



Universiteit
Leiden
The Netherlands

Compressed Σ -protocol theory

Attema, T.

Citation

Attema, T. (2023, June 1). *Compressed Σ -protocol theory*. Retrieved from <https://hdl.handle.net/1887/3619596>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3619596>

Note: To cite this publication please use the final published version (if applicable).

CHAPTER 4

Compressed Σ -Protocols: Higher Level Functionalities

4.1 Introduction

Instantiating compressed Σ -protocols with a homomorphic and compact vector commitment scheme establishes an honest-verifier zero-knowledge proof for opening linear forms L on committed vectors $\mathbf{x}$. More precisely, its most basic variant is a protocol for proving knowledge of a commitment opening $(\mathbf{x}; \gamma)$ satisfying the *linear* constraint $L(\mathbf{x}) = y$. This functionality might seem somewhat restrictive; in many practical scenarios the statement one wishes to prove cannot be captured by a linear constraint *directly*. In this chapter, we enhance this basic linear functionality by treating two (classes of) relations that cannot be captured by a homomorphism directly. In both cases our strategy is to reduce the relation to our desired starting point, i.e., a prover claiming to know a homomorphism preimage.

First, in Section 4.2, we consider proving the correctness of a large set of committed multiplication triples $(\alpha_i, \beta_i, \gamma_i = \alpha_i \beta_i) \in \mathbb{Z}_q^3$. The corresponding multiplicative relation is clearly nonlinear and therefore cannot be captured by a homomorphism directly. Our approach is to *linearize* this relation to bring about the desired starting point and, subsequently, apply a compressed Σ -protocol. This approach is based on the work of [CDP12] that shows how to prove *arbitrary constraints* on committed vectors by exploiting techniques from secure multi-party computation (MPC) based on arithmetic secret sharing. More concretely, our work is based on the ideas underlying the *commitment multiplication* protocol from [CDM00]. It is this combination of “compact commitments with linear openings” and arithmetic secret sharing that allows for “linearizing nonlinear relations.” This section is based on the article [AC20], co-authored by Ronald Cramer.

Second, in Section 4.3, we consider a prover that claims to know the homomorphism preimages for a *subset* of public elements $P_1, \dots, P_n$, i.e., a prover claiming to have *partial* knowledge about the preimages of these elements. Proofs of partial knowledge were introduced in [CDS94]. Their solution combines Σ -protocol theory and linear secret sharing, and achieves linear communication complexity. We present a Σ -protocol, inspired by the [CDS94]-approach, for linearizing the proof of partial knowledge relation. However, a careful re-design of the original protocol is necessary to allow for compression. After composing with the appropriate

compressed Σ -protocol, we establish a proof of partial knowledge with logarithmic communication complexity. This section is based on the article [ACF21], co-authored by Ronald Cramer and Serge Fehr.

4.2 An Arithmetic Secret-Sharing Based Linearization Technique

The main result of [CDP12] is a zero-knowledge protocol for proving the correctness of a large number of committed multiplication triples $(\alpha_i, \beta_i, \gamma_i = \alpha_i \beta_i) \in \mathbb{Z}_q^3$. Their technique requires some adaptations to make it work for us here. In Section 4.2.1, we first outline the technique from [CDP12] and then discuss the required adaptations. These adaptations allow us to linearize the nonlinear relations defined by multiplication triples. Combined with our compressed Σ -protocols for opening linear forms, we obtain an interactive proof that allows a prover to commit to a large vector of multiplication triples and prove that the committed vector is of the appropriate form.

In practice, it may happen that the prover is already committed to the vector of multiplication triples *before* being asked to prove its correctness. This is referred to as the “commit-and-prove” scenario. In order to deal with this scenario some further utility enhancements are needed. The required enhancements, based on the compactification techniques of Section 3.4.4, are described in Section 4.2.2.

Finally, the linearization technique of Section 4.2.1 requires $q > 3m$, i.e., the multiplication triples must be defined over a large enough field $\mathbb{Z}_q$. In Section 4.2.3, we show how to handle the case $q \leq 3m$.

4.2.1 Proving Correctness of Multiplication Triples

Let us first outline the technique from [CDP12] for proving the correctness of committed multiplication triples. Subsequently, we describe our adaptations to this technique. The work of [CDP12] considers homomorphic commitment schemes where the secret committed to is not a vector in $\mathbb{Z}_q^n$, but a *single element* of $\mathbb{Z}_q$ instead. Their primary result is a Σ -protocol showing the correctness of commitments to m multiplication triples $(\alpha_i, \beta_i, \gamma_i := \alpha_i \beta_i)$. In other words, each of the α_i ’s, β_i ’s and γ_i ’s is individually committed to, and the protocol verifies the multiplicative relations $\gamma_i = \alpha_i \cdot \beta_i$. The communication complexity of the [CDP12]-approach is linear in the number of multiplication triples m . Adaptations are required to make the protocol amenable for compression and reduce the communication complexity down to logarithmic.

Their solution employs multiplicative packed secret sharing (Section 2.10). For instance, consider Shamir’s scheme over $\mathbb{Z}_q$, with privacy parameter $p = 1$, but with secret-space dimension m . This scheme uses random polynomials of degree $\leq m$, subject to the evaluations on the points $1, \dots, m$ comprising the desired secret vector. Note that, for each sharing, a single random $\mathbb{Z}_q$ -element is required (which can be taken as the evaluation at 0). Moreover, this packed secret scheme can be instantiated with $q - m$ players, with shares corresponding to the $q - m$ evaluations outside the points $1, \dots, m$.

Remark 4.1. Actually, the above scheme can be instantiated with $q - m + 1$ players by taking the evaluation at infinity as an additional share. Because this adaptation only has a minor impact on the properties of the protocol, we will ignore the point at infinity in our analysis. For more details see [CDN15].

It is important to note that, given a secret vector $\mathbf{x} \in \mathbb{Z}_q^m$ and random element $r \in \mathbb{Z}_q$, it holds by *Lagrange Interpolation* that, for each $c \in \mathbb{Z}_q$, the evaluation $f(c)$ of such polynomial $f(X)$ is some public $\mathbb{Z}_q$ -linear combination over the coordinates of the secret vector and the random element. Namely, consider the map that, on input $m + 1$ arbitrary evaluations on the points $0, \dots, m$, outputs the (coefficients of the) unique polynomial $f(X)$ of degree $\leq m$ that maps the points $0, \dots, m$ to these given evaluations. A transformation matrix describing this map corresponds to the inverse of the Vandermonde-matrix

$$\begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 1 & 1 & 1 & \cdots & 1 \\ 1 & 2 & 4 & \cdots & 2^m \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 1 & m & m^2 & \cdots & m^m \end{pmatrix} \in \mathbb{Z}_q^{(m+1) \times (m+1)}.$$

Composed with the *linear* evaluation at c mapping, this transformation describes the desired $\mathbb{Z}_q$ -linear combination.

Now, assume that $3m < q$. In this case, the total number of shares $q - m$ is at least $2m + 1$ and the above instantiation of Shamir's secret-sharing scheme is multiplicative. More precisely, the secret-sharing scheme has $(2m + 1)$ -product-reconstruction. In Section 4.2.3, we describe how to handle the case $3m \geq q$. The [CDP12]-protocol goes as follows.

- The vectors of commitments to the multiplication triples are assumed to be part of the common input.
- The prover selects a random polynomial $f(X)$ of degree at most m that defines a packed secret sharing of the vector $(\alpha_1, \dots, \alpha_m)$. The prover also selects a random polynomial $g(X)$ of degree at most m that defines a packed secret sharing of the vector $(\beta_1, \dots, \beta_m)$. Finally, the prover computes the product polynomial $h(X) := f(X)g(X)$ of degree at most $2m < q$.
- The prover commits to the random $\mathbb{Z}_q$ -element for the sharing based on $f(X)$, i.e., $f(0)$, and commits to the random $\mathbb{Z}_q$ -element for the sharing based on $g(X)$, i.e., $g(0)$. The prover also commits to the evaluations of $h(X)$ on the points $0, m + 1, \dots, 2m$.¹ Note that the “absent” evaluations at $1, \dots, m$ comprise the γ_i 's and their commitments are already assumed to be part of the common input.
- The prover sends these $m + 3$ commitments to the verifier.
- The verifier selects a random challenge $c \in \mathbb{Z}_q$, distinct from $1, \dots, m$, and sends it to the prover.

¹By Lagrange interpolation these points, together with the γ_i 's, determine $h(X)$.

- By public linear combinations, both prover and verifier can compute three commitments: one to $u := f(c)$, one to $v := g(c)$ and one to $w := h(c)$. The prover opens each of these (assuming, of course, that $c \notin \{1, \dots, m\}$).
- The verifier checks the three openings and verifies that $u \cdot v = w$.

If the committed polynomials do not satisfy $f(X)g(X) = h(X)$, and under the assumption that the commitment scheme is binding, there are at most $2m$ values of c out of the $q - m$ possibilities such that the final check goes through. So a dishonest prover succeeds with probability at most $2m/(q - m)$, which is smaller than 1 since $3m < q$. More precisely, the protocol can be shown to be $(2m + 1)$ -out-of- $(q - m)$ special-sound under the assumption that the commitment scheme is binding. Honest-verifier zero-knowledge essentially follows from 1-privacy of the secret-sharing scheme.

We now make the following observation. In the above protocol, the prover may as well use our compressed Σ -protocol for opening linear forms as a black-box. Indeed, the entire $(4m + 3)$ -dimensional $\mathbb{Z}_q$ -vector

$$\mathbf{y} = (\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_m, \gamma_1, \dots, \gamma_m, f(0), g(0), h(0), h(m + 1), \dots, h(2m))$$

of data that the prover commits to in the protocol above can be committed to in a *single* compact commitment. Note that, by definition, $\gamma_i = h(i)$ for all $1 \leq i \leq m$, i.e., the γ_i 's comprise the “missing” evaluations of $h(X)$. Furthermore, all of the data *opened* to the verifier is some fixed linear form on the (long) secret committed vector $\mathbf{y}$. Indeed, each of the values u, v and w corresponds to an opening of a public linear form applied to $\mathbf{y}$. The linear form is determined by Lagrange interpolation as addressed above, under the convention that the form takes zeros on the portion of the coordinates of $\mathbf{y}$ not relevant to the computation, i.e., all three linear forms correspond to the evaluation of a polynomial whose coefficients are defined by a different part of $\mathbf{y}$.

Overall, in this adaptation of the [CDP12]-protocol, the prover sends a single compact commitment to $\mathbf{y}$ to the verifier and, after receiving a challenge $c \leftarrow_R \mathbb{Z}_q \setminus \{1, \dots, m\}$, the prover and verifier proceed by running a compressed Σ -protocol to open three different linear forms. This interactive proof for committing to m multiplication triples and proving the correctness of these triples only requires the prover to send $\mathcal{O}(\log(m))$ elements.

4.2.1.1 Generalizing to Arbitrary Packed Secret-Sharing Schemes

For concreteness, the linearization technique has thus far been instantiated with Shamir's packed secret-sharing scheme. This scheme allows an m -dimensional vector with coefficients in a finite field to be secret shared amongst $q - m$ players. Moreover, the deployed scheme has 1-privacy, $(m + 1)$ -reconstruction and $(2m + 1)$ -product-reconstruction. Hence, if $2m + 1 \leq q - m$, or equivalently $q > 3m$, this scheme is multiplicative.

More generally, as long as $n \geq R$, the linearization technique can be instantiated with any n -player linear secret-sharing scheme (LSSS) $\mathcal{S}$ for m -dimensional vectors that has R -product-reconstruction and $(p \geq 1)$ -privacy. As in Section 2.10, we

let $\widehat{\mathcal{S}}$ denote the LSSS such that every component-wise product $[\mathbf{a}; \mathbf{r}_\mathbf{a}]_{\mathcal{S}} * [\mathbf{b}; \mathbf{r}_\mathbf{b}]_{\mathcal{S}}$ of secret sharings is a secret sharing of the component-wise product $\mathbf{a} * \mathbf{b}$ with respect to $\widehat{\mathcal{S}}$. The linearization technique, now instantiated with $\mathcal{S}$, proceeds as follows.

- The prover samples $\mathbf{r}_\alpha, \mathbf{r}_\beta \leftarrow_R \mathbb{Z}_q^t$ uniformly at random and computes $\mathbf{r}_\gamma$ such that

$$[\alpha_1, \dots, \alpha_m; \mathbf{r}_\alpha]_{\mathcal{S}} * [\beta_1, \dots, \beta_m; \mathbf{r}_\beta]_{\mathcal{S}} = [\gamma_1, \dots, \gamma_m; \mathbf{r}_\gamma]_{\widehat{\mathcal{S}}}.$$

- The prover commits to the long vector

$$\mathbf{y} = (\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_m, \gamma_1, \dots, \gamma_m, \mathbf{r}_\alpha, \mathbf{r}_\beta, \mathbf{r}_\gamma).$$

in a single compact commitment.

- The verifier selects an index $i \leftarrow_R \{1, \dots, n\}$ uniformly at random and sends it to the prover.
- By linearity of the secret-sharing scheme, both prover and verifier can determine three linear forms L_1 , L_2 and L_3 : one corresponding to the i -th share $L_1(\mathbf{y}) = u$ of $[\alpha_1, \dots, \alpha_m; \mathbf{r}_\alpha]_{\mathcal{S}}$ when evaluated in $\mathbf{y}$, one to the i -th share $L_2(\mathbf{y}) = v$ of $[\beta_1, \dots, \beta_m; \mathbf{r}_\beta]_{\mathcal{S}}$ and one to the i -th share $L_3(\mathbf{y}) = w$ of $[\gamma_1, \dots, \gamma_m; \mathbf{r}_\gamma]_{\widehat{\mathcal{S}}}$.
- The prover uses a compressed Σ -protocol to open the three linear forms L_1 , L_2 and L_3 on the compactly committed vector $\mathbf{y}$.
- The verifier checks the three openings and checks whether $u \cdot v = w$.

The following lemma shows that, if $\alpha_i \cdot \beta_i \neq \gamma_i$ for some i , then $u \cdot v = w$ with probability at most $(R-1)/n$. Hence, assuming that the commitment scheme is binding, a dishonest prover succeeds with probability at most $(R-1)/n$ in convincing the verifier.

As before, honest-verifier zero-knowledge essentially follows from $(p \geq 1)$ -privacy of the secret-sharing scheme. In fact, the verifier can ask the prover to open the shares of p different players instead of only one. For $p > 1$, this reduces the success probability of a dishonest prover from $(R-1)/n$ down to $\binom{R-1}{p} / \binom{n}{p}$.

Lemma 4.1 (Arithmetic Secret Sharing Based Linearization). *Let $m, n, t, R \in \mathbb{N}$ with $R \leq n$, q prime and $\mathcal{S}$ the linear secret-sharing scheme (LSSS) defined by $M \in \mathbb{Z}_q^{n \times (m+t)}$. Further, let $\widehat{\mathcal{S}}$ be the LSSS defined by*

$$\widehat{M} = \begin{pmatrix} M_1 \otimes M_1 \\ \vdots \\ M_n \otimes M_n \end{pmatrix} \in \mathbb{Z}_q^{n \times (m+t)^2},$$

where M_i denotes the i -th row of M . Suppose that $\widehat{\mathcal{S}}$ has R -reconstruction or, equivalently, that $\mathcal{S}$ has R -product-reconstruction.

Then, for all $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{Z}_q^m$ with $\mathbf{a} * \mathbf{b} \neq \mathbf{c}$ and for all $\mathbf{r}_\mathbf{a}, \mathbf{r}_\mathbf{b}$ and $\mathbf{r}_\mathbf{c}$, it holds that the vectors

$$[\mathbf{a}; \mathbf{r}_\mathbf{a}]_S * [\mathbf{b}; \mathbf{r}_\mathbf{b}]_S \in \mathbb{Z}_q^n \quad \text{and} \quad [\mathbf{c}; \mathbf{r}_\mathbf{c}]_{\hat{S}} \in \mathbb{Z}_q^n$$

coincide in at most $R - 1$ coefficients.

Proof. First, recall that the component-wise product $[\mathbf{a}; \mathbf{r}_\mathbf{a}]_S * [\mathbf{b}; \mathbf{r}_\mathbf{b}]_S$ of $\mathcal{S}$ -sharings is a secret sharing of $\mathbf{a} * \mathbf{b}$ with respect to $\hat{S}$ (see Section 2.10), i.e.,

$$[\mathbf{a}; \mathbf{r}_\mathbf{a}]_S * [\mathbf{b}; \mathbf{r}_\mathbf{b}]_S = [\mathbf{a} * \mathbf{b}; \mathbf{r}]_{\hat{S}} \in \mathbb{Z}_q^n$$

for some vector $\mathbf{r}$.

Since $\hat{S}$ has R -reconstruction, any A of cardinality R of the secret sharing $[\mathbf{a} * \mathbf{b}; \mathbf{r}]_{\hat{S}}$ uniquely determines $\mathbf{a} * \mathbf{b}$. Hence, if there exists an R -subset A for which the shares of $[\mathbf{a} * \mathbf{b}; \mathbf{r}]_{\hat{S}}$ and $[\mathbf{c}; \mathbf{r}_\mathbf{c}]_{\hat{S}}$ coincide, it follows that $\mathbf{a} * \mathbf{b} = \mathbf{c}$. This contradicts the assumption $\mathbf{a} * \mathbf{b} \neq \mathbf{c}$ and therefore such an R -subset A cannot exist. This completes the proof of the lemma. $\square$

4.2.2 A Commit-and-Prove Variant

The compressed Σ -protocol for proving the correctness of m multiplication triples, described in Section 4.2.1, requires the prover to commit to the vector of triples

$$\mathbf{x} = (\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_m, \gamma_1, \dots, \gamma_m) \in \mathbb{Z}_q^{3m}$$

and the auxiliary information

$$\mathbf{aux} = (f(0), g(0), h(0), h(m+1), \dots, h(2m)) \in \mathbb{Z}_q^{m+3}$$

in a single compact commitment. By contrast, the original [CDP12]-protocol allows the prover to generate this auxiliary information *after* it has committed to the multiplication triples. A protocol where the prover can first commit to the secret input data and at a later point in time prove that the committed input satisfies some constraint, unknown at the time of committing to the input data, is called a *commit-and-prove* protocol. Hence, whereas the original [CDP12]-protocol is commit-and-prove, the compressed Σ -protocol of Section 4.2.1 is not.

In particular, note that the [CDP12]-protocol can be repeated arbitrarily many times, e.g., to prove to multiple verifiers that a fixed set of commitments comprises a set of committed multiplication triples. In each repetition the protocol generates fresh auxiliary information. By contrast, the compressed Σ -protocol variant outputs a commitment to multiplication triples together with the auxiliary information. Hence, repeating this protocol would output different commitments, allowing a dishonest prover to commit to different sets of multiplication triples in each invocation. Moreover, since the deployed secret-sharing scheme only has 1-privacy, a prover can not reuse a compact commitment to the long vector $\mathbf{y} = (\mathbf{x}, \mathbf{aux}) \in \mathbb{Z}_q^{4m+3}$. More precisely, it is crucial that the prover only opens the evaluations $f(c)$, $g(c)$ and $h(c)$ for a single challenge $c \in \mathbb{Z}_q \setminus \{1, \dots, m\}$. This prevents the prover from reusing a given commitment to the long vector $\mathbf{y}$.

In many practical scenarios, commit-and-prove functionality is crucial. Fortunately, enhancing the compressed Σ -protocol for multiplication triples with this

functionality turns out to be merely a matter of plug-and-play with the basic theory.

To see this, suppose $P \in \mathbb{H}$ is a fixed compact commitment to the vector of m multiplication triples $\mathbf{x} \in \mathbb{Z}_q^{3m}$. We aim to bring about the desired starting point, i.e., a single compact commitment to multiplication triples and freshly generated auxiliary information $\mathbf{aux} \in \mathbb{Z}_q^{m+3}$. Let now Q be a commitment to the vector $\mathbf{aux}$ prepended with $3m$ zeros, i.e., to $(0, \mathbf{aux}) \in \mathbb{Z}_q^{4m+3}$ (here 0 denotes a vector of $3m$ zeros). Then, since the commitment scheme is assumed to be homomorphic, $P \cdot Q$ is the required compact commitment to the vector $\mathbf{y} = (\mathbf{x}, \mathbf{aux})$ containing both the multiplication triples and the auxiliary information. What remains is for the prover to convince the verifier that the commitment Q is of the appropriate form. More precisely, it must prove that Q is a commitment to a vector $(0, \mathbf{aux})$ starting with at least $3m$ zeros. This simply amounts to opening the $3m$ linear forms

$$L_i: \mathbb{Z}_q^{4m+3} \rightarrow \mathbb{Z}_q, \quad \mathbf{x} \mapsto x_i,$$

for $1 \leq i \leq 3m$. Namely, opening L_i to 0 on a compactly committed vector shows that the i -th coordinate of this vector equals 0.

The commit-and-prove variant thus runs two amortized compressed Σ -protocols for opening linear forms as subroutines. The first one opens the three linear forms, corresponding to the polynomial evaluations $f(c)$, $g(c)$ and $h(c)$, on the commitment $P \cdot Q$. The second one opens the n linear forms L_i , corresponding to the required nullity checks, on commitment Q . Recall that the costs of opening n different linear forms on a single compact commitment can be amortized (Section 3.4.2). Therefore, the naive commit-and-prove approach incurs roughly a factor two loss in communication efficiency. By deploying the compactification techniques of Section 3.4.4, this factor two loss can be avoided.

The foregoing describes how to handle the scenario where the prover is already committed to all multiplication triples in a single compact commitment P . Section 3.4.4 also shows how to handle the case where the prover is committed to all coefficients of the vector $\mathbf{x}$ of multiplication triples individually, i.e., in separate 1-dimensional commitments. Further, in Section 7.2.2, we deploy similar techniques to achieve a commit-and-prove circuit satisfiability protocol.

4.2.3 Correctness of Multiplication Triples in Small Fields

The linearization technique of Section 4.2.1 requires $q > 3m$, where m is the number of multiplication triples in the finite field $\mathbb{Z}_q$. In fact, linearization is well defined as long as $q > 2m$; the prover must commit to $2m + 1$ distinct evaluations of the polynomial $h(X) \in \mathbb{Z}_q[X]$. However, since the challenges are sampled from $\mathbb{Z}_q \setminus \{1, \dots, m\}$, the linearization step itself is a $(2m + 1)$ -out-of- $(q - m)$ special-sound Σ -protocol. In Chapter 6, we show that this implies a knowledge error $\kappa \geq 2m/(q - m)$. Hence, to ensure a nontrivial knowledge error $\kappa < 1$, we must even require $q > 3m$. In the discrete logarithm instantiation the prime q is exponential in the security parameter and, with m polynomial in the security parameter, this gives a negligible knowledge error. However, when the multiplication triples are defined over smaller fields, possibly even with cardinality $q \leq 3m$, the approach above does not suffice. In this section, we show that only minor adaptations are required when considering multiplication triples in small fields $\mathbb{Z}_q$.

So let, as before,

$$\mathbf{x} = (\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_m, \gamma_1, \dots, \gamma_m) \in \mathbb{Z}_q^{3m}$$

be a vector of m multiplication triples $(\alpha_i, \beta_i, \gamma_i)$, but now defined over a small field with cardinality $q \leq 3m$. To handle the fact that $q \leq 3m$, we simply define Shamir's secret-sharing scheme over a field extension $\mathbb{F}/\mathbb{Z}_q$ with cardinality at least $3m + 1$. More precisely, in this case $f(X) \in \mathbb{F}[X]$, $g(X) \in \mathbb{F}[X]$ and $h(X) = f(X)g(X)$ define packed secret sharings, with 1-privacy, of the α_i 's, the β_i 's and the γ_i 's. Hence, the difference with before is that the polynomials are now defined over the field extension $\mathbb{F}$ instead of the base field $\mathbb{Z}_q$. Let d be the degree of the extension $\mathbb{F}/\mathbb{Z}_q$, i.e., $|\mathbb{F}| = q^d > 3m$. For simplicity, we enumerate the elements of $\mathbb{F}$, i.e., every $0 \leq i < q^d$ uniquely corresponds to a field element. Then, by choosing an appropriate basis of the extension $\mathbb{F}/\mathbb{Z}_q$, the vector

$$\mathbf{y} = (\mathbf{x}, f(0), g(0), h(0), h(m+1), \dots, h(2m)) \in \mathbb{F}^{4m+3},$$

containing all relevant information, can be viewed as a vector in $\mathbb{Z}_q^{3m+dm+3d}$. Here, we use that the coefficients of $\mathbf{x}$ are elements of the base field $\mathbb{Z}_q$. What remains is to observe that, also in this generalization, all evaluations of $f(X)$, $g(X)$ and $h(X)$ are accessible as $\mathbb{Z}_q$ -affine combinations of the coefficients of $\mathbf{y}$.

Taking d such that $|\mathbb{F}| = q^d > 2m$ results in a well-defined linearization technique for multiplication triples in $\mathbb{Z}_q$. It is a perfectly complete and $(2m+1)$ -out-of- $(q^d - m)$ special-sound Σ -protocol with knowledge error

$$\kappa = \frac{2m+1}{q^d - m},$$

i.e., $q^d > 3m$ is required for the knowledge error to be nontrivial. Moreover, when composed with a compressed Σ -protocol for opening linear forms, the prover only needs to send $\mathcal{O}(\log(n + dm))$ elements to the verifier.

Exactly the same approach applies to multiplication triples defined over a ring $\mathcal{R}$. In this case, the evaluation points of the Shamir polynomials $f(X), g(X), h(X) \in \mathcal{R}[X]$ should be chosen from an *exceptional* subset of $\mathcal{R}$. If the maximal size of such an exceptional subset is too small, i.e., at most $3m$, one simply defines the secret-sharing scheme over an appropriate ring extension of $\mathcal{R}$.

4.3 Proofs of Partial Knowledge

In a k -out-of- n proof of partial knowledge [CDS94] a prover knowing witnesses for some k -subset, i.e., subset of cardinality k , of n given public statements can convince the verifier of this claim without revealing which k -subset. Typically, the secrets are solutions to public instances of intractable problems, such as the discrete logarithm problem. The work of [CDS94] gives an elegant solution with linear communication complexity that combines Σ -protocol theory with linear secret sharing. Especially the “1-out-of- n ” case $k = 1$ has seen myriad applications during the last decades, e.g., in electronic voting, ring signatures, and confidential transaction systems. Our goal is to construct proofs of partial knowledge with *logarithmic* communication complexity in both k and n .

4.3.1 Knowledge of k -out-of- n Homomorphism Preimages

Before we present our solution, let us formalize the k -out-of- n proof of partial knowledge problem. To this end, for a prime q and a group $(\mathbb{G}_T, +)$ with exponent q , let

$$\psi: \mathbb{Z}_q \rightarrow \mathbb{G}_T$$

be a homomorphism. The prover now claims to know the preimages for some k -subset of a set of n public group elements $y_1, \dots, y_n \in \mathbb{G}_T$. We aim to construct an interactive proof for convincing a verifier of the veracity of this claim, without revealing the preimages or the k -subset. More precisely, we aim to construct an interactive proof for relation

$$\mathfrak{R}_{k\text{-out-of-}n} = \{ (y_1, \dots, y_n; S, \mathbf{x}) : |S| = k, y_i = \psi(x_i) \ \forall i \in S \subseteq \{1, \dots, n\} \}.$$

Note that, for notational convenience, the secret $\mathbf{x}$ is defined as a vector in $\mathbb{Z}_q^n$, while only the k coefficients $(x_i)_{i \in S}$ are relevant in this relation. Further, for simplicity we assume the domain of the homomorphism Ψ to be $\mathbb{Z}_q$. Our techniques are easily generalized to arbitrary domains $\mathbb{G}$. However, this would require a vector commitment scheme for *mixed* vectors with coefficients in both $\mathbb{Z}_q$ and $\mathbb{G}$, such as the commitment schemes presented in Section 5.3.

Inspired by the design principle of [CDS94], we reduce the k -out-of- n scenario to the n -out-of- n scenario by having the prover “eliminate” the preimages $(x_i)_{i \notin S}$ that it does not know, and then we apply an amortized compressed Σ -protocol to prove the n instances in one go. However, the original solution of [CDS94] to reduce the k -out-of- n to the n -out-of- n scenario, achieved by secret sharing the challenge, does not work for us, as the resulting protocol is not in the shape for the compression mechanism to apply. More precisely, the third message in the [CDS94]-protocol includes a consistent secret sharing of the challenge, which cannot be compressed.

Instead, we use the following solution. The prover first chooses an $(n - k + 1)$ -out-of- n Shamir secret sharing $(s_1 = p(1), \dots, s_n = p(n)) \in \mathbb{Z}_q^n$ of the default secret $s = 1$, where it selects the non-constant “random” coefficients $a_1, \dots, a_{n-k}$ of the sharing polynomial

$$p(X) = 1 + a_1X + \dots + a_{n-k}X^{n-k} \in \mathbb{Z}_q[X]$$

such that $s_i = 0$ for $i \notin S$. Hence, $p(X)$ is the unique polynomial of degree at most $n - k$ such that $p(0) = 1$ and $p(i) = 0$ for all $i \notin S$.

Now let $t_i = s_i x_i$ for any i , i.e., $t_i = 0$ for all $i \notin S$. The prover then commits to the vector

$$(\mathbf{a}, \mathbf{t}) = (a_1, \dots, a_{n-k}, t_1, \dots, t_n) \in \mathbb{Z}_q^{2n-k}$$

in a single compact commitment $P = \text{COM}(\mathbf{x}; \gamma)$. We assume the commitment scheme $\text{COM}: \mathbb{Z}_q^{2n-k} \times \text{Rand} \rightarrow \mathbb{H}$ to be homomorphic.

What remains is for the prover to show that

$$\psi(t_i) = s_i y_i \tag{4.1}$$

for all $i \in \{1, \dots, n\}$. Recall that $s_i = p(i) = 1 + \sum_{j=1}^{n-k} a_j i^j$. Thus, Equation 4.1 can be rewritten as

$$\phi_i(\mathbf{a}, \mathbf{t}) := \psi(t_i) - y_i \cdot \sum_{j=1}^{n-k} a_j i^j = y_i,$$

where the left hand side is a group homomorphism $\phi_i: \mathbb{Z}_q^{2n-k} \rightarrow \mathbb{G}_T$ evaluated in the committed vector $(\mathbf{a}, \mathbf{t})$. Hence, proving knowledge of an opening of commitment P that satisfies Equation 4.1 for all $1 \leq i \leq n$, is reduced to proving knowledge of a Ψ -preimage of $(P, y_1, \dots, y_n)$, where

$$\Psi: \mathbb{Z}_q^{2n-k} \times \mathbf{Rand} \rightarrow \mathbb{H} \times \mathbb{G}_T^n, \quad (\mathbf{a}, \mathbf{t}; \gamma) \mapsto (\text{COM}(\mathbf{a}, \mathbf{t}; \gamma), \phi_1(\mathbf{a}, \mathbf{t}), \dots, \phi_n(\mathbf{a}, \mathbf{t})).$$

In other words, in the final step of the k -out-of- n proof of partial knowledge, the prover opens n homomorphisms ϕ_i on the compactly committed vector $(\mathbf{a}, \mathbf{t})$.

For efficiency reasons, the costs of opening these n homomorphism can be amortized. More precisely, instead of opening the homomorphisms ϕ_i individually, the prover opens a single homomorphism $\Phi_c = \sum_{i=1}^n c^{i-1} \phi_i$ for a challenge $c \leftarrow_R \mathbb{Z}_q$ sampled uniformly at random by the verifier. This approach is a minor adaptation of the amortization technique presented in Section 3.4.2. The difference is that here the coefficients of the committed vector $(\mathbf{a}, \mathbf{t}) \in \mathbb{Z}_q^t$ are of a different type than the homomorphism openings $y_1, \dots, y_n \in \mathbb{G}_T$. For this reason, the evaluations y_i can not be incorporated into the commitment.

Opening the homomorphism Φ_c with a standard Σ -protocol gives a novel secret-sharing based realization of [CDS94], with linear communication complexity. However, in contrast to the original [CDS94]-approach, this novel realization is now amenable to the compression techniques of Chapter 3, allowing us to reduce the communication complexity from linear down to logarithmic.

The resulting interactive proof, denoted $\Pi_{k\text{-out-of-}n}$, is formalized in Protocol 13. Its main properties are summarized in Theorem 4.1. In particular, note that the communication costs are logarithmic in both k and n . In this theorem, we minimize the communication costs by applying the compression mechanism $\log_2(2n - k) - 2$ times to reduce the dimension of the secret vector $(\mathbf{a}, \mathbf{t})$ from $2n - k$ down to 4. Namely, since every factor two reduction of the dimension comes at the cost of sending two $\mathbb{H}$ -elements and two $\mathbb{G}_T$ -elements, it is suboptimal to continue further and reduce the dimension of the witness down to two or even one. This also means that we implicitly assume that $n \geq 4$.

Theorem 4.1 (*k -out-of- n Proof of Partial Knowledge*). *Let q be a prime and $k, n, \mu \in \mathbb{N}$ such that $k \leq n$ and $2n - k = 2^\mu$. Further, let $\psi: \mathbb{Z}_q \rightarrow \mathbb{G}_T$ be a homomorphism and $\text{COM}: \mathbb{Z}_q^n \times \mathbf{Rand} \rightarrow \mathbb{H}$ a homomorphic vector commitment scheme.*

Then the compressed Σ -protocol $\Pi_{k\text{-out-of-}n}$ for relation $\mathfrak{R}_{k\text{-out-of-}n}$, described in Protocol 13, is perfectly complete, $(n, 2, 3, \dots, 3)$ -out-of- $(q, \dots, q)$ special-sound, under assumption that the commitment scheme is binding, and special honest-verifier zero-knowledge (SHVZK), under the assumption that the commitment scheme is hiding. Moreover, it has $2\mu + 1$ communication rounds and the communication costs are:

- $\mathcal{P} \rightarrow \mathcal{V}$: 4 elements of $\mathbb{Z}_q$, $2\mu - 3$ elements of $\mathbb{G}_T$, $2\mu - 2$ elements of $\mathbb{H}$ and 1 element of $\mathbf{Rand}$;
- $\mathcal{V} \rightarrow \mathcal{P}$: μ elements of $\mathbb{Z}_q$.

Protocol 13 k -out-of- n Proof of Partial Knowledge $\Pi_{k\text{-out-of-}n}$.

PARAMETERS:	$k, n \in \mathbb{N}$, prime q , groups $(\mathbb{G}_T, +)$ and $(\mathbb{H}, \cdot)$ with exponent q , $\psi \in \text{Hom}(\mathbb{Z}_q, \mathbb{G}_T)$ and COM: $\mathbb{Z}_q^{2n-k} \times \mathbf{Rand} \rightarrow \mathbb{H}$ (homomorphic)
PUBLIC INPUT:	$y_1, \dots, y_n \in \mathbb{G}_T$
PROVER'S PRIVATE INPUT:	$S \subseteq \{1, \dots, n\}$ with $ S = k$, $x_1, \dots, x_n \in \mathbb{Z}_q^n$
PROVER'S CLAIM:	$\psi(x_i) = y_i$ for all $i \in S$

Prover $\mathcal{P}$ Verifier $\mathcal{V}$

$$p(X) = 1 + \sum_{i=1}^{n-k} a_i X^i$$

$$\text{s.t. } p(i) = 0 \quad \forall i \notin S$$

$$\mathbf{t} = (p(1)x_1, \dots, p(n)x_n)$$

$$\gamma \leftarrow_R \mathbf{Rand}$$

$$P = \text{COM}(\mathbf{a}, \mathbf{t}; \gamma)$$

$$\xrightarrow{P}$$

$$\xleftarrow{c}$$

$$c \leftarrow_R \mathbb{Z}_q$$

Run the compressed Σ -protocol Σ_{comp} of Section 3.2.3 to prove knowledge of a preimage of $(P, \sum_{i=1}^n c^{i-1} y_i)$ with respect to homomorphism

$$\Psi: \mathbb{Z}_q^{2n-k} \times \mathbf{Rand} \rightarrow \mathbb{H} \times \mathbb{G}_T, \quad (\mathbf{a}, \mathbf{t}; \gamma) \mapsto (\text{COM}(\mathbf{a}, \mathbf{t}; \gamma), \sum_{i=1}^n c^{i-1} \phi_i(\mathbf{a}, \mathbf{t})),$$

where $\phi_i(\mathbf{a}, \mathbf{t}) := \psi(t_i) - y_i \cdot \sum_{j=1}^{n-k} a_j i^j$ for all $1 \leq i \leq n$.

Proof. Completeness: This property follows from the completeness of the compressed Σ -protocol Σ_{comp} .

SHVZK: This property follows from the fact that the commitment P is hiding and from the corresponding zero-knowledge property of Σ_{comp} .

Special-Soundness: Similar to the proof of Theorem 3.12 it follows that, under the assumption that the commitment scheme is binding, there exists an extractor that, on input an $(n, 2, 3, \dots, 3)$ -tree of accepting transcripts, outputs an opening $(\mathbf{a}, \mathbf{t}; \gamma)$ to the commitment P such that $\phi_i(\mathbf{a}, \mathbf{t}) := y_i$ for all $1 \leq i \leq n$.

Let $p(X) = 1 + \sum_{j=1}^{n-k} a_j X^j$. Then, $\phi_i(\mathbf{a}, \mathbf{t}) := y_i$ can be rewritten as $\psi(t_i) = p(i)y_i$. Given the bounded degree of p and the non-zero constant coefficient, $p(i) = 0$ for at most $n-k$ choices of $i \in \{1, \dots, n\}$. Thus, setting $S = \{i : p(i) \neq 0\}$, we have $|S| \geq k$, and for any $i \in S$ we can set $x_i := t_i/p(i)$. This then implies that $\psi(x_i) = y_i$ for all $i \in S$, which completes the proof. $\square$

Example 4.1 (Discrete Logarithm Instantiations). Taking $\psi: \mathbb{Z}_q \rightarrow \mathbb{H}$, $x \mapsto h^x$ allows one to prove knowledge of the discrete logarithms of a k -subset of public group elements $P_1, \dots, P_n \in (\mathbb{H}, \cdot)$. Moreover, it is easily seen that the proofs of partial knowledge immediately generalize to homomorphism $\psi: \mathbb{Z}_q^s \rightarrow \mathbb{G}_T$ with arbitrary input dimension s . This observation allows one to instantiate ψ as the Pedersen (vector) commitment function and prove knowledge of k -out-of- n commitment openings.

Remark 4.2. Similar to the linearization technique of Section 4.2, the proof of partial knowledge deploys Shamir's linear secret-sharing scheme (LSSS). However, the linearization technique for multiplication triples crucially depends on the multiplicativity of the LSSS. By contrast, the k -out-of- n proof of partial knowledge does not require multiplicativity and can be instantiated with any n player linear secret-sharing scheme that has $(n-k+1)$ -reconstruction and $(n-k)$ -privacy.

4.3.2 Pairing-Based Reduction of the Communication Costs

The amortized communication costs for opening the homomorphisms $\phi_i: \mathbb{Z}_q^{2n-k} \rightarrow \mathbb{G}_T$ are roughly $4 \log_2(2n-k)$ elements. This is approximately a factor two larger than the communication costs for opening n linear forms $L_i: \mathbb{Z}_q^{2n-k} \rightarrow \mathbb{Z}_q$ (Section 3.4.2). The reason is that, for a linear form, the input and output coefficients are of the same type; both are $\mathbb{Z}_q$ elements. Therefore, using the techniques of Section 3.4.2, the linear form evaluations can be “incorporated” into the commitment. More precisely, opening n linear forms L_i on a compact commitment $P = \text{COM}(\mathbf{x}; \gamma)$ can be reduced to proving knowledge of a preimage for the homomorphism

$$\Psi_c: \mathbb{Z}_q^{2n-k} \times \text{Rand} \rightarrow \mathbb{H}, \quad (\mathbf{x}; \gamma) \mapsto \text{COM}\left(\mathbf{x}, \sum_{i=1}^n c^i L_i(\mathbf{x}); \gamma\right),$$

where $c \leftarrow_R \mathbb{Z}_q$ is a challenge sampled uniformly at random by the verifier. Applying the same technique for the homomorphisms $\phi_i: \mathbb{Z}_q^{2n-k} \rightarrow \mathbb{G}_T$ requires a compact commitment scheme for *mixed* vectors $(\mathbf{x}, \sum_{i=1}^n c^i L_i(\mathbf{x})) \in \mathbb{Z}_q^{2n-k} \times \mathbb{G}_T$ containing both $\mathbb{Z}_q$ and $\mathbb{G}_T$ coefficients. In some settings, e.g., when proving knowledge of k -out-of- n discrete logarithms or Pedersen commitment openings, pairing-based commitment schemes with the required properties exist (Section 5.3). These commitment schemes allow the communication costs of the corresponding k -out-of- n proof of partial knowledge protocol to be reduced with a factor two, down to roughly $2 \log_2(2n-k)$ elements. For more details we refer to [ACF21].

4.3.3 General Access Structures

Thus far, we have restricted ourselves to provers that claim to know the preimages of some (secret) subset S , of cardinality at least k , of n (public) elements $P_1, \dots, P_n$, i.e., the secret subset S is an element of a *threshold* access structure

$$\Gamma_{k,n} = \{A \subseteq \{1, \dots, n\} : |A| \geq k\} \subseteq 2^{\{1, \dots, n\}}.$$

Here, we describe how the proof of partial knowledge can easily be generalized to arbitrary monotone access structures $\Gamma \subseteq 2^{\{1, \dots, n\}}$, i.e., to provers that claim to know the preimages of some subset of $S \in \Gamma$ of n public elements. Recall that Γ is called a monotone access structure if for all $A \in \Gamma$ and for all $B \subseteq 2^{\{1, \dots, n\}}$ with $A \subseteq B$ it holds that $B \in \Gamma$. The proofs of partial knowledge of [CDS94] already considered arbitrary access structures and we adapt their techniques by combining them with our compression framework.

Our k -out-of- n proofs of partial knowledge implicitly deploy a linear secret-sharing scheme (LSSS) for access structure $\Gamma_{k,n}^* = \Gamma_{n-k,n}$. Here, Γ^* denotes the *dual* of access structure Γ , generally given by

$$\Gamma^* = \{A \subseteq \{1, \dots, n\} : A^c \notin \Gamma\}.$$

More concretely the protocol of Section 4.3.1 uses Shamir's secret-sharing scheme and the polynomial $p(X) = 1 + \sum_{j=1}^{n-k} a_j X^j$ defines a secret sharing of the field element 1.

Now let Γ be a monotone access structure and $\mathcal{S}$ an LSSS for sharing field elements for access structure Γ^* . This implies that the adversary structure of $\mathcal{S}$ equals $\{S : S \notin \Gamma^*\}$, i.e., all player subsets are either qualified or unqualified [CDN15]. Depending on the access structure Γ^* , it might be required that shares are allowed to consist of several field elements.

Then, to construct a proof of partial knowledge for Γ , we simply replace $p(i)$ by the i -th share of a secret sharing of 1, with the randomness chosen so that the “right” shares (i.e., those corresponding to the x_i 's that the prover does not know) vanish. Since the adversary structure of $\mathcal{S}$ equals $\{S : S \notin \Gamma^*\}$, the randomness can always be chosen such that the appropriate shares vanish, showing completeness of the generalized proof of partial knowledge. Special-soundness follows from the following observation. Let $A \subseteq \{1, \dots, n\}$ be the subset for which all the corresponding shares vanish. Then, by linearity of the secret-sharing scheme and since the secret sharing reconstructs to 1, it follows that $A \notin \Gamma^*$. Hence, $A^c \in \Gamma$ and special-soundness follows as before.

The communication complexity of the resulting protocol depends logarithmically on the size of the LSSS for Γ^* , which is given by the monotone-span-program complexity of Γ^* [SJM91] and which coincides with the monotone-span-program complexity of Γ [Gál95].

