Multivariate Commitments and Signatures with Efficient Protocols

Charles Bouillaguet[1], Thibauld Feneuil[2], Jules Maire[1], Matthieu Rivain[2], Julia Sauvage[1], and Damien Vergnaud[1,3]

[1] Sorbonne Université, CNRS, LIP6, F-75005 Paris, France
[2] CryptoExperts, Paris, France

Abstract. We revisit multivariate commitments based on the hardness of solving systems of multivariate quadratic (MQ) equations over finite fields. We analyze a simple construction where a message μ is committed as $c = (\mu + \mathcal{F}(r), \mathcal{G}(r))$, with $\mathcal{F}$ and $\mathcal{G}$ random quadratic maps. We prove that the scheme is computationally hiding assuming the intractability of the MQ problem. Its binding property reduces to solving random bilinear systems. We prove that this problem is NP-complete and study the performance of existing algebraic and hybrid attacks.
We show that this commitment is well-suited for integration with zero-knowledge proofs. Using the *Threshold-Computation-in-the-Head* framework, we construct zero-knowledge efficient arguments of knowledge for the opening and arguments for relations on committed values. We apply this to construct an efficient blind signature scheme *à la* Fischlin, and we demonstrate that our techniques yield a fully multivariate construction of signatures with efficient protocols, enabling practical post-quantum anonymous credentials.

1 Introduction

Mitigating the quantum threat in information security requires cryptographic schemes based on problems intractable for quantum computers. One promising approach, *multivariate cryptography*, relies on the hardness of solving systems of multivariate polynomial equations over finite fields. A central problem is the (NP-hard) Multivariate Quadratic (MQ) problem, which consists in solving systems of quadratic equations in multiple variables over a finite field. The history of multivariate cryptography has seen many schemes broken due to the use of structured MQ variants that introduced exploitable algebraic weaknesses (e.g., [KS98,GJS06,DFS07,Beu22a,BS24]). However, recent advancements have led to more robust constructions, particularly for signature applications, several of which are being considered in the NIST post-quantum cryptography standardization process (*e.g.* [KPG99,Beu22b,FIKT21,BFR24,BBFR25]).

A commitment scheme [Blu82] is a cryptographic primitive that allows a party to commit to a hidden value while ensuring its immutability. It is essential for protocols like zero-knowledge proofs and secure multi-party computation. A commitment must be both *hiding* (concealing the value until revealed) and

binding (preventing alterations after commitment). Since the MQ problem is assumed to be intractable on average, it is natural to leverage its hardness for constructing hash functions [BRP07,DY07], stream ciphers [BGP06] and commitment schemes [PSM17].

The observation that quadratic maps are not necessarily collision-resistant seems to belong to the folklore—in any case, it has often be rediscovered in works that involve multivariate polynomials [BRP07,DY07,AHI+17,Beu24]. When basing security on the one-wayness of quadratic functions, this observation calls for particular caution when designing hash functions (that may fail to be collision-resistant) and commitment schemes (that may fail to be binding).

On the basis of this simple observation, Beullens [Beu24] broke a *blind signature scheme* [Cha82] proposed in [PSM17]. The root of the problem was that an MQ-based commitment was not binding.

In a blind signature scheme, a user gets a message signed without the signer seeing the content. The broader notion of *signatures with efficient protocols* (SEP), introduced by Camenisch and Lysyanskaya in [CL03], supports efficient zero-knowledge proofs, allowing a prover to demonstrate possession of a valid signature without revealing the signature or signed message (but allowing proving statements about it). This property is crucial for applications such as anonymous credentials and privacy-preserving authentication.

Beulens suggested that, while *overdetermined* systems of quadratic equations might still serve as a basis for binding commitments, their weaker resistance to inversion algorithms could compromise the scheme's hiding property. He left open the security analysis of such schemes and the derivation of concrete parameters to achieve a desired security level. He suggested replacing them with symmetric primitives and proof systems based on the MPC-in-the-head paradigm [IKOS09] to repair the blind signature scheme from [PSM17]. While this approach is likely feasible, it would not lead to a practical protocol for SEP. The aim of our paper is to provide a thorough security analysis of MQ-based commitments and derive an efficient construction for post-quantum SEP.

Contributions of the paper. We consider the following simple and natural commitment scheme whose security relies on the hardness of the MQ problem. Let q be a prime power and $\mathbb{F}_q$ denote the finite field with q elements. Given integers (k, n, m), with $m > k$, we consider two random quadratic maps $\mathcal{F} : \mathbb{F}_q^n \longrightarrow \mathbb{F}_q^k$ and $\mathcal{G} : \mathbb{F}_q^n \longrightarrow \mathbb{F}_q^{m-k}$. The commitment of a message $\mu \in \mathbb{F}_q^k$ is computed as $c = (\mu + \mathcal{F}(r), \mathcal{G}(r)) \in \mathbb{F}_q^m$ where r is picked uniformly at random in $\mathbb{F}_q^n$.

Hiding property. One can see that if the quadratic map $x \mapsto (\mathcal{F}(x), G(x))$ is *contracting* (i.e. $m < n$), the output $(\mathcal{F}(r), \mathcal{G}(r))$ becomes *statistically* close to a uniform distribution and the commitment is *statistically hiding*. Conversely, if $(\mathcal{F}, \mathcal{G})$ is *expanding*, the situation changes: with high probability, the output $(\mathcal{F}(r), \mathcal{G}(r))$ has a unique preimage r, and if an adversary can efficiently test for the existence of such an r corresponding to candidate messages μ_0, μ_1, they

can distinguish which message is committed. Therefore, in order for the commitment to be hiding, the MQ problem has to be computationally intractable. We show that, in such cases, $x \mapsto (\mathcal{F}(x), G(x))$ may still be *computationally* indistinguishable from a uniform vector, assuming the hardness of MQ. This was proven for the case $q = 2$ by Berbain, Gilbert, and Patarin [BGP06], and we extend this result to arbitrary finite fields $\mathbb{F}_q$ with $q > 2$. We rely on a generalization of the Goldreich–Levin theorem [GRS00] over larger finite fields due to Dodis, Goldwasser, Tauman Kalai, Peikert, and Vaikuntanathan [DGK+10].

Binding property. If an adversary is able to find two different values $r, r' \in \mathbb{F}_q^n$ with $r \neq r'$ such that $\mathcal{G}(r) = \mathcal{G}(r')$, then they can break the binding property of the commitment scheme since a commitment $c = (\mu + \mathcal{F}(r), \mathcal{G}(r))$ is also equal to $(\mu' + \mathcal{F}(r'), \mathcal{G}(r'))$ with $\mu' = \mu - \mathcal{F}(r') + \mathcal{F}(r)$, and can be opened to μ and μ' (and $\mu' \neq \mu$ with high probability). As mentioned above, it is folklore that multivariate quadratic maps are not binding (or collision-resistant) in general. We analyze the computational hardness of finding collisions in quadratic maps over finite fields. We first prove that the general problem is NP-complete by providing a reduction to MQ. The problem of finding a collision can be formulated as finding a solution to a system of *bilinear* multivariate quadratic maps. It is interesting to study the concrete hardness of solving such problems using known techniques. Classical algorithms for solving MQ systems include exhaustive search, which has exponential complexity, and algebraic techniques such as Gröbner basis methods or the XL algorithm proposed by Courtois, Klimov, Patarin, and Shamir in 2000 [CKPS00], which offer more efficient resolution but still have exponential complexity in the number of variables. The hybrid method [YC04,BFP09,BFP12] that combines exhaustive search and algebraic techniques is currently the best technique to solve quadratic systems over finite fields.

Claw-freeness of quadratic maps. We also analyze the hardness of finding *claws* between two independent random quadratic maps $\mathcal{F}, \mathcal{G} : \mathbb{F}_q^n \to \mathbb{F}_q^m$, that is, pairs (x, y) such that $\mathcal{F}(x) = \mathcal{G}(y)$. We first formalize the corresponding Quadratic Map Claw problem and observe that it is NP-complete, as it directly generalizes the classical MQ problem. We then study its average-case properties and analyze concrete algorithms for finding claws, focusing on adaptations of Gröbner basis and XL-type methods to the specific algebraic structure of claw systems. Our analysis indicates that exploiting the claw structure yields only marginal algorithmic gains, supporting the assumption that random quadratic maps can be safely treated as claw-free for practical parameters. As a by-product of this analysis, we propose another bit commitment scheme that achieves statistical hiding and whose binding property directly reduces to the claw-freeness assumption.

Commitments with efficient zero-knowledge proofs. When commitments serve as foundational components in complex cryptographic protocols, it is often crucial to demonstrate specific properties of a committed message without disclosing any further information about its content. This requirement is typically addressed using zero-knowledge proofs (or arguments) of knowledge [GMR89]. Many works

were devoted to designing quantum-resilient commitment schemes together with efficient zero-knowledge protocols: e.g., based on the learning parity with noise assumption [JKPT12], on lattice-based assumptions [BKLP15,BDL+18,ALS20], or on the subset-sum problem [MV23]. Our MQ commitment protocol is particularly well-suited for integration with zero-knowledge arguments, as it enables not only efficient arguments of *knowledge of the opening*, but also more general arguments of *relations over the committed values*. Our zero-knowledge arguments rely on recent variants of the MPC-in-the-head paradigm [IKOS09], a technique in which the prover simulates a multi-party computation and reveals selective views to convince the verifier. In particular, we adopt the *Threshold-Computation-in-the-Head* (TCitH) framework due to Feneuil and Rivain [FR25] that is currently one of the most efficient MPCitH-based instantiations. Proving the knowledge of opening is very similar to the MQOM signature proposed in [BFR24,BBFR25] and currently a candidate in the NIST post-quantum standardization process. We show how to adapt it to prove that committed values satisfy a given (public) arithmetic circuit over $\mathbb{F}_q$ (without revealing further information). Table 1 summarizes existing post-quantum commitments with zero-knowledge arguments of knowledge of opening achieving 128-bit security, using metrics from the respective papers (though message lengths and security properties may vary).

Scheme	Assumptions	Commitment size	Argument size
[BKLP15]	MLWE	54.5 KB	30 KB
[BDL+18]	MLWE	8.1 - 9 KB	6.6 - 29 KB
[ALS20]	MLWE	5.5 - 9 KB	5.1 - 31.3 KB*
[MV23]	Subset-Sum	32 - 64 B	33 - 35 KB
Ours	**MQ**	**131 - 257 B**	**3.44 - 3.79 KB**

Table 1. Comparison of existing post-quantum commitments with zero-knowledge arguments of knowledge of opening reaching 128 bits of security.
* The provided sizes are actually for a more complex "product proof"

Blind signatures with efficient protocols. Fischlin's approach [Fis06] to blind signatures based on commitments involves having the signer generate a signature on a committed message μ and the user's final signature on a message μ is a non-interactive zero-knowledge proof that they know a commitment c of μ and a signature on c. This elegant approach ensures both blindness and unforgeability. It is natural to try to adapt it to our MQ commitment scheme. In order to stay within the realm of multivariate cryptography and achieve good efficiency, we choose to use signature schemes that rely on (structured) MQ problems [KPG99,Beu22b,FIKT21]. We consider another (trapdoor) quadratic map $\mathcal{H} : \mathbb{F}_q^\ell \longrightarrow \mathbb{F}_q^m$, for some integer ℓ, and, given a commitment $c \in \mathbb{F}_q^m$, the signer computes—using their trapdoor—a preimage $x \in \mathbb{F}_q^\ell$ such that $\mathcal{H}(x) = c = ((\mu + \mathcal{F}(r), \mathcal{G}(r))$. Proving the knowledge of (μ, r, x) can then simply be done

efficiently using the TCitH framework. The unforgeability of this scheme under a no-message attack relies on another variant of the MQ that we also analyze in detail in the paper (finding claws between quadratic functions). Unfortunately, the one-more unforgeability of the resulting scheme cannot be proven under a well-established, non-interactive, and falsifiable variant of the MQ assumption. Instead, as in Chaum's original RSA-based blind signature scheme [Cha82], our security proof must rely on an interactive assumption of a *one-more-inversion* flavor [BNPS03], where an adversary is challenged to invert a quadratic map more times than it is allowed to query a trapdoor oracle. While this assumption appears reasonable in light of the current hardness of MQ inversion problems, it remains relatively unexplored and likely requires further cryptanalytic scrutiny. We therefore present our scheme as a promising *candidate* for blind signatures in the multivariate setting.

Beyond blind signatures, privacy-preserving authentication mechanisms aim to limit information disclosure to what is functionally necessary. A common requirement in these schemes is proving knowledge of a signature on hidden attributes, often achieved using SEP [CL03,CL04,BB08,PS16]. Transitioning these mechanisms to the quantum-resilient setting remains a challenge, and only lattice-based constructions are known [LLM+16,JRS23,AGJ+24] with still large proof sizes, in spite of the recent improvements. Recently, Bouillaguet, Feneuil, Maire, Rivain, Sauvage and Vergnaud [BFM+25] introduced the notion of *commit-append-and-prove* systems, which generalizes traditional commit-and-prove systems. They used this notion and the TCitH framework to construct blind versions of structured multivariate signatures (for UOV variants, [KPG99,Beu22b,FIKT21]) and code-based signatures (Wave, [BCC+23]). Their blind UOV signatures range from 3.8 KB to 11 KB, but are not suitable for designing SEP.

Using the TCitH framework, we demonstrate that our approach yields highly competitive multivariate SEP. As an illustration, applying our approach to a use case suggested by [JRS23] yields proof sizes of 13.6 KB, whereas the lattice-based constructions required 724 KB in [JRS23] and 80 KB in [AGJ+24]. Table 2 summarizes existing post-quantum SEP achieving 128-bit security.

Scheme	Assumptions	Signature size	Credential size
[JRS23]	MSIS/MLWE	261 - 418 KB	724 KB
[BLNS23b]	SIS + LWE + NTRU	-	62 - 243 KB
[AGJ+24]	MSIS/MLWE	6.8 KB	80 KB
[BFM+25]	Wave	39 - 138 KB	—
[BFM+25]	UOV	3.8 - 11 KB	—
Ours	**One-more Quadratic Claw**	**6.8 KB**	**13.6 KB**

Table 2. Comparison of existing post-quantum blind signatures with efficient protocols reaching 128 bits of security. The credential size is given of the use case suggested in [JRS23], namely 10 128-bit attributes, of which 4 are disclosed.

2 Preliminaries

Multivariate quadratic systems. Let $\mathbb{F}_q$ denote the finite field with q elements and $\mathbb{F}_q[x_1, \ldots, x_n]$ denote the ring of multivariate polynomials in n variables. We usually consider a sequence of m polynomials $f_1, \ldots, f_m$ in n variables with coefficients in $\mathbb{F}_q$.

Definition 1 (Quadratic Function). *A* quadratic function $\mathcal{F} : \mathbb{F}_q^n \to \mathbb{F}_q^m$ *is such that the i-th coordinate of $\mathcal{F}(x)$ is a quadratic polynomial in the n coordinates of x.*

A quadratic function is fully described by the $m(n+1)(n+2)/2$ coefficients of the m quadratic polynomials. We denote by $\mathsf{MQ}(n, m, q)$ the set of all quadratic functions from $\mathbb{F}_q^n$ to $\mathbb{F}_q^m$ with coefficients in $\mathbb{F}_q$. In other terms, the letters m and n denote the number of polynomials and variables, respectively.

Indistinguishability. Two distributions $\{D_\kappa\}_\kappa$ and $\{E_\kappa\}_\kappa$ indexed by a security parameter κ are (t, ε)-*indistinguishable* (where t and ε are $\mathbb{N} \to \mathbb{R}$ functions) if, for any algorithm $\mathcal{A}$ running in time at most $t(\kappa)$ we have

$$\left| \Pr[\mathcal{A}^{D_\kappa}() = 1] - \Pr[\mathcal{A}^{E_\kappa}() = 1] \right| \leq \varepsilon(\kappa) \ ,$$

with $\mathcal{A}^{Dist}$ meaning that $\mathcal{A}$ has access to a sampling oracle of distribution $Dist$.

Commitments. In the following, a *commitment scheme* is a triple of algorithms $(\mathsf{Setup}, \mathsf{Com}, \mathsf{Verif})$ where for a security parameter κ:

- $\mathsf{Setup}(1^\kappa)$ generates public parameters pp;
- $\mathsf{Com}(\mathsf{pp}, \mu, \mathsf{aux})$ takes a message μ and randomness aux, and returns a commitment c;
- $\mathsf{Verif}(\mathsf{pp}, \mu, c, \mathsf{aux})$ checks whether c is a valid commitment to μ.

with the correctness property that, for any μ and aux, we have:

$$\mathsf{Verif}(\mathsf{pp}, \mu, \mathsf{Com}(\mathsf{pp}, \mu, \mathsf{aux}), \mathsf{aux}) = 1.$$

The scheme is:

- *Hiding* if no efficient adversary can distinguish which of two messages was committed, beyond a negligible advantage over 1/2;
- *Binding* if no efficient adversary can find two valid openings (μ_1, aux_1) and (μ_2, aux_2) to the same commitment with $\mu_1 \neq \mu_2$.

These notions come in different flavors, typically defined either in a computational or statistical setting, depending on the desired level of security. The formal definitions are provided in Appendix A. The main goal of this article is to analyze the commitment scheme shown in Protocol 1. The sizes of the integers q, k, n, m output by $\mathsf{Setup}(1^\kappa)$ to achieve a given security level κ, will be derived in the following. Intuitively, the pseudo-randomness of $\mathcal{F}(r)$ hides the committed value, while the collision-resistance of $\mathcal{G}$ makes the commitment binding. We analyze these properties in the following two sections.

$\mathsf{Setup}(1^\kappa)$ outputs a prime power q, positive integers k, n, m and two random quadratic maps $\mathcal{F} \in \mathsf{MQ}(n, k, q)$ and $\mathcal{G} \in \mathsf{MQ}(n, m-k, q)$ (i.e., $\mathsf{pp} \leftarrow (q, k, n, m, \mathcal{F}, \mathcal{G})$).

$\mathsf{Com}(\mathsf{pp}, \mu, \mathsf{aux})$ is defined for messages $\mu \in \mathbb{F}_q^k$ and randomness $\mathsf{aux} = r \xleftarrow{\$} \mathbb{F}_q^n$. A commitment c is computed as $c \leftarrow (\mu + \mathcal{F}(r), \mathcal{G}(r))$.

$\mathsf{Verif}(\mathsf{pp}, \mu, c, \mathsf{aux})$ is simply the canonical verification procedure that recomputes the commitment using (μ, aux) and checks equality with c.

Protocol 1: Computationally hiding MQ-based commitment scheme.

3 Randomness properties of quadratic maps

3.1 Universal hashing

The set of all quadratic maps $\mathsf{MQ}(n, m, q)$ forms a family of hash functions (as defined in [Gol01, §3.5.1.1]). First of all, when $\mathcal{F} \xleftarrow{\$} \mathsf{MQ}(n, m, q)$, then the value of $\mathcal{F}(x)$ is uniformly random over $\mathbb{F}_q^m$ for any $x \in \mathbb{F}_q^n$. This follows from the random choice of the constant terms of $\mathcal{F}$. But in fact, we have a bit more: $\mathsf{MQ}(n, m, q)$ is a *universal hash function family*. We believe that this was already known; in any case [AHI+17] shows that they are universal one-way hash functions. We now state two simple results that apply to all universal hash functions families, and in particular to quadratic maps. All proofs have been postponed to Appendix B.

Lemma 1. *A random quadratic map $\mathcal{F} \in \mathsf{MQ}(n, m, q)$ is:*

i) Injective with probability at least $1 - q^{2n-m}$.
ii) Surjective with probability at least $1 - q^{2m-n}$.

When two random *contracting* quadratic maps are evaluated on a random input, the output distributions are statistically close with overwhelming probability (over the choice of the functions).

Lemma 2. *For any pair $(\mathcal{F}, \mathcal{G})$ of quadratic maps in $\mathsf{MQ}(n, m, q)$, define*

$$\Delta(\mathcal{F}, \mathcal{G}) = \sum_{\alpha \in \mathbb{F}_q^m} \left| \mathbb{P}_x[\mathcal{F}(x) = \alpha] - \mathbb{P}_x[\mathcal{G}(x) = \alpha] \right|.$$

Then

$$\mathbb{P}_{\mathcal{F},\mathcal{G}} \left[\Delta(\mathcal{F}, \mathcal{G}) \geq 2q^{\frac{m-n}{3}} \right] \leq 2q^{\frac{m-n}{3}}.$$

3.2 Pseudorandomness

Intuitively, if $\mathcal{F}$ is a contracting quadratic map, then $\mathcal{F}(x)$ is *statistically* indistinguishable from a uniformly random vector. On the other hand, if $\mathcal{F}$ is expanding, then this cannot hold, because x is the single preimage of $\mathcal{F}(x)$ with high probability, and it can be found by inverting the quadratic map. In this

1: **function** DOT-PRODUCT-DISTINGUISHER($\mathcal{F}, y, r, \sigma$)
2: Sample a random vector $a \xleftarrow{\$} \mathbb{F}_q^m$.
3: Define $\mathcal{F}'(z) := \mathcal{F}(z) + \langle r, z\rangle \cdot a$ and $y' := y + \sigma \cdot a$.
4: Set $b \leftarrow \mathcal{B}(\mathcal{F}', y')$. ▷ if $\sigma = \langle r, x\rangle$, then $y' = \mathcal{F}'(x)$
5: **return** b.

Algorithm 1: Pseudo-code of the dot-product distinguisher needed by Theorem 2.

case, $\mathcal{F}(x)$ is *computationally* indistinguishable from a uniformly random vector, when x is randomly chosen. This has been proved when $q = 2$ by Berbain, Gilbert and Patarin [BGP06]. We extend the result to arbitrary finite fields; this requires the use of a generalization of the Goldreich-Levin theorem to finite fields with more than two elements. Two such results are readily available in the literature: one in [GRS00] that uses a *predictor* and one in [DGK+10] that uses a *distinguisher*. We use the latter to establish the following.

Theorem 1. *If there is an algorithm $\mathcal{D}$ that (τ, ϵ)-distinguishes $\mathcal{F}(x)$ from uniform, where $\mathcal{F} \in \mathsf{MQ}(n, m, q)$ and $x \xleftarrow{\$} \mathbb{F}_q^n$, then there is an algorithm that (τ', ϵ')-inverts $\mathcal{F}$ with*

$$\tau' = 128\frac{q^2n^2}{\epsilon^2}\tau + [\textit{negligible terms}], \qquad \textit{and} \qquad \epsilon' \geq \frac{\epsilon^3}{512nq^2}.$$

The generalized Goldreich-Levin theorem from [DGK+10] requires a dot-product distinguisher, namely an algorithm that distinguishes $\langle r, x\rangle$ from random, given $\mathcal{F}$, $y := \mathcal{F}(x)$ and r. In exchange, it provides an algorithm that solves $\mathcal{F}(x) = y$.

Lemma 3. *Let $\mathcal{F} \in \mathsf{MQ}(n, m, q)$ be a quadratic map, let $x \xleftarrow{\$} \mathbb{F}_q^n$ and define $y := \mathcal{F}(x)$. Suppose that there is an algorithm $\mathcal{B}$ that (τ, ϵ)-distinguishes $\mathcal{F}(x)$ from random. Then the* DOT-PRODUCT-DISTINGUISHER *function shown in Algorithm 1, (τ', ϵ)-distinguishes $\langle r, x\rangle$ from random given $\mathcal{F}$, $y := \mathcal{F}(x)$ and a random $r \xleftarrow{\$} \mathbb{F}_q^n$, with $\tau' = \tau + \mathcal{O}(n^2m)$.*

Proof. If $\sigma = \langle r, x\rangle$, then $y' = \mathcal{F}'(x)$. Otherwise $y' = \mathcal{F}'(x) + (\sigma - \langle r, x\rangle)) \cdot a$. Because a is uniformly random, then so is y'. It follows that $\mathcal{D}$ has the same distinguishing advantage as $\mathcal{B}$. □

We combine this with the following.

Theorem 2. ***[Theorem 1 from [DGK+10], slighlty modified]*** *Let q be a prime power and let $\mathcal{F} \in \mathsf{MQ}(n, m, q)$ denote a quadratic map. If there is a dot-product distinguisher $\mathcal{D}$ that runs in time τ such that*

$$\left|\Pr\left[x \xleftarrow{\$} \mathbb{F}_q^n, y \leftarrow \mathcal{F}(x), r \xleftarrow{\$} \mathbb{F}_q^n : \mathcal{D}(\mathcal{F}, y, r, \langle r, x\rangle) = 1\right]\right.$$
$$\left. - \Pr\left[x \xleftarrow{\$} \mathbb{F}_q^n, y \leftarrow \mathcal{F}(x), r \xleftarrow{\$} \mathbb{F}_q^n, u \xleftarrow{\$} \mathbb{F}_q : \mathcal{D}(\mathcal{F}, y, r, u) = 1\right]\right| = \epsilon$$

then INVERTER *(described in Algorithm 2)* (τ', ϵ')*-inverts* $\mathcal{F}$*. More precisely, it runs in time*

$$\tau' = \left(192\frac{n}{\epsilon^2} + 128\frac{q^2n^2}{\epsilon^2}\right)\tau + \mathcal{O}\left(\frac{q^2n^2}{\epsilon^2}\right)$$

and is such that

$$\Pr\left[x \xleftarrow{\$} \mathbb{F}_q^n, y \leftarrow \mathcal{F}(x) : \mathcal{F}(\text{INVERTER}(\mathcal{F}, y, n, \epsilon)) = y\right] \geq \frac{\epsilon^3}{512nq^2}.$$

There are three important differences between this statement and the original in [DGK+10]:

1. The original is stated for a prime q. However, a careful review of the proof shows that this restriction is inessential; the proof is valid over any finite field.
2. The original assumes that the coefficients of x may be restricted to belong to a subset of $\mathbb{F}_q$. This restriction is not useful to us.
3. The complexity of the inverter is given asymptotically in the original, and more concretely here.

Here are a few details that are necessary to read the pseudo-code of the inverter promised by Theorem 2, shown in Algorithm 2. The running time of the function is fairly easy to establish. A quick examination of the pseudo-code shows that the distinguisher $\mathcal{D}$ is invoked by the inverter $192n/\epsilon^2$ times on line 6 and $128n^2q^2/\epsilon^2$ times on line 19. The total number of other operations performed by the algorithm is dominated by the number of invocations of the distinguisher. The success probability of the algorithm is proved in [DGK+10]. Write again $y := \mathcal{F}(x)$, and consider the probabilities:

$$\alpha_{x,y} = \Pr\left[r \xleftarrow{\$} \mathbb{F}_q^n : \mathcal{D}(\mathcal{F}, y, r, \langle r, x\rangle) = 1\right],$$
$$\beta_{x,y} = \Pr\left[r \xleftarrow{\$} \mathbb{F}_q^n, u \leftarrow \mathbb{F}_q : \mathcal{D}(\mathcal{F}, y, r, u) = 1\right].$$

By hypothesis, the dot-product distinguisher $\mathcal{D}$ is such that $\mathbb{E}_x(\alpha_{x,y} - \beta_{x,y}) \geq \epsilon$ — without loss of generality, we assume that $\alpha_{x,y} \geq \beta_{x,y}$, by flipping the decisions of $\mathcal{D}$ if needed. The pair (x, y) is said to be *good* if $\alpha_{x,y} - \beta_{x,y} \geq \epsilon/2$. It follows that $\Pr[x \leftarrow \mathbb{F}_q^n, y \leftarrow \mathcal{F}(x) : (x, y) \text{ is good}] \geq \epsilon/2$. As in line 2 of the algorithm, let c denote the smallest integer such that $q^c \geq 128qn/\epsilon^2$. Algorithm 2 recovers x from y probability greater than $\epsilon/4q^c$, whenever the pair (x, y) is good.

Proof (of Theorem 1). Suppose that an adversary $\mathcal{A}$ is capable of (τ, ϵ)-distinguishing $\mathcal{F}(x)$ from random, where $\mathcal{F} : \mathbb{F}_q^n \to \mathbb{F}_q^m$ is a quadratic map and x is randomly chosen in $\mathbb{F}_q^n$. Then Lemma 3 provides the distinguisher needed by Theorem 2. □

function INVERTER(n, ϵ, y)
  Set c to the smallest integer such that $q^c \geq 128qn/\epsilon^2$. ▷ note that $c \geq 2$
  Set $N \leftarrow 0$. ▷ Estimate $\beta_{x,y}$
  for $i = 1, \ldots, 192n/\epsilon^2$ **do**
    Sample a vector $r \xleftarrow{\$} \mathbb{F}_q^n$ and a scalar $u \xleftarrow{\$} \mathbb{F}_q$.
    if $\mathcal{D}(y, r, u) = 1$ **then** increment N.
  Set $e \leftarrow N\epsilon^2/n/192$. ▷ $\Pr[|e - \beta| \geq \epsilon/8] \leq 2^{-n}$ as soon as $n \geq 3$
  Set $\gamma \leftarrow e + \epsilon/4$. ▷ $\beta_{x,y} + \epsilon/8 \leq \gamma \leq \alpha_{x,y} - \epsilon/8$.
  Choose c random vectors $z_1, \ldots, z_c$ in $\mathbb{F}_q^n$ and c random scalars $g_1, \ldots, g_c$ from $\mathbb{F}_q$.
    ▷ g_i is a "guess" for the value of $\langle z_i, x \rangle$
  Set $m \leftarrow 128qn/\epsilon^2$. Sample a random m-subset S of $\mathbb{F}_q^c/\{0\}$.
  for every $\rho = (\rho_1, \ldots, \rho_c) \in S$ **do**
    Compute $r_\rho = \sum_{j=1}^{c} \rho_j z_j$ and $h_\rho = \sum_{j=1}^{c} \rho_j g_j$.
      ▷ if the "guesses" g_i are all correct, then $h_\rho = \langle r_\rho, x \rangle$ for every ρ
  for $i = 1, 2, \ldots, n$ **do**
    Set $x_i \leftarrow \perp$. ▷ the following tries to determine x_i
    for $a \in \mathbb{F}_q$ **do** ▷ assume that $x_i = a$; the following checks if this is correct
      Set $N \leftarrow 0$.
      for $\rho \in S$ **do**
        Pick a random scalar $\tau \xleftarrow{\$} \mathbb{F}_q$.
        if $\mathcal{D}(y, r_\rho + \tau \cdot e_i, h_\rho + \tau \cdot a) = 1$ **then** increment N.
      if $N \geq m\gamma$ **then** set $x_i \leftarrow a$.
    if $x_i = \perp$ **then abort** the algorithm. ▷ could not determine x_i
  return x

Algorithm 2: Pseudo-code of the inverter promised by Theorem 2.

4 Collision-resistance and claw-freeness of quadratic maps

4.1 Collision-resistance

Quadratic maps are not necessarily collision-resistant. We believe that this result belongs to folklore; in any case, it is explicitly stated in [BRP07,AHI+17,Beu24]. More precisely, the "difference equation"

$$\mathcal{F}(x + \Delta^{in}) = \mathcal{F}(x) + \Delta^{out} \qquad (\Delta^{in} \neq 0) \tag{1}$$

can potentially be easy to solve when Δ^{in} and Δ^{out} are fixed. Write the i-th coordinate of $\mathcal{F}(x)$ as $f_i(x) = {}^t x A_i x + b_i x + c_i$, where A_i is an $n \times n$ matrix. On

the i-th coordinate, (1) becomes:

$$\begin{aligned} f_i(x) + \Delta_i^{out} &= f_i(x + \Delta^{in}) \\ &= {}^t(x + \Delta^{in}) A_i (x + \Delta^{in}) + b_i x + b_i \Delta^{in} + c_i \\ &= f_i(x) + {}^t x A_i \Delta^{in} + {}^t \Delta^{in} A_i x + {}^t \Delta^{in} A_i \Delta^{in} + b_i \Delta^{in} \\ \Delta_i^{out} &= {}^t x (A_i + {}^t A_i) \Delta^{in} + {}^t \Delta^{in} A_i \Delta^{in} + b_i \Delta^{in} \end{aligned} \tag{2}$$

If Δ^{in} is fixed to a non-zero value, then this is a system of m *linear* (in fact, affine) equations in the n coefficients of x. If it admits a solution, then this solution is easy to find. This is in particular the case when $m \leq n$.

In a recent article [Beu24], Beullens observed that, for this reason, the tentative commitment scheme, proposed by Petzoldt, Szepieniec and Mohamed in [PSM17], and defined by $\mathsf{Com}(\mu, r) := H(\mu) + \mathcal{F}(r)$ is not binding when $\mathcal{F} \in \mathsf{MQ}(n, m, q)$ with $m \approx n$ and H is a random function with values in $\mathbb{F}_q^m$. Indeed, given two arbitrary values μ_0, μ_1, an adversary can produce a commitment c that can be opened to either μ_0 or μ_1. To do so, set $\Delta^{out} = H(\mu_0) - H(\mu_1)$, choose an arbitrary non-zero Δ^{in} and solve (2). This yields an x such that $\mathcal{F}(x + \Delta^{in}) + H(\mu_1) = \mathcal{F}(x) + H(\mu_0)$. As a consequence, $\mathsf{Com}(\mu_0, x) = \mathsf{Com}(\mu_1, x + \Delta^{in})$.

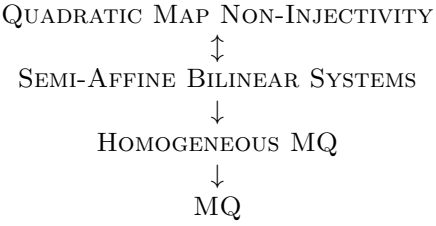
4.2 NP-hardness

Given the coefficients of a quadratic function $\mathcal{F}$, it is well-known that the problem of deciding whether $\mathcal{F}(x) = 0$ has a solution is NP-hard [GJ79]. This is the "decisional MQ problem" (or simply MQ). We show in the following that the problem of deciding if (1) has a solution with $\Delta^{out} = 0$ is hard in general, at least in odd characteristic. Note that this does not contradict the fact that this is easy in special cases, namely when $m \approx n$. This is the Quadratic Map Non-Injectivity problem in Fig. 1.

We strongly believe that it is NP-hard in the usual sense (using Karp reductions), but a proof of this fact eluded us. Nevertheless, we are able to prove the following slightly weaker statement (using Cook reductions).

Theorem 3. *If a polynomial-time algorithm decides* Quadratic Map Non-Injectivity*, then* P = NP*.*

The proof goes through a series of reductions between the problems described in Fig. 1:



Quadratic Map Non-Injectivity

INSTANCE Integers n, m, q; a quadratic map $\mathcal{F} \in \mathsf{MQ}(n, m, q)$.
QUESTION Do there exist $x \neq y \in \mathbb{F}_q^n$ such that $\mathcal{F}(x) = \mathcal{F}(y)$?

Homogeneous MQ

INSTANCE Integers n, m, q; matrices $A_1, \ldots, A_m$ in $\mathbb{F}_q^{n \times n}$.
QUESTION Is there a vector $x \in \mathbb{F}_q^n$ such that $x \neq 0$ and ${}^t x A_i x = 0$ for all $1 \leq i \leq m$?

Semi-Affine Bilinear Systems

INSTANCE Integers n, m, q; matrices $A_1, \ldots, A_m$ in $\mathbb{F}_q^{n \times n}$; vectors $b_1, \ldots, b_m$ in $\mathbb{F}_q^n$.
QUESTION Do there exist $x, y \in \mathbb{F}_q^n$ such that $y \neq 0$, and ${}^t x A_i y + b_i y = 0$ for all $1 \leq i \leq m$?

Quadratic Map Claw

INSTANCE Integers n, m, q; two quadratic maps $\mathcal{F}, \mathcal{G} \in \mathsf{MQ}(n, m, q)$.
QUESTION Do there exist $x, y \in \mathbb{F}_q^n$ such that $\mathcal{F}(x) = \mathcal{G}(y)$?

Fig. 1. Related hard computational problems.

The starting point is the fact that Homogeneous Quadratic Systems is NP-complete. This is shown by a reduction to MQ in [Koi00,GKP13].

If q is odd, then any quadratic polynomial $f(x) = {}^t x A x + bx + c$ can be written with a *symmetric* matrix A: observe that ${}^t x A x = {}^t x A^t x$, so that $2f(x) = {}^t x (A + A^t) x + 2bx + 2c$, and $f(x) = {}^t x \frac{1}{2}(A + A^t) x + bx + c$. The matrix $\frac{1}{2}(A + A^t)$ describing the quadratic terms is symmetric. If A_i is symmetric, then (2) becomes $\Delta_i^{out} = \left(2x + \Delta_i^{in}\right)^t A_i \Delta_i^{in} + b_i \Delta^{in}$. It follows that there is a collision $\mathcal{F}(x) = \mathcal{F}(x'), x \neq x'$ iff the system $y A_i z + b_i z = 0$ $(1 \leq i \leq m)$ admits a solution with $z \neq 0$. The collision can then be recovered with $\Delta^{in} \leftarrow z$ and $x \leftarrow (y - \Delta^{in})/2$. This is exactly the Semi-Affine Bilinear Systems problem in Fig. 1.

Lemma 4. *There is a polynomial-time reduction $\mathcal{R}$ that takes an instance I of* Homogeneous MQ *and produces n instances of* Semi-Affine Bilinear Systems*, such that:*

i) If I has no solution, then none of the instances in $\mathcal{R}(I)$ has a solution.
ii) If I has a solution, then at least one of the instances in $\mathcal{R}(I)$ has a solution.

Proof. Let $I = (q, n, m, A_1, \ldots, A_m)$ denote an instance of Homogeneous MQ. Let us have two sets of $n+1$ variables $(v, x_1, \ldots, x_n)$ and $(u, y_1, \ldots, y_n)$. We define a sequence of n semi-affine bilinear systems, denoted by $\mathcal{R}_1, \ldots, \mathcal{R}_n$, where $\mathcal{R}_k(I)$ is defined as:

$$
\begin{aligned}
0 &= {}^t x A_i y && 1 \leq i \leq m && (a)\\
0 &= x_j v - y_j && 1 \leq j \leq n && (b)\\
0 &= u y_j - y_j && 1 \leq j \leq n && (c)\\
0 &= uv - v && && (d)\\
0 &= uv - y_k && && (e)
\end{aligned}
$$

i) If (u, x, v, y) is a solution of $\mathcal{R}_k(I)$ with $(v, y) \neq 0$, then x is a non-trivial solution of I. Indeed, if $v = 0$; then (b) would imply $y = 0$, a contradiction. Therefore, $v \neq 0$. Then (d) implies $u = 1$, and (e) implies $y_k \neq 0$. Because of (b) again, $y_j = x_j v$, and $x_k \neq 0$. Finally, (a) shows that ${}^t x A_i x = 0$ for all $1 \leq i \leq m$. Since, $v \neq 0$, then x is a solution of I, and $x \neq 0$.

ii) Suppose that I has a solution x with $x_k \neq 0$. Then $x' := x_k^{-1} x$ is also a solution of I with $x'_k = 1$, and $(1, x', 1, x')$ is a non-trivial solution of $\mathcal{R}_k(I)$.

□

Proof (of Theorem 3). Suppose such an algorithm exists. It has been shown above that when q is odd, this algorithm can be used to solve Semi-Affine Bilinear Systems in polynomial time. Therefore, denote by $\mathcal{A}$ a polynomial-time algorithm for Semi-Affine Bilinear Systems.

We now build a polynomial-time algorithm for Homogeneous MQ, an NP-complete problem, thus showing that P = NP. Let $\mathcal{R}$ denote the one-to-many reduction given by Lemma 4. Take a homogeneous quadratic system; apply $\mathcal{R}$ to obtain n semi-affine bilinear systems; decide each of them in polynomial time using $\mathcal{A}$; if at least one of them has a solution, return "yes"; otherwise return "no". □

4.3 Efficient algorithms for finding collisions

We now consider the most efficient ways to find a collision in a given quadratic map. There is a simple algorithm: choose a random Δ^{in}, try to solve the linear system given by (2) with $\Delta^{out} = 0$, retry if there is no solution. Let us heuristically assume that this is a random linear system of m equations in n unknowns. If $m \geq n$, then the probability that it admits a solution is approximately q^{n-m} (we verified this assumption experimentally). Then the above procedure requires q^{m-n} iterations on average. It follows that $m = n + \lambda / \log_2 q$ is the minimum output size that may potentially yield λ bits of collision-resistance. In other terms, *sufficiently expanding* quadratic maps may be collision-resistant.

As already argued above, when q is odd and A_i is assumed to be symmetric, a collision $\mathcal{F}(x) = \mathcal{F}(x')$ can be found by solving the semi-affine bilinear system:

$$ {}^t y A_i z + b_i z = 0, \qquad z \neq 0. \tag{3}$$

Let $\alpha \in \mathbb{F}_q$, observe that if $(\hat{y}, \hat{z})$ is a solution of (3), then $(\hat{y}, \alpha \hat{z})$ is too. It follows that if there is a solution where $\hat{z}_k \neq 0$ (for some index k), then there is a solution where $\hat{z}_k = 1$. Therefore, to solve the system, we pick $j \xleftarrow{\$} \{1, n\}$ and add the extra equation $z_j = 1$. If $\hat{z}_j \neq 0$ (this is true with high probability), then the system we have written has a solution. This reduces the problem to finding a solution to a random affine bilinear system of m equations in $(n, n-1)$ variables. We study the computational cost of this problem in Section 5.2.

4.4 Claw-freeness of quadratic maps

Let us be given two random quadratic maps $\mathcal{F}, \mathcal{G} : \mathbb{F}_q^n \to \mathbb{F}_q^m$. Finding a *claw* between $\mathcal{F}$ and $\mathcal{G}$ means finding (x, y) such that $\mathcal{F}(x) = \mathcal{G}(y)$. This is the

Quadratic Map Claw problem in Fig. 1. It is clearly NP-complete: just set $\mathcal{G} = 0$, then there exists a claw between $\mathcal{F}$ and $\mathcal{G}$ if and only $\mathcal{F}(x) = 0$ has a solution. This is exactly the MQ problem.

Intuitively, if the two quadratic maps are sufficiently expanding, then the probability that there is a claw is very small. For any $\alpha \in \mathbb{F}_q^m$, let $N_\alpha(\mathcal{F})$ denote the number of preimages of α for $\mathcal{F}$. We have $N = \sum_{z \in \mathbb{F}_q^m} N_z(\mathcal{F}) N_z(\mathcal{G})$. Because $\mathcal{F}$ and $\mathcal{G}$ are drawn independently, and by linearity of expectation, we find that $\mathrm{E}\, N = q^{2n-m}$. By Markov's inequality, $\Pr[N > 0] \leq \mathrm{E}\, N = q^{2n-m}$.

As opposed to the problem of finding a collision on $\mathcal{F}$, finding a claw on $(\mathcal{F}, \mathcal{G})$ seems hard even when both $\mathcal{F}$ and $\mathcal{G}$ are contracting (when the expected number of such claws is high). We conjecture that it is indeed the case, asymptotically, for $n = \alpha m$, for any constant α.

4.5 Efficient algorithms for finding claws

Finding a claw directly reduces to solving the quadratic system in $2n$ variables and m equations given by $\mathcal{F}(x) = \mathcal{G}(y)$. The expected number of solutions is q^{2n-m}, therefore we can arbitrarily fix the value of $2n-m$ variables while heuristically preserving the existence of a solution with high probability. This reduces the problem to a "square" quadratic system of m equations in m variables. If $n \geq m$, it is possible to "guess" y entirely: this leaves us with $\mathcal{F}(x) = cst$, a truly random instance of MQ. In this case, the complexity analysis is straightforward. The algorithm of Thomae-Wolf [TW12] may improves this if the system is sufficiently underdetermined (i.e., if $n \geq m$).

The downside of this approach is that it completely gives up on the structure of the input system. An alternative solution would be to guess $n - m/2$ variables in x and the same quantity in y, preserving the structure "$\mathcal{F}(x) = \mathcal{G}(y)$" (there are no cross-terms $x_i y_j$). We empirically observed that such systems, even though they are structured, are not easier to solve using the usual Gröbner basis algorithm than completely generic, unstructured systems (they are *regular*). Ad hoc algorithms that exploit the structure of the system are discussed in Section 5.3.

5 Complexity of state-of-the-art algorithms for MQ over "large" finite fields

Our purpose here is to obtain *lower bounds* on the number of operations required to solve random or sufficiently generic polynomial systems using state-of-the-art algorithms.

5.1 Quadratic systems

For random quadratic systems over "large" finite fields (with 16 or 256 elements), we use the same estimates as the designers of the UOV, MAYO and MQOM signature schemes submitted to the NIST call for additional signature

schemes [KPG99,Beu22b,FIKT21,BFR24,BBFR25]: we assume that the "hybrid method" that combines exhaustive search and the XL-Wiedemann [CCNY12a] algorithm is the best way to solve unstructured quadratic systems.

Given a collection of polynomials $F = (f_1, \ldots, f_m) \subset \mathcal{R}$, the *Macaulay matrix* of degree d of F is the matrix whose rows are the μf_j, where μ ranges across monomial such that $\deg \mu f_j \leq d$. The columns of the matrix correspond to all possible monomials of degree at most d.

The XL algorithm was proposed by Courtois, Klimov, Patarin and Shamir in 2000 [CKPS00]. In fact, it turned out to be a reinvention of a technique due to Lazard in 1983 [Laz83], and is more-or-less equivalent to modern Gröbner basis algorithms. Implementations of XL [CCNY12b] are currently the state of the art for solving (overdetermined) polynomial systems over "large" finite fields.

The underlying idea of the XL algorithm is simple. Pick a degree d and build the affine Macaulay matrix of degree d. Cut the column corresponding to the constant monomial. Call the resulting vector b and the truncated matrix A. Consider the linear system $A\hat{x} + b = 0$. Suppose that the polynomial system has a solution $a = (a_1, \ldots, a_n) \in \mathbb{F}_q^n$. Then the linear system $A\hat{x} + b = 0$ has a solution $\hat{x}$ that can be formed by assembling the evaluations of all monomials of degree at most d on a. In that case, the solution of the polynomial system can be read in $\hat{x}$, for instance, by looking at the coefficients that correspond to the degree-1 monomials x_i's. The trick is that if the degree of the Macaulay matrix is large enough, then this always works. The XL algorithm then proceeds as follows: Find a solution to $A\hat{x} + b = 0$; read from $\hat{x}$ the values of $x_1, \ldots, x_n$; test if all polynomials vanish on these values; if so, the system has been solved, and otherwise the algorithm has failed. If the input system spans a *maximal ideal*— this should happen with high probability if it is overdetermined and sufficiently generic— then XL succeeds if the degree of the Macaulay matrix is at least the *solving degree* of the input polynomials, namely the smallest d such that a Gröbner basis of the ideal they span is contained in the rowspan of a degree-d Macaulay matrix. If a function $f(x)$ admits a series expansion $f(x) = \sum_i a_i x^i$, then we use the notation $[x^i]f(x)$ to denote a_i. If the input quadratic system is sufficiently generic, then the solving degree is the smallest index i such that

$$[z^i] \frac{\left(1 - z^2\right)^m}{(1 - z)^{n+1}}$$

is nonpositive [Bar04,BFS15].

It is well-known that Macaulay matrices are rank-defficient. However, their rank can be easily predicted under the assumption that the polynomials are sufficiently generic. Consider the sequence $R_{d,j}$ derived from the behavior of the

matrix-F5 algorithm [Bar04,BFS15]:

$$\begin{aligned} R_{0,j} &= 0 \\ R_{1,j} &= 0 \\ R_{d,0} &= 0 \\ R_{d,j+1} &= R_{d,j} + \binom{n+d-2}{d-2} - R_{d-2,j} \end{aligned}$$

As long as d is strictly less than the solving degree, then the rank of the degree-d Macaulay matrix of $(f_1, \ldots, f_j)$ is less than $R_{d,j}$. There is equality if the f_i's form a *regular sequence* or a *semi-regular sequence*. This essentially means that the f_i's are not bound by unexpected algebraic relations. The solving degree of the system is the smallest degree such that $R_{d,m}$ is at least the number of monomials of degree at most d (eventually minus one). In fact, it is possible to show that

$$R_{d,m} = \binom{n+d}{d} - \left[z^d\right] \frac{\left(1-z^2\right)^m}{(1-z)^{n+1}},$$

so that both definitions of the solving degree coincide.

In order to solve $Ax = b$ with a sparse matrix A, the Block Wiedemann algorithm is usually the solution of choice. We refer the reader to [CCNY12a,BGG+20] for details about the algorithm and practical results. We rely on a consensual lower-bound on the number of arithmetic operations it requires:

$$T \geq 2\binom{n+2}{2}\binom{n+d}{d}^2,$$

where d is the solving degree defined above.

5.2 Bilinear systems

Let $m = n_x + n_y$ and $F = (f_1, \ldots, f_m) \subset \mathbb{F}_q[\mathbf{x}, \mathbf{y}]$ be an affine bilinear system. The complexity of solving such bilinear systems was first investigated in [FDS11], under the assumption that $m \leq n_x + n_y$. Some incorrect results in [FDS11] are rectified in [Spa12]. The *bidegree* of a monomial is the pair of integers that gives its degree in $\mathbf{x}$ and in $\mathbf{y}$, respectively. For instance, $x_1 x_3 y_2^5$ has bidegree $(2, 5)$. A bilinear monomial has bidegree $(1, 1)$. Bidegrees are compared as follows: $(a, b) \leq (c, d)$ when $a \leq c$ and $b \leq d$. The set of monomials of total degree d is partitioned into $d + 1$ categories according to the bidegree. Let $N_{u,v}$ denote the total number of monomials of bidegree at most (u, v). Then:

$$N_{u,v} := \sum_{i=0}^{u} \sum_{j=0}^{v} \binom{n_x+i-1}{i} \binom{n_y+j-1}{j}.$$

It is useful to consider the "bilinear Macaulay" matrix of a given bidegree (u, v): its column are indexed by monomials of bidegree at most (u, v) and its

rows correspond to products $\mu_j f_i$ where f_i is one of the input bilinear polynomials and μ_i is a monomial of bidegree at most $(u-1, v-1)$.

[BCV23] presents an algorithm that essentially consists in applying the XL algorithm to a Macaulay matrix of bidegree $(1, d)$, in other terms the input polynomials are multiplied by monomials in $\mathbf{y}$ only. This often yields smaller Macaulay matrices than the original XL algorithm.

[PST20] presents a generalization of this idea: a bidegree (u, v) is a *solving bidegree* if the rank of the Macaulay matrix of bidegree (u, v) is equal to its number of (non-zero) columns, eventually minus one. In this case, applying the XL algorithm to these matrices is going to work. The same article argues that a solving bidegree has to be larger than the lowest bidegrees (u, v) such that

$$[s^u t^v] \frac{(1-st)^m}{(1-s)^{n_x}(1-t)^{n_y}}$$

is non-positive.

This result can be obtained in a different way, by simulating the execution of a "bilinear variant" of the matrix-F5 algorithm. Write $S_{u,v,j}$ the sequence defined by:

$$\begin{aligned} S_{0,v,j} &= 0 \\ S_{u,0,j} &= 0 \\ S_{u,v,0} &= 0 \\ S_{u,v,j+1} &= S_{u,v,j} + N_{u-1,v-1} - S_{u-1,v-1,j} \end{aligned}$$

When (u, v) is strictly less than a solving bidegree, the rank of the bidegree-(u, v) bilinear Macaulay matrix of $(f_1, \ldots, f_j)$ is always less than or equal to $S_{u,v,j}$, with equality if the input polynomials are overdetermined and sufficiently generic (this is an empirical observation; in [PST20] this is derived from non-standard algebraic assumptions). In any case, the smallest bidegrees (u, v) such that $S_{u,v,m}$ is larger than or equal $N_{u,v} - 1$ are solving bidegrees.

This enables us to lower-bound the size of a Macaulay matrix such that a "bilinear" variant of the XL algorithm works. Indeed, the bilinear Macaulay matrix of bidegree (u, v) has $N_{u,v}$ columns, and the cost of running the XL algorithm is then

$$T \geq 2N_{u,v}^2 (n_x + 1)(n_y + 1).$$

5.3 Claw systems "$\mathcal{F}(x) = \mathcal{G}(y)$"

Suppose that $\mathcal{F}, \mathcal{G} : \mathbb{F}_q^n \to \mathbb{F}_q^m$ are two random quadratic maps. The m polynomials obtained by writing $\mathcal{F}(x) = \mathcal{G}(y)$ are linear combinations of monomials belonging to the set $M_1 := \{1, x_i, y_i, x_i x_j, y_i y_j : 1 \leq i, j \leq n\}$, where $x_i y_i$ is absent. Define $M_{k+1} = \{ab : a \in M_1, b \in M_k\}$. Then $M := \cup_k M_k$ is a (multiplicative) semigroup, and $\mathbb{F}_q[M]$ is a (graded) *semigroup algebra*. It is possible to design a variant of the XL algorithm that operates in this algebra, borrowing ideas from [FSS14] (note that some complexity results in this reference have later

been shown to be incorrect in [BFT19]). Given a collection of such polynomials $(f_1, \ldots, f_m)$, define the the order-k *semigroup Macaulay matrix* (for $k \geq 1$): its rows are the μf_j where μ ranges accross M_{k-1} and its columns are indexed by monomials in M_k. Note that it is a submatrix of the degree-$2k$ regular Macaulay matrix.

M_k contains all possible monomials of degree at most $2k$, except those of total degree $2k$ with an odd degree in x (and y). This shows that

$$|M_k| = \binom{2n+2k}{2k} - \sum_{i=0}^{k-1} \binom{n+2i}{2i+1} \binom{n+2(k-i)-2}{2(k-i)-1}.$$

We also observed empirically that the ranks of these semigroup Macaulay matrices can be easily predicted using the F5 criterion [BFS15]. Consider the sequence $T_{k,j}$ defined by:

$$\begin{aligned} T_{0,j} &= 0 \\ T_{k,0} &= 0 \\ T_{k,j+1} &= T_{k,j} + |M_{k-1}| - T_{k-1,j} \end{aligned}$$

The smallest degree where this rank becomes greater than the number of columns, is the "semigroup solving degree", and it is sufficient for the XL algorithm to operate. Under the assumption that the input polynomials are sufficiently generic (not bound by unexpected algebraic relations), and as long as d is strictly less than the solving degree, the rank of the order-k semigroup Macaulay matrix is $T_{k,m}$.

$\mathsf{Setup}(1^\kappa)$ outputs a prime power q, positive integers n, m and two random quadratic maps $\mathcal{F}_0, \mathcal{F}_1 \in \mathsf{MQ}(n, m, q)$—i.e., $\mathsf{pp} \leftarrow (q, n, m, \mathcal{F}_0, \mathcal{F}_1)$.

$\mathsf{Com}(\mathsf{pp}, b, \mathsf{aux})$ is defined for bits $b \in \{0, 1\}$ and randomness $\mathsf{aux} = r \xleftarrow{\$} \mathbb{F}_q^n$. A commitment c is computed as $c \leftarrow \mathcal{F}_b(r)$.

$\mathsf{Verif}(\mathsf{pp}, b, c, \mathsf{aux})$ is simply the canonical verification procedure that recomputes the commitment using (b, aux) and checks equality with c.

Protocol 2: Statistically hiding MQ-based commitment scheme.

Application: statistically hiding bit commitment. Consider the (classic) bit-commitment scheme [GK96] instanciated with random quadratic maps, and shown in Protocol 2. It is statistically hiding if, after $\mathcal{F}$ and $\mathcal{G}$ have been randomly chosen, the random variables $\mathcal{F}(r)$ and $\mathcal{G}(r)$ are statistically close (when r is randomly chosen). Using Lemma 2, we set $n = 2m$, and we find that, except with probability less than $2q^{-m/3}$, the statistical distance is less than $2q^{-m/3}$. It seems tempting to choose $q = 256$ and $m = 48$ (thus $n = 96$), as this is sufficient

to ensure the hiding property. Breaking the binding property is strictly equivalent to finding a claw between $\mathcal{F}$ and $\mathcal{G}$. As argued above, this problem seems hard.

Guessing all the y's and half of x leaves a random quadratic system with $m = 48$. The best way to solve it is to use guess 3 additional variables, then run the XL-Wiedemann algorithm (there are going to be q^3 iterations). This builds a degree-22 Macaulay matrix of 48 polynomials in 45 variables, that has $2^{57.9}$ columns and has full row rank. Completing each iteration requires $2^{126.9}$ arithmetic operations. The total cost is $2^{150.9}$ arihtmetic operations.

Now, consider the "semigroup" variation described above. Guess 26 variables on each side (26 in x, 26 in y). There are going to be q^4 iterations. Then assemble the order-10 semigroup Macaulay matrix. It has $2^{53.5}$ columns and has full row rank. Solving the linear system then requires $2^{116.1}$ arithmetic operations, and the total cost is $2^{148.1}$. So the semigroup variation gains a factor of about 7.

The semigroup variant guesses 4 extra variables, while the traditional solution guesses only 3. Guessing 4 variables reduces the solving degree to 20. The regular Macaulay matrix of degree 20 has $2^{54.1}$ columns: it is about 50% larger than its semigroup counterpart.

Lastly, using the Thomae-Wolf algorithm yields a square system in 45 variables, and the XL-Wiedemann algorithm runs in $2^{143.5}$ operations.

We conclude that, while the systems that arise naturally when searching a claw are structured, the ways we have found to exploit this structure only bring marginal gains. In any case, we believe that the parameters suggested above offer 128 bits of security for the statistically hiding bit-commitment.

6 Multivariate commitment with efficient protocols

This section is devoted to our main MQ-based commitment shown in Protocol 1.

Lemma 5. *The MQ-based commitment scheme of Protocol 1 is statistically binding if* $m \geq 2k + \kappa/\log_2 q$.

Proof. If $\mathcal{G}$ is injective, then only a single value of r allows the correct opening; in turn, the only possible value of μ is then fixed by c and r. The scheme is then clearly binding. The result then follows from Lemma 1. □

We believe that the best attack against its hiding property is the following: choose two arbitrary messages μ_0 and μ_1, then given the commitment c of either one, try to solve $c = (\mu_0 + \mathcal{F}(r), \mathcal{G}(r))$ for r. If a solution exists, then μ_0 was committed, otherwise it was μ_1. When $m \gg n$, i.e. when this polynomial system is overdetermined, it is unlikely that a solution exists simply out of chance, so the success probability of this procedure is very high. Beyond this heuristic argument, it is possible to show that an adversary against hiding can be turned into an MQ solver:

Theorem 4. *Suppose that there is a two-part algorithm $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ that (τ, ϵ)-breaks the hiding property of the commitment. Then there is an algorithm that (τ', ϵ')-inverts random quadratic maps in $\mathsf{MQ}(n, m, q)$, with*

$$\tau' = 512\frac{q^2n^2}{\epsilon^2}\tau \qquad \textit{and} \qquad \epsilon' = \frac{\epsilon^3}{4096nq^2}.$$

1: **function** PRG-DISTINGUISHER($\mathcal{H}, \sigma$)
2: Split $\mathcal{H}$ into two quadratic maps $\mathcal{F} : \mathbb{F}_q^n \to \mathbb{F}_q^k$ and $\mathcal{G} : \mathbb{F}_q^n \to \mathbb{F}_q^{m-k}$ such that $\mathcal{H} = \mathcal{F} \parallel \mathcal{G}$
3: $\mathsf{pp} \leftarrow (\mathcal{F}, \mathcal{G})$
4: $(\mu_0, \mu_1, s) \xleftarrow{\$} \mathcal{A}_1(\mathsf{pp})$
5: $b \xleftarrow{\$} \{0, 1\}$
6: $c \leftarrow (\mu_b \parallel 0^{m-k}) + \sigma$ ▷ if $\sigma = \mathcal{H}(\tilde{x})$, then $\mathsf{Verif}(\mathsf{pp}, \mu_b, c, \tilde{x}) = 1$ and $\mathcal{A}$ is presented a proper input
7: $b' \xleftarrow{\$} \mathcal{A}_2(c, s)$
8: **return** $b = b'$

Algorithm 3: Adversary against the pseudo-randomness of a random quadratic map using an adversary against the hiding property of the commitment.

Proof. Let $\mathcal{H}$ denote a quadratic function. We claim that the function PRG-DISTINGUISHER shown in algorithm 3 $(\tau, \epsilon/2)$-distinguishes $\mathcal{H}(x)$ from random, when $\mathcal{H} \xleftarrow{\$} \mathsf{MQ}(n, m, q)$ and $x \xleftarrow{\$} \mathbb{F}_q^n$.

PRG-DISTINGUISHER receives a quadratic function $\mathcal{H}$ and an input vector $\sigma \in \mathbb{F}_q^m$; it behaves differently when $\sigma = \mathcal{H}(x)$ than when σ is random. To this end, it simulates the environment of $\mathcal{A}$, with two modifications: 1) the Setup procedure is not invoked; instead $\mathcal{H}$ is used directly to provide the public parameters of the commitment and 2) σ is used to build the committed value.

If σ is uniformly random, then so is c. In this case, $\mathcal{A}_2$ receives strictly no information about b, and $\Pr[b' = 1 \mid b = 1] = \Pr[b' = 1 \mid b = 0]$. It follows that

$$\Pr[\mathcal{H} \xleftarrow{\$} \mathsf{MQ}(n, m, q), \sigma \xleftarrow{\$} \mathbb{F}_q^m : \text{PRG-DISTINGUISHER}(\mathcal{H}, \sigma) = 1] = 1/2.$$

Otherwise, if $\sigma = \mathcal{H}(\tilde{x})$ for some $\tilde{x}$, then PRG-DISTINGUISHER faithfully simulates $\mathcal{A}$'s expected environment. The distinguishing advantage of $\mathcal{A}$ then kicks in:

$$\left|\Pr[b' = 1 \mid b = 1] - \Pr[b' = 1 \mid b = 0]\right| = \epsilon.$$

We will drop the absolute value without loss of generality, reversing $\mathcal{A}$'s decision if needed. It follows that

$$\begin{aligned}\Pr[b = b'] &= \frac{1}{2}\Pr[b' = 1 \mid b = 1] + \frac{1}{2}\Pr[b' = 0 \mid b = 0] \\ &= \frac{1}{2}\Pr[b' = 1 \mid b = 1] + \frac{1}{2}(1 - \Pr[b' = 1 \mid b = 0]) \\ &= \frac{1}{2} + \frac{1}{2}\epsilon.\end{aligned}$$

Therefore, when $\mathcal{H} \xleftarrow{\$} \mathsf{MQ}(n, m, q)$,

$$\begin{aligned}&\Pr\left[\sigma \xleftarrow{\$} \mathbb{F}_q^m : \text{PRG-DISTINGUISHER}(\mathcal{H}, \sigma) = 1\right] - \\ &\quad \Pr\left[r \xleftarrow{\$} \mathbb{F}_q^n : \text{PRG-DISTINGUISHER}(\mathcal{H}, \mathcal{H}(r)) = 1\right] = \frac{\epsilon}{2}.\end{aligned}$$

Applying Theorem 1 concludes the proof. □

6.1 Parameters selection

We propose three sets of parameters. They all use a finite field of size $q = 256$, so that m gives the size of the commitment in bytes, with k the size of the committed value in bytes. They are shown in Table 3.

To make the commitment statistically binding, we rely on Lemma 5 and choose $m = 2n + \kappa$. It remains to choose the value of n to ensure the hiding property. Our first parameter "provably" offers the hiding property using Theorem 4. However, because it ultimately relies on a variant of the Goldreich-Levin theorem, the security proof is not tight at all and there is a large security loss. We therefore propose a second, more aggressive set of parameters by giving up on "provable security" and instead relying on the heuristic assumption that the best attack against the hiding property is the one described above. These parameters are chosen to make this attack intractable. Finally, we propose an even more aggressive set of parameters by choosing the smallest possible value of m, giving up on making the commitment statistically binding. It is then only computationally binding, based on the collision-resistance of the quadratic maps.

"Provable" computational hiding / statistical binding. To guarantee the hiding property, we conservatively commit to Theorem 4 and choose a value of n such that if an adversary (τ, ϵ)-breaks the hiding property, then $\tau \geq 2^\kappa$ or $\epsilon \leq 2^{-\kappa}$, for some security level κ, under the assumption that the XL-Wiedemann algorithm is the fastest way to invert random quadratic maps.

Suppose that there is an adversary $\mathcal{A}$ against the hiding property with $\tau \leq 2^\kappa$ and $\epsilon \geq 2^{-\kappa}$. Theorem 4 then asserts the existence of an algorithm $\mathcal{B}$ that (τ', ϵ')-solves random instances of $\mathsf{MQ}_{n,m}$, with $\tau' \leq 512q^2n^2\tau/\epsilon^2$ and $\epsilon' \geq \epsilon^3/4096nq^2$. Because XL-Wiedemann always succeeds, it seems fair to compare the running time of XL-Wiedemann with τ'/ϵ', namely the running time of $\mathcal{B}$ weighted by its

success probability, which is the expected time one has to wait before $\mathcal{B}$ succeeds in solving a quadratic system when tried repeatedly. A quick computation shows that $\tau'/\epsilon' \leq 2^{21}q^4n^3\tau/\epsilon^5$. With our choice of q and the bounds imposed on (τ, ϵ), this becomes $\tau'/\epsilon' \leq n^3 2^{53+6\lambda}$. This cannot be less than the running time of XL-Wiedemann, otherwise the assumption that XL-Wiedemann is the fastest algorithm would be contradicted.

Using the results of Section 5, we therefore choose a value of n that makes the running time of XL-Wiedemann larger than this bound. The inverter promised by Theorem 4 is thus "too efficient" and the adversary $\mathcal{A}$ cannot exist.

Hiding		Provable		Heuristic			
Binding		Statistical				Computational	
κ	k	n	m	n	m	n	m
128	32	912	1872	115	278	83	131
	128	949	2042	137	418	113	257
192	32	1344	2744	190	436	122	178
	128	1394	2940	217	586	162	314
256	32	1785	3634	260	584	161	225
	128	1831	3822	292	744	208	368

Table 3. Parameter sets for the MQ-commitment with $q = 256$.

Heuristic computational hiding / statistical binding. We still choose $m = 2n+\kappa$, so that the commitment is statistically binding. However, we assume that the best attack against the hiding property is to choose two arbitrary messages μ_0 and μ_1, then given the commitment c, solve $c = (\mu_0 + \mathcal{F}(r), \mathcal{G}(r))$. We choose n so that solving this polynomial system requires sufficiently many operations using the XL-Wiedemann algorithm.

Heuristic computational hiding / computational binding. We choose parameters such that solving $c = (\mu_0 + \mathcal{F}(r), \mathcal{G}(r))$ and finding a collision on $\mathcal{G}$ are both intractable. We are looking for the smallest m possible. For the latter, we rely on the analysis presented in Section 4.3.

6.2 Efficient arguments from *Threshold-Computation-in-the-Head*

We present two types of arguments for our commitment scheme:

- a proof of opening to convince someone that one knows the message and the randomness used for a given (public) commitment digest without revealing them;
- a proof that the committed message satisfies some public relations.

To build such arguments, we rely on the non-interactive zero-knowledge arguments described in Appendix C. Essentially, everything comes down to quadratic systems, and we follow the MQOM scheme [BFR24,BBFR25].

We consider a committed message satisfying some circuit C whose number of multiplicative gates (denoted $|C|$) is chosen for the following concrete application: *a person commits to their personal information. Later, they want to prove in zero-knowledge that she is older than 17. They can show that the committed value satisfies an arithmetic circuit C of size 18* (see Equation (5) in Appendix C). We introduce a parameter M to generalize our proof to M independently committed messages. The choice of the other parameters—N (the number of parties involved in the MPC protocol), L (the number of revealed parties), ω (the proof-of-work parameter), and t_{open} (the maximal number of nodes opened)—is detailed in Appendix C, along with security proofs.

We obtain the proof sizes described in Table 4 (using the parameters for $\kappa = 128$ security level and computational security in Table 3).

	Commitment parameters			NIZK parameters					Proved statement		Proof size
κ	k	n	m	τ	L	N	ω	t_{open}	M	$\lvert C\rvert$	(in KB)
Proof of opening											
128	128	113	257	12	1	2^{11}	9	118	-	-	3.79
128	32	83	131	12	1	2^{11}	9	118	-	-	3.44
Proof for circuit satisfiability											
128	128	113	257	12	1	2^{11}	9	118	1	18	5.50
128	32	83	131	11	1	2^{12}	9	119	1	18	4.22
128	32	83	131	12	1	2^{11}	9	118	10	18	16.16

Table 4. Performances of our MQ-based commitment scheme with $q = 256$.

7 Candidate blind signatures with efficient protocols

A blind signature scheme allows a user to obtain a signature on a message without revealing the message to the signer, ensuring both authenticity and privacy. A blind signature scheme is *round-optimal* if the interaction between the user and the signer consists of only one message from the user, one reply from the signer, and a final local computation by the user to extract the signature. In the following, it is defined as a tuple of probabilistic algorithms $(\mathsf{KeyGen}, \mathsf{Sign}_1, \mathsf{Sign}_2, \mathsf{Sign}_3, \mathsf{Verif})$ where:

- $\mathsf{KeyGen}(1^\kappa)$ outputs a key pair $(\mathsf{sk}, \mathsf{pk})$;
- $\mathsf{Sign}_1(\mathsf{pk}, \mu)$ (user side) produces a state $S_{\mathcal{U}}$ and a first message ρ_1;
- $\mathsf{Sign}_2(\mathsf{sk}, \rho_1)$ (signer side) produces a response ρ_2;
- $\mathsf{Sign}_3(S_{\mathcal{U}}, \rho_2)$ (user side) outputs a signature sig or $\perp$;
- $\mathsf{Verif}(\mathsf{sig}, \mathsf{pk}, \mu)$ checks the validity of the signature.

Such a scheme must satisfy:

- **Correctness:** honest executions produce valid signatures;
- **Blindness:** the signer cannot tell which message was signed;
- **One-more unforgeability:** a user cannot obtain more valid signatures than interactions with the signer.

The formal definitions are provided in Appendix A.

We propose a blind signature scheme based on a family of trapdoor quadratic functions. It is a tuple $\mathsf{H} = (\mathsf{KeyGen}, \mathsf{Eval}, \mathsf{Invert})$ where:

- $\mathsf{KeyGen}(1^\kappa)$ outputs a function $\mathcal{H} \in \mathsf{MQ}(q, \ell, m)$, an evaluation key $\mathfrak{H}$, and a trapdoor $\mathfrak{H}^{-1}$;
- $\mathsf{Eval}(\mathfrak{H}, x)$ evaluates $\mathcal{H}(x)$ for $x \in \mathbb{F}_q^\ell$;
- $\mathsf{Invert}(\mathfrak{H}^{-1}, y)$ returns some x such that $\mathcal{H}(x) = y$.

To simplify notation, we may write $\mathsf{Eval}(\mathfrak{H}, x) = \mathcal{H}(x)$, $\mathsf{Invert}(\mathfrak{H}^{-1}, y) = \mathcal{H}^{-1}(y)$ and $\mathsf{H} = (\mathcal{H}, \mathcal{H}^{-1})$. The family H is one-way if, given $\mathfrak{H}$ and $y = \mathcal{H}(x)$ for random x, no efficient adversary can recover any x' such that $\mathcal{H}(x') = y$. The formal definitions are provided in Appendix A.

This primitive gives rise naturally to *hash-and-sign* signatures (in the random oracle model), and this was used in [PSM17] to design their multivariate blind signature scheme following the general construction due to Fischlin [Fis06]. To sign a message μ, the user derives a commitment $t := \mathrm{Hash}(\mu) - \mathcal{F}(r)$ to μ, where Hash is a cryptographic hash function, $\mathcal{F}$ is a random quadratic function, and r is the randomness. The signer then computes s such that $\mathcal{H}(s) = t$, where $\mathcal{H}$ is a trapdoor quadratic function. The signature is a zero-knowledge proof of knowledge of (r, s) that satisfies $\mathcal{F}(r) + \mathcal{H}(s) = \mathrm{Hash}(\mu)$. Petzoldt *et al.* [PSM17] used this commitment scheme (which was broken in [Beu24]) and Rainbow[DS05] as the underlying trapdoor quadratic function (which was also broken in [Beu22a]).

Bouillaguet, Feneuil, Maire, Rivain, Sauvage and Vergnaud [BFM+25] recently introduced another hash-and-sign blind signature scheme based on the same family of trapdoor functions. Interestingly, it could be made more efficient by leveraging our MQ-commitments. More precisely, it relies on the *commit-append-and-prove system*. They first derive a commitment c_r to the randomness r, which is input to the hash function with the message μ. This defines a value $y = \mathcal{F}(r) + \mathrm{Hash}(\mu, c_r)$ sent by the User. After receiving $s = \mathcal{H}^{-1}(y)$ from the Signer, the User produces a proof π of knowledge of (r, s) for the statement that $\mathcal{H}(s) = \mathcal{F}(r) + \mathrm{Hash}(\mu, c_r)$ and that c_r is a commitment to r. The signature is the pair (π, c_r). This proof size could be reduced if the second part of the proof—namely, proving the commitment c_r—were improved. To this end, we replace the value y using our commitment of μ with randomness r as $y = (\mu + \mathcal{F}(r), \mathcal{G}(r))$. The message to be signed is then concealed by the hiding property of the commitment scheme and the User only has to generate

a proof that $\mathcal{H}(s) = (\mu + \mathcal{F}(r), \mathcal{G}(r))$ which becomes more efficient compared to [BFM$^+$25]. Producing a forgery for this signature scheme (under a key-only attack) consists essentially in finding a claw for the underlying quadratic maps.

Moreover, due to the use of the hash function Hash modeled as a random oracle in the schemes from [PSM17] and [BFM$^+$25], they are not suitable to achieve the broader notion of *signatures with efficient protocols* (SEP) [CL03]. These schemes support efficient zero-knowledge proofs, allowing a prover to demonstrate possession of a valid signature without revealing the signature or signed message. It allows proving statements about it, including proofs that the signed message satisfies a given public circuit (e.g., ensuring that the message contains specific attributes). Our blind signature scheme gives rise to an efficient multivariate blind SEP with improved efficiency and believable security.

7.1 Description of our signatures with efficient protocols

As mentioned above, to build a SEP, we replace the hash function in the previous constructions with two additional random quadratic functions $\mathcal{F}$ and $\mathcal{G}$, which serve to commit to the message using our MQ commitment. A signature on a message μ is then a pair $(s, r) \in \mathcal{F}_q^\ell \times \mathbb{F}_q^n$ such that $\mathcal{H}(s) = (\mu + \mathcal{F}(r), \mathcal{G}(r))$.

Let us define H a family of trapdoor quadratic functions over $\mathsf{MQ}(q, \ell, m)$ and random quadratic functions over $\mathsf{MQ}(q, n, k)$ and $\mathsf{MQ}(q, n, m-k)$ for parameters ℓ, m, n, k. The set $\mathbb{F}_q^k$ is the message space and $\mathbb{F}_q^n$ the randomness space.

Our construction employs non-interactive zero-knowledge argument systems (NIZK, see Section C) that are straight-line extractable. We denote as $\mathsf{NIZK.P}$ and $\mathsf{NIZK.V}$ the generation and verification algorithms of such a proof. Our construction is described in Protocols 3 and 4.

$\underline{\mathsf{KeyGen}(1^\kappa)}$

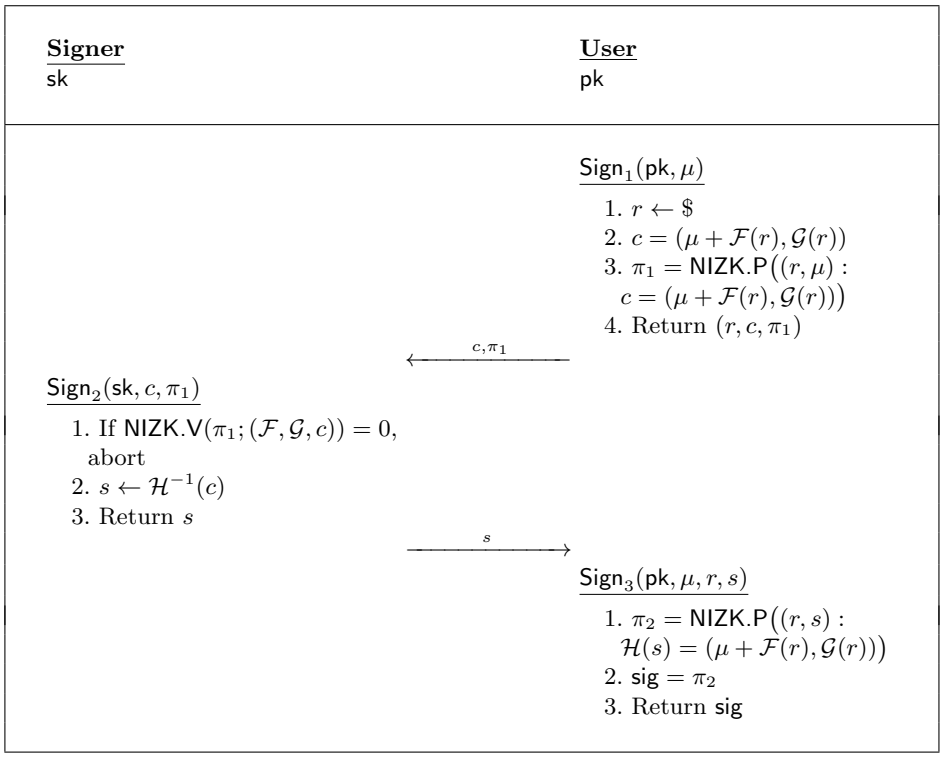
1. Sample $(\mathcal{H}, \mathcal{H}^{-1}) \in \mathsf{H}$
2. Sample $\mathcal{F} \in \mathsf{MQ}(q, n, k)$, $\mathcal{G} \in \mathsf{MQ}(q, n, m-k)$
3. $\mathsf{pk} = (\mathcal{F}, \mathcal{G}, \mathcal{H})$, $\mathsf{sk} = (\mathcal{F}, \mathcal{H}, \mathcal{H}^{-1}, \mathcal{G})$

$\underline{\mathsf{Verif}(\mathsf{sig}, \mathsf{pk})}$

1. Parse $\pi_2 = \mathsf{sig}$, $(\mathcal{F}, \mathcal{G}, \mathcal{H}) = \mathsf{pk}$
2. If $\mathsf{NIZK.V}(\pi_2; (\mathcal{F}, \mathcal{G}, \mathcal{H})) = 0$, return 0
3. Return 1

Protocol 3: Key generation and verification algorithms.

From our vanilla blind signature to a SEP. Turning our protocol 4 into a SEP relies on the favorable properties of our blind signature, which supports efficient zero-knowledge proofs. The standard requirement for a SEP is to prove



Protocol 4: Blind signature protocol.

the message-signature possession. For this purpose, the prover can compute $\mathsf{NIZK.P}\big((\mu, r, s) : \mathcal{H}(s) = (\mu + \mathcal{F}(r), \mathcal{G}(r))\big)$. The prover can additionally prove statements about the signed message, e.g. μ satisfies a public circuit C. Using R1CS to deal with C, this proof follows the one of our MQ-based commitment for circuit satisfiability: $\mathsf{NIZK.P}\big((\mu, r, s) : \mathcal{H}(s) = (\mu + \mathcal{F}(r), \mathcal{G}(r)) \wedge$ Eq. 4 is satisfied$\big)$.

7.2 Security

The security analysis provided by Petzoldt *et al.* in [PSM17] is (at best) heuristic, and it is unclear if their scheme can be proven secure (in the random oracle model) even if we replace internal components by a secure commitment scheme and a secure signature scheme. The flaw in their unforgeability security claim appears in [PSM17, Proof of Theorem 3, Game 4], where the authors claim that *"the ability to query inverses $\langle \cdots \rangle$ does not help the adversary at all $\langle \ldots$ and $\ldots \rangle$ should not help the adversary to either recover the secret key or find his own inverses. Otherwise, it would be possible to mount an attack exploiting this fact"*. This is not true since the signer in the blind signature uses their trapdoor on arbitrary inputs, whereas in the hash-and-sign scheme, it is only applied to hash values (programmed in the random oracle model). This situation is frequent in the design of round-optimal blind signatures (in particular, if one does not use zero-knowledge proofs of "well-formedness"), and it is often necessary (see, e.g. [BMV08]) to consider interactive assumptions of a *one-more-inversion* flavor as defined by Bellare, Namprepre, Pointcheval, and Semanko [BNPS03] to analyze the security of Chaum's original RSA-based blind signature scheme [Cha82].

We therefore consider a new variant of one-wayness for a trapdoor quadratic function family H. It is called *one-more claw-free* if, given oracle access to $\mathsf{Invert}(\mathfrak{H}^{-1}, \cdot)$ (at most q times) and a random quadratic function $\mathcal{J}$, no efficient adversary can find more than q distinct pairs (x_i, z_i) such that $\mathcal{H}(x_i) = \mathcal{J}(z_i)$ (see Appendix A for a formal definition). This assumption appears reasonable in light of the current hardness of MQ inversion problems, but it requires further cryptanalytic scrutiny. This is the reason why we present our scheme as a *candidate* for blind signatures in the multivariate setting.

We demonstrate that the unforgeability of the blind signature scheme reduces to this assumption (and the straight-line extractable soundness of the proof systems). Furthermore, the blindness of the scheme relies on the pseudorandomness of the MQ function, the blindness of the MQ-commitment, and the zero-knowledge property of the proof systems.

Theorem 5. *The blind signature scheme in 3 and 4 is ℓ-one-more unforgeable based on the hardness of the $\ell+1$-claw-free problem and the knowledge soundness of the NIZK. More precisely, for every adversary $\mathcal{A}$ with advantage $\mathsf{Adv}_{\mathsf{OMUG}}(\mathcal{A})$ against the ℓ-one-more unforgeability game, there exist an adversary $\mathcal{B}$ with advantage $\mathsf{Adv}_{\mathsf{KS\text{-}NIZK}}(\mathcal{B})$ against the knowledge soundness of the NIZK and an*

adversary $\mathcal{C}$ with advantage $\mathsf{Adv}_{\mathsf{OM\text{-}CF}}(\mathcal{C})$ against the one-more claw-free game, such that

$$\mathsf{Adv}_{\mathsf{OMUG}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{KS\text{-}NIZK}}(\mathcal{B}) + \mathsf{Adv}_{\mathsf{OM\text{-}CF}}(\mathcal{C}).$$

Theorem 6. *The blind signature scheme in 3 and 4 is anonymous based on the zero-knowledge property of the NIZK, the blindness property of the MQ-based commitment, and the security of the MQ-based PRG. More precisely, for every adversary $\mathcal{A}$ with advantage $\mathsf{Adv}_{\mathsf{ANON}}(\mathcal{A})$ against the anonymity game, there exist an adversary $\mathcal{B}$ that can distinguish simulated NIZK proofs from real ones with advantage $\mathsf{Adv}_{\mathsf{ZK\text{-}NIZK}}(\mathcal{B})$, an adversary $\mathcal{C}$ that can break the blindness simulation game with advantage $\mathsf{Adv}_{\mathsf{MQ\text{-}Com}}(\mathcal{C})$, and an adversary $\mathcal{D}$ that can distinguish MQ-based PRG outputs from true randomness with advantage $\mathsf{Adv}_{\mathsf{PRG}}(\mathcal{D})$ such that*

$$\mathsf{Adv}_{\mathsf{ANON}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{ZK\text{-}NISK}}(\mathcal{B}) + \mathsf{Adv}_{\mathsf{MQ\text{-}Com}}(\mathcal{C}) + \mathsf{Adv}_{\mathsf{MQ\text{-}PRG}}(\mathcal{D}) \ .$$

7.3 Instances

We instantiate π_2 from our blind signature with the TCitH framework (described in Appendix C.2), which is straight-line extractable in the random oracle model. To apply this construction, we need to express the statement we want to prove as polynomial constraints. The binary relation is

$$\mathcal{R} = \{(\mathcal{F}, \mathcal{G}, \mathcal{H}), (\mu, r, s) \mid \mathcal{H}(s) = (\mu + \mathcal{F}(r), \mathcal{G}(r))\} \ ,$$

where $\mathcal{F} \in \mathbb{F}_q^{k\times n}, \mathcal{H} \in \mathbb{F}_q^{(k+m)\times \ell}, \mathcal{G} \in \mathbb{F}_q^{m\times n}$. Let us assume that the message μ owns attributes and not all of them need to stay hidden, and call $\delta \in [0,1]$ the ratio of the message that is not revealed (when $\delta = 1$ one gets a signature-message possession proof). $\mathcal{R}$ directly describes a system of quadratic constraints, so we can apply the proof system given in Appendix C with a witness of length $|w| = n + \ell + \delta k$. In addition, one may want to prove that the signed message satisfies a public circuit C. Then, dealing with C by using R1CS (see Eq. 4, Appendix C), the proof size of our SEP is (in bits):

$$\textsc{Size} = 4\lambda + \lambda t_{open} + \tau\left((n + \ell + \delta k + |C|)\log_2 q + \log_2 q^{\rho e} + 2\lambda\right),$$

where parameters are chosen following Appendix C.2. In particular, (ρ, e) are entirely determined once (κ, N, q) have been fixed.

Since π_2 is built with the TCitH framework and is therefore computationally secure, we choose MQ-based commitment parameters achieving computational security. Moreover, we rely on our family of trapdoor quadratic functions H on the UOV scheme [KPG99]. Other parameters have been introduced and detailed in Appendix C.2 (the overall proof-of-work of $\log_2 \theta + \omega$ is fixed to around 11 bits).

As a use-case, we follow [AGJ+24] and considers a message of 10 attributes of 128 bits each, where 4 of them are revealed (i.e. $\delta = 0.6$). In [JRS23], a

	Commitment parameters			NIZK parameters						Proof size
κ	n	l	k	δ	τ	N	L	ω	t_{open}	(in KB)
128	118	733	160	0.6	12	2^{11}	1	9	118	13.57
128	73	300	16	0	12	2^{11}	1	9	118	6.84

Table 5. Our SEP parameters and performances with $q = 257$

credential size of 724 KB is obtained for the common metric of message-signature possession (for a message of 128 bits). This size was significantly improved to 80 KB in [AGJ+24]. Table 5 gives our parameters to achieve the performances given in the introduction (with a blind signature size of 6.84 KB and a credential size of 13.6 KB for this use case).

Acknowledgements. The authors would like to thank Olivier Sanders and Olivier Blazy for helpful discussions about potential applications of SEP, as well as Florette Martinez for spotting an error into an earlier proof of the NP-hardness of finding collisions in quadratic maps. The authors also thank the anonymous reviewers of Asiacrypt'25 who provided very interesting feedback. This work is supported in part by the French ANR SANGRIA and GORILLA project (ANR-21-CE39-0006 and ANR-20-CE39-0002).

References

AGJ+24. Sven Argo, Tim Güneysu, Corentin Jeudy, Georg Land, Adeline Roux-Langlois, and Olivier Sanders. Practical post-quantum signatures for privacy. In Bo Luo, Xiaojing Liao, Jun Xu, Engin Kirda, and David Lie, editors, *ACM CCS 2024*, pages 1523–1537. ACM Press, October 2024.

AHI+17. Benny Applebaum, Naama Haramaty, Yuval Ishai, Eyal Kushilevitz, and Vinod Vaikuntanathan. Low-complexity cryptographic hash functions. In Christos H. Papadimitriou, editor, *ITCS 2017*, volume 4266, pages 7:1–7:31, 67, January 2017. LIPIcs.

ALS20. Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. Practical product proofs for lattice commitments. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 470–499. Springer, Cham, August 2020.

Bar04. Magali Bardet. *Étude des systèmes algébriques surdéterminés. Applications aux codes correcteurs et à la cryptographie.* PhD thesis, Pierre and Marie Curie University, Paris, France, 2004.

BB08. Dan Boneh and Xavier Boyen. Short signatures without random oracles and the SDH assumption in bilinear groups. *Journal of Cryptology*, 21(2):149–177, April 2008.

BBFR25. Ryad Benadjila, Charles Bouillaguet, Thibauld Feneuil, and Matthieu Rivain. Mqom: Mq on my mind. Algorithm Specifications and Supporting Documentation (Version 2.1), 2025.

BBM+24. Carsten Baum, Ward Beullens, Shibam Mukherjee, Emmanuela Orsini, Sebastian Ramacher, Christian Rechberger, Lawrence Roy, and Peter Scholl.

One tree to rule them all: Optimizing GGM trees and owfs for post-quantum signatures. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part I*, volume 15484 of *Lecture Notes in Computer Science*, pages 463–493. Springer, 2024.

BCC+23. Gustavo Banegas, Kévin Carrier, André Chailloux, Alain Couvreur, Thomas Debris-Alazard, Philippe Gaborit, Pierre Karpman, Johanna Loyer, Ruben Niederhagen, Nicolas Sendrier, Benjamin Smith, and Jean-Pierre Tillich. Wave. Technical report, National Institute of Standards and Technology, 2023. available at `https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures`.

BCV23. John B. Baena, Daniel Cabarcas, and Javier A. Verbel. On the complexity of solving generic overdetermined bilinear systems. *Adv. Math. Commun.*, 17(5):1210–1234, 2023.

BDL+18. Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In Dario Catalano and Roberto De Prisco, editors, *SCN 18*, volume 11035 of *LNCS*, pages 368–385. Springer, Cham, September 2018.

Beu22a. Ward Beullens. Breaking rainbow takes a weekend on a laptop. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 464–479. Springer, Cham, August 2022.

Beu22b. Ward Beullens. MAYO: Practical post-quantum signatures from oil-and-vinegar maps. In Riham AlTawy and Andreas Hülsing, editors, *SAC 2021*, volume 13203 of *LNCS*, pages 355–376. Springer, Cham, September / October 2022.

Beu24. Ward Beullens. Multivariate blind signatures revisited. In Maria Eichlseder and Sébastien Gambs, editors, *Selected Areas in Cryptography - SAC 2024 - 31st International Conference, Montreal, QC, Canada, August 28-30, 2024, Revised Selected Papers, Part I*, volume 15516 of *Lecture Notes in Computer Science*, pages 230–236. Springer, 2024.

BFM+25. Charles Bouillaguet, Thibauld Feneuil, Jules Maire, Matthieu Rivain, Julia Sauvage, and Damien Vergnaud. Blinding post-quantum hash-and-sign signatures. *IACR Cryptol. ePrint Arch.*, (895), 2025. (to appear in IEEE Security & Privacy 2026).

BFP09. Luk Bettale, Jean-Charles Faugère, and Ludovic Perret. Hybrid approach for solving multivariate systems over finite fields. *J. Math. Cryptol.*, 3(3):177–197, 2009.

BFP12. Luk Bettale, Jean-Charles Faugère, and Ludovic Perret. Solving polynomial systems over finite fields: improved analysis of the hybrid approach. In Joris van der Hoeven and Mark van Hoeij, editors, *International Symposium on Symbolic and Algebraic Computation, ISSAC'12, Grenoble, France - July 22 - 25, 2012*, pages 67–74. ACM, 2012.

BFR24. Ryad Benadjila, Thibauld Feneuil, and Matthieu Rivain. MQ on my mind: Post-quantum signatures from the non-structured multivariate quadratic problem. In *2024 IEEE European Symposium on Security and Privacy*, pages 468–485. IEEE Computer Society Press, July 2024.

BFS15. Magali Bardet, Jean-Charles Faugère, and Bruno Salvy. On the complexity of the F5 gröbner basis algorithm. *J. Symb. Comput.*, 70:49–70, 2015.

BFT19. Matías R. Bender, Jean-Charles Faugère, and Elias P. Tsigaridas. Gröbner basis over semigroup algebras: Algorithms and applications for sparse polynomial systems. In James H. Davenport, Dongming Wang, Manuel Kauers, and Russell J. Bradford, editors, *Proceedings of the 2019 on International Symposium on Symbolic and Algebraic Computation, ISSAC 2019, Beijing, China, July 15-18, 2019*, pages 42–49. ACM, 2019.

BGG+20. Fabrice Boudot, Pierrick Gaudry, Aurore Guillevic, Nadia Heninger, Emmanuel Thomé, and Paul Zimmermann. Comparing the difficulty of factorization and discrete logarithm: A 240-digit experiment. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 62–91. Springer, Cham, August 2020.

BGP06. Côme Berbain, Henri Gilbert, and Jacques Patarin. QUAD: A practical stream cipher with provable security. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 109–128. Springer, Berlin, Heidelberg, May / June 2006.

BKLP15. Fabrice Benhamouda, Stephan Krenn, Vadim Lyubashevsky, and Krzysztof Pietrzak. Efficient zero-knowledge proofs for commitments from learning with errors over rings. In Günther Pernul, Peter Y. A. Ryan, and Edgar R. Weippl, editors, *ESORICS 2015, Part I*, volume 9326 of *LNCS*, pages 305–325. Springer, Cham, September 2015.

BLNS23a. Ward Beullens, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Lattice-based blind signatures: Short, efficient, and round-optimal. In Weizhi Meng, Christian Damsgaard Jensen, Cas Cremers, and Engin Kirda, editors, *ACM CCS 2023*, pages 16–29. ACM Press, November 2023.

BLNS23b. Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Alessandro Sorniotti. A framework for practical anonymous credentials from lattices. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 384–417. Springer, Cham, August 2023.

Blu82. Manuel Blum. Coin flipping by telephone - A protocol for solving impossible problems. In *COMPCON'82, Digest of Papers, Twenty-Fourth IEEE Computer Society International Conference, San Francisco, California, USA, February 22-25, 1982*, pages 133–137. IEEE Computer Society, 1982.

BMV08. Emmanuel Bresson, Jean Monnerat, and Damien Vergnaud. Separation results on the "one-more" computational problems. In Tal Malkin, editor, *CT-RSA 2008*, volume 4964 of *LNCS*, pages 71–87. Springer, Berlin, Heidelberg, April 2008.

BNPS03. Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The one-more-RSA-inversion problems and the security of Chaum's blind signature scheme. *Journal of Cryptology*, 16(3):185–215, June 2003.

BRP07. Olivier Billet, Matthew J. B. Robshaw, and Thomas Peyrin. On building hash functions from multivariate quadratic equations. In Josef Pieprzyk, Hossein Ghodosi, and Ed Dawson, editors, *ACISP 07*, volume 4586 of *LNCS*, pages 82–95. Springer, Berlin, Heidelberg, July 2007.

BS24. Charles Bouillaguet and Julia Sauvage. Preliminary cryptanalysis of the biscuit signature scheme. *CiC*, 1(1):30, 2024.

CCNY12a. Chen-Mou Cheng, Tung Chou, Ruben Niederhagen, and Bo-Yin Yang. Solving quadratic equations with XL on parallel architectures. In Emmanuel Prouff and Patrick Schaumont, editors, *Cryptographic Hardware*

and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings, volume 7428 of *Lecture Notes in Computer Science*, pages 356–373. Springer, 2012.

CCNY12b. Chen-Mou Cheng, Tung Chou, Ruben Niederhagen, and Bo-Yin Yang. Solving quadratic equations with XL on parallel architectures. In Emmanuel Prouff and Patrick Schaumont, editors, *CHES 2012*, volume 7428 of *LNCS*, pages 356–373. Springer, Berlin, Heidelberg, September 2012.

Cha82. David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *CRYPTO'82*, pages 199–203. Plenum Press, New York, USA, 1982.

CKPS00. Nicolas Courtois, Alexander Klimov, Jacques Patarin, and Adi Shamir. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 392–407. Springer, Berlin, Heidelberg, May 2000.

CL03. Jan Camenisch and Anna Lysyanskaya. A signature scheme with efficient protocols. In Stelvio Cimato, Clemente Galdi, and Giuseppe Persiano, editors, *SCN 02*, volume 2576 of *LNCS*, pages 268–289. Springer, Berlin, Heidelberg, September 2003.

CL04. Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 56–72. Springer, Berlin, Heidelberg, August 2004.

DFS07. Vivien Dubois, Pierre-Alain Fouque, and Jacques Stern. Cryptanalysis of SFLASH with slightly modified parameters. In Moni Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 264–275. Springer, Berlin, Heidelberg, May 2007.

DGK+10. Yevgeniy Dodis, Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Public-key encryption schemes with auxiliary inputs. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 361–381. Springer, Berlin, Heidelberg, February 2010.

DS05. Jintai Ding and Dieter Schmidt. Rainbow, a new multivariable polynomial signature scheme. In John Ioannidis, Angelos Keromytis, and Moti Yung, editors, *ACNS 2005*, volume 3531 of *LNCS*, pages 164–175. Springer, Berlin, Heidelberg, June 2005.

DY07. Jintai Ding and Bo-Yin Yang. Multivariates polynomials for hashing. In Dingyi Pei, Moti Yung, Dongdai Lin, and Chuankun Wu, editors, *Information Security and Cryptology, Third SKLOIS Conference, Inscrypt 2007, Xining, China, August 31 - September 5, 2007, Revised Selected Papers*, volume 4990 of *Lecture Notes in Computer Science*, pages 358–371. Springer, 2007.

FDS11. Jean-Charles Faugère, Mohab Safey El Din, and Pierre-Jean Spaenlehauer. Gröbner bases of bihomogeneous ideals generated by polynomials of bidegree (1, 1): Algorithms and complexity. *J. Symb. Comput.*, 46(4):406–437, 2011.

FIKT21. Hiroki Furue, Yasuhiko Ikematsu, Yutaro Kiyomura, and Tsuyoshi Takagi. A new variant of unbalanced Oil and Vinegar using quotient ring: QR-UOV. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 187–217. Springer, Cham, December 2021.

Fis06. Marc Fischlin. Round-optimal composable blind signatures in the common reference string model. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 60–77. Springer, Berlin, Heidelberg, August 2006.

FR25. Thibauld Feneuil and Matthieu Rivain. Threshold computation in the head: Improved framework for post-quantum signatures and zero-knowledge arguments. *J. Cryptol.*, 38(3):28, 2025.

FS86. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology - CRYPTO '86, Santa Barbara, California, USA, 1986, Proceedings*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, 1986.

FSS14. Jean-Charles Faugère, Pierre-Jean Spaenlehauer, and Jules Svartz. Sparse gröbner bases: the unmixed case. In Katsusuke Nabeshima, Kosaku Nagasaka, Franz Winkler, and Ágnes Szántó, editors, *International Symposium on Symbolic and Algebraic Computation, ISSAC '14, Kobe, Japan, July 23-25, 2014*, pages 178–185. ACM, 2014.

GJ79. M. R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

GJS06. Louis Granboulan, Antoine Joux, and Jacques Stern. Inverting HFE is quasipolynomial. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 345–356. Springer, Berlin, Heidelberg, August 2006.

GK96. Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *Journal of Cryptology*, 9(3):167–190, June 1996.

GKP13. Bruno Grenet, Pascal Koiran, and Natacha Portier. On the complexity of the multivariate resultant. *J. Complexity*, 29(2):142–157, 2013.

GMR89. Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.

Gol01. Oded Goldreich. *The Foundations of Cryptography - Volume 1: Basic Techniques*. Cambridge University Press, 2001.

GRS00. Oded Goldreich, Ronitt Rubinfeld, and Madhu Sudan. Learning polynomials with queries: The highly noisy case. *SIAM J. Discret. Math.*, 13(4):535–570, 2000.

IKOS09. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge proofs from secure multiparty computation. *SIAM J. Comput.*, 39(3):1121–1152, 2009.

JKPT12. Abhishek Jain, Stephan Krenn, Krzysztof Pietrzak, and Aris Tentes. Commitments and efficient zero-knowledge proofs from learning parity with noise. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 663–680. Springer, Berlin, Heidelberg, December 2012.

JRS23. Corentin Jeudy, Adeline Roux-Langlois, and Olivier Sanders. Lattice signature with efficient protocols, application to anonymous credentials. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 351–383. Springer, Cham, August 2023.

Koi00. Pascal Koiran. The complexity of local dimensions for constructible sets. *J. Complex.*, 16(1):311–323, 2000.

KPG99. Aviad Kipnis, Jacques Patarin, and Louis Goubin. Unbalanced Oil and Vinegar signature schemes. In Jacques Stern, editor, *EUROCRYPT'99*,

volume 1592 of *LNCS*, pages 206–222. Springer, Berlin, Heidelberg, May 1999.

KS98. Aviad Kipnis and Adi Shamir. Cryptanalysis of the Oil & Vinegar signature scheme. In Hugo Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 257–266. Springer, Berlin, Heidelberg, August 1998.

Laz83. Daniel Lazard. Gröbner-bases, gaussian elimination and resolution of systems of algebraic equations. In J. A. van Hulzen, editor, *EUROCAL*, volume 162 of *Lecture Notes in Computer Science*, pages 146–156. Springer, 1983.

LLM+16. Benoît Libert, San Ling, Fabrice Mouhartem, Khoa Nguyen, and Huaxiong Wang. Signature schemes with efficient protocols and dynamic group signatures from lattice assumptions. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 373–403. Springer, Berlin, Heidelberg, December 2016.

MV23. Jules Maire and Damien Vergnaud. Efficient zero-knowledge arguments and digital signatures via sharing conversion in the head. In Gene Tsudik, Mauro Conti, Kaitai Liang, and Georgios Smaragdakis, editors, *ESORICS 2023, Part I*, volume 14344 of *LNCS*, pages 435–454. Springer, Cham, September 2023.

PS16. David Pointcheval and Olivier Sanders. Short randomizable signatures. In Kazue Sako, editor, *CT-RSA 2016*, volume 9610 of *LNCS*, pages 111–126. Springer, Cham, February / March 2016.

PSM17. Albrecht Petzoldt, Alan Szepieniec, and Mohamed Saied Emam Mohamed. A practical multivariate blind signature scheme. In Aggelos Kiayias, editor, *FC 2017*, volume 10322 of *LNCS*, pages 437–454. Springer, Cham, April 2017.

PST20. Ray Perlner and Daniel Smith-Tone. Rainbow band separation is better than we thought. Cryptology ePrint Archive, Paper 2020/702, 2020.

Spa12. Pierre-Jean Spaenlehauer. *Résolution de systèmes multi-homogènes et déterminantiels algorithmes - complexité - applications*. PhD thesis, 2012. Thèse de doctorat dirigée par Faugère, Jean-Charles Informatique Paris 6 2012.

TW12. Enrico Thomae and Christopher Wolf. Solving underdetermined systems of multivariate quadratic equations revisited. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 156–171. Springer, Berlin, Heidelberg, May 2012.

YC04. Bo-Yin Yang and Jiun-Ming Chen. Theoretical analysis of XL over small fields. In Huaxiong Wang, Josef Pieprzyk, and Vijay Varadharajan, editors, *ACISP 04*, volume 3108 of *LNCS*, pages 277–288. Springer, Berlin, Heidelberg, July 2004.

Supplementary material

A Security definitions

A.1 Commitment schemes

Definition 2 (Commitment Scheme). *A commitment scheme is a triplet of algorithms* $(\mathsf{Setup}, \mathsf{Com}, \mathsf{Verif})$ *such that*

Setup *A PPT algorithm that, on input* 1^κ*, outputs some public parameters* $\mathsf{pp} \in \{0,1\}^{\mathsf{poly}(\kappa)}$ *containing a definition of the message space, the randomness space, and the commitment space.*

Com *A deterministic polynomial-time algorithm that, on input the public parameters* pp*, a message* μ *and randomness* aux*, outputs a commitment* c*.*

Verif *A deterministic polynomial-time algorithm that, on input the public parameters* pp*, a message* μ*, a commitment* c *and randomness* aux*, outputs a bit* $b \in \{0,1\}$*.*

Definition 3 (Correctness Property). *A commitment scheme achieves correctness if for any public parameters* pp*, any message* μ *and any randomness* aux*:*

$$\Pr[\mathsf{Verif}(\mathsf{pp}, \mu, c, \mathsf{aux}) = 1 \mid c \leftarrow \mathsf{Com}(\mathsf{pp}, \mu, \mathsf{aux})] = 1 \ .$$

Definition 4 (Security Properties). *Let* $t : \mathbb{N} \to \mathbb{N}$ *and* $\epsilon : \mathbb{N} \to [0,1]$*. A commitment scheme* $(\mathsf{Setup}, \mathsf{Com}, \mathsf{Verif})$ *is said:*

- (t,ϵ)-computationally hiding *if for all two-phase algorithm* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, we have for all* $\kappa \in \mathbb{N}$*:*

$$\Pr\left[b = b' \;\middle|\; \begin{array}{l} \mathsf{pp} \xleftarrow{\$} \mathsf{Setup}(1^\kappa), (\mu_0, \mu_1, \mathsf{aux}, s) \xleftarrow{\$} \mathcal{A}_1(\mathsf{pp}), \\ b \xleftarrow{\$} \{0,1\}, c \xleftarrow{\$} \mathsf{Com}(\mathsf{pp}, \mu_b, \mathsf{aux}), b' \xleftarrow{\$} \mathcal{A}_2(c, s) \end{array}\right] \leq \frac{1}{2} + \epsilon(\kappa)$$

  when $\mathcal{A}$ *runs in time at most* $t(\kappa)$ *in this probabilistic computational game.*

- (t,ϵ)-computationally binding *if for all algorithm* $\mathcal{A}$*, we have for all* $\kappa \in \mathbb{N}$*:*

$$\Pr\left[\begin{array}{c} \mu_1 \neq \mu_2, \\ \mathsf{Verif}(\mathsf{pp}, \mu_1, c, \mathsf{aux}_1) = 1, \\ \mathsf{Verif}(\mathsf{pp}, \mu_2, c, \mathsf{aux}_2) = 1 \end{array} \;\middle|\; \begin{array}{l} \mathsf{pp} \xleftarrow{\$} \mathsf{Setup}(1^\kappa), \\ (\mu_1, \mu_2, \mathsf{aux}_1, \mathsf{aux}_2, c) \xleftarrow{\$} \mathcal{A}(1^\kappa, \mathsf{pp}) \end{array}\right] \leq \epsilon(\kappa)$$

  when $\mathcal{A}$ *runs in time at most* $t(\kappa)$ *in this probabilistic computational game.*

These two security notions can be strengthened to *statistical security* – where the adversary is unbounded (i.e., t is not restricted) and the advantage ϵ is negligible – and further to *perfect security* – where the adversary is unbounded and $\epsilon = 0$.

A.2 Blind signatures

There exist several (essentially equivalent) syntactic definitions for (round-optimal) blind signatures. In the present paper, we follow the definition from [BLNS23a].

Definition 5 (Blind Signatures). *A round-optimal blind signature scheme is a tuple of five (probabilistic) PPT algorithms* $(\mathsf{KeyGen}, \mathsf{Sign}_1, \mathsf{Sign}_2, \mathsf{Sign}_3, \mathsf{Verif})$*:*

- $\mathsf{KeyGen}(1^\kappa) \to (\mathsf{sk}, \mathsf{pk})$*: this probabilistic algorithm takes the security parameter as input and outputs a signing secret key* sk *and a verifying public key* pk*;*
- $\mathsf{Sign}_1(\mathsf{pk}, \mu) \to (S_{\mathcal{U}}, \rho_1)$*: this probabilistic algorithm is run by a user* $\mathcal{U}$*; it takes as input a public key* pk *belonging to a signer* $\mathcal{S}$ *and a message* $\mu \in \{0,1\}^*$ *and outputs a state* $S_{\mathcal{U}}$ *and a first message* ρ_1*, which they send to* $\mathcal{S}$*;*
- $\mathsf{Sign}_2(\mathsf{sk}, \rho_1) \to \rho_2$*: this probabilistic algorithm is run by a signer* $\mathcal{S}$ *and takes as input a secret key* sk *and a first message* ρ_1 *and outputs a second message* ρ_2*, which they send to* $\mathcal{U}$ *;*
- $\mathsf{Sign}_3(S_{\mathcal{U}}, \rho_2) \to \mathsf{sig}$*: this probabilistic algorithm is run by a user* $\mathcal{U}$*; it takes as input a second message* ρ_2 *and a state* $S_{\mathcal{U}}$ *and outputs a signature* sig *or an error message* $\perp$*;*
- $\mathsf{Verif}(\mathsf{sig}, \mathsf{pk}, \mu) \to \{0,1\}$*: this deterministic algorithm takes as input a public key* pk*, a message* μ *and a (putative) signature* sig *and outputs a bit to indicate whether the signature is deemed valid (1) or invalid (0).*

A round-optimal blind signature scheme must satisfy the three following properties:

1. **correctness**, ensuring that signatures generated honestly will verify with overwhelming probability;
2. **blindness**, ensuring that the signer cannot link signatures to the interactions in which they were created;
3. and **one-more unforgeability**, ensuring that if a user is runs the signing protocol q times, they cannot obtain more than q valid signatures.

Definition 6 (Correctness). *A round-optimal blind signature scheme* $(\mathsf{KeyGen}, \mathsf{Sign}_1, \mathsf{Sign}_2, \mathsf{Sign}_3, \mathsf{Verif})$ *is deemed* correct *if, given there exists a negligible function* negl *such that, we have:*

$$\Pr\left[\mathsf{Verif}(\mathsf{sig}, \mathsf{pk}, \mu) = 0 \;\middle|\; \begin{array}{l} (\mathsf{sk}, \mathsf{pk}) \xleftarrow{\$} \mathsf{KeyGen}(1^\kappa) \\ (S_{\mathcal{U}}, \rho_1) \xleftarrow{\$} \mathsf{Sign}_1(\mathsf{pk}, \mu) \\ \rho_2 \xleftarrow{\$} \mathsf{Sign}_2(\mathsf{sk}, \rho_1) \\ \mathsf{sig} \xleftarrow{\$} \mathsf{Sign}_3(S_{\mathcal{U}}, \rho_2) \end{array}\right] \leq \mathsf{negl}(\kappa) \ .$$

Definition 7 (Blindness). *A round-optimal blind signature scheme* $(\mathsf{KeyGen}, \mathsf{Sign}_1, \mathsf{Sign}_2, \mathsf{Sign}_3, \mathsf{Verif})$ *is deemed* blind *if for every three-part stateful PPT*

adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$*, there exists a negligible function* negl *such that, for any two messages* $\mu_0, \mu_1 \in \{0,1\}^*$*, we have:*

$$\mathsf{Adv}_{\mathsf{BLIND}}(\mathcal{A}) := \left| \Pr\left[\mathcal{A}_3(\mathsf{sig}_0, \mathsf{sig}_1) = b \;\middle|\; \begin{array}{l} \mathsf{pk} \xleftarrow{\$} \mathcal{A}_1(1^\kappa) \\ b \xleftarrow{\$} \{0,1\} (S^0_{\mathcal{U}}, \rho^0_1) \leftarrow \mathsf{Sign}_1(\mathsf{pk}, \mu_0) \\ (S^1_{\mathcal{U}}, \rho^1_1) \leftarrow \mathsf{Sign}_1(\mathsf{pk}, \mu_1) \\ (\rho^b_2, \rho^{1-b}_2) \leftarrow \mathcal{A}_2(\rho^b_1, \rho^{1-b}_1) \\ \mathsf{sig}^0 \leftarrow \mathsf{Sign}_3(S^0_{\mathcal{U}}, \rho^0_2) \\ \mathsf{sig}^1 \leftarrow \mathsf{Sign}_3(S^1_{\mathcal{U}}, \rho^1_2) \\ \mathit{if}\ \ \bot \in \{\mathsf{sig}_0, \mathsf{sig}_1\}\ \mathit{then}\ (\mathsf{sig}_0, \mathsf{sig}_1) \leftarrow (\bot, \bot) \end{array} \right] - \frac{1}{2} \right| \leq \mathsf{negl}(\kappa)\ .$$

Definition 8 (One-more unforgeability). *A round-optimal blind signature scheme* $(\mathsf{KeyGen}, \mathsf{Sign}_1, \mathsf{Sign}_2, \mathsf{Sign}_3, \mathsf{Verif})$ *is deemed* one-more unforgeable *if for every PPT adversary* $\mathcal{A}$ *that makes at most* q *queries to a signing oracle (where* q *is a function of* κ*), there exists a negligible function* negl *such that, for any two messages* $\mu_0, \mu_1 \in \{0,1\}^*$*, we have:*

$$\mathsf{Adv}_{\mathsf{BSIG}}(\mathcal{A}) := \left| \Pr\left[\begin{array}{l} \forall i \in \{1, \ldots, q+1\}, \mathsf{Verif}(\mathsf{sig}_i, \mathsf{pk}, \mu_i) = 1 \\ \forall i, j \in \{1, \ldots, q+1\}, i \neq j \Rightarrow \mu_i \neq \mu_j \end{array} \;\middle|\; \begin{array}{c} (\mathsf{sk}, \mathsf{pk}) \xleftarrow{\$} \mathsf{KeyGen}(1^\kappa) \\ \{(\mu_i, \mathsf{sig}_i)\}_{i \in \{1,\ldots,q+1\}} \\ \xleftarrow{\$} \mathcal{A}^{\mathsf{Sign}_2(\mathsf{sk}, \cdot)}(\mathsf{pk}) \end{array} \right] \right| \leq \mathsf{negl}(\kappa)\ .$$

A.3 Trapdoor quadratic functions

We formally define the notion of trapdoor quadratic function family, which is central to our blind signature construction.

Definition 9 (Trapdoor Quadratic Function Family). *A* trapdoor quadratic function family *is a tuple of PPT algorithms* $\mathsf{H} = (\mathsf{KeyGen}, \mathsf{Eval}, \mathsf{Invert})$ *defined as follows:*

- $\mathsf{KeyGen}(1^\kappa) \to (q, \ell, m, \mathcal{H}, \mathfrak{H}, \mathfrak{H}^{-1})$*: on input a security parameter* κ*, outputs a prime power* q*, integers* $\ell, m \in \mathbb{N}$*, a function* $\mathcal{H} \in \mathcal{MQ}(q, \ell, m)$*, an evaluation key* $\mathfrak{H}$ *and a trapdoor* $\mathfrak{H}^{-1}$*;*
- $\mathsf{Eval}(\mathfrak{H}, x) \to y$*: given the evaluation key* $\mathfrak{H}$ *and an input* $x \in \mathbb{F}_q^\ell$*, outputs* $y = \mathcal{H}(x) \in \mathbb{F}_q^m$*;*
- $\mathsf{Invert}(\mathfrak{H}^{-1}, y) \to x$*: given the trapdoor* $\mathfrak{H}^{-1}$ *and a point* $y \in \mathbb{F}_q^m$*, outputs some* $x \in \mathbb{F}_q^\ell$ *such that* $\mathcal{H}(x) = y$*.*

To simplify notation, we sometimes write $\mathsf{Eval}(\mathfrak{H}, x) = \mathcal{H}(x)$ and $\mathsf{Invert}(\mathfrak{H}^{-1}, y) = \mathcal{H}^{-1}(y)$, while keeping in mind that these are algorithms, not just mathematical functions.

Definition 10 (One-Wayness). *Let $t : \mathbb{N} \to \mathbb{N}$ and $\epsilon : \mathbb{N} \to [0,1]$. The function family $\mathsf{H} = (\mathsf{KeyGen}, \mathsf{Eval}, \mathsf{Invert})$ is (t, ϵ)-*one-way *if for all adversaries $\mathcal{A}$ running in time at most $t(\kappa)$, the probability $\mathsf{Adv}_{\mathsf{OW}}(\mathcal{A})$*

$$\Pr\left[\mathcal{H}(x') = \mathcal{H}(x) \;\middle|\; \begin{array}{l}(q, \ell, m, \mathcal{H}, \mathfrak{H}, \mathfrak{H}^{-1}) \xleftarrow{\$} \mathsf{KeyGen}(1^\kappa) \\ x \xleftarrow{\$} \mathbb{F}_q^\ell \\ y \leftarrow \mathsf{Eval}(\mathfrak{H}, x) \\ x' \xleftarrow{\$} \mathcal{A}(q, \ell, m, \mathfrak{H}, y)\end{array}\right] \leq \epsilon(\kappa).$$

The notion of one-more one-wayness was introduced by Bellare, Namprepre, Pointcheval, and Semanko [BNPS03].

Definition 11 (One-More One-Wayness). *Let $t : \mathbb{N} \to \mathbb{N}$, $q : \mathbb{N} \to \mathbb{N}$ and $\epsilon : \mathbb{N} \to [0,1]$. The function family $\mathsf{H} = (\mathsf{KeyGen}, \mathsf{Eval}, \mathsf{Invert})$ is (t, q, ϵ)-*one-more one-way *if for all adversary $\mathcal{A}$ running in time at most $t(\kappa)$ and queries to two oracles:*

- *an inversion oracle $\mathcal{O}_{\mathsf{Invert}}(\mathfrak{H}^{-1}, \cdot)$ which on input $y \in \mathbb{F}_q^m$ returns some $x \in \mathbb{F}_q^\ell$ such that $\mathcal{H}(x) = y$ or $\perp$ if none exists;*
- *a random challenge oracle $\mathcal{O}_{\mathsf{ch}}$ which returns a uniformly random element $y \leftarrow \mathbb{F}_q^m$ and adds it to a list $\mathcal{L}$*

The probability $\mathsf{Adv}_{\mathsf{OM\text{-}OW}}(\mathcal{A})$ that $\mathcal{A}$ outputs more than $q(\kappa)+1$ distinct valid preimages (for distinct images) while making at most $q(\kappa)$ queries to $\mathcal{O}_{\mathsf{Invert}}(\mathfrak{H}^{-1}, \cdot)$ is at most $\epsilon(\kappa)$:

$$\Pr\left[\begin{array}{c}y_i = \mathcal{H}(x_i) \in \mathcal{L}, \\ \text{for all } i \in \{1, \ldots, q(\kappa)+1\} \\ \text{and all } y_i \text{ distinct}\end{array} \;\middle|\; \begin{array}{c}(q, \ell, m, \mathcal{H}, \mathfrak{H}, \mathfrak{H}^{-1}) \leftarrow \mathsf{KeyGen}(1^\kappa) \\ \mathcal{L} \leftarrow \emptyset \\ \{x_i\}_{i=1}^{q(\kappa)+1} \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Invert}}(\mathfrak{H}^{-1}, \cdot), \mathcal{O}_{\mathsf{ch}}}(q, \ell, m, \mathfrak{H})\end{array}\right] \leq \epsilon(\kappa)$$

Our signature scheme relies on a related variant that we call *one-more claw-freeness*:

Definition 12 (One-More Claw-Freeness). *Let $t, n : \mathbb{N} \to \mathbb{N}$, $q : \mathbb{N} \to \mathbb{N}$ and $\epsilon : \mathbb{N} \to [0,1]$. The function family $\mathsf{H} = (\mathsf{KeyGen}, \mathsf{Eval}, \mathsf{Invert})$ is (t, q, ϵ)-*one-more claw-free *if for every adversary $\mathcal{A}$ running in time at most $t(\kappa)$ and making at most $q(\kappa)$ queries to the inversion oracle*

$$\mathcal{O}_{\mathsf{Invert}}(\mathfrak{H}^{-1}, \cdot) : y \in \mathbb{F}_q^m \mapsto x \in \mathbb{F}_q^\ell \text{ such that } \mathcal{H}(x) = y, \text{ or } \perp \text{ if none exists,}$$

the probability $\mathsf{Adv}_{\mathsf{OM\text{-}CF}}(\mathcal{A})$ that $\mathcal{A}$, for a randomly chosen quadratic function $\mathcal{J} \in \mathsf{MQ}(q, n(\kappa), m)$, outputs more than $q(\kappa)$ pairs $\{(x_i, z_i)\}_{i=1}^{q(\kappa)+1} \subseteq \mathbb{F}_q^\ell \times \mathbb{F}_q^n$ satisfying $\mathcal{H}(x_i) = \mathcal{J}(z_i)$, for all $i \in \{1, \ldots, q(\kappa)+1\}$, and such that all these values $\mathcal{H}(x_i) = \mathcal{J}(z_i)$ are distinct, is at most $\epsilon(\kappa)$:

$$\Pr\left[\begin{array}{c}(x_i, z_i) \text{ pairwise distinct} \\ \mathcal{H}(x_i) = \mathcal{J}(z_i) \\ \text{for } i \in \{1, \ldots, q(\kappa)+1\}\end{array} \;\middle|\; \begin{array}{c}(q, \ell, m, \mathcal{H}, \mathfrak{H}, \mathfrak{H}^{-1}) \leftarrow \mathsf{KeyGen}(1^\kappa), \\ \mathcal{J} \leftarrow \mathsf{MQ}(q, n, m), \\ \{(x_i, z_i)\}_{i=1}^{q(\kappa)+1} \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Invert}}(\mathfrak{H}^{-1}, \cdot)}(q, \ell, m, n, \mathfrak{H}, \mathcal{J})\end{array}\right] \leq \epsilon(\kappa).$$

It is worth mentioning that the analysis of the *one-more RSA problem* [BNPS03] is highly complex, as proving its hardness would, among other things, imply a separation between the RSA problem and integer factorization. The analogous problems we defined for multivariate quadratic systems are also open and appear to be very difficult. However, one can establish results showing that, even with access to a certain number of preimages of zero (or claws for two MQ systems), it remains NP-hard to decide whether the system admits another solution (or another claw).

B Proofs of randomness properties of quadratic maps

Lemma 6. *Let $\mathcal{F}$ be drawn uniformly at random from $\mathsf{MQ}(n, m, q)$. For any $x \neq y \in \mathbb{F}_q^n$, $\mathcal{F}(x)$ and $\mathcal{F}(y)$ are independent and uniformly distributed in $\mathbb{F}_q^m$.*

Proof. Let us fix two values $x \neq y \in \mathbb{F}_q^n$ and two scalars $\alpha, \beta \in \mathbb{F}_q$. Let us have a random quadratic polynomial $(x_1, \ldots, x_n) := \sum_{i,j} f_{i,j} x_i x_j + \sum_i f'_i x_i + f''$. What is the probability that $\mathcal{F}(x) = \alpha$ and $\mathcal{F}(y) = \beta$ at the same time? Denote by k the total number of terms of $\mathcal{F}$, so that there are possible q^k ways to choose $\mathcal{F}$.

Since $x \neq y$, there exists i^* such that $x_{i^*} \neq y_{i^*}$. Write $f(x) = f'' + f'_{i^*} x_{i^*} + g(x)$, where $g(x)$ contains all the other $k-2$ terms of f. It is not difficult to check that if $f(x) = \alpha$ and $f(y) = \beta$, then the two following relations have to hold:

$$f'_{i^*} = (\alpha - \beta + g(y) - g(x))/(x_{i^*} - y_{i^*}),$$
$$f'' = \alpha - (\alpha - \beta + g(y) - g(x))/(x_{i^*} - y_{i^*}) x_{i^*} - g(x).$$

It follows that there are exactly q^{k-2} ways to choose an f satisfying $f(x) = \alpha$ and $f(y) = \beta$. In turn, this shows that, when f is randomly chosen, $\Pr[f(x) = \alpha \text{ and } f(y) = \beta] = 1/q^2 = \Pr[f(x) = \alpha] \Pr[f(y) = \beta]$. The result easily generalizes from quadratic polynomials to quadratic maps (all coordinates are independent). □

It follows directly that $\mathsf{MQ}(n, m, q)$ is a *universal hash function family*. Indeed, for any $x \neq y \in \mathbb{F}_q^n$, the probability $\Pr[\mathcal{F}(x) = \mathcal{F}(y)]$ is equal to $\sum_{\alpha \in \mathbb{F}_q^m} \Pr[\mathcal{F}(x) = \alpha \text{ and } \mathcal{F}(y) = \alpha] = 1/q^m$. We believe that this was already known; in any case [AHI+17] shows that they are universal one-way hash functions.

We now state two simple results that apply to all function families satisfying the conclusion of Lemma 6, i.e., all hash function families, and in particular to quadratic maps.

Lemma 7. *For any $\mathcal{F} \in \mathsf{MQ}(n, m, q)$ and $\alpha \in \mathbb{F}_q^m$, let $N_\alpha(\mathcal{F})$ denote the number of preimages of α for $\mathcal{F}$. Then $\mathrm{E}[N_\alpha(\mathcal{F})] = q^{n-m}$ and $\mathrm{Var}[N_\alpha(\mathcal{F})] \leq q^{n-m}$ (taken over the random choice of $\mathcal{F}$).*

Proof. For any $x \in \mathbb{F}_q^n$, denote by $[\mathcal{F}(x) = \alpha]$ the indicator random variable of the corresponding event. The expectation of these random variables (over the

random choice of $\mathcal{F}$) is clearly q^{-m} and their variance is $q^{-m}(1-q^{-m})$, because they are Bernoulli random variables. Lemma 6 implies that they are pairwise independent. The number of preimages of α by $\mathcal{F}$ is:

$$N_\alpha = \sum_{x \in \mathbb{F}_q^n} [\mathcal{F}(x) = \alpha]$$

By linearity of expectation, $\mathrm{E}[N_\alpha] = q^{n-m}$, and by pairwise independence $\mathrm{Var}[N_\alpha] \leq q^{n-m}$. □

This implies that a random quadratic map is somewhat regular.

B.1 Proof of Lemma 1

Lemma 1. *A random quadratic map $\mathcal{F} \in \mathsf{MQ}(n, m, q)$ is:*

i) Injective with probability at least $1 - q^{2n-m}$.
ii) Surjective with probability at least $1 - q^{2m-n}$.

Proof. **Injectivity.** Because quadratic maps are universal hash functions, a simple union bound shows that:

$$\Pr[\mathcal{F} \text{ not injective}] = \Pr[\exists x \neq y, \mathcal{F}(x) = \mathcal{F}(y)] \leq \sum_{x \neq y} \Pr[\mathcal{F}(x) = \mathcal{F}(y)] = \sum_{x \neq y} q^{-m} \leq q^{2n-m}.$$

Surjectivity. Lemma 7 enables the use of Chebyshev's inequality, which shows that for any $t > 0$,

$$\mathbb{P}\big[\big|N_\alpha(\mathcal{F}) - q^{n-m}\big| \geq t\big] \leq \frac{q^{n-m}}{t^2}.$$

In particular, with $t = q^{n-m}$, we obtain that $\mathbb{P}[N_\alpha(\mathcal{F}) = 0] \leq q^{m-n}$. A simple union bound enables us to conclude that

$$\mathbb{P}[\mathcal{F} \text{ is not surjective}] \leq \sum_{\alpha \in \mathbb{F}_q^m} \mathbb{P}[N_\alpha(\mathcal{F}) = 0] \leq q^{2m-n}.$$

□

B.2 Proof of Lemma 2

Lemma 2. *For any pair $(\mathcal{F}, \mathcal{G})$ of quadratic maps in $\mathsf{MQ}(n, m, q)$, define*

$$\Delta(\mathcal{F}, \mathcal{G}) = \sum_{\alpha \in \mathbb{F}_q^m} \big|\mathbb{P}_x[\mathcal{F}(x) = \alpha] - \mathbb{P}_x[\mathcal{G}(x) = \alpha]\big|.$$

Then

$$\mathbb{P}_{\mathcal{F},\mathcal{G}}\left[\Delta(\mathcal{F}, \mathcal{G}) \geq 2q^{\frac{m-n}{3}}\right] \leq 2q^{\frac{m-n}{3}}.$$

Proof. Let $\mathcal{F}, \mathcal{G}$ be fixed. Using the same notations as above, we have $\mathbb{P}_x[\mathcal{F}(x) = \alpha] = q^{-n} N_\alpha(\mathcal{F})$, so that $\Delta(\mathcal{F}, \mathcal{G}) = q^{-n} \sum_{\alpha \in \mathbb{F}_q^m} \left| N_\alpha(\mathcal{F}) - N_\alpha(\mathcal{G}) \right|$.

Assume that both $|N_\alpha(\mathcal{F}) - q^{n-m}| \leq t$ and $|N_\alpha(\mathcal{G}) - q^{n-m}| \leq t$. Then $\Delta(\mathcal{F}, \mathcal{G}) \leq 2tq^{m-n}$ by the triangular inequality. Using Lemma 7 and Chebyshev's inequality again, we find that this assumption holds for $\mathcal{F}$ with probability $1 - q^{n-m}/t^2$ (over the random choice of $\mathcal{F}$). It therefore holds with probabilty $1 - 2q^{n-m}/t^2$ for both $\mathcal{F}$ and $\mathcal{G}$. Setting $t = q^{\frac{2}{3}(n-m)}$ yields the statement of the lemma. □

C Efficient arguments for our MQ commitment

C.1 Non-interactive zero-knowledge arguments (NIZK)

A *zero-knowledge protocol* for some polynomial-time decidable binary relation $\mathcal{R}$ (i.e., a relation that defines a language in NP) is defined by two PPT interactive algorithms, a prover P and a verifier V: both P and V are given a common input x and P is given in addition a witness w such that $(x, w) \in \mathcal{R}$. Then, P and V exchange a sequence of messages alternatively until V outputs a bit b (with $b = 1$ indicating that V accepts P's claim and $b = 0$ indicating that V rejects the claim). The entire sequence of messages exchanged by P and V, along with the answer b, is called a *transcript*, and is denoted as $\langle \mathsf{P}(x, w), \mathsf{V}(x) \rangle$.

A zero-knowledge argument for $\mathcal{R}$ with ϵ-knowledge soundness and (t, ξ)-zero-knowledge satisfies the following properties:

1. **perfect completeness**: if $(x, w) \in \mathcal{R}$, and P knows a witness w for x, they will succeed in convincing V, i.e.,

$$\Pr[\langle \mathsf{P}(x, w), \mathsf{V}(x) \rangle = 1] = 1.$$

2. **ϵ-knowledge soundness**: if there exists a PPT algorithm $\tilde{\mathsf{P}}$ such that

$$\tilde{\epsilon} := \Pr[\langle \tilde{\mathsf{P}}(x), \mathsf{V}(x) \rangle = 1] > \epsilon,$$

   then there exists a PPT algorithm E (called the *extractor*) which, given rewindable black-box access to $\tilde{\mathsf{P}}$ outputs a witness w' for x in time $\mathsf{poly}(\kappa, (\tilde{\epsilon} - \epsilon)^{-1})$ with probability at least $1/2$.
3. **(t, ξ)-zero-knowledge**: for every PPT algorithm $\tilde{\mathsf{V}}$, there exists a PPT algorithm S (called the *simulator*) which, given the input statement x and rewindable black-box access to $\tilde{\mathsf{V}}$, outputs a simulated transcript which is (t, ξ)-indistinguishable from $\langle \mathsf{P}(x, w), \tilde{\mathsf{V}}(x) \rangle$.

Zero-knowledge proofs are an important tool for many cryptographic protocols and applications. To build them, many techniques exist. In the setting of advanced protocols (such as blind signatures), we often need to build a proof of knowledge for a small circuit/statement. In such a context, the most efficient and versatile zero-knowledge proofs are obtained thanks to the MPC-in-the-Head (MPCitH) framework, which provides a generic construction using techniques

from secure multiparty computation. After sharing the secret witness w, the prover emulates "in her head" an MPC protocol with N parties and commits each party's view independently. The verifier then challenges the prover to reveal the views of a random subset of parties. By the privacy of the MPC protocol, nothing is revealed about the witness, which implies the zero-knowledge property. On the other hand, a malicious prover needs to cheat for at least one party, which shall be discovered by the verifier with high probability, hence ensuring the soundness property.

Then, this interactive protocol can be turned into a non-interactive one by applying standard techniques like the Fiat-Shamir heuristic [FS86].

C.2 Threshold-Computation-in-the-Head

The Threshold-Computation-in-the-Head (TCitH) framework [FR25] is currently the most efficient MPCitH-based instantiation (when considering small statements to prove). A proof system derived from TCitH is parameterized by the following parameters:

- The security level κ in bits;
- The number τ of parallel repetition of our proof;
- The number N of parties involved in the MPC protocol;
- The number L of revealed parties.

One additionaly considers some optimizations from [BBM+24] for MPCitH-based schemes relying on PRG trees:

1. Instead of considering τ independent PRG trees of N leaves in parallel, one relies on a unique large PRG tree of τN leaves where the i-th share of the e-th proof repetition is associated to the $((i-1)\tau + e)$-th leaf of the large PRG tree. Essentially, opening τL leaves of the big tree is more efficient than opening L leaves in each of the τ smaller trees, because with high probability some of the sibling paths in the tree will merge relatively close to the leaves, reducing the number of nodes that need to be revealed.
2. The last hash Fiat-Shamir challenge is recomputed with an incremented counter until the number of revealed nodes in the revealed sibling paths does not exceed some threshold t_{open}. This induces a rejection probability $p_{rej} := 1 - 1/\theta$, and the signer hence needs to perform an average of θ hash computations for the challenge (instead of 1). This strategy can be thought of as losing $\log_2 \theta$ bits of security (because of a smaller challenge space), which are regained thanks to a proof-of-work (performing an average of θ hash computations before getting a valid challenge).
3. One adds another explicit proof-of-work to the Fiat-Shamir hash computation of the last challenge. The signer must get a hash digest for which the ω last bits are zero, for some parameter ω. The same counter as for the previous improvement is used as a nonce in this hash and increased until the ω-zeros property is satisfied. This strategy increases the cost of hashing

the last challenge by a factor 2^ω and hence increases the security of ω bits. This thus allows to take smaller parameters (N, τ) achieving $\kappa - \omega$ bits of security.

Hence, our TCitH proof system has the following additional parameters:

- The proof-of-work parameter ω;
- The maximal number t_{open} of nodes opened in the tree used to commit the input shares.

The global proof-of-work, composed of the grinding process on the seed trees and the explicit proof-of-work on the Fiat-Shamir hash computation, has an average cost $\theta 2^\omega$ Fiat-Shamir hash computations. Note that the gain in size comes from the smaller parameters (N, τ) on the one hand, and the smaller sibling paths (of size $\leq t_{open}$ instead of $\approx \tau \log_2 N$) on the other hand. This gain in size is traded for an increased number of Fiat-Shamir hash attempts ($\theta 2^\omega$ on average instead of 1).

Modeling and performances. When relying on the MPC protocol Π_{pc} described in [FR25, Section 5.2] and when using the Fiat-Shamir heuristic, we obtain a non-interactive zero-knowledge proof of knowledge of a witness $w \in \mathbb{F}_q^{|w|}$ satisfying

$$f_j(w) = 0, \ \forall j \in \{1, \ldots, m\},$$

where $f_1, \ldots, f_m$ are some public degree-d polynomials from $\mathbb{F}_q[X_1, \ldots, X_{|w|}]$. For a given ω, one picks parameters (N, L, τ) ensuring $\kappa - \omega$ bits of security, i.e.,

$$\left(\frac{\binom{dL}{L}}{\binom{N}{L}} \right)^\tau \leq 2^{\kappa - \omega}.$$

Then one fixes t_{open} for these (N, L, τ) such that the global proof-of-work costs on average $\theta 2^\omega$ hash computations. The transcript size of this proof system when considering $L = 1$ is (in bits):

$$\textsc{Size} = 4\kappa + \kappa t_{open} + \tau \left(|w| \log_2 |\mathbb{F}_q| + (d-1)\rho e \log_2 |\mathbb{F}_q| + 2\kappa \right),$$

where e satisfies $|\mathbb{F}_q|^e \geq N$ (Shamir secret sharing requirement) and ρ satisfies $(|\mathbb{F}_q|^e)^\rho \geq 2^\kappa$ (the number of MPC repetition for each of the τ proof repetition).

C.3 Efficient arguments for our MQ commitment

To apply the previous TCitH proof system, we just need to express the statement we want to prove from Subsection 6.2 into (quadratic) polynomial constraints.

These MPCitH-like proofs have already been used to prove properties of commitments. [MV23] examined those based on subset-sum problems.

Argument of opening. The binary relation we want to prove is

$$\mathcal{R} = \{(\mathcal{F}, \mathcal{G}, c), (\mu, r) \mid c = (\mu + \mathcal{F}(r), \mathcal{G}(r)\}.$$

Let us denote c_1 the k first coefficients of c and c_2 the $m-k$ last coefficients. We can remark that wan can equivalently work on the binary relation

$$\mathcal{R}' = \{(\mathcal{G}, c_2), r \mid c_2 = \mathcal{G}(r)\},$$

since the message can be derived as $\mu := c_1 - \mathcal{F}(r)$, once we know r such that $c_2 = \mathcal{G}(r)$. $\mathcal{R}'$ directly describes a system of quadratic constraints, so we can apply the proof system in C.2 with $d = 2, |w| = |r| = n$.

The proof size is then (in bits):

$$\textsc{Size} = 4\kappa + \kappa t_{open} + \tau \left(n \log_2 q + \log_2 q^{\rho e} + 2\kappa\right),$$

where parameters are chosen following C.2. In particular, (ρ, e) are entirely determined once (κ, N, q) have been fixed.

Since the proof of the commitment's content is based on the TCitH framework and is therefore computational, one chooses commitment parameters achieving computational security in Table 3. Then we obtain the proof sizes described in Table 4, with an overall proof-of-work of $\log_2 \theta + \omega$ bits that is fixed around 11 bits, where $1/\theta$ is the success probability that the number of revealed nodes is lower than t_{open}.

Theorem 7 (Security proofs). *Let the PRG* PRG *used in our protocol be* $(t, \epsilon_{\mathsf{PRG}})$*-secure and the commitment scheme* Com *instantiated with a hash function be* $(t, \epsilon_{\mathsf{Com}})$*-hiding. Our interactive protocol is a* $(t, \epsilon_{\mathsf{PRG}} + \epsilon_{\mathsf{Com}})$*-zero-knowledge argument of knowledge for the language*

$$\{\mathcal{G} \in \mathbb{F}_q^{(m-k)\times n} \text{ a quadratic map, } c_2 \in \mathbb{F}_q^{m-k} \mid \exists\, r \in \mathbb{F}_q^n \text{ s.t. } \mathcal{G}(r) = c_2\},$$

with perfect completeness, and $\left[\frac{1}{q^{e\rho}} + \left(1 - \frac{1}{q^{e\rho}}\right) \frac{\binom{2L}{L}}{\binom{N}{L}}\right]$*-knowledge soundness, where* N *is the number of parties emulated by the prover,* L *the number of opened parties,* e *such that* $q^e \geq N$*, and* ρ *such that* $q^{e\rho} \geq 2^\kappa$*.*

Proof. Completeness. For any sampling of the random coins of P and V, if the computation described in the proof system of Appendix C adapted to our argument of opening is honestly performed, all the checks of V pass. The completeness is therefore perfect.

Soundness. For simplicity, we assume that the commitment scheme used in the proof system of Appendix C is perfectly binding (if it were computationally binding, we would need to handle cases where the extractor produces a commitment collision). For any two successful transcripts T_1, T_2 with the same commitment and two different challenges $i^*_{T_1} \neq i^*_{T_2}$ for the unopened leaf (among $\binom{N}{L}$ leaves) of the PRG tree, either the opened shares of $[\![r]\!]$ are inconsistent, leading to

a hash collision, or the openings are unique, defining the witness r (and this witness can be recovered). Let us call $[\![r]\!]$ a *valid* witness if $\mathcal{G}(r) = c_2$. For the rest, we assume that the extractor only processes consistent shares (otherwise, it would detect a hash collision). Let R_h denote the randomness used by $\tilde{\mathsf{P}}$ to generate the initial commitment h, and r_h a possible realization of R_h. The extractor procedure is described as follows:

1. Repeat $+\infty$ times:
2. Run $\tilde{\mathsf{P}}$ with honest $\mathsf{V}$to get transcript T_1
3. If T_1 is not a successful transcript, go to the next iteration
4. Do N_1 times:
5. Run $\tilde{\mathsf{P}}$ with honest $\mathsf{V}$and same r_h as T_1 to get transcript T_2
6. If T_2 is a successful transcript, $i^*_{T_1} \neq i^*_{T_2}$ and (T_1, T_2) reveals a valid witness,
7. Return (T_1, T_2)

Protocol 5: Extractor E

We estimate the extraction complexity, *i.e.*, the average number of times the extractor calls $\tilde{\mathsf{P}}$.

Let $\mathsf{succ}_{\tilde{\mathsf{P}}}$ be the event that $\tilde{\mathsf{P}}$ convinces V, with $\mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}}] = \tilde{\epsilon}$. Fix $\alpha \in (0,1)$ such that $(1-\alpha)\tilde{\epsilon} > \epsilon$. Call r_h *good* if $\mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}} \mid R_h = r_h] \geq (1-\alpha)\tilde{\epsilon}$. By the Splitting Lemma, $\mathbb{P}[r_h \text{ good} \mid \mathsf{succ}_{\tilde{\mathsf{P}}}] \geq \alpha$. Assume that we sample a successful transcript T_1 with a good r_h.

Let us prove by contraposition that if r_h is good, then the witness is valid. Assume that the witness in T_1 is invalid. This means that, if the sharings are genuine, the polynomial constraints check does not pass. Let FP be the event of a false positive in the polynomial constraints check, *i.e.* $\mathbb{P}[\mathsf{FP}] \leq \frac{1}{q^\rho}$ (see Lemma 2, [FR25]). The probability that the inner loop finds a successful transcript is:

$$\mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}} \mid R_h = r_h] \leq \mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}}, \mathsf{FP} \mid R_h = r_h] + \mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}}, \neg\mathsf{FP} \mid R_h = r_h]$$

$$\leq \frac{1}{q^\rho} + \left(1 - \frac{1}{q^\rho}\right) \mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}} \mid \neg\mathsf{FP}, R_h = r_h] \ .$$

Having a successful transcript when $\neg\mathsf{FP}$ occurs implies the prover must cheat for at least one subset of L parties. Cheating for multiple subsets fails, while cheating for exactly one subset succeeds with probability at most $1/\binom{N}{L}$ (all the sets of $L+1$ shares among the $[\![r]\!]_i$'s encode the same witness r). Since one relies on the proof system for polynomial constraints (with polynomial of degree 2) described in Appendix C (see Subsection 6.2, [FR25]), $\mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}} \mid R_h = r_h, \neg\mathsf{FP}] \leq \frac{\binom{2L}{L}}{\binom{N}{L}}$. Then one has $\mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}} \mid R_h = r_h] \leq \frac{1}{q^\rho} + \left(1 - \frac{1}{q^\rho}\right) \frac{\binom{2L}{L}}{\binom{N}{L}} = \epsilon < (1-\alpha)\tilde{\epsilon}$, meaning that r_h is *not* good. By contraposition, we get that if r_h is good, then the witness is valid. Now, let us lower bound the probability that the i-th iteration of the inner loop finds a successful transcript T_2 (whose event is denoted as $\mathsf{succ}^{T_2}_{\tilde{\mathsf{P}}}$)

such that $i^*_{T_1} \neq i^*_{T_2}$ given a good r_h.

$$\begin{aligned}
&\mathbb{P}[\mathsf{succ}^{T_2}_{\tilde{\mathsf{P}}} \cap (i^*_{T_1} \neq i^*_{T_2}) \mid r_h \text{ good}] \\
&\qquad = \mathbb{P}[\mathsf{succ}^{T_2}_{\tilde{\mathsf{P}}} \mid r_h \text{ good}] - \mathbb{P}[\mathsf{succ}^{T_2}_{\tilde{\mathsf{P}}} \cap (i^*_{T_1} = i^*_{T_2}) \mid r_h \text{ good}] \\
&\qquad \geq (1-\alpha)\tilde{\epsilon} - \mathbb{P}[i^*_{T_1} = i^*_{T_2} \mid r_h \text{ good}] \\
&\qquad = (1-\alpha)\tilde{\epsilon} - \frac{1}{\binom{N}{L}} \\
&\qquad \geq (1-\alpha)\tilde{\epsilon} - \epsilon \ .
\end{aligned}$$

Let define $p_0 := (1-\alpha)\tilde{\epsilon} - \epsilon$. By running η times $\tilde{\mathsf{P}}$ with the same r_h as for T_1, we hence obtain a second non-colliding transcript T_2 with probability at least $1/2$ when $\eta \approx \frac{\ln(2)}{\ln\left(\frac{1}{1-p_0}\right)} \leq \frac{\ln(2)}{p_0}$.

Let C denote the number of calls to $\tilde{\mathsf{P}}$ made by the extractor before finishing. In each iteration, the extractor makes one call to $\tilde{\mathsf{P}}$ to obtain T_1. If T_1 is not successful, which occurs with probability at least $(1 - \mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}}])$, the extractor continues to the next iteration and makes an average of $\mathbb{E}[C]$ calls to $\tilde{\mathsf{P}}$. If T_1 is successful, which occurs with probability $\mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}}]$, either r_h is good which happens with probability α and the extractor makes at most η calls to $\tilde{\mathsf{P}}$ in the inner loop of E and outputs a pair (T_1, T_2) with probability $1/2$, or the extractor makes η calls to $\tilde{\mathsf{P}}$ in the inner loop of E without stopping, with probability at most $(1 - \frac{\alpha}{2})$. The expected number of calls to $\tilde{\mathsf{P}}$ hence satisfies:

$$\begin{aligned}
\mathbb{E}[C] &\leq 1 + \underbrace{(1 - \mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}}])\mathbb{E}[C]}_{T_1 \text{ unsuccessful}} + \underbrace{\mathbb{P}[\mathsf{succ}_{\tilde{\mathsf{P}}}]\Big(\eta + \big(1 - \frac{\alpha}{2}\big)\mathbb{E}[C]\Big)}_{T_1 \text{ successful}} \\
&\leq 1 + \tilde{\epsilon}\eta + \mathbb{E}[C]\Big(1 - \frac{\tilde{\epsilon}\alpha}{2}\Big) \\
&\leq \frac{2}{\alpha\tilde{\epsilon}}(1 + \tilde{\epsilon}\eta) \\
&\leq \frac{2}{\alpha\tilde{\epsilon}}\left(1 + \tilde{\epsilon}\frac{\ln(2)}{(1-\alpha)\tilde{\epsilon} - \epsilon}\right)
\end{aligned}$$

To obtain an α-free formula, let us take α such that $(1-\alpha)\tilde{\epsilon} = \frac{1}{2}(\tilde{\epsilon} + \epsilon)$. We have $\alpha = \frac{1}{2}\left(1 - \frac{\epsilon}{\tilde{\epsilon}}\right)$ and the average number of calls to $\tilde{\mathsf{P}}$ is upper bounded by $\frac{4}{\tilde{\epsilon} - \epsilon}\left(1 + \tilde{\epsilon}\frac{2\ln(2)}{\tilde{\epsilon} - \epsilon}\right)$ which concludes the proof.

Zero-knowledge. We describe the simulator S (which produces transcripts indistinguishable from real ones without knowing the secret), with oracle access to some probabilistic-polynomial time $\tilde{\mathsf{V}}$, as follows:

1. Sample a challenge $i^* \xleftarrow{\$} [1, \binom{N}{L}]$.
2. Sample a master seed $\mathsf{mseed} \xleftarrow{\$} \{0,1\}^\lambda$.

3. Compute seeds $\{(\mathsf{seed}_i, \rho_i)\}_{i\in[1,\binom{N}{L}]}$ with $\mathsf{TreePRG}(\mathsf{mseed})$.
4. For each $i \in [1, \binom{N}{L}]\backslash\{i^*\}$, set $r_i \leftarrow \mathsf{PRG}(\mathsf{seed}_i)$ and $\mathsf{com}_i = \mathsf{Com}(\mathsf{seed}_i; \rho_i)$.
5. Sample $r, r_{i^*} \xleftarrow{\$} \mathbb{F}_q^n$ and a random commitment com_{i^*}, and define $\Delta r = r - \sum_i r_i$.
6. Given $\{r_i\}_{i\in\{1,\dots,\binom{N}{L}\}\backslash\{i^*\}}$, compute $\{[\![r]\!]_j\}_{j\in S:|S|=L}$ using the technique described in [FR25], where S is a subset of indexes determined by i^*.
7. Call $\tilde{\mathsf{V}}$ with the hash digest h of Δr (and of the commitments of the seed and associated randomness of each party) and get batch challenges.
8. Sample $\alpha \xleftarrow{\$} \mathbb{F}_q^e$.
9. Simulate computation of parties $j \in S$ to get $\{[\![\alpha]\!]_j\}_{j\in S}$.
10. Adapt the messages and the outputs of a random party $j^* \notin S$ such that the Lagrange interpolation of $[\![\alpha]\!]_{j^*}$ and $\{[\![\alpha]\!]_j\}_{j\in S}$ equals α.
11. Call $\tilde{\mathsf{V}}$ with the hash digest h' of $[\![\alpha]\!]$ and gets a challenge $\tilde{i}^*$. If $\tilde{i}^* \neq i^*$, then S restarts the simulation from scratch.
12. Outputs the transcript $\big(h, h', (\mathsf{seed}_i, \rho_i)_{i\neq i^*}, \mathsf{com}_{i^*}, \Delta r, [\![\alpha]\!]_{j^*}\big)$.

The output transcript is identically distributed to the genuine transcript except for the commitment of the party i^*. Distinguishing them means breaking the commitment hiding property or the PRG security. The above simulator S is a probabilistic polynomial-time algorithm since the challenge set $[1, \binom{N}{L}]$ (from which i^* is sampled) has a polynomial size in the security parameter.

Argument for arithmetic circuits. Let us design a zero-knowledge argument proving that some committed messages $\mu_1, \dots, \mu_M$ (which have been committed independently) satisfy some public relation

$$C(\mu_1, \dots, \mu_M) = 0,$$

where C is an arithmetic circuit. For the sake of simplicity, we assume that the messages have all the same size, i.e. live in $\mathbb{F}_q^k$. Let us denote o the vector where o_i is the output of the i^{th} multiplication gate in $C(\mu_1, \dots, \mu_M)$. To deal with C, we can use a Rank-1 Constraint System (R1CS): there exists three matrices A, B and D such that

$$C(\mu_1, \dots, \mu_M) = 0 \iff \exists o \text{ s.t. } A \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_M \\ o \\ 1 \end{pmatrix} \circ B \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_M \\ o \\ 1 \end{pmatrix} + D \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_M \\ o \\ 1 \end{pmatrix} = 0, \quad (4)$$

where $\circ$ denotes the component-wise multiplication. The R1CS constraints are quadratic ones. We can use the proof system from C.2 to prove that we know (μ_1, r_1), $\dots$, (μ_M, r_M) and $o \in \mathbb{F}_q^{|C|}$ such that $c_i = (\mu_i + \mathcal{F}_i(r_i), \mathcal{G}_i(r_i))$ $\forall i \in \{1, \dots, M\}$ and all the R1CS constraints are satisfied. The proof size is then (in bits):

$$\textsc{Size} = 4\kappa + \kappa t_{open} + \tau\left([M(n+k) + |C|]\log_2 q + \log_2 q^{\rho e} + 2\kappa\right).$$

If we apply it on the commitment parameters previously proposed in this section, we obtain the proof sizes described in Table 4.

Theorem 8 (Security proofs). *Let the PRG* PRG *used in our protocol be* $(t, \epsilon_{\mathsf{PRG}})$*-secure and the commitment scheme* Com *instantiated with a hash function be* $(t, \epsilon_{\mathsf{Com}})$*-hiding. Our interactive protocol is a* $(t, \epsilon_{\mathsf{PRG}}+\epsilon_{\mathsf{Com}})$*-zero-knowledge argument of knowledge for the language*

$$\{(c_i \in \mathbb{F}_q^m, \mathcal{G}_i \in \mathbb{F}_q^{(m-k)\times n}, \mathcal{F}_i \in \mathbb{F}_q^{k\times n})_{i\in\{1,\dots,M\}} \mid \exists\, (r_i \in \mathbb{F}_q^n, \mu_i \in \mathbb{F}_q^k)_{i\in\{1,\dots,M\}}$$
$$\text{and } o \in \mathbb{F}_q^{|C|} \text{ satisfying Eq. 4 and } c_i = (\mu_i + \mathcal{F}_i(r_i), \mathcal{G}_i(r_i))\ \forall i \in \{1,\dots,M\}\},$$

with perfect completeness, and $\left[\frac{1}{q^{e\rho}} + \left(1 - \frac{1}{q^{e\rho}}\right)\frac{\binom{2L}{L}}{\binom{N}{L}}\right]$*-knowledge soundness, where* N *is the number of parties emulated by the prover,* L *the number of opened parties,* e *such that* $q^e \geq N$*, and* ρ *such that* $q^{e\rho} \geq 2^\kappa$*.*

Proof. Completeness. For any sampling of the random coins of P and V, if the computation described in the proof system of Appendix C adapted to our argument of circuit satisfiability is honestly performed, all the checks of V pass. The completeness is therefore perfect.

Soundness. As in proof of Theorem 7, we assume that the commitment scheme used in the proof system of Appendix C is perfectly binding. For any two successful transcripts T_1, T_2 with the same commitment and two different challenges $i^*_{T_1} \neq i^*_{T_2}$ for the unopened leaf of the PRG tree, either the opened shares of $\{[\![r_j]\!], [\![\mu_j]\!]\}_{j\in\{1,\dots,M\}}, [\![o]\!]$ are inconsistent, leading to a hash collision, or the openings are unique, defining the witnesses $\{r_j, \mu_j\}_{j\in\{1,\dots,M\}}, o$ (and these witnesses can be recovered). Let us call $\{[\![r_j]\!], [\![\mu_j]\!]\}_{j\in\{1,\dots,M\}}, [\![o]\!]$ *valid* witnesses if $c_j = (\mu_j + \mathcal{F}_j(r_j), \mathcal{G}_j(r_j))\ \forall j \in \{1,\dots,M\}$ and Equation 4 is satisfied. The extractor procedure is identical to the one for Theorem 7 (see Algorithm 5). Then, the extraction complexity follows from the proof of Theorem 7. In particular, cheating for multiple subsets fails, while cheating for exactly one subset succeeds with probability at most $1/\binom{N}{L}$ (all the sets of $L+1$ shares among the $\{[\![r_j]\!], [\![\mu_j]\!]\}_{j\in\{1,\dots,M\}}, [\![o]\!]$'s encode the same witnesses $\{r_j, \mu_j\}_{j\in\{1,\dots,M\}}, o$). Since one also relies on the proof system for polynomial constraints (with polynomial of degree 2) described in Appendix C, the average number of calls to $\tilde{\mathsf{P}}$ and thus the soundness error is identical to Theorem 7.

Zero-knowledge. We describe the simulator S with oracle access to some PPT $\tilde{\mathsf{V}}$, as follows:

1. Sample a challenge $i^* \xleftarrow{\$} [1, \binom{N}{L}]$.
2. Sample a master seed $\mathsf{mseed} \xleftarrow{\$} \{0,1\}^\lambda$.

3. Compute seeds $\{(\mathsf{seed}_i, \rho_i)\}_{i\in[1,\binom{N}{L}]}$ with $\mathsf{TreePRG}(\mathsf{mseed})$.
4. For each $i \in [1, \binom{N}{L}]\backslash\{i^*\}$, set $\{r_{\ell,i}, \mu_{\ell,i}\}_{\ell\in\{1,\dots,M\}}, o_i \leftarrow \mathsf{PRG}(\mathsf{seed}_i)$ and $\mathsf{com}_i = \mathsf{Com}(\mathsf{seed}_i; \rho_i)$.
5. Sample $\{r_\ell, r_{\ell,i^*}\}_{\ell\in\{1,\dots,M\}} \xleftarrow{\$} \mathbb{F}_q^n$, $\{\mu_\ell, \mu_{\ell,i^*}\}_{\ell\in\{1,\dots,M\}} \xleftarrow{\$} \mathbb{F}_q^k$, $o, o_{i^*} \xleftarrow{\$} \mathbb{F}_q^{|C|}$ and a random commitment com_{i^*}. Define $\{\Delta r_\ell = r_\ell - \sum_i r_{\ell,i}, \Delta\mu_\ell = \mu_\ell - \sum_i \mu_{\ell,i}\}_{\ell\in\{1,\dots,M\}}, \Delta o = o - \sum_i o_i$.
6. Given $\{r_{\ell,i}, \mu_{\ell,i}, o_i\}_{\ell\in\{1,\dots,M\},i\in\{1,\dots,\binom{N}{L}\}\backslash\{i^*\}}$, compute $\{[\![r_\ell]\!]_j, [\![\mu_\ell]\!]_j, [\![o]\!]_j\}_{\ell\in\{1,\dots,M\},j\in S:|S|=L}$ using the technique described in [FR25], where S is a subset of indexes determined by i^*.
7. Call $\tilde{\mathsf{V}}$ with the hash digest h of $\{\Delta r_\ell, \Delta\mu_\ell\}_{\ell\in\{1,\dots,M\}}, \Delta o$ (and of the commitments of the seed and associated randomness of each party) and get batch challenges.
8. Sample $\alpha \xleftarrow{\$} \mathbb{F}_q^e$.
9. Simulate computation of parties $j \in S$ to get $\{[\![\alpha]\!]_j\}_{j\in S}$ and $\{\Delta\alpha_\ell\}_{\ell\in\{1,\dots,\rho\}}$.
10. Adapt the messages and the outputs of a random party $j^* \notin S$ such that the Lagrange interpolation of $[\![\alpha]\!]_{j^*}$ and $\{[\![\alpha]\!]_j\}_{j\in S}$ equals α.
11. Call $\tilde{\mathsf{V}}$ with the hash digest h' of $[\![\alpha]\!]$ and gets a challenge $\tilde{i}^*$. If $\tilde{i}^* \neq i^*$, then S restarts the simulation from scratch.
12. Outputs the transcript

$$\big(h, h', (\mathsf{seed}_i, \rho_i)_{i\neq i^*}, \mathsf{com}_{i^*}, \{\Delta r_\ell, \Delta\mu_\ell\}_{\ell\in\{1,\dots,M\}}, \Delta o, [\![\alpha]\!]_{j^*}\big).$$

The output transcript is identically distributed to the genuine transcript except for the commitment of the party i^*. Distinguishing them means breaking the commitment hiding property or the PRG security. The above simulator S is a probabilistic polynomial-time algorithm since the challenge set $[1, \binom{N}{L}]$ (from which i^* is sampled) has a polynomial size in the security parameter.

Motivation. Let us develop our motivation for this proof with an example. Consider a set of M people, and for simplicity, let us take $M = 1$. This person commits to their personal information (e.g., age, birthday, social insurance number, etc.), and the committed vector of information is denoted as $\mu \in \mathbb{F}_q^k$. Assume that the first element μ_1 of μ corresponds to the age information (recall that $q = 256$). If this person later wants to prove in zero-knowledge that they are older than 17, they can show that the committed value satisfies the arithmetic circuit computing:

$$\prod_{i=0}^{17}(\mu_1 - i)s \neq 0, \tag{5}$$

where s is randomly chosen by this person (the prover) and is an input of the circuit. This circuit has size $|C| = 18$, and by applying our proof system, we get a proof size below 6 KB (cf Table 4).

D Security proofs of the blind signature scheme

D.1 Proof of Theorem 5

Theorem 5. *The blind signature scheme in 3 and 4 is ℓ-one-more unforgeable based on the hardness of the $\ell+1$-claw-free problem and the knowledge soundness of the NIZK. More precisely, for every adversary $\mathcal{A}$ with advantage $\mathsf{Adv}_{\mathsf{OMUG}}(\mathcal{A})$ against the ℓ-one-more unforgeability game, there exist an adversary $\mathcal{B}$ with advantage $\mathsf{Adv}_{\mathsf{KS\text{-}NIZK}}(\mathcal{B})$ against the knowledge soundness of the NIZK and an adversary $\mathcal{C}$ with advantage $\mathsf{Adv}_{\mathsf{OM\text{-}CF}}(\mathcal{C})$ against the one-more claw-free game, such that*

$$\mathsf{Adv}_{\mathsf{OMUG}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{KS\text{-}NIZK}}(\mathcal{B}) + \mathsf{Adv}_{\mathsf{OM\text{-}CF}}(\mathcal{C}).$$

Proof. Assuming the existence of an adversary $\mathcal{A}$ with advantage $\mathsf{Adv}_{\mathsf{OMUG}}(\mathcal{A})$ against the one-more unforgeability game, we exhibit a reduction such that $\mathsf{Adv}_{\mathsf{KS\text{-}NIZK}}(\mathcal{B}) + \mathsf{Adv}_{\mathsf{OM\text{-}CF}}(\mathcal{C})$ is greater than or equal to $\mathsf{Adv}_{\mathsf{OMUG}}(\mathcal{A})$.

Let $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}$ be the public and secret keys from our blind signature with $\mathsf{pk} = (\mathcal{F}, \mathcal{G}, \mathcal{H}), \mathsf{sk} = \mathcal{G}^{-1}$. Given these parameters, the ℓ-one-more unforgeability game for our blind signature asks, given ℓ couples of distinct message and valid signature $(\mu_j, \mathsf{Sign}_j)_{j\in[j]}$, to produce a valid signature Sign' for a message $\mu' \notin \{\mu_j\}_{j\in[\ell]}$.

The inputs of the reduction $\mathcal{C}$ are $\mathcal{F}, \mathcal{G}, \mathcal{H}$, which are the parameters of the $\ell + 1$-claw-free game that the reduction has to win: the reduction has to find $\ell + 1$ claws for the functions $x \to \mathcal{G}(x)$ and $(x, y) \to (x + \mathcal{F}(y), \mathcal{H}(y))$.

The adversary $\mathcal{A}$ chooses ℓ distinct messages $\{\mu_j\}_{j\in[\ell]}$ and makes ℓ signing queries for these messages (with corresponding randomness $\{r_j\}_{j\in[\ell]}$). $\mathcal{A}$ receives ℓ responses $\{s_j\}_{j\in[\ell]}$ and produces the corresponding signatures $\{\mathsf{Sign}_j = \mathrm{NIZK}.P\big((r_j, s_j) : \mathcal{G}(s_j) = (\mu_j + \mathcal{F}(r_j), \mathcal{H}(r_j))\big)\}_{j\in[\ell]}$. Now, let us consider some $\mu \notin \{\mu_j\}_{j\in[\ell]}$. By the knowledge soundness of the NIZK, the adversary cannot come up with a valid proof π_2 for which the reduction fails to extract a valid couple (r, s), except with advantage $\mathsf{Adv}_{\mathsf{KS\text{-}NIZK}}(\mathcal{B})$ (in that case the OMUG wins but not the OM-CF). So assume the statement $\mathcal{G}(s) = (\mu + \mathcal{F}(r), \mathcal{H}(r))$ is true. By assumption, $\mathcal{A}$ comes up with such a valid statement with advantage $\mathsf{Adv}_{\mathsf{OMUG}}(\mathcal{A})$. Then, $\mathcal{A}$ sends these $\ell + 1$ claws $\{(\mu_j, r_j, s_j)\}_{j\in[\ell]}$ and (μ, r, s) to the reduction $\mathcal{C}$, which wins the $\ell + 1$-claw-free game with advantage at least $\mathsf{Adv}_{\mathsf{OMUG}}(\mathcal{A}) - \mathsf{Adv}_{\mathsf{KS\text{-}NIZK}}(\mathcal{B})$.

D.2 Proof of Theorem 6

Theorem 6. *The blind signature scheme in 3 and 4 is anonymous based on the zero-knowledge property of the NIZK, the blindness property of the MQ-based commitment, and the security of the MQ-based PRG. More precisely, for every adversary $\mathcal{A}$ with advantage $\mathsf{Adv}_{\mathsf{ANON}}(\mathcal{A})$ against the anonymity game, there exist an adversary $\mathcal{B}$ that can distinguish simulated NIZK proofs from real ones with advantage $\mathsf{Adv}_{\mathsf{ZK\text{-}NIZK}}(\mathcal{B})$, an adversary $\mathcal{C}$ that can break the blindness simulation game with advantage $\mathsf{Adv}_{\mathsf{MQ\text{-}Com}}(\mathcal{C})$, and an adversary $\mathcal{D}$ that can distinguish*

MQ-based PRG outputs from true randomness with advantage $\mathsf{Adv}_{PRG}(\mathcal{D})$ *such that*

$$\mathsf{Adv}_{ANON}(\mathcal{A}) \leq \mathsf{Adv}_{ZK\text{-}NISK}(\mathcal{B}) + \mathsf{Adv}_{MQ\text{-}Com}(\mathcal{C}) + \mathsf{Adv}_{MQ\text{-}PRG}(\mathcal{D}) \ .$$

Proof. The adversary first generates $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}$, and calls the random oracle with pk. The random oracle chooses μ_b and simulates the interaction for the first signing session with the corresponding signature Sign_b. Then the random oracle produces the interaction for the second signing session with some μ_{1-b} and corresponding signature Sign_{1-b}. Finally, the adversary receives μ_0, μ_1. The anonymity game for the adversary is to determine b, *i.e.*, trying to link signatures with signing sessions. Let us fix notations with $c_b = (\mu_b + \mathcal{F}(r_b), \mathcal{H}(r_b))$ and $\pi_{1,b}$ (resp. $c_b = (\mu_b + \mathcal{F}(r_b), \mathcal{H}(r_b))$ and $\pi_{1,1-b}$) the signing query sent by the random oracle for the message μ_b (resp. for the message μ_{1-b}).

By the zero-knowledge property of the NIZK, the adversary cannot extract information on the witness and then cannot distinguish between proofs except with advantage $\mathsf{Adv}_{\mathsf{ZK\text{-}NIZK}}(\mathcal{B})$. So assume that the proofs are indistinguishable. By the blindness property of the MQ-commitment, the adversary cannot extract information on the randomness r_b, r_{1-b} in c_b, c_{1-b} except with advantage $\mathsf{Adv}_{\mathsf{MQ\text{-}PRG}}(\mathcal{D})$. So assume that the commitments $\mathcal{H}(r_b), \mathcal{H}(r_{1-b})$ do not leak information on r_b, r_{1-b}. By the security of the MQ-PRG, the adversary cannot distinguish $\mu_b + \mathcal{F}(r_b), \mu_{1-b} + \mathcal{F}(r_{1-b})$, except with advantage $\mathsf{Adv}_{\mathsf{MQ\text{-}PRG}}(\mathcal{D})$. So assume that the PRG outputs are indistinguishable. At this stage, all the adversary can see is independent of the messages μ_0, μ_1, and of the random bit b drawn by the anonymity game, making the adversary's advantage equal to 0. Hence, the advantage of the adversary $\mathcal{A}$ against the anonymity game is at most

$$\mathsf{Adv}_{\mathsf{ZK\text{-}NISK}}(\mathcal{B}) + \mathsf{Adv}_{\mathsf{MQ\text{-}Com}}(\mathcal{C}) + \mathsf{Adv}_{\mathsf{MQ\text{-}PRG}}(\mathcal{D}) \ .$$