y) &= \sum_{0 \leq i, j \leq N} v_{\kappa,i} \cdot g_{\kappa,j} \cdot Y^{N+i-j} \\ &= \sum_{0 \leq i \leq N} v_{\kappa,i} \cdot g_{\kappa,i} \cdot Y^N + \sum_{0 \leq i \neq j \leq N} v_{\kappa,i} \cdot g_{\kappa,j} \cdot Y^{N+i-j} \end{aligned}$$

and

$$\begin{aligned} \sum_{0 \leq i \leq N} v_{\kappa,i} \cdot g_{\kappa,i} &= \sum_{0 \leq i \leq N} u_i(\beta_\kappa) \cdot f_i(\beta_\kappa) &= \tilde{\mathbf{d}}(\beta_\kappa) - \mathbf{z}^*(\beta_\kappa) + \mathbf{q}^*(\beta_\kappa) \cdot \mathbf{r}(\beta_\kappa) \\ &= \tilde{d}_\kappa - z_\kappa^* + q_\kappa^* \cdot r_\kappa \end{aligned}$$

where $\tilde{d}_\kappa = \tilde{\mathbf{d}}(\beta_\kappa)$ and $r_\kappa = \mathbf{r}(\beta_\kappa)$ can be publicly computed, and so one proves

$$v_\kappa(Y) \cdot \mathbf{g}_\kappa^*(Y) - (\tilde{d}_\kappa - z_\kappa^* + q_\kappa^* \cdot r_\kappa) \cdot Y^N = \mathbf{s}_\kappa^*(Y) \quad (11)$$

with a hiding commitment of \mathbf{s}_κ^* from $\mathcal{R}_5 = \mathbb{Z}_q[Y^{2N \setminus N}]$ and a public commitment of Y^N .

Again, from Corollary 9, when all the tests pass, the above equations (8), (9), (10), and (11) are all satisfied with error probability less than $3\Lambda \cdot 2^{-(S+1)K/4}$.

They altogether prove that, for each $\kappa \in \llbracket 1; \Lambda \rrbracket$, $d_\kappa - z_\kappa^* + q_\kappa^* \cdot r_\kappa$ is the coefficient of Y^N in $\mathbf{v}_\kappa(Y) \cdot \mathbf{g}_\kappa^*(Y)$, so

$$\tilde{\mathbf{d}}(\beta_\kappa) = \sum_j \mathbf{u}_j(\beta_\kappa) \cdot \mathbf{f}_j(\beta_\kappa) + \mathbf{z}^*(\beta_\kappa) - \mathbf{q}^*(\beta_\kappa) \cdot \mathbf{r}(\beta_\kappa) \pmod q$$

for the random β_κ , on polynomials committed beforehand. Hence,

$$\tilde{\mathbf{d}} = \sum_j \mathbf{u}_j \cdot \mathbf{f}_j + \mathbf{z}^* - \mathbf{q}^* \cdot \mathbf{r},$$

excepted with error probability bounded by $2\ell n/p \leq n/2^{S+1}N$, for each κ . Again, these relations hold for both $\tilde{\mathbf{d}}$ and $\tilde{\mathbf{d}}'$ excepted with error probability bounded by $2 \times (3\Lambda \cdot 2^{-(S+1)K/4} + (n/2^{S+1}N)^\Lambda)$.

Globally, the error probability is bounded by $2 \times (5\Lambda \cdot 2^{-(S+1)K/4} + 2 \times (n/2^{S+1}N)^\Lambda)$. We thus need to take $\Lambda = \lceil \log_2(2\nu_e/\varepsilon_s) / \log_2(2^{S+1}N/n) \rceil$ different values for β_κ so that $4 \times (n/2^{S+1}N)^\Lambda \leq \varepsilon_s/6$, and $K > 4 \log_2(60\Lambda/\varepsilon_s)/(S+1)$ so that $2 \times 5\Lambda \cdot 2^{-(S+1)K/4} \leq \varepsilon_s/4$ too, so that the global error on the equations be bounded by $\varepsilon_s/2$.

On the other hand, as shown on Figure 6, the global number of commitments is bounded by $15\Lambda + 14$. We thus need to take $\varepsilon_c \leq \varepsilon_s/(2 \times \nu_c) = \varepsilon_s/(2(15\Lambda + 14))$.

D.4 SND-Security of our Scheme

For each commitment generated by the prover and sent to the receiver, Theorem 8 grants that if it passes validity and quadratic root checks, then at least $4K/5$ of the twin encodings encode the polynomial \mathbf{u} in the appropriate subspace, that can be extracted by extractability of valid encodings, excepted with probability less than $2^{-(S+1)K/5} + 2K \cdot 2^{-(S+1)M/3}$, where $q = \prod_{i=1}^\ell p_i$ with all $p_i \geq 2^S \times 2\ell \times 2N$. We choose K and M such that

$$2^{-(S+1)K/5} + 2K \cdot 2^{-(S+1)M/3} \leq \frac{\varepsilon_s}{2\nu_c},$$

where ε_s is the soundness security parameter and ν_c the total number of prover-generated commitments in the protocol. So, as ν_c commitments are going to be checked, this holds for all of them except with probability less than $\varepsilon_s/2$.

Then, in our protocol, a total of ν_e quadratic equations on committed polynomials need to be checked thanks to their commitments. The checks for one of these relations grants it is valid except for a probability smaller than $2^{-SK/5}$, where K is chosen so that this is less than $\frac{\varepsilon_s}{2\nu_e}$. For ν_e checked equations, they all hold with probability less than $\varepsilon_s/2$.

Then, the proof also uses the decryption of 2 FHE ciphertxts, and their soundness is ensured by the FHE security on more than 128 bits.

In the end, our soundness is granted except for a probability less than our soundness security parameter ε_s .

D.5 R-Privacy-Security of our Scheme

The scheme is R-Privacy-secure under the semantic security of the FHE FV scheme. Sending the public key containing the FHE public information and one FHE ciphertxts of \mathbf{m} , with an FHE scheme with at least 128 bits of security does not reveal any information on the cleartxts, for any malicious adversary performing polynomial time calculations.

D.6 S-Privacy-Security of our Scheme

Under the statistically hiding properties of our commitment scheme given in Theorem 10 and the statistical security provided by the noise-flooding, no sender's information leaks to the receiver.

E Applications

Our protocol can be deployed for several applications. Hereafter we detail the cases of PSI (using OPE), and of SPIR.

E.1 Application to Private Set Intersection

As already explained, we can apply this technique to PSI, where the receiver owns a set $\mathcal{X} = \{x_1, \dots, x_a\}$ of cardinality $\#\mathcal{X}$, the sender a set $\mathcal{Y} = \{y_1, \dots, y_b\}$ of cardinality $\#\mathcal{Y}$, and the receiver wants to learn the intersection. We consider the case where $\#\mathcal{Y}$ is much larger than $\#\mathcal{X}$, and hope to get a protocol essentially linear in $\#\mathcal{X}$ only. To this aim, we follow the basic approach from [FNP04], with the polynomial $P(Y) = \prod_{i=1}^N (Y - h_i)$ committed once for all in P^* , where the values $h_i \in \mathcal{R}_t$ encode the elements y_i in the sender's set, and the degree $N = \#\mathcal{Y}$. Then the receiver wants to check whether $P(x_i) = 0$ on every input x_i .

Thanks to the verifiability of our proof, and even the knowledge-soundness of our commitments, our PSI protocol is secure against malicious senders, as they cannot use distinct polynomials as P in each execution.

The protocol can be adapted to support adding elements to the sender's set, with a new check to make sure the old sender's polynomial divides the new one which only has additional roots.

E.2 Application to Symmetric Private Information Retrieval

In the Symmetric PIR setting, the sender owns a set $\mathcal{Y} = \{y_1, \dots, y_N\}$ of cardinality N , and the receiver wants to retrieve an element of the sender's set with a private index μ . To this aim, the sender defines $\mathbf{F} = (f_j)_j$ as a vector of representations of their set elements in \mathcal{R}_t , the receiver has a representation $\mathbf{m} = \tau(\mu)$ of their index μ in \mathcal{R}_t . But using our previous notations, we want to define $\mathbf{M} = (m_j)_j$, with $m_j = \delta_{\mu,j}$, where δ is the Kroenecker symbol ($\delta_{i,j} = 1$ if $i = j$ and 0 otherwise): $\langle \mathbf{F}, \mathbf{M} \rangle = \sum f_j m_j = \sum \delta_{\mu,j} f_j = f_\mu$, which is the μ -th element in the sender's set.

F Parameter Calculations

F.1 RNS Compatible Application

In order to use the SEAL FV implementation for our scheme, we need to be able to take a modulus q that is a product of primes on less than 60 bits. This implies higher false positive probabilities given by the Schwartz-Zippel lemma, hence we need repetitions in our proofs to obtain the required soundness. This changes our parameter values, with $q = p_1 \times \dots \times p_\ell$, $p = \min\{p_1, \dots, p_\ell\}$, and $p \geq 2^S \cdot 4N\ell$.

The size of the proof depends a lot on Λ . The larger N/n is, the shorter the proof, but S also has a huge impact, when S makes Λ decrease, with the number of encodings. However, when S grows, so does p , which can make the FHE ciphertexts heavier.

F.2 Summary of Commitments and Checks

In Figures 4, 5, and 6, we present a summary of all the parameters according to the expected security levels.

#Prover-Generated Commitments: c	2	3
$\varepsilon_c = 2^{-38}$ K	26	35
M (and binding case #Scalars)	20	20
#Equations (hiding case)	2184	2940
#Secrets (hiding case)	4448	5960
#Encodings for $1v$ Pol.	572	770
#Encodings for $2v$ Pol.	1092	1470
$\varepsilon_c = 2^{-136}$ K	75	100
M (and binding case #Scalars)	62	62
#Equations (hiding case)	18900	25200
#Secrets (hiding case)	38048	50648
#Encodings for $1v$ Pol.	4800	6400
#Encodings for $2v$ Pol.	9450	12600

Fig. 4. Parameters for Commitments with $p \geq 64\ell\mathcal{D}$, four repetitions in the zero-knowledge proofs, $\varepsilon_s = 2^8 \cdot \varepsilon_c$, and $\nu_e = 21$.

Equation	c
$\mathbf{u}(X, Y) - \mathbf{v}_\kappa(Y) = (X - \beta_\kappa) \cdot \mathbf{w}_\kappa(X, Y)$	3
$\mathbf{u}'(X, Y) - \mathbf{v}'_\kappa(Y) = (X - \beta_\kappa) \cdot \mathbf{w}'_\kappa(X, Y)$	3
$l(X) - l_\kappa = (X - \beta_\kappa) \cdot l_\kappa(X)$	2
$l'(X) - l'_\kappa = (X - \beta_\kappa) \cdot l'_\kappa(X)$	2
$\mathbf{v}_\kappa(Y) \cdot \mathbf{t}_\kappa(Y) - (b_\kappa + l_\kappa \cdot r_\kappa) \cdot Y^N = \mathbf{y}_\kappa(Y)$	2
$\mathbf{v}'_\kappa(Y) \cdot \mathbf{t}'_\kappa(Y) - (b'_\kappa + l'_\kappa \cdot r_\kappa) \cdot Y^N = \mathbf{y}'_\kappa(Y)$	2
$\mathbf{f}^*(X, Y) - \mathbf{g}_\kappa^*(Y) = (X - \beta_\kappa) \cdot \mathbf{h}_\kappa^*(X, Y)$	3
$\mathbf{q}^*(X) - \mathbf{q}'_\kappa = (X - \beta_\kappa) \cdot \mathbf{q}_\kappa^*(X)$	2
$\mathbf{q}'^*(X) - \mathbf{q}''_\kappa = (X - \beta_\kappa) \cdot \mathbf{q}''_\kappa(X)$	2
$\mathbf{v}_\kappa(Y) \cdot \mathbf{g}_\kappa^*(Y) - (\mathbf{d}_\kappa - z_\kappa^* + \mathbf{q}_\kappa^* \cdot r_\kappa) \cdot Y^N = \mathbf{s}_\kappa^*(Y)$	3
$\mathbf{v}'_\kappa(Y) \cdot \mathbf{g}'_\kappa^*(Y) - (\mathbf{d}'_\kappa - z'_\kappa + \mathbf{q}'_\kappa \cdot r_\kappa) \cdot Y^N = \mathbf{s}'_\kappa^*(Y)$	3
$\bar{\mathbf{u}}^*(X) \cdot \mathbf{u}^*(X) - n_u^* \cdot X^{n-1} = \mathbf{x}_u^*(X)$	3
$\bar{\mathbf{e}}_1^*(X) \cdot \mathbf{e}_1^*(X) - n_{e,1}^* \cdot X^{n-1} = \mathbf{x}_{e,1}^*(X)$	3
$\bar{\mathbf{e}}_2^*(X) \cdot \mathbf{e}_2^*(X) - n_{e,2}^* \cdot X^{n-1} = \mathbf{x}_{e,2}^*(X)$	3
$\mathbf{u}^*(X) - \mathbf{u}'_\kappa = (X - 1/\beta_\kappa) \cdot \mathbf{u}_\kappa^*(X)$	2
$\mathbf{e}_1^*(X) - \mathbf{e}_{1,\kappa} = (X - 1/\beta_\kappa) \cdot \mathbf{e}_{1,\kappa}^*(X)$	2
$\mathbf{e}_2^*(X) - \mathbf{e}_{2,\kappa} = (X - 1/\beta_\kappa) \cdot \mathbf{e}_{2,\kappa}^*(X)$	2
$\bar{\mathbf{u}}^*(X) - \beta_\kappa^{n-1} \cdot \mathbf{u}_\kappa^* = (X - \beta_\kappa) \cdot \bar{\mathbf{u}}_\kappa^*(X)$	2
$\bar{\mathbf{e}}_1^*(X) - \beta_\kappa^{n-1} \cdot \mathbf{e}_{1,\kappa}^* = (X - \beta_\kappa) \cdot \bar{\mathbf{e}}_{1,\kappa}^*(X)$	2
$\bar{\mathbf{e}}_2^*(X) - \beta_\kappa^{n-1} \cdot \mathbf{e}_{2,\kappa}^* = (X - \beta_\kappa) \cdot \bar{\mathbf{e}}_{2,\kappa}^*(X)$	2

Total number of checked equations: $\nu_e = 17\Lambda + 3$

Fig. 5. Equations to verify with quadratic checks in our full protocol

#variables	c	Binding or Hiding	Polynomials	Domain	#commitments
$1v$	2	B	$l, l', l_\kappa, l'_\kappa$	\mathcal{R}_1	$2(\Lambda + 1)$
			y_κ, y'_κ	\mathcal{R}_5	2Λ
		H	$\mathbf{q}^*, \mathbf{q}'_\kappa, \mathbf{q}_\kappa^*, \mathbf{q}''_\kappa$	\mathcal{R}_1	$5\Lambda + 2$
			$\mathbf{u}_\kappa^*, \mathbf{e}_{1,\kappa}^*, \mathbf{e}_{2,\kappa}^*$		
	3	B	$\mathbf{v}_\kappa, \mathbf{v}'_\kappa$	\mathcal{R}_4	2Λ
			$\mathbf{u}^*, \mathbf{e}_1^*, \mathbf{e}_2^*, \bar{\mathbf{u}}^*, \bar{\mathbf{e}}_1^*, \bar{\mathbf{e}}_2^*$	\mathcal{R}_1	6
		H	$\mathbf{x}_u^*, \mathbf{x}_{e,1}^*, \mathbf{x}_{e,2}^*$	\mathcal{R}_2	3
			\mathbf{g}_κ^*	\mathcal{R}_4	Λ
		$\mathbf{s}_\kappa^*, \mathbf{s}'_\kappa$	\mathcal{R}_5	2Λ	
$2v$	3	B	$\mathbf{u}, \mathbf{u}', \mathbf{w}_\kappa, \mathbf{w}'_\kappa$	\mathcal{R}_3	$2(\Lambda + 1)$
		H	$\mathbf{f}^*, \mathbf{h}_\kappa^*$	\mathcal{R}_3	$\Lambda + 1$
Total					$17\Lambda + 16$

Fig. 6. Number of Commitments in the Global Proof, and Size in Number of Encodings, with $\Lambda = \lceil (3 + \log_2(3/\varepsilon_s)) / (S + 1 + \log_2(N/n)) \rceil$