

# *aPlonK*: Aggregated *PlonK* from Multi-Polynomial Commitment Schemes

Miguel Ambrona, Marc Beunardeau, Anne-Laure Schmitt, and Raphaël R. Toledo

Nomadic Labs, Paris, France  
name.surname@nomadic-labs.com

**Abstract.** *PlonK* is a prominent universal and updatable zk-SNARK for general circuit satisfiability. We present *aPlonK*, a variant of *PlonK* that reduces the proof size and verification time when multiple statements are proven in a batch. Both the aggregated proof size and the verification complexity of *aPlonK* are logarithmic in the number of aggregated statements. Our main building block, inspired by the techniques developed in SnarkPack (Gailly, Maller, Nitulescu, FC 2022), is a multi-polynomial commitment scheme, a new primitive that generalizes polynomial commitment schemes. Our techniques also include a mechanism for involving committed data into *PlonK* statements very efficiently, which can be of independent interest.

We also implement an open-source industrial-grade library for zero-knowledge *PlonK* proofs with support for *aPlonK*. Our experimental results show that our techniques are suitable for real-world applications (such as blockchain rollups), achieving significant performance improvements in proof size and verification time.

## 1 Introduction

In 1985 [GMR85], Goldwasser, Micali and Rackoff introduced the notion of zero-knowledge arguments. They allow a *prover* to convince a *verifier* of the validity of a certain statement without revealing any other information, e.g. why the statement is true. A few years later, Blum, Feldman and Micali [BFM88], extended this notion and considered *non-interactive* zero-knowledge arguments (NIZK), where the communication between the two parties is unilateral: the prover produces a “certificate” that can be verified by everyone.

Existing generic protocols that implement zero-knowledge argument systems [BFM88, DMP90, FLS90] for any NP relation have been perceived as mainly theoretical results for many years: they used to involve expensive NP reductions and repetitions of the same routine in order to achieve reasonable soundness. Only special purpose protocols for specific NP languages [Sch91, Cra97, CDS94, GS08] were considered efficient enough for practical deployment, and they have been widely used for building digital signatures and anonymous credentials.

Recently, the research community has witnessed significant improvements on the design of efficient general-purpose zero-knowledge proof systems which offer various degrees of practicality [Gro16, BCG<sup>+</sup>17, BBB<sup>+</sup>18, MBKM19, LMR19]. Such improvements have been driven by the increasing development of blockchain systems that make use of zero-knowledge arguments to achieve privacy and scalability [BCG<sup>+</sup>14, DFKP13]. In these systems, communication complexity is one of the most important performance factors, which has led to an increasing interest and remarkable progress in so-called *succinct non-interactive arguments of knowledge* (SNARKs) [GGPR13, BCG<sup>+</sup>13, PHGR13, Gro16], a class of non-interactive arguments of knowledge with sublinear (if not constant) communication and verification complexity. This comes at the cost of a significantly slower prover, compared to other zero-knowledge proof systems with higher communication complexity [JKO13, GMO16, CGM16].

*PlonK*, which stands for *Permutations over Lagrange-bases for Oecumenical Non-interactive arguments of Knowledge*, is a universal and updatable zero-knowledge SNARK for general circuit satisfiability. Given its significant improvements with respect to its predecessor Sonic [MBKM19], especially on prover efficiency, *PlonK* has become very popular and has been adopted by several state-of-the-art blockchain projects such as Zcash [HBHW], Mina [BMRS20], the Dusk Network [MKF21] or Anoma [GYB21].

In this work we present *aPlonK*, a new affluent of the *PlonK* family which focuses on reducing the proof size and verification time when multiple statements are proven in a batch. The aggregated proof size and the verification complexity of *aPlonK* are logarithmic in the number of aggregated statements, making it an appealing building block for many blockchain applications, where having low verification complexity is paramount.

## 1.1 Blockchain applications of SNARKs

Blockchain developers were among the first to deploy large-scale real-life applications of general purpose zero-knowledge proof systems, starting different lines of research in this area. We can cite Virgo [ZXZS20] and its successors, used for Overeality [Ove22]; STARK [BBHR19], used by Starkware [Sta21]; Halo [BGH19a], used by Zcash [HBHW]; or *PlonK* [GWC19], created by Aztec [Wil18].

*Privacy.* Blockchains revolve around the property of public verifiability, since anyone must be able to verify the transition between successive blockchain states. As such, all the information on a blockchain must be public, which makes it difficult to support privacy-friendly applications.

Using zk-SNARKs is a natural approach to keep public verifiability while maintaining privacy. This was first illustrated theoretically [MGGR13, DFKP13, BCG<sup>+</sup>14] but also in practice [HBHW]. These systems leverage zero-knowledge proofs to allow users to generate private transactions that hide the sender, the recipient, and (potentially) the transferred amount. Although zero-knowledge proofs are practical, they incur a considerable overhead on both the prover (the user) and the verifier (the blockchain). Note that a SNARK proof verification typically involves heavier computations than a simple signature verification.

*Scalability and validity rollups.* Scalability is an inherent issue in blockchain systems. The blockchain throughput cannot be simply increased with additional computing power, since every node should be able to validate state transitions. SNARKs can be of help here when they are used to certify expensive computations (e.g., the validity of multiple transactions), since the SNARK verification can become cheaper than the direct validation of the statement being proven. This idea has been explored in so-called validity rollups. (Note that in this context, zero-knowledge is not necessarily relevant.)

A validity rollup is an alternative chain that runs in parallel to the main chain, but stores a small amount of data on the main chain, e.g., a commitment to the rollup state. Transactions can be sent to a rollup *operator*, who knows the exhaustive rollup state and can update it accordingly. Periodically, the rollup operator will communicate to the main chain a commitment to the most updated version of the rollup state together with a proof that ensures its validity (ergo the name).<sup>1</sup> The commitment to the new rollup state and such proof are published on the main blockchain and the nodes only need to check this single proof (instead of validating all the operations performed between rollup states).<sup>2</sup> The blockchain (layer 1) becomes more scalable at the cost of having to produce such proof, which is generated by an independent operator (in layer 2). Unlike in layer 1, the operator can make use of extra computing power and parallelization to speedup the process of creating proofs, thus reducing the rollup latency.

Despite such promising properties and even if the rollup operator can use large computing power, producing proofs is a major bottleneck. A possible idea to reduce the proving cost (and thus the rollup latency) is to split the statement into smaller ones. For example, instead of proving the validity of 10,000 rollup transactions with one proof, one could produce 100 proofs of 100 transactions each. Dealing with smaller proofs can significantly simplify the prover cost, whose complexity is linearithmic in the circuit size. Unfortunately, this would require that the blockchain nodes receive and verify 100 proofs instead of 1.

Our techniques in this work are particularly suitable for the above scenario. They allow the prover to combine the batch proofs, producing an aggregated proof that can be verified very efficiently by the blockchain nodes. An alternative solution would be to use incrementally verifiable computation (IVC) [Val08]. We discuss the differences between these two approaches in Section 1.3.

<sup>1</sup>A proof that the new committed state has been achieved by applying legitimate operations to the previous committed state.

<sup>2</sup>Remarkably, the blockchain nodes do not even need access to the rollup operations that were involved.

*Privacy-preserving rollups.* To make rollups privacy-friendly, we can leverage the zero-knowledge property of SNARKs. For example, by having users create zero-knowledge proofs which are then aggregated by the rollup operator, their private information could remain hidden. Our techniques are also applicable to this scenario but require some coordination between users. We need the users to synchronize a few times during the proving process to achieve correctness, as all parties must use the same Fiat-Shamir randomness (see Section 4.1). On the other hand, our distributed version of *aPlonK* can be adapted to prevent DoS attacks: if a user aborts the execution of their proof, or misbehaves, the aggregation of the rest of proofs can still be completed. Again, recursion and IVC are an alternative for implementing privacy-preserving rollups (see Section 1.3).

## 1.2 Our contributions

We pursue the study of the *PlonK* proving system and establish several general techniques that reduce the proof size and verification time when multiple statements are proven in a batch.

*aPlonK.* Our main contribution is a multi-statement proving system coined *aggregated PlonK* or *aPlonK* for short, which allows one to combine  $k$  proofs into a single aggregated proof of  $\mathcal{O}(\log k)$  size that can be verified in  $\mathcal{O}(\log k)$  time. The aggregated proofs must be created coordinately, but their computation is highly parallelizable. *aPlonK* is the result of extending the techniques of Gailly, Maller and Nitulescu (SnarkPack) [GMN20], designed over Groth16 [Gro16], to the framework of *PlonK*.

*Multi-polynomial commitments.* We introduce the notion of multi-polynomial commitment schemes, a generalization of polynomial commitment schemes designed to commit to several polynomials at the same time, while achieving sublinear commitment and proof sizes and sublinear verification complexity in the number of committed polynomials.

We then present a generic construction of a multi-polynomial commitment scheme from any homomorphic polynomial commitment scheme whose commitment space is one of the source groups of a set of bilinear groups. Our construction is inspired by the techniques of SnarkPack for building an inner-product argument with logarithmic verification time by combining a modified version of the inner-product argument [BBB+18, BGH19b, BCL+21, DRZ20] with a KZG-like [KZG10] commitment scheme whose commitment space is the target group of a set of bilinear groups.

Our new notion of multi-polynomial commitments captures the essence of SnarkPack, hardcoded in their *ad hoc* construction for aggregating Groth16 proofs. We consider this an important contribution as it provides clarity, intuition and continues the modularity of *PlonK*-based systems.

*Improvements over SnarkPack.* While the verification of SnarkPack is presented as sublinear<sup>3</sup>, their verifier needs to perform a linear number of scalar operations for dealing with public inputs. (This is inherent for verifiable computations.) We observe that for many applications (e.g. a validity rollup) most public inputs can be hidden from the verifier as long as some relation on them is ensured (e.g., they form a chain). Our system can exploit this fact, to achieve actual sublinear verification time, when the use case allows for it.

Furthermore, we double the efficiency of the main subroutine of SnarkPack by observing that their *pair group commitments* [GMN20, Section 3.2] do not need to be binding in order to achieve the desired security properties, if the underlying polynomial commitment scheme is *inner-product binding* and *inner-product extractable* (see Section 3).

*Commitments in PlonK relations.* En route, we present a mechanism that allows a *PlonK* statement to refer to the data inside a public commitment. Such link does not require a high number of constraints to model the commitment opening, as it is performed *outside of the PlonK circuit*. This building block, necessary to instantiate *aPlonK* efficiently, can be of independent interest, as it can be used for building hybrid proving systems or for proving statements modeled with non-deterministic circuits (see Section 4.2).

---

<sup>3</sup>It is in terms of elliptic curve operations.

*Implementation and evaluation.* We implement a general library for (zero-knowledge) *PlonK* proofs with support for *aPlonK*. Our library is implemented over the BLS12-381 elliptic curve [Bow17] and uses bindings to the *blst* library [Sup21]. Our experiments show that the techniques described in this work are suitable for real-world applications, providing significant performance improvements in proof size and verification time, while introducing a light overhead on prover complexity. Our code is publicly available as open-source [Nom22].

### 1.3 Related work

In this section, we compare our techniques with other approaches for combining zero-knowledge proofs and present the main advantages of *aPlonK*.

**IVC and recursion.** Incrementally verifiable computation (IVC) [Val08], conceived by Valiant, is a framework that provides proof composability: with IVC one can conjunctively combine two proofs of size  $k$  into a proof of size  $k$  as well. This is a powerful technique that can be used to implement recursion. In the context of SNARKS, recursion allows one to prove statements like the following (parametrized by a state):

*“I know a previous state from which the current state can be reached and I also have a proof of this very statement for such previous state.”*

This can be achieved by expressing a SNARK verifier in a SNARK circuit. One real-world application of this technique is the Mina blockchain [BMRS20], which provides its user with a constant-size proof of validity of its most updated state. In particular, the proof ensures that one transition of the blockchain has been performed correctly and that there exists another proof for the preceding state. This allows the blockchain state to be constant.

Incidentally, recursion can also be used to aggregate proofs together by proving that one has seen valid proofs. This allows for natural parallelization by splitting a complex computation into smaller ones that are then aggregated.

However, the strength of recursive SNARKs comes with high costs. Expressing a SNARK verifier in a SNARK circuit is very expensive. The current known techniques are (i) using cycles of pairing friendly elliptic curves [CCW19] which require very big group elements, (ii) or implementing non-native operations such as modular arithmetic over a modulus (e.g. the SNARK’s base field order) that does not coincide with the SNARK’s scalar field order. This typically leads to a decrease in performance of several orders of magnitude.

This performance issue has led to new lines of research exploring alternative techniques for achieving weaker versions of IVC. We can cite, Halo [BGH19a], and its successor Halo2 (which uses *PlonK* instead of Sonic [MBKM19]), Fractal [COS19], Bünz et al. work [BCMS20], or Nova [KST21], a novel construction based on folding schemes. These works explore the idea of performing a weaker version of recursion by not modeling some expensive parts of the SNARK verification in the circuit. These excluded verification steps can be accumulated and carried out for future verification. These techniques achieve IVC by using a cycle of (not necessarily pairing-friendly) elliptic curves, leading to better performance.

**Proof aggregation without recursion.** Proof aggregation can be achieved more efficiently without recursion and still be suitable for many applications such as validity rollups.

*Aztec.* The company Aztec [Wil18], creator of *PlonK*, achieves a form of proof aggregation which can be seen as a weak version of IVC. Thanks to this simplification, they do not require cycles of elliptic curves. However, they still need to model elliptic curves in a SNARK circuit, which involves simulating non-native field operations. The expensive pairing checks are accumulated as in Halo, by using standard batching techniques.

*SnarkPack*. Gailly, Maller and Nitulescu [GMN20] provide a framework for aggregating Groth16 proofs. As we explained in Section 1.2, their techniques are the starting point of this paper and combine a homomorphic pair group commitment schemes with an inner-product argument to achieve logarithmic-size proofs and logarithmic verification complexity (in the number of aggregated proofs).

Our work achieves very efficient proof aggregation without cycles of elliptic curves and without simulating non-native operations. This is an improvement over Halo and Aztec, which brings us at the level of SnarkPack. However, unlike SnarkPack, *aPlonK* is defined over a universal SNARK. An immediate consequence is that we can aggregate different circuits. Furthermore, we can perform proof aggregation that connects the proven statements in an arbitrary fashion (e.g. Section 4.3).

#### 1.4 Technical overview

In a nutshell, a *PlonK* proof consists of a set of commitments to secret polynomials together with evaluations of such polynomials at a random point (sampled after the polynomials have been committed). A *PlonK* verifier simply checks that the evaluations are valid with respect to the corresponding polynomial commitments and that they satisfy a series of equations.

In the multi-statement setting, in order to achieve sublinear verification time in the number of aggregated proofs, the verifier will need to delegate some computations to the prover and independently verify that they were performed honestly. As we will see, such verification can be performed significantly faster than the delegated computation. In our construction, such delegation occurs twice: (i) a multi-polynomial commitment scheme is used to achieve sublinear commitments size and sublinear verification complexity on checking the commitment evaluations; (ii) we use *meta-verification* (which we describe below in more detail) to achieve constant verification complexity on checking the evaluation equations.

**Multi-polynomial commitments.** The main challenge of building a multi-polynomial commitment scheme is achieving sublinear commitment size (and sublinear verification) in the number of committed polynomials. We follow the techniques of Gailly, Maller and Nitulescu [GMN20], and start from the KZG polynomial commitment scheme [KZG10] defined over a set of bilinear groups  $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_t)$  of prime order  $p$  equipped with a pairing  $e$ . We use the implicit notation  $[x]_i$  to denote  $xG_i \in \mathbb{G}_i$  for every  $i \in \{1, 2, t\}$ , where  $G_i$  is the designated generator of  $\mathbb{G}_i$  and  $x \in \mathbb{Z}_p$ . The KZG commitment scheme is defined as follows:

- The commitment key is  $\text{ck} := [s, \dots, s^{d-1}]_1$  for a uniformly sampled  $s \in \mathbb{Z}_p$ .
- The verification key is  $\text{vk} := [s]_2$ .
- A polynomial  $f \in \mathbb{Z}_p^{<d}[X]$  is committed as  $\mu := [f(s)]_1$ , using  $\text{ck}$ .
- A proof that  $f(z) = v$  is  $\pi := [h(s)]_1$ , where  $h(X) := (f(X) - f(z))/(X - z)$ .
- Verification that  $f(z) = v$  is done by checking  $e(\mu - [v]_1, [1]_2) = e(\pi, \text{vk} - [z]_2)$ .

Observe that such scheme is homomorphic in the sense that if  $\mu_1$  and  $\mu_2$  are commitments to polynomials  $f_1$  and  $f_2$  respectively, then  $\mu_1 + \mu_2$  is a commitment to polynomial  $f_1 + f_2$ .

The homomorphic property allows for the following optimization when verifying that several commitments  $\mu_1, \dots, \mu_k$  evaluate to claimed evaluations  $v_1, \dots, v_k$  on a common point  $z \in \mathbb{Z}_p$ . First, compute a random linear combination of the commitments,  $\hat{\mu} := \sum_i r^i \mu_i$  for a uniformly sampled  $r \in \mathbb{Z}_p$ . Then, verify that  $\hat{\mu}$  opens to  $\hat{v} := \sum_i r^i v_i$  on  $z$ . This trick allows the verifier to only check one pairing equation instead of  $k$ , at the cost of a negligible statistical error. Indeed, it could occur that the aggregated commitment opens to the aggregated evaluation, whereas some of the commitments  $\mu_i$  do not open to the claimed  $v_i$  on  $z$ , but the probability of this event can be upper-bounded by  $k/p$  if  $r$  is chosen uniformly and independently.

*Delegating the computation of  $\hat{\mu}$ .* In order to achieve sublinear commitment size, the prover will *commit to the commitments*  $\mu_1, \dots, \mu_k$ . A value  $r$  will be sampled after this meta-commitment, either by the verifier (in an interactive protocol) or through the Fiat-Shamir heuristic. The computation of the aggregated commitment

$\hat{\mu}$  will be delegated to the prover, who will include a proof ensuring that such computation is correct with respect to the meta-commitment.

What would be a suitable scheme for committing to the commitments? The computation that we need to assert on  $\hat{\mu}$  can be seen as a polynomial evaluation on  $r$  of a polynomial whose “coefficients” are  $\mu_i$ . Thus, a good candidate for our meta-commitment scheme is again a polynomial commitment scheme. This is precisely what Gailly et al. suggest in [GMN20]. We can use a variant of the KZG commitment scheme whose committing space is  $\mathbb{G}_t$ , by committing to  $\mu_1, \dots, \mu_k$  as  $M := \sum_i e(\mu_i, [\tau^i]_2)$ , where  $\tau$  is a new SRS secret, independent of  $s$ .

A difficulty arises: how can we generalize the KZG proof of opening strategy in that case? If we define  $f(X) := \sum_i \text{dlog}(\mu_i) X^i$ , we could provide the verifier with  $f(r)$  and  $\pi := [h(\tau)]_2$ , where  $h(X)$  is defined as  $(f(X) - v)/(X - r)$ . The verifier would then check that  $M - [f(r)]_t = e(\pi, [\tau]_1 - [r]_1)$  and that  $\hat{\mu} = [f(r)]_1$ , which equals  $\sum_i r^i \mu_i$ , as desired. Unfortunately, this method requires explicitly knowing the coefficients of polynomial  $f$ . Given that group elements  $\mu_i$  are the result of committing to certain non-constant polynomial, their discrete logarithm will not be known to the prover.

A possible solution [GMN20] is to implement the opening of commitment  $M$  at  $r$  via an inner-product argument [BBB<sup>+</sup>18]. In particular, a modified version similar to those in [BGH19b, BCL<sup>+</sup>21], that we describe in detail in Figure 3, adjusted to support relation  $\text{PoK}\{\boldsymbol{\mu} : \langle \boldsymbol{\mu}, \boldsymbol{\tau} \rangle = M \wedge \langle \boldsymbol{r}, \boldsymbol{\mu} \rangle = \hat{\mu}\}$ . However, inner-product arguments are known to have linear verification. More concretely, the verification complexity is logarithmic in  $k$  except for one final check that a certain  $M' \in \mathbb{G}_t$  corresponds to the commitment of a polynomial  $g(x) := \prod_{j=1}^{\kappa} (u_j^{-1} + u_j X^{2^{\kappa-j}})$ , for some known coefficients  $u_j$ , where  $\kappa = \lceil \log_2(k) \rceil$ . Polynomial  $g$ , given its nice factored form, can be evaluated in logarithmic time. This opens the possibility of, instead of performing the (expensive) linear check that  $M'$  is the commitment to  $g$ , verifying a proof of opening of  $M'$  at a random point  $\rho \in \mathbb{Z}_p$ , and checking that it opens to  $g(\rho)$ . This can be done precisely as we described in the previous paragraph. Intuitively, the inner-product argument has allowed us to replace the KZG-like proof of opening of unknown polynomial  $f$ , by a KZG-like proof of opening of a *known* polynomial  $g$  (at the cost of some other logarithmic complexity checks).

*Remark 1.* The meta-committing function  $\boldsymbol{\mu} \mapsto \sum_i e(\mu_i, [\tau^i]_2)$  is not binding.<sup>4</sup> This is because  $\tau$  is also available in  $\mathbb{G}_1$ , which is necessary for the KZG-like verification. The authors of [GMN20] suggest to make the commitment binding by computing it twice with respect to two independent structured reference string, namely:  $\boldsymbol{\mu} \mapsto (\sum_i e(\mu_i, [\tau^i]_2), \sum_i e(\mu_i, [\hat{\tau}^i]_2))$ .

Interestingly, we show that such duplication is not strictly necessary if the underlying polynomial commitment scheme satisfies two additional properties which we coin the *inner-product binding property* and *inner-product extractability* (see Section 3.1). We then show that the KZG polynomial commitment scheme satisfies both (Lemmas 2 and 3). This observation reduces the number of  $\mathbb{G}_t$  elements and  $\mathbb{G}_t$  operations involved in the aggregated proofs and the inner-product argument by a factor of 2 compared to the protocol from [GMN20].

*Achieving sublinear verification complexity.* The above techniques allowed us to delegate the computation of  $\hat{\mu}$ . In order to get a complete multi-polynomial commitment scheme, we need to design a sublinear verification algorithm that takes as input a commitment to the evaluations instead of the evaluations themselves. This can be achieved generically through a proof for relation  $\text{PoK}\{\boldsymbol{v} : \text{Commit-Evals}(\boldsymbol{v}) = \text{com}_{\boldsymbol{v}} \wedge \sum_i r^i v_i = \hat{v}\}$ . We refer to Section 3 for more details and we note that such relation will be proven with a *PlonK* circuit in what we call *meta-verification* (Section 4.1).

**Meta-verification.** Using a multi-polynomial commitment is not enough to achieve sublinear verification. There is a linear number of equations/identities in the number of aggregated proofs that need to be verified. We exploit the fact that these identities involve scalar operations over  $\mathbb{Z}_p$ , the native field of *PlonK* circuits.

This observation allows us to delegate the verification of the identities to the prover, who will compute a *PlonK* proof of the fact that the identities are satisfied on the evaluations inside  $\text{com}_{\boldsymbol{v}}$ , the commitment to

<sup>4</sup>For example, the commitments of vectors  $([\tau]_1, [0]_1)$  and  $([0]_1, [\tau]_1)$  are identical.

the evaluations, whose validity has been ensured by the multi-polynomial commitment scheme (see Figure 5 for a precise description of the statement).

We then show that if function `Commit-Evals`, used for committing to the evaluations, is chosen adequately, it can be linked very naturally to a `PlonK` proof, without having to model the commitment opening with `PlonK` constraints, but outside of the circuit. This technique, necessary to have a small meta-verification circuit and thus maximize the number of proofs  $k$  that can be aggregated, can be of independent interest (see Section 4.2).

## 2 Preliminaries

### 2.1 Notation

For a finite set  $S$ , we write  $a \leftarrow S$  to denote that  $a$  is uniformly sampled from  $S$ . We denote the security parameter by  $\lambda \in \mathbb{N}$ . Given two functions  $f, g : \mathbb{N} \rightarrow [0, 1]$ , we write  $f \approx g$  if the difference  $|f(\lambda) - g(\lambda)|$  is asymptotically smaller than the inverse of any polynomial. A function  $f$  is said to be *negligible* if  $f \approx 0$ , whereas it is said to be *overwhelming* when  $f \approx 1$ . For integers  $m, n$ , such that  $m \leq n$ , we denote by  $[m, n]$  the range  $\{m, m+1, \dots, n\}$ . We denote by  $[n]$  the range  $[1, n]$ . Given  $d \in \mathbb{N}$  and a ring  $R$ , we denote by  $R^{<d}[X]$  the set of univariate polynomials over  $X$  with coefficients in  $R$  and degree strictly smaller than  $d$ . For  $n \in \mathbb{N}$ , we denote by  $\mathbf{v} \in R^n$  a vector length  $n$  over  $R$ , and for every  $i \in [n]$ , we denote by  $v_i$  its  $i$ -th component. Furthermore, for any  $k \leq n$ ,  $\mathbf{v}[k]$  denotes the vector formed by the first  $k$  components of  $\mathbf{v}$ .

We consider a bilinear group generator  $\mathcal{G}$  that on input  $1^\lambda$ , produces a set of bilinear groups  $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_t)$  of order  $p$  (a  $\lambda$ -bits prime), equipped with a non-degenerate bilinear pairing  $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_t$ , satisfying  $e(aG, bH) = ab \cdot e(G, H)$  for all  $G \in \mathbb{G}_1, H \in \mathbb{G}_2$  and  $a, b \in \mathbb{Z}_p$ . We use additive notation for all three groups.<sup>5</sup> Unless specified otherwise, we implicitly assume that all algorithms share the same common set of bilinear groups, sampled from the appropriate security parameter. For  $n \in \mathbb{N}$ , such that  $n | p-1$ , let  $\mathcal{H}_n$  be the subgroup generated by  $\omega_n \in \mathbb{Z}_p$ , a designated primitive  $n$ -th root of unity over  $\mathbb{Z}_p$ , and let  $Z_{\mathcal{H}_n}(X) := X^n - 1$ , which vanishes over  $\mathcal{H}_n$ . For every  $i \in [n]$ , let  $L_{i,n}$  be the Lagrange polynomial such that  $L_{i,n}(\omega_n^i) = 1$  and  $L_{i,n}(h) = 0$  for all  $h \in \mathcal{H}_n \setminus \{\omega_n^i\}$ . Throughout the paper, such  $n$  will denote the number of constraints in the constraint system of interest. To speed up polynomial operations through the discrete (I)FFT algorithm, it is convenient that  $n$  be a power of two.<sup>6</sup>

### 2.2 Succinct non-interactive arguments of knowledge

SNARKs are a class of arguments of knowledge that allow a prover to convince a verifier of the validity of a certain statement. They have the important property that proofs and verification time must be polylogarithmic in the length of the statement and the witness.

**Definition 1 (SNARKs).** *A succinct non-interactive argument of knowledge (SNARK) for a binary relation  $\mathcal{R}$  is a triple of PPT algorithms*

- $\text{Setup}(1^\lambda, \mathcal{R}) \rightarrow pp$ , on input the security parameter  $\lambda$  and relation  $\mathcal{R}$ , outputs a set of public parameters  $pp$ , also known as a common reference string.
- $\text{Prove}(pp, x, w) \rightarrow \pi$ , on input  $pp$ , statement  $x$  and witness  $w$ , outputs a proof.
- $\text{Verify}(pp, x, \pi) \rightarrow 1/0$ , on input  $pp$ , statement  $x$  and proof  $\pi$ , outputs a bit.

<sup>5</sup>It is more common to express  $\mathbb{G}_t$  in multiplicative notation, since its group operation is typically implemented through a polynomial multiplication.

<sup>6</sup>Some elliptic-curves are designed so that a big power of 2 divides  $p-1$ , e.g.,  $2^{32}$  divides the order of the multiplicative subgroup of BLS12-381 [Bow17] scalar field.

**Completeness.** A SNARK is complete if for every  $\lambda \in \mathbb{N}$  and every  $(x, w) \in \mathcal{R}$ :

$$\Pr [pp \leftarrow \text{Setup}(1^\lambda, \mathcal{R}); \pi \leftarrow \text{Prove}(pp, x, w) : \text{Verify}(pp, x, \pi) = 1] = 1 .$$

**Knowledge Soundness.** A SNARK is knowledge sound if for every PPT algorithm  $\mathcal{A}$ , there exists an expected polynomial-time extractor  $\mathcal{E}$  (with access to  $\mathcal{A}$ 's random tape) such that the following probability is negligible in  $\lambda$ :

$$\Pr [pp \leftarrow \text{Setup}(1^\lambda, \mathcal{R}); (\pi, x) \leftarrow \mathcal{A}(pp) : w \leftarrow \mathcal{E}(pp); \text{Verify}(pp, x, \pi) = 1 \wedge (x, w) \notin \mathcal{R}] .$$

**Succinctness.** A SNARK is succinct if  $|\pi| = \text{poly}(\lambda, \log(|x| + |w|))$ , for every  $(x, w) \in \mathcal{R}$ .

**Zero-knowledge.** A SNARK is zero-knowledge if there exists a PPT (stateful) simulator  $\mathcal{S}$  such that for every PPT (stateful) algorithm  $\mathcal{A}$ , the following probabilities are negligibly close (in  $\lambda$ ):

$$\begin{aligned} \Pr [pp \leftarrow \text{Setup}(1^\lambda, \mathcal{R}); (x, w) \leftarrow \mathcal{A}(pp); \pi \leftarrow \text{Prove}(pp, x, w) : (x, w) \in \mathcal{R} \wedge \mathcal{A}(\pi) = 1] , \\ \Pr [pp \leftarrow \mathcal{S}(1^\lambda); (x, w) \leftarrow \mathcal{A}(pp); \pi \leftarrow \mathcal{S}(pp, x) : (x, w) \in \mathcal{R} \wedge \mathcal{A}(\pi) = 1] . \end{aligned}$$

### 2.3 Polynomial commitment schemes

A polynomial commitment scheme (PCS) [KZG10] is a commitment scheme where the objects being committed are univariate polynomials (of bounded degree). These systems are also equipped with a mechanism for proving (not necessarily in zero-knowledge) that the polynomial “inside” a certain commitment evaluates to a claimed value at a given evaluation point.

**Definition 2 (Polynomial Commitment).** A polynomial commitment scheme over a ring  $R$  consists of four PPT algorithms:

- $\text{Setup}(1^\lambda, d) \rightarrow (\text{ck}, \text{vk})$ , on input the security parameter  $\lambda$  and a degree bound  $d \in \mathbb{N}$ , outputs a commitment key  $\text{ck}$  and a verification key  $\text{vk}$ .
- $\text{Commit}(\text{ck}, f) \rightarrow \text{com}$ , given  $\text{ck}$  and a polynomial  $f \in R^{<d}[X]$ , outputs a commitment  $\text{com}$ .
- $\text{Open}(\text{ck}, \text{com}, z, f) \rightarrow \pi$ , given a commitment key, a commitment  $\text{com}$ , an evaluation point  $z \in R$  a polynomial  $f$  (that was committed in  $\text{com}$ ), outputs a proof  $\pi$ .
- $\text{Check}(\text{vk}, \text{com}, z, v, \pi) \rightarrow 1/0$ , given a verification key  $\text{vk}$ , a commitment  $\text{com}$ , an evaluation point  $z$ , a claimed evaluation  $v$  and a proof  $\pi$ , outputs a bit (1 representing acceptance, 0 representing rejection).

For the sake of simplicity in the next definitions, we require that all algorithms except  $\text{Setup}$  be deterministic. Note that most instantiations from the literature are deterministic [KZG10, BGH19b, BDFG20]. We also require a polynomial commitment scheme to satisfy the following properties.

**Completeness.** A multi-polynomial commitment scheme is *complete* if for every  $\lambda, d$  and every  $(\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d)$ , any  $f \in R^{<d}[X]$ ,  $z \in R$ , it holds:

$$\text{Check}(\text{vk}, \text{com}, z, f(z), \text{Open}(\text{ck}, \text{com}, z, f)) = 1 ,$$

where  $\text{com} := \text{Commit}(\text{ck}, f)$ .

**Binding Property.** A polynomial commitment scheme is *binding* if for every polynomial  $d \in \mathbb{N}$  and every PPT adversary  $\mathcal{A}$ , the following probability is negligible in  $\lambda$ :

$$\Pr \left[ \begin{array}{l} (\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d) \\ (f, f') \leftarrow \mathcal{A}(\text{ck}) \end{array} : \text{Commit}(\text{ck}, f) = \text{Commit}(\text{ck}, f') \wedge f \neq f' \right] .$$

**Knowledge Soundness.** A polynomial commitment scheme is *knowledge sound* if for every polynomial  $d \in \mathbb{N}$  and every PPT adversary  $\mathcal{A}$ , there exists an (expected polynomial time) extractor  $\mathcal{E}$  such that the following probability is negligible in  $\lambda$ :

$$\Pr \left[ \begin{array}{l} (\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d) \\ (\text{com}, z, v, \pi) \leftarrow \mathcal{A}(\text{ck}) \\ f \leftarrow \mathcal{E}(\text{ck}) \end{array} : \begin{array}{l} \text{Check}(\text{vk}, \text{com}, z, v, \pi) = 1 \\ \wedge (\text{com} \neq \text{Commit}(\text{ck}, f) \vee (f(z) \neq v)) \end{array} \right] .$$



## 2.4 Constraint systems

A constraint system is a list of polynomial equations over  $\mathbb{Z}_p[X_1, \dots, X_m]$ , of restricted form. For simplicity in our exposition, in this work we consider polynomials of the form

$$\mathbf{q}_L X_i + \mathbf{q}_R X_j + \mathbf{q}_O X_k + \mathbf{q}_M X_i X_j + \mathbf{q}_C ,$$

for certain scalar coefficients  $\mathbf{q}_L, \mathbf{q}_R, \mathbf{q}_O, \mathbf{q}_M, \mathbf{q}_C \in \mathbb{Z}_p$ . This corresponds to the classical identity considered in the original *PlonK* paper [GWC19]. All our results extend to other versions of *PlonK*, that involve additional identities such as [GW19, PFM<sup>+</sup>22] and even to implementations that use a different number of wires per gate (instead of 3).

**Definition 3 (Constraint System).** *A constraint system on  $m$  variables is a list of tuples  $(a, b, c, \mathbf{q}_L, \mathbf{q}_R, \mathbf{q}_O, \mathbf{q}_M, \mathbf{q}_C)$  with  $a, b, c \in [m]$ ,  $\mathbf{q}_L, \mathbf{q}_R, \mathbf{q}_O, \mathbf{q}_M, \mathbf{q}_C \in \mathbb{Z}_p$ . We say a vector  $\mathbf{x} \in \mathbb{Z}_p^m$  satisfies constraint system  $\mathcal{C} = \{(a_i, b_i, c_i, \mathbf{q}_{L_i}, \mathbf{q}_{R_i}, \mathbf{q}_{O_i}, \mathbf{q}_{M_i}, \mathbf{q}_{C_i})\}_{i \in [n]}$  if for every  $i \in [n]$ :*

$$\mathbf{q}_{L_i} x_{a_i} + \mathbf{q}_{R_i} x_{b_i} + \mathbf{q}_{O_i} x_{c_i} + \mathbf{q}_{M_i} x_{a_i} x_{b_i} + \mathbf{q}_{C_i} = 0 .$$

The *PlonK* proving system is a zk-SNARK for the following relation, defined over so-called *public inputs*  $\mathbf{x} \in \mathbb{Z}_p^\ell$  and *witness*  $\mathbf{w} \in \mathbb{Z}_p^{m-\ell}$ :

$$\text{PoK} \{ \mathbf{w} \in \mathbb{Z}_p^{m-\ell} : (\mathbf{x}, \mathbf{w}) \in \mathbb{Z}_p^m \text{ satisfies } \mathcal{C} \} . \quad (1)$$

The statement being proved is thus parametrized by both  $\mathcal{C}$  and  $\mathbf{x}$ .

## 2.5 The *PlonK* proving system

Let  $\mathcal{C} = \{(a_i, b_i, c_i, \mathbf{q}_{L_i}, \mathbf{q}_{R_i}, \mathbf{q}_{O_i}, \mathbf{q}_{M_i}, \mathbf{q}_{C_i})\}_{i \in [n]}$  be a constraint system on  $m$  variables. *PlonK* requires that the system be preprocessed by defining univariate polynomials  $\mathbf{q}_L(X), \mathbf{q}_R(X), \mathbf{q}_O(X), \mathbf{q}_M(X), \mathbf{q}_C(X)$  in  $\mathbb{Z}_p[X]$  satisfying:

$$\mathbf{q}_L(\omega_n^i) = \mathbf{q}_{L_i} \quad \mathbf{q}_R(\omega_n^i) = \mathbf{q}_{R_i} \quad \mathbf{q}_O(\omega_n^i) = \mathbf{q}_{O_i} \quad \mathbf{q}_M(\omega_n^i) = \mathbf{q}_{M_i} \quad \mathbf{q}_C(\omega_n^i) = \mathbf{q}_{C_i} ,$$

for every  $i \in [n]$ . We recall that  $\omega_n$  is a designated  $n$ -th primitive root of unity. Furthermore the relations between indices  $\{a_i, b_i, c_i\}_{i \in [n]}$  are captured through a permutation  $\sigma : [3n] \rightarrow [3n]$ , which decomposes in exactly  $m$  cycles: the  $j$ -th cycle involving all positions where the  $j$ -th variable is used. Such permutation is then transformed into a list of 3 polynomials  $S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3}$ , which are involved in the definition of the so-called *permutation identities*, parametrized by two scalars  $\beta, \gamma \in \mathbb{Z}_p$ :

$$\text{perm-ids}_{\beta, \gamma}^\sigma(\mathbf{A}(X), \mathbf{B}(X), \mathbf{C}(X), \mathbf{Z}(X)) ,$$

defining two polynomials which must vanish over the whole subgroup  $\mathcal{H}_n$ . We describe in detail this identity, how it depends on polynomials  $S_{\sigma_i}$ , and how these polynomials are created in Appendix A.1. Here, we just assume there exists an efficient mechanism to compute a polynomial  $\mathbf{Z}$  (of degree at most  $n$ ) that satisfies the *perm-ids* <sup>$\sigma$</sup> , from  $\beta, \gamma$ , if polynomials  $\mathbf{A}, \mathbf{B}, \mathbf{C}$  were honestly generated from a satisfying assignment to the constraint system  $\mathcal{C}$ .

Let  $\Psi$  be a polynomial commitment scheme, and let  $(\text{ck}, \text{vk}) \leftarrow \Psi.\text{Setup}(1^\lambda, n)$ . *PlonK*'s preprocessing phase concludes by committing to the above polynomials using  $\Psi$ . That is,  $\mu_{\mathbf{q}_L} \leftarrow \Psi.\text{Commit}(\text{ck}, \mathbf{q}_L)$ , and similarly for  $\mu_{\mathbf{q}_R}, \mu_{\mathbf{q}_O}, \mu_{\mathbf{q}_M}, \mu_{\mathbf{q}_C}$  and  $\mu_{S_{\sigma_1}}, \mu_{S_{\sigma_2}}, \mu_{S_{\sigma_3}}$ . These polynomial commitments, together with  $(\text{ck}, \text{vk})$  form *PlonK*'s public parameters  $\text{pp}$ .

We describe *PlonK*'s prover and verifier in Figure 1. In a nutshell, the prover commits to certain polynomials  $\mathbf{A}, \mathbf{B}, \mathbf{C}$ , that represent a valid trace witness. The prover then argues that such polynomials satisfy the identities over the whole subgroup  $\mathcal{H}_n$  by showing that the identities, instantiated with the witness polynomials, lead to polynomials which are divisible by  $Z_{\mathcal{H}_n}$ . This is done by committing to the quotient  $\mathbf{T}$  of such division and evaluating all (committed) polynomials on a uniformly sampled point  $\xi$ . The verifier then

$\mathcal{PlonK}.\text{Prove}(\mathcal{C}, \text{pp}, \mathbf{x}, \mathbf{w})$ :

**Inputs:** constraint system  $\mathcal{C}$  on  $m$  variables, preprocessed parameters  $\text{pp}$  for  $\mathcal{C}$ , instance  $\mathbf{x}$ , witness trace  $\mathbf{w}$

**Output:** PoK  $\{ \mathbf{w} \in \mathbb{Z}_p^{m-\ell} : (\mathbf{x}, \mathbf{w}) \in \mathbb{Z}_p^m \text{ satisfies } \mathcal{C} \}$

- 1:  $\tilde{\mathbf{w}} := (\mathbf{x}, \mathbf{w})$ ;  $\text{A}(X) := \sum_{i=1}^n \tilde{w}_{a_i} \text{L}_i(X)$ ;  $\text{B}(X) := \sum_{i=1}^n \tilde{w}_{b_i} \text{L}_i(X)$ ;  $\text{C}(X) := \sum_{i=1}^n \tilde{w}_{c_i} \text{L}_i(X)$ <sup>a</sup>
- 2:  $[\text{A}] := \Psi.\text{Commit}(\text{ck}, \text{A})$ ;  $[\text{B}] := \Psi.\text{Commit}(\text{ck}, \text{B})$ ;  $[\text{C}] := \Psi.\text{Commit}(\text{ck}, \text{C})$
- 3:  $\beta \leftarrow \text{Hash}([\text{A}], [\text{B}], [\text{C}])$ ;  $\gamma \leftarrow \text{Hash}(\beta)$
- 4: compute polynomial  $\text{Z}(X)$ , satisfying  $\text{perm-ids}_{\beta, \gamma}^\sigma(\text{A}, \text{B}, \text{C}, \text{Z})$  on  $\mathcal{H}_n$  ▷ See Appendix A.1
- 5:  $[\text{Z}] \leftarrow \Psi.\text{Commit}(\text{ck}, \text{Z})$
- 6:  $F(X) = (\text{q}_L \text{A} + \text{q}_R \text{B} + \text{q}_O \text{C} + \text{q}_M \text{A} \text{B} + \text{q}_C)(X)$
- 7:  $\text{ids}(X) = \{F(X)\} \cup \text{perm-ids}_{\beta, \gamma}^\sigma(\text{A}, \text{B}, \text{C}, \text{Z})$ ;  $\alpha \leftarrow \text{Hash}(\gamma, [\text{Z}])$
- 8:  $\text{T}(X) := (\sum_{i \in \|\text{ids}\|} \alpha^i \text{ids}_i(X)) / Z_{\mathcal{H}_n}(X)$
- 9:  $[\text{T}] \leftarrow \Psi.\text{Commit}(\text{ck}, \text{T})$
- 10:  $\xi \leftarrow \text{Hash}(\alpha, [\text{T}])$
- 11: **return**  $([\text{A}], [\text{B}], [\text{C}], [\text{Z}], [\text{T}])$  together with evaluations on  $\xi$  and proofs of their correctness, of these commitments as well as all polynomials commitments in  $\text{pp}$ ; evaluate  $\text{Z}$  also on  $\xi \omega$  ▷ Using  $\Psi.\text{Open}$

$\mathcal{PlonK}.\text{Verify}(\mathcal{C}, \text{pp}, \mathbf{x}, \pi)$ :

**Inputs:** constraint system  $\mathcal{C}$  on  $m$  variables, preprocessed parameters  $\text{pp}$  for  $\mathcal{C}$ , instance  $\mathbf{x}$ , proof  $\pi$

**Output:** bool (*true* if the proof is accepted, *false* otherwise)

- 1: parse  $\pi$  as  $([\text{A}], [\text{B}], [\text{C}], [\text{Z}], [\text{T}], \pi_{\text{evals}})$
- 2: assert that the claimed evaluations in  $\pi_{\text{evals}}$  are correct ▷ Using  $\Psi.\text{Check}$
- 3:  $\beta \leftarrow \text{Hash}([\text{A}], [\text{B}], [\text{C}])$ ;  $\gamma \leftarrow \text{Hash}(\beta)$ ;  $\alpha \leftarrow \text{Hash}(\gamma, [\text{Z}])$ ;  $\xi \leftarrow \text{Hash}(\alpha, [\text{T}])$
- 4: evaluate  $\{\text{ids}\}_i$  on  $\xi$ , leveraging the claimed evaluations, obtaining  $\{\overline{\text{ids}}\}_i$  ▷  $\text{ids}$  depend on  $\beta, \gamma$
- 5: **return**  $\sum_{i \in \|\text{ids}\|} \alpha^i \overline{\text{ids}}_i = Z_{\mathcal{H}_n}(\xi) \bar{t}$  ▷  $\bar{t}$  is the claimed evaluation of  $T$  on  $\xi$

**Fig. 1.** The  $\mathcal{PlonK}$  proving system for  $\mathcal{C} := \{(a_i, b_i, c_i, \text{q}_L, \text{q}_R, \text{q}_O, \text{q}_M, \text{q}_C)\}_{i \in [n]}$ , with preprocessed parameters  $\text{pp}$  that include commitment keys  $(\text{ck}, \text{vk})$  of a polynomial commitment scheme  $\Psi$  and polynomial commitments  $[\text{q}_L], [\text{q}_R], [\text{q}_O], [\text{q}_M], [\text{q}_C]$  and  $[\text{S}_{\sigma,1}], [\text{S}_{\sigma,2}], [\text{S}_{\sigma,3}]$ .

<sup>a</sup>Additional multiples of polynomial  $Z_{\mathcal{H}_n}$  may be optionally added to each  $\text{A}, \text{B}, \text{C}$ , to achieve zero-knowledge.

checks that  $Z_{\mathcal{H}_n}(\xi) \text{T}(\xi)$  equals the evaluation of the identities on  $\xi$ , which ensures that the previous division (over polynomials) was exact thanks to the knowledge soundness of  $\Psi$  and the Schwartz-Zippel Lemma.

If instantiated with a secure polynomial commitment  $\Psi$  which has logarithmic verification on the degree bound of polynomials<sup>7</sup>, the protocol from Figure 1 constitutes a SNARK for relation (1) by virtue of [GWC19, Theorem 7.1 & Corollary 7.2].

<sup>7</sup>The authors propose to instantiate  $\Psi$  with the KZG polynomial commitment scheme [KZG10], which leads to a SNARK whose knowledge soundness holds in the algebraic group model.

### 3 Multi-polynomial commitment schemes

We introduce the notion of *multi-polynomial commitment schemes*, a generalization of polynomial commitment schemes designed to commit to several polynomials at the same time. We require the commitment size be sublinear in the number of committed polynomials. Furthermore, we require that verification can be performed from a succinct (standard) commitment to the polynomial evaluations. That way, the verifier does not need to obtain the actual evaluations, which allows its running time to be sublinear in the number of polynomials involved.

**Definition 4 (Multi-polynomial commitment).** *A multi-polynomial commitment scheme over a ring  $R$  consists of five polynomial-time algorithms:*

- $\text{Setup}(1^\lambda, d, K) \rightarrow (\text{ck}, \text{vk})$ , on input the security parameter  $\lambda$ , a degree bound  $d \in \mathbb{N}$ , and a vector length bound  $K \in \mathbb{N}$ , outputs a commitment key  $\text{ck}$  and a verification key  $\text{vk}$ .<sup>8</sup>
- $\text{Commit-Polys}(\text{ck}, \mathbf{f}) \rightarrow \text{com}_{\mathbf{f}}$ , given a commitment key  $\text{ck}$  and a vector of  $k$  polynomials  $\mathbf{f} \in R^{<d}[X]^k$ , with  $k \leq K$ , outputs a commitment  $\text{com}_{\mathbf{f}}$ . We require that the size of  $\text{com}_{\mathbf{f}}$  be sublinear in  $k$ .
- $\text{Commit-Evals}(\mathbf{v}) \rightarrow \text{com}_{\mathbf{v}}$ , given a vector  $\mathbf{v} \in R^k$ , with  $k \leq K$ , outputs a commitment  $\text{com}_{\mathbf{v}}$ . We require that the size of  $\text{com}_{\mathbf{v}}$  be sublinear in  $k$ .
- $\text{Open}(\text{ck}, \text{com}_{\mathbf{f}}, z, \mathbf{f}) \rightarrow \pi$ , given a commitment key, a commitment  $\text{com}_{\mathbf{f}}$ , an evaluation point  $z \in R$  a vector of  $k$  polynomials in  $R^{<d}[X]$  (that were committed in  $\text{com}_{\mathbf{f}}$ ) with  $k \leq K$ , outputs a proof  $\pi$ .
- $\text{Check}(\text{vk}, \text{com}_{\mathbf{f}}, z, \text{com}_{\mathbf{v}}, \pi) \rightarrow 1/0$ , given a verification key  $\text{vk}$ , a commitment to polynomials  $\text{com}_{\mathbf{f}}$ , an evaluation point  $z$ , a commitment to evaluations  $\text{com}_{\mathbf{v}}$ , and a proof  $\pi$ , outputs a bit. We require that the verification complexity be sublinear in  $K$ .

For the sake of simplicity, we require that all algorithms except  $\text{Setup}$  be deterministic. Our definitions could be adjusted to support non-determinism, but this is not necessary for our use case. (Commitments do not need to be hiding, as zero-knowledge can be enforced by other mechanisms.)

**Completeness.** A multi-polynomial commitment scheme is *complete* if for every  $\lambda, d, K$  and all  $(\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d, K)$ , for any  $k \leq K$ , any vector  $\mathbf{f} \in R^{<d}[X]^k$ , and any  $z \in R$ , it holds:

$$\text{Check}(\text{vk}, \text{com}_{\mathbf{f}}, z, \text{com}_{\mathbf{v}}, \text{Open}(\text{ck}, \text{com}_{\mathbf{f}}, z, \mathbf{f})) = 1 \quad ,$$

where  $\text{com}_{\mathbf{f}} := \text{Commit-Polys}(\text{ck}, \mathbf{f})$  and  $\text{com}_{\mathbf{v}} := \text{Commit-Evals}(\mathbf{f}(z))$ .

**Binding Property.** A multi-polynomial commitment scheme is *binding* if for every polynomial  $d, K \in \mathbb{N}$  and every PPT adversary  $\mathcal{A}$ , the following probabilities are negligible in  $\lambda$ :

$$\Pr \left[ (\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d, K); (\mathbf{f}, \mathbf{f}') \leftarrow \mathcal{A}(\text{ck}) : \begin{array}{l} \mathbf{f} \neq \mathbf{f}' \wedge \mathbf{f} \in R^{<d}[X]^k, \mathbf{f}' \in R^{<d}[X]^{k'}, k, k' \leq K \\ \text{Commit-Polys}(\text{ck}, \mathbf{f}) = \text{Commit-Polys}(\text{ck}, \mathbf{f}') \end{array} \right] ,$$

$$\Pr \left[ (\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d, K); (\mathbf{v}, \mathbf{v}') \leftarrow \mathcal{A}(\text{ck}) : \begin{array}{l} \mathbf{v} \neq \mathbf{v}' \wedge \mathbf{v} \in R^k, \mathbf{v}' \in R^{k'}, k, k' \leq K \\ \text{Commit-Evals}(\mathbf{v}) = \text{Commit-Evals}(\mathbf{v}') \end{array} \right] .$$

**Knowledge Soundness.** A multi-polynomial commitment scheme is *knowledge sound* if for every polynomial  $d, K \in \mathbb{N}$  and every PPT adversary  $\mathcal{A}$ , there exists an (expected polynomial time) extractor  $\mathcal{E}$  such that the following probability is negligible in  $\lambda$ :

$$\Pr \left[ \begin{array}{l} (\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d, K) \\ (\text{com}_{\mathbf{f}}, z, \text{com}_{\mathbf{v}}, \pi) \leftarrow \mathcal{A}(\text{ck}) \\ \mathbf{f} \leftarrow \mathcal{E}(\text{ck}) \end{array} : \begin{array}{l} \text{Check}(\text{vk}, \text{com}_{\mathbf{f}}, z, \text{com}_{\mathbf{v}}, \pi) = 1 \\ \text{com}_{\mathbf{f}} \neq \text{Commit-Polys}(\text{ck}, \mathbf{f}) \\ \vee \text{com}_{\mathbf{v}} \neq \text{Commit-Evals}(\mathbf{f}(z)) \end{array} \right] .$$

<sup>8</sup>We assume both keys implicitly contain  $d, K$  and that  $\text{ck}$  implicitly contains  $\text{vk}$ .

### 3.1 A multi-polynomial commitment scheme from KZG and IPA

We present a generic construction of multi-polynomial commitments from:

- (i) a polynomial commitment scheme which is homomorphic over  $\mathbb{G}_1$ , inner-product binding and inner-product extractable (as defined below);
- (ii) a sublinear-verifier argument system for the following relation, parametrized by  $\mathbf{G} \in \mathbb{G}_2^k$ ,  $C \in \mathbb{G}_t$ ,  $P \in \mathbb{G}_1$  and  $r \in \mathbb{Z}_p$ , where  $\mathbf{r} = (1, r, \dots, r^{k-1})$ :

$$\text{PoK}\{\boldsymbol{\mu} \in \mathbb{G}_1^k : \langle \boldsymbol{\mu}, \mathbf{G} \rangle = C \wedge \langle \mathbf{r}, \boldsymbol{\mu} \rangle = P\} , \quad (2)$$

The first building block can be instantiated with the celebrated KZG commitment scheme [KZG10] (described in Appendix B). For the second building block, we propose a modified version of the inner-product argument [BBB<sup>+</sup>18], inspired by [GMN20] (see Figure 3).

**Homomorphic Property.** A polynomial commitment scheme over group  $\mathbb{G}$  of prime order  $p$  is *homomorphic* if for every  $\lambda, d \in \mathbb{N}$ ,  $(\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d)$  and all  $f, g \in \mathbb{Z}_p^{<d}[X]$ , it holds:

$$\text{Commit}(\text{ck}, f) +_{\mathbb{G}} \text{Commit}(\text{ck}, g) = \text{Commit}(\text{ck}, f + g) ,$$

**Inner-Product Binding Property.** A homomorphic polynomial commitment scheme (over  $\mathbb{G}_1$ ) is *inner-product binding* if for every polynomial  $d, K \in \mathbb{N}$  and every PPT (stateful) algorithm  $\mathcal{A}$ , the following probability is negligible in  $\lambda$ :

$$\Pr \left[ \begin{array}{l} (\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d) \\ \tau \leftarrow \mathbb{Z}_p \\ \mathbf{f}, \mathbf{f}' \leftarrow \mathcal{A}(\text{ck}, [\tau]_1, [\tau]_2) \end{array} : \begin{array}{l} \mathbf{f} \neq \mathbf{f}', \mathbf{f}, \mathbf{f}' \in \mathbb{Z}_p^{<d}[X]^k, \text{ with } k \leq K \\ \langle \text{Commit}(\text{ck}, \mathbf{f}), \boldsymbol{\tau}[:k] \rangle = \langle \text{Commit}(\text{ck}, \mathbf{f}'), \boldsymbol{\tau}[:k] \rangle \end{array} \right] ,$$

where  $\text{Commit}(\text{ck}, \mathbf{f})$  is a shorthand for  $(\text{Commit}(\text{ck}, f_1), \dots, \text{Commit}(\text{ck}, f_k))$  and  $\boldsymbol{\tau} := (1, \tau, \dots, \tau^{K-1})$ .

**Proposition 1.** *The KZG PCS (Figure 7) is inner-product binding.*

We refer to Appendix B for a proof.

**Inner-Product Extractability.** A homomorphic polynomial commitment scheme (over  $\mathbb{G}_1$ ) is *inner-product extractable* if for every polynomial  $d, K \in \mathbb{N}$  and every PPT (stateful) algorithm  $\mathcal{A}$ , there exists an (expected polynomial time) extractor  $\mathcal{E}$  such that the following probability is negligible in  $\lambda$ :

$$\Pr \left[ \begin{array}{l} (\text{ck}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, d) \\ \tau, r \leftarrow \mathbb{Z}_p \\ G, z, \mathbf{v} \leftarrow \mathcal{A}(\text{ck}, [\tau]_1, [\tau]_2) \\ (\boldsymbol{\mu}, \pi) \leftarrow \mathcal{A}(r) \\ \mathbf{f} \leftarrow \mathcal{E}(\text{ck}, [\tau]_1, [\tau]_2, r) \end{array} : \begin{array}{l} G \in \mathbb{G}_t, z \in \mathbb{Z}_p, \mathbf{v} \in \mathbb{Z}_p^k, \boldsymbol{\mu} \in \mathbb{G}_1^k, \text{ with } k \leq K \\ \langle \boldsymbol{\mu}, \boldsymbol{\tau}[:k] \rangle = G \\ \text{Check}(\text{vk}, \langle \boldsymbol{\mu}, \mathbf{r} \rangle, z, \langle \mathbf{v}, \mathbf{r} \rangle, \pi) = 1 \\ (\mathbf{f}(z) \neq \mathbf{v} \vee \langle \text{Commit}(\text{ck}, \mathbf{f}), \boldsymbol{\tau}[:k] \rangle \neq G) \end{array} \right] ,$$

where  $\text{Commit}(\text{ck}, \mathbf{f})$  is a shorthand for  $(\text{Commit}(\text{ck}, f_1), \dots, \text{Commit}(\text{ck}, f_k))$ ,  $\boldsymbol{\tau} := (1, \tau, \dots, \tau^{K-1})$  and  $\mathbf{r} := (1, r, \dots, r^{k-1})$ .

**Proposition 2.** *The KZG PCS (Figure 7) is inner-product extractable.*

We refer to Appendix B for a proof.

**Theorem 1.** *If  $\text{Hash} : \{0, 1\}^* \rightarrow \mathbb{Z}_p$  is a random oracle, and polynomial commitment scheme  $\Psi$  is complete, binding, knowledge sound, homomorphic, inner-product binding and inner-product extractable, then the scheme from Figure 2 is a complete, binding and knowledge sound multi-polynomial commitment scheme.*

We refer to Appendix C.1 for a proof.

Setup( $1^\lambda, d, K$ ):

- 1:  $(\text{ck}_\Psi, \text{vk}_\Psi) \leftarrow \Psi.\text{Setup}(1^\lambda, d)$
- 2:  $\tau \leftarrow \mathbb{Z}_p$ ;  $\text{ck}_\tau := [1, \tau, \tau^2, \dots, \tau^{K-1}]_2$
- 3: **return**  $(\text{ck} := (\text{ck}_\Psi, \text{ck}_\tau), \text{vk} := (\text{vk}_\Psi, [\tau]_1))$

Commit-Polys( $\text{ck} := (\text{ck}_\Psi, \text{ck}_\tau), \mathbf{f}$ ):

- 1:  $\mu_i \leftarrow \Psi.\text{Commit}(\text{ck}_\Psi, f_i) \forall i \in [k]$  ( $k := |\mathbf{f}| \leq K$ )
- 2: **return**  $\text{com}_\mathbf{f} := (k, \sum_{i=1}^k e(\mu_i, \text{ck}_{\tau_i}))$

Commit-Evals( $\mathbf{v}$ ):

Any function that is (sublinearly) *shrinking*, *binding* and admits a succinct and efficient proof for relation:

$$\text{PoK}\{\mathbf{v} : \text{Commit-Evals}(\mathbf{v}) = \text{com}_\mathbf{v} \wedge \sum_{i=1}^k r^{i-1} v_i = \hat{v}\} \quad (3)$$

Open( $\text{ck} := (\text{ck}_\Psi, \text{ck}_\tau), \text{com}_\mathbf{f} := (k, G), z, \mathbf{f}$ ):

- 1:  $\mathbf{v} = \mathbf{f}(z)$ ;  $\text{com}_\mathbf{v} := \text{Commit-Evals}(\mathbf{v})$ ;  $k = |\mathbf{f}|$ ;  $\kappa := \lceil \log_2(k) \rceil$
- 2:  $\mu_i \leftarrow \Psi.\text{Commit}(\text{ck}_\Psi, f_i) \forall i \in [k]$  ▷ Not necessary if  $\mu_i$  were stored on Commit-Polys
- 3:  $r := \text{Hash}(\text{com}_\mathbf{f}, z, \text{com}_\mathbf{v})$ ;  $\mathbf{r} := (1, r, \dots, r^{n-1})$ ;  $\hat{\mathbf{f}} := \langle \mathbf{f}, \mathbf{r} \rangle$ ;  $\hat{\boldsymbol{\mu}} := \langle \boldsymbol{\mu}, \mathbf{r} \rangle$ ;  $\hat{v} := \langle \mathbf{v}, \mathbf{r} \rangle$
- 4:  $\pi_\Psi \leftarrow \Psi.\text{Open}(\text{ck}_\Psi, \hat{\boldsymbol{\mu}}, z, \hat{\mathbf{f}}, \hat{v})$
- 5: produce a proof  $\pi_v$  of relation (3) w.r.t  $\text{com}_\mathbf{v}$ ,  $\hat{v}$  and  $r$
- 6:  $\pi_{\text{IPA}} \leftarrow \text{IPA}.\text{Prove}(k, \text{ck}_\tau, (G, r, \hat{\boldsymbol{\mu}}), \boldsymbol{\mu})$  (let  $\{u_j\}_{j=1}^\kappa$  be the sampled random challenges)
- 7:  $g(X) := \prod_{j=1}^\kappa (u_j^{-1} + u_j X^{2^{\kappa-j}})$ ;  $\rho := \text{Hash}(\pi_{\text{IPA}})$  and  $v_\rho := g(\rho)$
- 8:  $h(X) := (g(X) - v_\rho)/(X - \rho)$ ;  $\pi_\tau := [h(\tau)]_2$  ▷  $\pi_\tau$  is computed using  $\text{ck}_\tau$
- 9: **return**  $(\hat{\boldsymbol{\mu}}, \hat{v}, \pi_\Psi, \pi_v, \pi_{\text{IPA}}, \pi_\tau)$

Check( $\text{vk} := (\text{vk}_\Psi, [\tau]_1), \text{com}_\mathbf{f} := (k, G), z, \text{com}_\mathbf{v}, \pi := (\hat{\boldsymbol{\mu}}, \hat{v}, \pi_\Psi, \pi_v, \pi_{\text{IPA}}, \pi_\tau)$ ):

- 1:  $r := \text{Hash}(\text{com}_\mathbf{f}, z, \text{com}_\mathbf{v})$
- 2:  $b_\Psi \leftarrow \Psi.\text{Check}(\text{vk}_\Psi, \hat{\boldsymbol{\mu}}, z, \hat{v}, \pi_\Psi)$
- 3: let  $b_v$  be the result of verifying that  $\pi_v$  is a valid proof of relation (3) w.r.t  $\text{com}_\mathbf{v}$ ,  $\hat{v}$  and  $r$
- 4:  $b_{\text{IPA}} \leftarrow \text{IPA}.\text{Verify}'(k, [\tau]_1, (G, r, \hat{\boldsymbol{\mu}}), \pi_{\text{IPA}})$  ▷ Skip steps 6-7 of Figure 3
- 5:  $\rho := \text{Hash}(\pi_{\text{IPA}})$ ;  $v_\rho := \prod_{j=1}^\kappa (u_j^{-1} + u_j \rho^{2^{\kappa-j}})$  ▷  $u_j$  are the challenges computed during  $\text{IPA}.\text{Verify}'$
- 6:  $b_\tau := e([\tau]_1 - [\rho]_1, \pi_\tau) \stackrel{?}{=} e([1]_1, G_0 - [v_\rho]_2)$  ▷  $G_0 \in \mathbb{G}_2$  is the last element of  $\pi_{\text{IPA}}$
- 7: **return**  $b_\Psi \wedge b_v \wedge b_{\text{IPA}} \wedge b_\tau$

**Fig. 2.** Multi-polynomial commitment scheme based on an inner-product binding and inner-product extractable polynomial commitment scheme  $\Psi$  (over  $\mathbb{G}_1$ ) and inner-product argument IPA from Figure 3.

IPA.Prove( $k, \mathbf{G}, \mathbf{H}, (C_G, C_H, r, P), \boldsymbol{\mu}$ ):

**Inputs:**  $k := 2^\kappa$ , vectors  $\mathbf{G}, \mathbf{H} \in \mathbb{G}_2^k$ , instance  $(C_G, C_H, r, P) \in \mathbb{G}_t^2 \times \mathbb{Z}_p \times \mathbb{G}_1$ , witness  $\boldsymbol{\mu} \in \mathbb{G}_1^k$

**Output:**  $\text{PoK}\{\boldsymbol{\mu} \in \mathbb{G}_1^k : \langle \boldsymbol{\mu}, \mathbf{G} \rangle = C_G \wedge \langle \boldsymbol{\mu}, \mathbf{H} \rangle = C_H \wedge \langle \mathbf{r}, \boldsymbol{\mu} \rangle = P\}$ , where  $\mathbf{r} = (1, r, r^2, r^3, \dots, r^{k-1})$

- 1: set  $\boldsymbol{\mu}^{(\kappa)} := \boldsymbol{\mu}$ ,  $\mathbf{r}^{(\kappa)} := (1, r, r^2, r^3, \dots, r^{2^\kappa-1})$ ,  $\mathbf{G}^{(\kappa)} := \mathbf{G}$ ,  $\mathbf{H}^{(\kappa)} := \mathbf{H}$ , and  $\text{ts} := (C_G, C_H, r, P)$
- 2: **for every**  $j$  from  $\kappa$  down to 1 **do**
- 3:   set  $L_G^{(j)} := \langle \boldsymbol{\mu}_L^{(j)}, \mathbf{G}_R^{(j)} \rangle$ ,    $L_H^{(j)} := \langle \boldsymbol{\mu}_L^{(j)}, \mathbf{H}_R^{(j)} \rangle$    and    $L_r^{(j)} := \langle \boldsymbol{\mu}_L^{(j)}, \mathbf{r}_R^{(j)} \rangle$
- 4:   set  $R_G^{(j)} := \langle \boldsymbol{\mu}_R^{(j)}, \mathbf{G}_L^{(j)} \rangle$ ,    $R_H^{(j)} := \langle \boldsymbol{\mu}_R^{(j)}, \mathbf{H}_L^{(j)} \rangle$    and    $R_r^{(j)} := \langle \boldsymbol{\mu}_R^{(j)}, \mathbf{r}_L^{(j)} \rangle$ ,
- 5:   set  $u_j \leftarrow \text{Hash}(L_G^{(j)}, R_G^{(j)}, L_H^{(j)}, R_H^{(j)}, L_r^{(j)}, R_r^{(j)}, \text{ts}) \in \mathbb{Z}_p$  and  $\text{ts} := u_j$
- 6:   set  $\boldsymbol{\mu}^{(j-1)} := u_j \boldsymbol{\mu}_L^{(j)} + u_j^{-1} \boldsymbol{\mu}_R^{(j)}$
- 7:   set  $\mathbf{G}^{(j-1)} := u_j^{-1} \mathbf{G}_L^{(j)} + u_j \mathbf{G}_R^{(j)} \triangleright u_j$  multiplies the left half of  $\boldsymbol{\mu}^{(j)}$ , but the right half of  $\mathbf{G}^{(j)}$ ,  $\mathbf{H}^{(j)}$ ,  $\mathbf{r}^{(j)}$
- 8:   set  $\mathbf{H}^{(j-1)} := u_j^{-1} \mathbf{H}_L^{(j)} + u_j \mathbf{H}_R^{(j)}$
- 9:   set  $\mathbf{r}^{(j-1)} := u_j^{-1} \mathbf{r}_L^{(j)} + u_j \mathbf{r}_R^{(j)}$
- 10: **return**  $\pi := (\{L_G^{(j)}, R_G^{(j)}, L_H^{(j)}, R_H^{(j)}, L_r^{(j)}, R_r^{(j)}\}_{j \in [\kappa]}, \boldsymbol{\mu}^{(0)}, \mathbf{G}^{(0)}, \mathbf{H}^{(0)}) \in (\mathbb{G}_t^4 \times \mathbb{G}_1^2)^\kappa \times \mathbb{G}_1 \times \mathbb{G}_2^2$

IPA.Verify( $k, \mathbf{G}, \mathbf{H}, (C_G, C_H, r, P), \pi$ ):

**Inputs:**  $k := 2^\kappa$ , vectors  $\mathbf{G}, \mathbf{H} \in \mathbb{G}_2^k$ , instance  $(C_G, C_H, r, P) \in \mathbb{G}_t^2 \times \mathbb{Z}_p \times \mathbb{G}_1$ , proof  $\pi$

**Output:** bool (*true* if the proof is accepted, *false* otherwise)

- 1: parse  $\pi$  as  $(\{L_G^{(j)}, R_G^{(j)}, L_H^{(j)}, R_H^{(j)}, L_r^{(j)}, R_r^{(j)}\}_{j \in [\kappa]}, \mu_0, G_0, H_0) \in (\mathbb{G}_t^4 \times \mathbb{G}_1^2)^\kappa \times \mathbb{G}_1 \times \mathbb{G}_2^2$  or fail
- 2: **for every**  $j$  from  $\kappa$  down to 1 **do**
- 3:   set  $u_j \leftarrow \text{Hash}(L_G^{(j)}, R_G^{(j)}, L_H^{(j)}, R_H^{(j)}, L_r^{(j)}, R_r^{(j)}, \text{ts}) \in \mathbb{Z}_p$  and set  $\text{ts} := u_j$
- 4: define  $g(X) := \prod_{j=1}^\kappa (u_j^{-1} + u_j X^{2^{\kappa-j}})$   $\triangleright g$  is a polynomial of degree  $2^\kappa - 1$
- 5: set  $r_0 := g(r)$   $\triangleright \mathcal{O}(\kappa)$  if  $g(r)$  is computed as  $\prod_{j=1}^\kappa (u_j^{-1} + u_j r^{2^{\kappa-j}})$
- 6: let  $\mathbf{g} = (g_0, g_1, \dots, g_{2^\kappa-1})$  be the coefficients of  $g$  in increasing order of degree
- 7: assert  $\langle \mathbf{g}, \mathbf{G} \rangle = G_0$ , and  $\langle \mathbf{g}, \mathbf{H} \rangle = H_0$   $\triangleright \mathcal{O}(2^\kappa)$
- 8: **return true** iff  $r_0 \cdot \mu_0 = P + \sum_{j=1}^\kappa (u_j^2 L_r^{(j)} + u_j^{-2} R_r^{(j)})$  and the following hold:

$$e(\mu_0, G_0) = C_G + \sum_{j=1}^\kappa (u_j^2 L_G^{(j)} + u_j^{-2} R_G^{(j)}) \quad \wedge \quad e(\mu_0, H_0) = C_H + \sum_{j=1}^\kappa (u_j^2 L_H^{(j)} + u_j^{-2} R_H^{(j)})$$

**Fig. 3.** Inner-product argument (IPA) for relation  $\text{PoK}\{\boldsymbol{\mu} \in \mathbb{G}_1^{2^\kappa} : \langle \boldsymbol{\mu}, \mathbf{G} \rangle = C_G \wedge \langle \boldsymbol{\mu}, \mathbf{H} \rangle = C_H \wedge \langle \mathbf{r}, \boldsymbol{\mu} \rangle = P\}$ . The figure describes two schemes obtained by including or discarding the colored terms.

*Remark 2.* Our scheme from Figure 2 could also be instantiated with a homomorphic commitment scheme  $\Psi$  that is not inner-product binding nor inner-product extractable. In that case, the multi-polynomial commitment scheme would need to be modified by adding a second  $\mathbb{G}_t$  element  $\sum_{i=1}^k e(\mu_i, [\tilde{\tau}^i]_2)$  to  $\text{com}_f$  for a new  $\tilde{\tau}$  independent of  $\tau$ .<sup>9</sup> This modification would make the committing function binding, which would allow us to prove security without relying on the inner-product binding and inner-product extractability

<sup>9</sup>Such  $\tilde{\tau}$  could be the same secret used during the setup of  $\Psi$ , if  $\Psi$  is such that its structured reference string is formed by the powers of a secret scalar over  $\mathbb{G}_1$  and  $\mathbb{G}_2$ . KZG (Figure 7) is an example of such scheme.

properties of  $\Psi$  (see Appendix C.1). Note that proofs of opening would need to include an extra element  $\pi_{\tilde{\tau}}$ , computed as  $[h(\tilde{\tau})]_2$ , analogously to  $\pi_{\tau}$  in step 8 of the `Open` algorithm, which would be verified with a second pairing equation in step 6 of the `Check` algorithm. Furthermore, the IPA protocol would need to be adapted, as described in Figure 3 through extra colored terms.

## 4 *PlonK* proof aggregation from multi-polynomial commitments

We study the problem of designing a multi-statement proving system, that can handle several *PlonK* proofs more efficiently than the simple parallel execution on every statement of the traditional *PlonK* system. For the sake of simplicity in our exposition, we assume that all statements are parametrized by the same *PlonK* constraint system (although each statement has its own public inputs). However, most of our techniques apply to the case with different systems.

### 4.1 *aPlonK*

A simple but effective first optimization is to share the random challenges sampled with Fiat-Shamir across all proofs. This can be beneficial for several reasons. For example, having a common evaluation point  $\xi$  for all proofs means that all polynomial commitments are opened at the same point, which can typically lead to significant optimizations by the underlying polynomial commitment scheme. Note that sharing such random challenges across proofs does not harm security as long as the challenges are computed from the partial transcripts of all proofs. In that case, from an extractor for the aggregated proving system one could build an extractor for any of the individual statements by fixing all other statements. On the other hand, this trick, which is the basis of many of our optimizations, requires that the provers of every different statement run coordinately or at least synchronize at every point where random challenges are sampled. This limitation prevents us from strictly achieving IVC [Val08], but does not limit the distribution of the prover computation. Thus, our system is perfectly applicable to creating a validity rollup (see Section 1.1).

Another rather simple optimization is to have a common polynomial  $T$  for all proofs, computed from a linear combination of all the identities. In the rest of this section, we describe our more sophisticated optimization techniques. The resulting proving system, that we call *aPlonK*, is described in Figure 4.

**Theorem 2.** *If multi-polynomial commitment scheme  $\Psi$  is complete, binding and knowledge sound, and  $\text{Hash} : \{0, 1\}^* \rightarrow \mathbb{Z}_p$  is a random oracle, the protocol described in Figure 4 constitutes a SNARK for relation:*

$$\text{PoK} \left\{ \left\{ \mathbf{w}_i \in \mathbb{Z}_p^{m-\ell} \right\}_{i \in [k]} : (\mathbf{x}_i, \mathbf{w}_i) \in \mathbb{Z}_p^m \text{ satisfies } \mathcal{C} \quad \forall i \in [k] \right\} .$$

We refer to Appendix C.2 for a proof.

**Shared permutation argument.** Permutation arguments [BG12, BCC<sup>+</sup>16, MBKM19] can be used for arguing correctness of a shuffle  $\sigma$ . *PlonK*'s permutation argument [GWC19] is used for enforcing that the wires which are supposed to be equal have indeed been instantiated with the same value. This is done by committing to a polynomial  $Z$  (see step 4 of the *PlonK.Prove* routine from Figure 1), which satisfies the permutation identities (see Appendix A.1) with respect to the wire polynomials  $A, B, C$ .

We observe that it is possible to share the permutation argument across all proofs for the same circuit, since the permutation  $\sigma$  is common to all of them. To do so, we linearly batch the wire polynomials  $A_j, B_j, C_j$  for every proof  $j \in [k]$ , with some uniformly sampled coefficient  $\delta \in \mathbb{Z}_p$  as follows:

$$\widehat{A} := \sum_{j=1}^k \delta^j A_j \quad \widehat{B} := \sum_{j=1}^k \delta^j B_j \quad \widehat{C} := \sum_{j=1}^k \delta^j C_j .$$

If each of the polynomial triples  $(A_j, B_j, C_j)$  satisfies the copy-constraints induced by permutation  $\sigma$ , so will the batched triple  $(\widehat{A}, \widehat{B}, \widehat{C})$ , which guarantees correctness. The converse is also true with overwhelming probability over the choice of  $\delta$ , which assures soundness (this intuition can be formalized via a forking argument). Thus, the permutation argument for  $k$  proofs can be achieved with just one  $Z$  polynomial instead of  $k$  of them (see steps 5 and 6 of the *aPlonK.Prove* routine from Figure 4). The number of permutation identities is consequently reduced from  $2k$  to just 2.

$aPlonK.Setup(1^\lambda, \mathcal{C} := \{a_i, b_i, c_i, q_L, q_R, q_O, q_M, q_C\}_{i \in [n]}, k)$ :

- 1:  $q_L(X) := \sum_{i=1}^n q_L L_{i,n}(X)$ ; define  $q_R, q_O, q_M, q_C$  analogously
- 2:  $\sigma : [3n] \rightarrow [3n]$  be  $\mathcal{C}_\sigma$ ;  $pp.polys := (q_L, q_R, q_O, q_M, q_C, S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3})$
- 3:  $(ck, vk) \leftarrow \Psi.Setup(1^\lambda, n, k)$ ;  $\mu_{pp} \leftarrow \Psi.Commit-Polys(ck, pp.polys)$
- 4: **return**  $pp := (n, \sigma, ck, vk, \mu_{pp}, pp.polys)$

$aPlonK.Prove(pp := (n, \sigma, ck, \_, \mu_{pp}, pp.polys), \{\mathbf{x}_j\}_{j \in [k]}, \{\mathbf{w}_j\}_{j \in [k]})$ :

- 1:  $\tilde{w}_j := (\mathbf{x}_j, \mathbf{w}_j)$  for all  $j \in [k]$
- 2:  $A_j(X) := \sum_{i=1}^n \tilde{w}_j a_i L_{i,n}(X)$ ;  $B_j(X) := \sum_{i=1}^n \tilde{w}_j b_i L_{i,n}(X)$ ;  $C_j(X) := \sum_{i=1}^n \tilde{w}_j c_i L_{i,n}(X)$  for all  $j \in [k]$
- 3:  $\mathbf{W} := (A_1, B_1, C_1, \dots, A_k, B_k, C_k)$ ;  $\mu_w := \Psi.Commit-Polys(ck, \mathbf{W})$
- 4:  $\beta \leftarrow Hash(\mu_w)$ ;  $\gamma \leftarrow Hash(\beta)$ ;  $\delta \leftarrow Hash(\gamma)$
- 5:  $\hat{A}(X) := \sum_{j=1}^k \delta^j A_j(X)$ ;  $\hat{B}(X) := \sum_{j=1}^k \delta^j B_j(X)$ ;  $\hat{C}(X) := \sum_{j=1}^k \delta^j C_j(X)$
- 6: compute polynomial  $Z(X)$ , satisfying  $perm-ids_{\beta, \gamma}^\sigma(\hat{A}, \hat{B}, \hat{C}, Z)$  on  $\mathcal{H}_n$  ▷ See Appendix A.1
- 7:  $\mu_z \leftarrow \Psi.Commit-Polys(ck, Z)$
- 8:  $F_j(X) := (q_L A_j + q_R B_j + q_O C_j + q_M A_j B_j + q_C + Pl_{\mathbf{x}_j})(X)$  for all  $j \in [k]$
- 9:  $ids(X) := \bigcup_{j \in [k]} F_j(X) \cup perm-ids_{\beta, \gamma}^\sigma(\hat{A}, \hat{B}, \hat{C}, Z)$
- 10:  $\alpha \leftarrow Hash(\delta, \mu_z)$
- 11:  $T(X) := (\sum_{i \in [ids]} \alpha^i ids_i(X)) / Z_{\mathcal{H}_n}(X)$  ▷ This division is exact if the identities hold over  $\mathcal{H}_n$
- 12:  $\mu_t \leftarrow \Psi.Commit-Polys(ck, T)$
- 13:  $\xi \leftarrow Hash(\alpha, \mu_t)$
- 14:  $ev(\text{com}, \mathbf{f}, x) := (\Psi.Open(ck, \text{com}, x, \mathbf{f}), \Psi.Commit-Evals(ck, \mathbf{f}(x)))$
- 15:  $(\pi_w, \nu_w) \leftarrow ev(\mu_w, \mathbf{W}, \xi)$   $(\pi_z, \nu_z) \leftarrow ev(\mu_z, Z, \xi)$   $(\bar{\pi}_z, \bar{\nu}_z) \leftarrow ev(\mu_z, Z, \omega \xi)$   $(\pi_t, \nu_t) \leftarrow ev(\mu_t, T, \xi)$
- 16:  $(\pi_{pp}, \nu_{pp}) \leftarrow ev(\mu_{pp}, pp.polys, \xi)$
- 17: compute  $\pi_{meta}$ , a PoK  $\{\mathbf{w}_{meta} : \mathcal{R}_{n,k}((\alpha, \beta, \gamma, \delta, \xi, \nu_w, \nu_z, \bar{\nu}_z, \nu_t, \nu_{pp}, \{\mathbf{x}_j\}_{j \in [k]}), \mathbf{w}_{meta}) = 1\}$  ▷ See Section 4.1
- 18: **return**  $\pi := (\mu_w, \mu_z, \mu_t, \nu_w, \nu_z, \bar{\nu}_z, \nu_t, \nu_{pp}, \pi_w, \pi_z, \bar{\pi}_z, \pi_t, \pi_{pp}, \pi_{meta})$

$aPlonK.Verify(pp := (n, \_, \_, vk, \mu_{pp}, \_), \{\mathbf{x}_j\}_{j \in [k]}, \pi := (\mu_w, \mu_z, \mu_t, \nu_w, \nu_z, \bar{\nu}_z, \nu_t, \nu_{pp}, \pi_w, \pi_z, \bar{\pi}_z, \pi_t, \pi_{pp}, \pi_{meta}))$ :

- 1:  $\beta \leftarrow Hash(\mu_w)$ ;  $\gamma \leftarrow Hash(\beta)$ ;  $\delta \leftarrow Hash(\gamma)$ ;  $\alpha \leftarrow Hash(\delta, \mu_z)$ ;  $\xi \leftarrow Hash(\alpha, \mu_t)$
- 2:  $v(\text{com}, \mathbf{v}, x, \pi) := \Psi.Check(vk, \text{com}, x, \mathbf{v}, \pi)$
- 3:  $b_\mu := v(\mu_w, \nu_w, \xi, \pi_w) \wedge v(\mu_z, \nu_z, \xi, \pi_z) \wedge v(\mu_z, \bar{\nu}_z, \omega \xi, \bar{\pi}_z) \wedge v(\mu_t, \nu_t, \xi, \pi_t) \wedge v(\mu_{pp}, \nu_{pp}, \xi, \pi_{pp})$
- 4: let  $b_{meta}$  be the result of verifying  $\pi_{meta}$  w.r.t. relation  $\mathcal{R}_{n,k}((\alpha, \beta, \gamma, \delta, \xi, \nu_w, \nu_z, \bar{\nu}_z, \nu_t, \nu_{pp}, \{\mathbf{x}_j\}_{j \in [k]}), \cdot)$
- 5: **return**  $b_\mu \wedge b_{meta}$

**Fig. 4.** The  $aPlonK$  proving system, based on multi-polynomial commitment scheme  $\Psi$ .

**Using a multi-polynomial commitment.** Replacing the polynomial commitment scheme used by  $PlonK$  by a multi-polynomial commitment scheme can lead to major improvements in proof size and verification



$$\begin{aligned}
& \mathcal{R}_{n,k}((\alpha, \beta, \gamma, \delta, \xi, \nu_w, \nu_z, \nu_{\bar{z}}, \nu_t, \nu_{pp}, \{\mathbf{x}_j\}_{j \in [k]}), (\{\mathbf{a}_j, \mathbf{b}_j, \mathbf{c}_j\}_{j \in [k]}, \mathbf{z}, \bar{\mathbf{z}}, \mathbf{t}, \mathbf{e}_{q_L}, \mathbf{e}_{q_R}, \mathbf{e}_{q_0}, \mathbf{e}_{q_M}, \mathbf{e}_{q_C}, \mathbf{e}_{s_1}, \mathbf{e}_{s_2}, \mathbf{e}_{s_3})) := \\
& \text{id}_j := \mathbf{e}_{q_L} \mathbf{a}_j + \mathbf{e}_{q_R} \mathbf{b}_j + \mathbf{e}_{q_0} \mathbf{c}_j + \mathbf{e}_{q_M} \mathbf{a}_j \mathbf{b}_j + \mathbf{e}_{q_C} + \text{Pl}_{\mathbf{x}_j}(\xi) \quad \forall j \in [k] \\
& \hat{\mathbf{a}} := \sum_{j=1}^k \delta^j \mathbf{a}_j; \quad \hat{\mathbf{b}} := \sum_{j=1}^k \delta^j \mathbf{b}_j; \quad \hat{\mathbf{c}} := \sum_{j=1}^k \delta^j \mathbf{c}_j \\
& \text{perm-id}_1 := (\hat{\mathbf{a}} + \beta \xi + \gamma)(\hat{\mathbf{b}} + \beta \eta \xi + \gamma)(\hat{\mathbf{c}} + \beta \eta' \xi + \gamma) \mathbf{z} - (\hat{\mathbf{a}} + \beta \mathbf{e}_{s_1} + \gamma)(\hat{\mathbf{b}} + \beta \mathbf{e}_{s_2} + \gamma)(\hat{\mathbf{c}} + \beta \mathbf{e}_{s_3} + \gamma) \bar{\mathbf{z}} \\
& \text{perm-id}_2 := (\mathbf{z} - 1) \text{L}_{1,n}(\xi) \\
& b_{ids} := (Z_{\mathcal{H}_n}(\xi) \cdot \mathbf{t} = (\sum_{j=1}^k \alpha^{j-1} \text{id}_j) + \alpha^k \text{perm-id}_1 + \alpha^{k+1} \text{perm-id}_2) \\
& \mathbf{w} := (\mathbf{a}_1, \mathbf{b}_1, \mathbf{c}_1, \dots, \mathbf{a}_k, \mathbf{b}_k, \mathbf{c}_k) \text{ and } \mathbf{e}_{pp} := (\mathbf{e}_{q_L}, \mathbf{e}_{q_R}, \mathbf{e}_{q_0}, \mathbf{e}_{q_M}, \mathbf{e}_{q_C}, \mathbf{e}_{s_1}, \mathbf{e}_{s_2}, \mathbf{e}_{s_3}) \\
& \text{return } b_{ids} \wedge \nu_w = \text{Commit-Evals}(\mathbf{w}) \wedge \nu_z = \text{Commit-Evals}(\mathbf{z}) \wedge \nu_{\bar{z}} = \text{Commit-Evals}(\bar{\mathbf{z}}) \\
& \quad \wedge \nu_t = \text{Commit-Evals}(\mathbf{t}) \wedge \nu_{pp} = \text{Commit-Evals}(\mathbf{e}_{pp}) .
\end{aligned}$$

**Fig. 5.** Meta-verification relation for aggregating  $k$  proofs of  $n$ -constraints circuits.  $\eta_2$  and  $\eta_3$  are non-quadratic residues over  $\mathbb{Z}_p$ , see Section A.1.

time. It allows us to commit to all wire polynomials together in one single multi-polynomial commitment with sublinear size in the number of aggregated proofs  $k$ . With our multi-polynomial commitment scheme from Figure 2, the commitment size would be constant (1  $\mathbb{G}_t$  element) instead of linear in  $k$ , and the commitment verification complexity would be  $\mathcal{O}(\log k)$  instead of  $\mathcal{O}(k)$ .

This technique achieves sublinear complexity (in  $k$ ) on commitment verification operations. However, the verifier still needs to check all the identities, which involves a  $\mathcal{O}(k)$  number of scalar operations. For that, the verifier needs to receive all the evaluations of the committed polynomials (whose validity can be asserted through the already verified evaluation commitment) and use them to verify the identities. Our next technique addresses this issue.

**Meta-verification.** The verification of identities only involves scalar operations over  $\mathbb{Z}_p$ , but this is the native field of  $\mathcal{PlonK}$  circuits. This opens the possibility of, instead of verifying the identities directly, verifying a  $\mathcal{PlonK}$  proof that the identities are correct. Such proof would need to ensure that:

- the prover knows evaluations satisfying all the identities,
- such evaluations coincide with the evaluations verified during the multi-polynomial commitment check.

We formally describe the meta-verification equation in Figure 5. It is parametrized by the number of constraints in the circuit  $n$ , and the number of aggregated proofs  $k$ . The public inputs to the meta-verification circuit are  $(\alpha, \beta, \gamma, \delta, \xi, \nu_w, \nu_z, \nu_{\bar{z}}, \nu_t, \nu_{pp}, \{\mathbf{x}_j\}_{j \in [k]})$ , where  $\alpha, \beta, \gamma, \delta, \xi$  are Fiat-Shamir sampled scalars;  $\nu_w, \nu_z, \nu_t, \nu_{pp}$  are evaluation commitments of (respectively) the wire polynomials, Z polynomial, T polynomial and setup polynomials at  $\xi$ ;  $\nu_{\bar{z}}$  is (a commitment to) the evaluation of polynomial Z at  $\omega\xi$ ; and for every  $j \in [k]$ ,  $\mathbf{x}_j$  is the vector of public inputs to the  $j$ -th statement. On the other hand, the secret inputs to the meta-verification relation are the actual polynomial evaluations at  $\xi$  and  $\omega\xi$  of the committed polynomials.

By just verifying a single  $\mathcal{PlonK}$  proof, the verifier can assert the correctness of all identities without performing a  $\mathcal{O}(k)$  number of scalar operations. On the other hand, we make three observations that deserve attention:

- (i) The  $\text{Commit-Evals}$  algorithm needs to be modeled in a  $\mathcal{PlonK}$  circuit. There is flexibility for the choice of such algorithm, but modeling any commitment scheme that is binding will require a significant number of constraints.
- (ii) The verifier complexity is still  $\mathcal{O}(k)$  on scalar operations, given the public inputs  $\{\mathbf{x}_j\}_{j \in [k]}$  to the meta-verification circuit.

- (iii) This technique imposes a bound on  $k$ , the number of aggregated proofs, since the meta-verification circuit size is linear in  $k$  and there is an inherent upper-bound on the size of  $\mathcal{PlonK}$  circuits.

The first issue can be partially solved by instantiating **Commit-Evals** with a SNARK-friendly hash function like Poseidon [GKR+21] or Anemoi [BBC+22], which can be implemented with a moderate number of constraints. Alternatively, in Section 4.2, we show how to build a commitment scheme that can be involved very efficiently in a  $\mathcal{PlonK}$  statement. This can be of independent interest.

We refer to Section 4.3 for details on how the second issue can be addressed, depending on the circuit being proved.

Finally, if the first issue is satisfactorily solved, the upper-bound on  $k$  could be sufficient for most applications. Even if it is not, one could consider proving the meta-verification relation with several  $\mathcal{PlonK}$  proofs, what can lead to a second layer of aggregation and iterate this process if necessary.

## 4.2 Commitments in $\mathcal{PlonK}$ statements

We present a mechanism that allows the data committed in a public commitment to be involved in a  $\mathcal{PlonK}$  statement. More concretely, let **Com** be a commitment scheme for vectors over  $\mathbb{Z}_p$ . We enhance the  $\mathcal{PlonK}$  proving system to support the following relation:

$$\text{PoK} \left\{ (\mathbf{w}, \mathbf{w}') \in \mathbb{Z}_p^m \times \mathbb{Z}_p^{m'} : (\mathbf{x}, \mathbf{w}, \mathbf{w}') \in \mathbb{Z}_p^{\ell+m+m'} \text{ satisfies } \mathcal{C} \wedge \text{Com}(\mathbf{w}) = \text{com} \right\} . \quad (4)$$

The statement being proved is thus parametrized by  $\mathcal{C}$ ,  $\mathbf{x}$  and **com**.

For that, we define **Com** of a vector  $\mathbf{w} \in \mathbb{Z}_p^m$  as the Kate commitment [KZG10] to any polynomial  $f$  that evaluates to  $w_i$  on  $\omega^{i-1}$  for all  $i \in [m]$ .<sup>10</sup> Similarly to how public inputs are treated in  $\mathcal{PlonK}$ , we will dedicate a section of  $m$  constraints to the link between **Com** and the circuit wires. Say, constraints  $i^*$  to  $i^*+m-1$  for some  $i^*$ . For that, we define a selector  $\mathbf{q}_{\text{com}}$  such that  $\mathbf{q}_{\text{com}}(\omega^i) = -1$  if  $i \in [i^*, i^*+m-1]$  or 0 otherwise. Furthermore, we deactivate (set to zero) all other selectors in this range of constraints except  $\mathbf{q}_L$ , which evaluates to 1 in it, allowing us to “fetch” the values inside **com** into **a**-wires, which will be then used across the circuit. The  $\mathcal{PlonK}$  identity will consequently get added the new term  $\mathbf{q}_{\text{com}}(X) \cdot \text{com}(\omega^{-i^*} X)$ . Observe that by setting  $i^* = 0$ , one can avoid having a new evaluation point  $\omega^{-i^*} \xi$ . This can be achieved by shifting away the constraints dedicated to public inputs, which in the original  $\mathcal{PlonK}$  were defined to be the first  $\ell$ . Note that if we do not want to open the whole vector  $\mathbf{w}$ , we can activate  $\mathbf{q}_{\text{com}}$  on a smaller range; then the verifier can dynamically adapt by querying  $\text{com}(\omega^{-j^*} X)$  for a chosen  $j$ . This can prove useful if we want to use this technique in the context of vector commitments.

**Hybrid statements.** Proving statements about committed data is a powerful tool that can be used for constructing hybrid proving systems [CGM16, AGM18, CFtQ19]. Such commitments can be the meeting point between a SNARK proof and other systems, e.g., a sigma protocol asserting the validity of the committed data with respect to some algebraic statement. The binding property of the commitment ensures that both proofs “talk about the same data”.

In particular, we can choose the **Commit-Evals** algorithm from our multi-polynomial commitment scheme (Figure 2) to be the above **Com**. This provides a satisfying solution to issue (i) from the previous section. In that case, the meta-verification circuit (Figure 5) could be extended to also handle the PoK from relation (3), see Figure 2.

**Randomized circuits.** Another remarkable application of involving commitments in the statement is being able to express *randomized circuits*. Randomness can be used to simplify the verification of certain

<sup>10</sup>Note that  $f$  can be randomly chosen among the set of all polynomials bound to  $\mathbf{v}$  on  $1, \omega, \dots, \omega^{m-1}$ , this results in a potentially perfectly hiding commitment.

computations: given a Boolean circuit  $C$ , it is often possible to find a (smaller) circuit  $\widehat{C}$ , taking an extra input  $r$ , which is equivalent to  $C$  in the sense that, for every  $x$ :

$$\Pr_{r \leftarrow \mathbb{Z}_p} [C(x) = 1 \Rightarrow \widehat{C}(x; r) = 1] = 1 \quad \text{and} \quad \Pr_{r \leftarrow \mathbb{Z}_p} [\widehat{C}(x; r) = 1 \Rightarrow C(x) = 1] \approx 1 .$$

Verifying the more efficient  $\widehat{C}(x; r) = 1$  for a uniformly sampled  $r$  is an overwhelming evidence that  $C(x) = 1$ . However, it is important to choose  $r$  uniformly and independently of  $x$ . When proving a randomized circuit in  $\mathcal{PlonK}$ ,  $r$  will be treated as an additional public input. We must guarantee that  $r$  cannot be biased by the prover and that the prover cannot change its secret input based on  $r$ . On the other hand, the circuit trace that the prover commits to depends on  $r$ , this is natural, since  $r$  is involved in the constraint system. The solution is to have the prover commit to the witness seed<sup>11</sup> using  $\text{Com}$ , then derive  $r$  from such commitment through the Fiat-Shamir heuristic and finally, completing the rest of the trace (which now depends on  $r$ ).

### 4.3 Hiding public inputs

Our techniques allowed us to significantly reduce the verifier complexity, for the most part now being logarithmic in the number of aggregated proofs  $k$ . However, as in SnarkPack [GMN20], our verifier complexity is still linear in  $k$ , due to the  $\ell$  public inputs per proof to process. This seems an inherent limitation.

However, depending on the application, such limitation could be relaxed. For example, the verifier may have access to a commitment to the relevant public inputs and may be interested in simply checking that verification passes with respect to some opening of the commitment. In other scenarios, the circuit public inputs themselves could be irrelevant to the verifier, who only wants to assert some relation between them. For example, in some incremental computations like a transactional rollup of a blockchain: where each of the  $k$  proofs takes 2 public inputs, an initial rollup state and a final (modified) state, and the verifier is simply interested in asserting that the final state of a given proof matches the initial state of the next proof, but not on the actual value of such intermediate states. This *rollup-like* configuration is used for our benchmarks (Section 5).

In those cases, our techniques can be very naturally extended to achieve actual sublinear verification complexity by performing the relevant checks on public inputs in the meta-verification circuit. This can be seen as new mechanism to implement a weak form of IVC.

## 5 Implementation and evaluation

We have implemented the algorithms described in this work and evaluated their performance in a series of benchmarks presented in Section 5.2. Our source code is written in OCaml with bindings to C implementations of the heaviest cryptographic functions. We use the BLS12-381 elliptic curve [Bow17] for pairings through bindings to the *blst* library [Sup21]. Our implementation is publicly available as open-source [Nom22].

Our comparisons are performed between  $\mathcal{PlonK}$  and  $a\mathcal{PlonK}$  for aggregating a batch of  $k$  different proofs of a constraint system of  $n$  constraints. In the case of  $\mathcal{PlonK}$ , we use the KZG polynomial commitment and implement a simple proof aggregation, but which uses our shared permutation argument and the batch verification optimizations applicable to KZG [GWC19]. In the case of  $a\mathcal{PlonK}$ , we implement the scheme described in Figure 4 instantiated with a multi-polynomial commitment scheme constructed from the KZG polynomial commitment scheme (Figure 7) and the IPA argument from Figure 3. Our experimental results from Section 5.2 show that the performance and proof size improvements of  $a\mathcal{PlonK}$  are significant even against the optimized version of  $\mathcal{PlonK}$  that we compare it with, which in turn is a lot more performant than the naïve parallel execution of standard  $\mathcal{PlonK}$ .

<sup>11</sup>In  $\mathcal{PlonK}$ , the term “witness” usually refers the whole trace of the circuit being verified. Such trace is typically derivable from a succinct witness that we call the witness seed: a value that determines (and from which one can efficiently compute) the rest of the trace.

**Table 1.** Proof size comparison.  $k$  is the number of aggregated proofs (formulas valid for  $k \geq 3$ ).

	$\mathbb{Z}_p$	$\mathbb{G}_1$	$\mathbb{G}_2$	$\mathbb{G}_t$
$\mathcal{PlonK}$	$3k + 12$	$3k + 4$	0	0
$a\mathcal{PlonK}$	19	$2\lceil \log_2 3k \rceil + 10$	2	$2\lceil \log_2 3k \rceil + 3$

**Table 2.** Verifier complexity comparison.  $k$  is the number of aggregated proofs,  $n$  is the circuit size,  $\ell$  is the number of public inputs per proof,  $n' = \mathcal{O}(k + \ell)$  is the meta-verification circuit size (see Table 4),  $\ell' (\leq n')$  is the number of public inputs to it (in our implementation, it is set to 15), and  $\kappa = \lceil \log_2(3k) \rceil$ .

Operation	$\mathcal{PlonK}$	$a\mathcal{PlonK}$
inv	$k\ell + 2k$	$\kappa + \ell' + 2$
$\mathbb{Z}_p$ mul	$k \log_2(n) + 2k\ell + 11k + 20$	$5\kappa + \log_2(n') + 2\ell' + 35$
add	$2k\ell + 11k + 14$	$\kappa + 2\ell' + 28$
$\mathbb{G}_1$ mul	$3k + 15$	$2\kappa + 13$
add	$3k + 16$	$2\kappa + 17$
$\mathbb{G}_t$ mul	0	$2\kappa$
add	0	$2\kappa + 3$
Pairing	2	4
Hash	9	$\kappa + 10$

*Remark 3.* We leverage the homomorphic properties of the IPA statement, see relation (2), to perform only one IPA instead of four: for the evaluation of the commitments to (i) public parameter polynomials, (ii) wire polynomials, (iii) the permutation polynomial and (iv) the  $\mathbb{T}$  polynomial. This is possible if all inner-product arguments use the same  $r$ , since in that case:

$$\langle \boldsymbol{\mu}, \mathbf{G} \rangle = C \quad \wedge \quad \langle \mathbf{r}, \boldsymbol{\mu} \rangle = P \quad \wedge \quad \langle \boldsymbol{\mu}', \mathbf{G} \rangle = C' \quad \wedge \quad \langle \mathbf{r}, \boldsymbol{\mu}' \rangle = P'$$

if and only if, with overwhelming probability over the choice of  $\zeta$ :

$$\langle \boldsymbol{\mu} + \zeta \boldsymbol{\mu}', \mathbf{G} \rangle = C + \zeta C' \quad \wedge \quad \langle \mathbf{r}, \boldsymbol{\mu} + \zeta \boldsymbol{\mu}' \rangle = P + \zeta P' .$$

## 5.1 Theoretical results

We present a detailed comparison between  $\mathcal{PlonK}$  and  $a\mathcal{PlonK}$  on their proof size, verifier complexity and prover complexity in terms of scalar field and group elements/operations. In the rest of this section  $k$  represents the number of aggregated statements,  $n$  is the common circuit size measured in number of constraints and  $\ell$  is the number of public inputs for each atomic statement.

*Proof size (Table 1).* Observe that the aggregated proof sizes of both  $\mathcal{PlonK}$  and  $a\mathcal{PlonK}$  are independent on the circuit size  $n$  and only depend on the number of aggregated statements  $k$ . With  $\mathcal{PlonK}$ , the proof size is linear in  $k$ , whereas with  $a\mathcal{PlonK}$ , it is logarithmic.

*Verifier complexity (Table 2).* The verifier complexity of  $\mathcal{PlonK}$  is  $\mathcal{O}(k \log n + k\ell)$ , whereas the verifier complexity of  $a\mathcal{PlonK}$  is  $\mathcal{O}(\log k + \ell')$ , where  $\ell' = \mathcal{O}(k\ell)$  is the number of public inputs to the meta-verification circuit. Note that, when applicable (e.g. for a rollup), our technique for hiding public inputs can lead to a constant  $\ell'$ . In that case, the verifier complexity of  $a\mathcal{PlonK}$  would be  $\mathcal{O}(\log k)$ . This is an important difference with respect to SnarkPack [GMN20], where the verification of aggregated proofs is linear in the number of public inputs.

**Table 3.** Prover complexity comparison.  $k$  is the number of aggregated proofs,  $n$  is the circuit size,  $n' = \mathcal{O}(k + \ell)$  is the meta-verification circuit size (see Table 4), and  $K$  is the first power of two over  $3k$ .

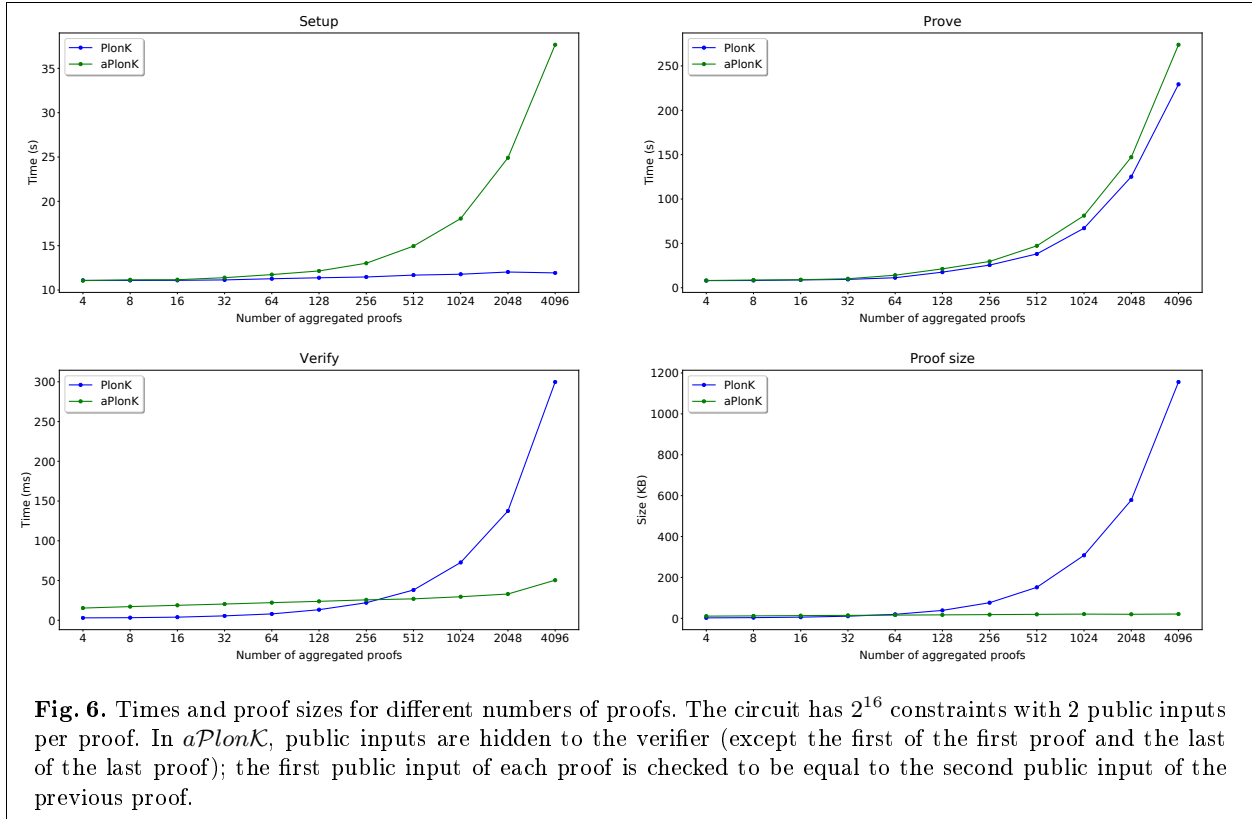
Operation	$\mathcal{PlonK}$	$a\mathcal{PlonK}$
inv	$n + 3k + 3$	$\mathcal{PlonK} + \log_2(K) + 2n' + 6$
$\mathbb{Z}_p$ mul	$\frac{15}{2}kn \log_2(n) + 45kn + 7n \log_2(n) + 105n + 3k + 9$	$\mathcal{PlonK} + 3K + 3 \log_2(K) + \frac{29}{2}n' \log_2(n') + 155n' + 16$
add	$15kn \log_2(n) + 42kn + 14n \log_2(n) + 95n + 3k + 11$	$\mathcal{PlonK} + 2K + \log_2(K) + 29n' \log_2(n') + 140n' + 13$
$\mathbb{G}_1$ mul	$3kn + 6n + 3k + 12$	$\mathcal{PlonK} + 9n' + 8$
add	$3kn + 6n + 3k + 12$	$\mathcal{PlonK} + 9n' + 9k + 27$
$\mathbb{G}_2$ mul	0	$3K - 2$
add	0	$2K - 1$
$\mathbb{G}_t$ add	0	$2K + k - 2$
Pairing	0	$2K + 3k + 2$
Hash	9	$\log_2(K) + 10$

**Table 4.** Meta-verification circuit size (measured in number of constraints). For circuits that aggregate  $k$  proofs and for two different alternatives implementations of **Commit-Evals** and two different schedules of public inputs.

	Without public inputs	Rollup-like public inputs
Base cost (ignoring <b>Commit-Evals</b> )	$25k + 344$	$28k + 339$
Extra cost if <b>Commit-Evals</b> = Poseidon	$200k + 1061$	$210k + 1061$
Extra cost if <b>Commit-Evals</b> = (KZG.Commit $\circ$ IFFT)	$3k + 12$	$3k + 12$

*Prover complexity (Table 3).* The prover complexity of  $\mathcal{PlonK}$  is  $\mathcal{O}(kn \log n)$ . The  $a\mathcal{PlonK}$  prover requires more operations but stays in the same order of complexity with respect to operations in  $\mathbb{Z}_p$  and  $\mathbb{G}_1$ . It additionally requires  $\mathcal{O}(\log k)$  hashes, pairings, and operations in  $\mathbb{G}_2$ ,  $\mathbb{G}_t$ . However, as evidenced by our experimental results from Section 5.2, the overhead of these additional computations is not very significant compared to the complexity of the rest of computations.

*Meta-verification circuit size (Table 4).* We present the size measured in number of constraints of the meta-verification circuit which models  $\mathcal{R}_{n,k}$  (described in Figure 5). Such numbers are helpful to interpret the verifier and prover complexities given in Tables 2 and 3 respectively. Table 4 describes the base cost of implementing relation  $\mathcal{R}_{n,k}$  when ignoring the logic related to **Commit-Evals**. We consider two different schedules of public inputs: (i) a circuit without public inputs, which serves as a lower-bound on the size of the meta-verification circuit; (ii) a rollup-like public input schedule, where each circuit being proved has two public inputs and the meta-verification circuit checks that they are linked in a chain (i.e., that the second public input of one circuit coincides with the first public input of the next). We also present the additional cost of modeling **Commit-Evals** through the Poseidon (implemented following the results of [ASTW22]) and the cost of modeling it with a Kate commitment as described in Section 4.2. Note how, even though Poseidon is a SNARK-friendly hash function, implementing **Commit-Evals** with it would be very costly and the main factor on the meta-verification circuit size. This would introduce a relatively small limit on the number of proofs that can be aggregated with this technique. On the other hand, when modeling **Commit-Evals** as described in Section 4.2, the overhead is minimal compared to the rest of the verification circuit, improving by an order of magnitude the number of proofs that can be aggregated.



**Fig. 6.** Times and proof sizes for different numbers of proofs. The circuit has  $2^{16}$  constraints with 2 public inputs per proof. In *aPlonK*, public inputs are hidden to the verifier (except the first of the first proof and the last of the last proof); the first public input of each proof is checked to be equal to the second public input of the previous proof.

## 5.2 Experimental results

We present our experimental results in Figure 6, which includes a comparison of the setup, proving and verification times, as well as proof sizes of *PlonK* and *aPlonK*, for different aggregation sizes. All experiments were performed on a 2.9GHz Intel Xeon Platinum 8375C vCPU with 1 TB of RAM and 128 processors.

We use a circuit of  $n = 2^{16}$  constraints. Our circuit performs a computation which involves several additions and multiplications across various inputs, two of which are considered public. That way we guarantee that all selectors  $q_L, q_R, q_O, q_M, q_C$  are non-trivial. Nevertheless, note that the complexity of all algorithms is independent of the actual architecture of the circuit and their performance only depends on the number of constraints  $n$ . We choose a rollup-like schedule of public inputs, joined in a chain as described in the previous section.

The logic associated to *Commit-Evals* in the meta-verification circuit should be implemented through our method for involving commitments outside of *PlonK* (Section 4.2). While we counted these extra constraints in Table 4, such method has not been implemented for simplicity. We expect the overhead of having this logic into account to be negligible given that it only increases the meta-verification circuit by  $3k + 12$  constraints.

*Setup.* The setup of *PlonK* is constant since the circuit of interest always has  $n = 2^{16}$  constraints. However, in *aPlonK*, it is linear in  $k$ . This is because the size of the meta-verification circuit grows linearly with the number of aggregated proofs. Fortunately, the impact of *aPlonK*'s setup is in the order of seconds for aggregating thousands or proofs. Furthermore, note that the setup performance is not critical, as it is precomputed only once.

*Proving.* All experiments have used parallelization over all the 128 available cores. When  $k \leq 128$  a core is assigned for each proof. After that threshold we can expect a linear growth since each core will need to

produce more than one proof. One CPU was also in charge of orchestrating the distribution and computing the meta-verification proof, a step which was performed sequentially. The difference between *PlonK* and *aPlonK* proving times comes from the proving time of the meta-verification circuit. We can see a overhead of approximately 20% (45 seconds) for *aPlonK* proving time for  $2^{12}$  aggregated proofs. This represents 1% of the total machine time. Furthermore, such overhead would be even less important if circuits were larger, as the complexity of our aggregation routines is independent of  $n$ .

*Verification and proof sizes.* Our experimental results on verification corroborate the fact that *PlonK* is linear while *aPlonK* is logarithmic in the number of aggregated proofs. *aPlonK* becomes more efficient after a threshold of about  $k = 300$  proofs. On the other hand, the proof size of *aPlonK* becomes smaller starting from  $k = 64$  proofs.

## Acknowledgments

We are very thankful to Antonio Locascio, Danny Willems, Julien Coolen, Marco Stronati, Marina Polubelova and Victor Dumitrescu, developers and co-authors of our implementation [Nom22], for very fruitful discussions and all their help and feedback.

We would also like to thank Mary Maller, for her feedback and clarifications about SnarkPack in the early stages of this project.

## References

- AGM18. Shashank Agrawal, Chaya Ganesh, and Payman Mohassel. Non-interactive zero-knowledge proofs for composite statements. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 643–673. Springer, Heidelberg, August 2018.
- ASTW22. Miguel Ambrona, Anne-Laure Schmitt, Raphael R. Toledo, and Danny Willems. New optimization techniques for PlonK’s arithmetization. Cryptology ePrint Archive, Report 2022/462, 2022. <https://eprint.iacr.org/2022/462>.
- BBB<sup>+</sup>18. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018.
- BBC<sup>+</sup>22. Clémence Bouvier, Pierre Briaud, Pyrros Chaidos, Léo Perrin, and Vesselin Velichkov. Anemoi: Exploiting the link between arithmetization-orientation and czz-equivalence. Cryptology ePrint Archive, Paper 2022/840, 2022. <https://eprint.iacr.org/2022/840>.
- BBHR19. Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable zero knowledge with no trusted setup. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 701–732. Springer, 2019.
- BCC<sup>+</sup>16. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 327–357. Springer, Heidelberg, May 2016.
- BCG<sup>+</sup>13. Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108. Springer, Heidelberg, August 2013.
- BCG<sup>+</sup>14. Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014.
- BCG<sup>+</sup>17. Jonathan Bootle, Andrea Cerulli, Essam Ghadafi, Jens Groth, Mohammad Hajiabadi, and Sune K. Jakobsen. Linear-time zero-knowledge proofs for arithmetic circuit satisfiability. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part III*, volume 10626 of *LNCS*, pages 336–365. Springer, Heidelberg, December 2017.

- BCL<sup>+</sup>21. Benedikt Bünz, Alessandro Chiesa, William Lin, Pratyush Mishra, and Nicholas Spooner. Proof-carrying data without succinct arguments. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 681–710, Virtual Event, August 2021. Springer, Heidelberg.
- BCMS20. Benedikt Bünz, Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. Proof-carrying data from accumulation schemes. Cryptology ePrint Archive, Report 2020/499, 2020. <https://eprint.iacr.org/2020/499>.
- BDFG20. Dan Boneh, Justin Drake, Ben Fisch, and Ariel Gabizon. Efficient polynomial commitment schemes for multiple points and polynomials. Cryptology ePrint Archive, Report 2020/081, 2020. <https://eprint.iacr.org/2020/081>.
- BFM88. Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th ACM STOC*, pages 103–112. ACM Press, May 1988.
- BG12. Stephanie Bayer and Jens Groth. Efficient zero-knowledge argument for correctness of a shuffle. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 263–280. Springer, Heidelberg, April 2012.
- BGH19a. Sean Bowe, Jack Grigg, and Daira Hopwood. Halo: Recursive proof composition without a trusted setup. *IACR Cryptol. ePrint Arch.*, page 1021, 2019.
- BGH19b. Sean Bowe, Jack Grigg, and Daira Hopwood. Halo: Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Report 2019/1021, 2019. <https://eprint.iacr.org/2019/1021>.
- BMRS20. Joseph Bonneau, Izaak Meckler, Vanishree Rao, and Evan Shapiro. Mina: Decentralized cryptocurrency at scale, 2020. Whitepaper. <https://docs.minaprotocol.com/static/pdf/technicalWhitepaper.pdf>.
- BN06. Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 390–399. ACM Press, October / November 2006.
- Bow17. Sean Bowe. Bls12-381: New zk-snark elliptic curve construction, 2017. ECC Posts. <https://electriccoin.co/blog/new-snark-curve/>.
- CCW19. Alessandro Chiesa, Lynn Chua, and Matthew Weidner. On cycles of pairing-friendly elliptic curves. *SIAM J. Appl. Algebra Geom.*, 3(2):175–192, 2019.
- CDS94. Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Yvo Desmedt, editor, *CRYPTO'94*, volume 839 of *LNCS*, pages 174–187. Springer, Heidelberg, August 1994.
- CFtQ19. Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: Modular design and composition of succinct zero-knowledge proofs. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2075–2092. ACM Press, November 2019.
- CGM16. Melissa Chase, Chaya Ganesh, and Payman Mohassel. Efficient zero-knowledge proof of algebraic and non-algebraic statements with applications to privacy preserving credentials. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 499–530. Springer, Heidelberg, August 2016.
- COS19. Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. Fractal: Post-quantum and transparent recursive proofs from holography. Cryptology ePrint Archive, Report 2019/1076, 2019. <https://eprint.iacr.org/2019/1076>.
- Cra97. Ronald Cramer. Modular design of secure yet practical cryptographic protocols. 1997.
- DFKP13. George Danezis, Cedric Fournet, Markulf Kohlweiss, and Bryan Parno. Pinocchio coin: Building zero coin from a succinct pairing-based proof system. In *Proceedings of the First ACM Workshop on Language Support for Privacy-Enhancing Technologies*, PETShop '13, page 27–30, New York, NY, USA, 2013. Association for Computing Machinery.
- DMP90. Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Non-interactive zero-knowledge with preprocessing. In Shafi Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 269–282. Springer, Heidelberg, August 1990.
- DRZ20. Vanesa Daza, Carla Ràfols, and Alexandros Zacharakis. Updateable inner product argument with logarithmic verifier and applications. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 527–557. Springer, Heidelberg, May 2020.
- FLS90. Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st FOCS*, pages 308–317. IEEE Computer Society Press, October 1990.
- GGPR13. Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.



- GKR<sup>+</sup>21. Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. Poseidon: A new hash function for zero-knowledge proof systems. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 519–535. USENIX Association, August 2021.
- GMN20. Nicolas Gailly, Mary Maller, and Anca Nitulescu. Snarkpack: Practical snark aggregation. In Ittay Eyal and Juan Garay, editors, *FC 2022: 26th International Conference on Financial Cryptography and Data Security*, volume 13411 of *LNCS*, St George's, Grenada, 2020. Springer, Heidelberg, Germany. <https://ia.cr/2021/529>.
- GMO16. Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. ZKBoo: Faster zero-knowledge for Boolean circuits. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016*, pages 1069–1083. USENIX Association, August 2016.
- GMR85. Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems (extended abstract). In *17th ACM STOC*, pages 291–304. ACM Press, May 1985.
- Gro16. Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.
- GS08. Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 415–432. Springer, Heidelberg, April 2008.
- GW19. Ariel Gabizon and Zachary J. Williamson. The turbo-plonk program syntax for specifying snark programs, 2019. Preprint. [https://docs.zkproof.org/pages/standards/accepted-workshop3/proposal-turbo\\_plonk.pdf](https://docs.zkproof.org/pages/standards/accepted-workshop3/proposal-turbo_plonk.pdf).
- GWC19. Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. <https://eprint.iacr.org/2019/953>.
- GYB21. Christopher Goes, Awa Sun Yin, and Adrian Brink. Anoma: Undefined money: A protocol for private, asset-agnostic digital cash and n-party bartering, 2021. <https://anoma.network/papers/whitepaper.pdf>.
- HBHW. Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash protocol specification. <https://zips.z.cash/protocol/protocol.pdf>.
- JKO13. Marek Jawurek, Florian Kerschbaum, and Claudio Orlandi. Zero-knowledge using garbled circuits: how to prove non-algebraic statements efficiently. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 955–966. ACM Press, November 2013.
- KST21. Abhiram Kothapalli, Srinath Setty, and Ioanna Tzialla. Nova: Recursive zero-knowledge arguments from folding schemes. Cryptology ePrint Archive, Report 2021/370, 2021. <https://eprint.iacr.org/2021/370>.
- KZG10. Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Heidelberg, December 2010.
- LMR19. Russell W. F. Lai, Giulio Malavolta, and Viktoria Ronge. Succinct arguments for bilinear group arithmetic: Practical structure-preserving cryptography. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2057–2074. ACM Press, November 2019.
- MBKM19. Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2111–2128. ACM Press, November 2019.
- MGGR13. Ian Miers, Christina Garman, Matthew Green, and Aviel D. Rubin. Zerocoin: Anonymous distributed E-cash from Bitcoin. In *2013 IEEE Symposium on Security and Privacy*, pages 397–411. IEEE Computer Society Press, May 2013.
- MKF21. Toghrul Maharrov, Dmitry Khovratovich, and Emanuele Francioni. The dusk network whitepaper, 2021. Whitepaper. [https://dusk.network/uploads/The\\_Dusk\\_Network\\_Whitepaper\\_v3\\_0\\_0.pdf](https://dusk.network/uploads/The_Dusk_Network_Whitepaper_v3_0_0.pdf).
- Nom22. Nomadic Labs' Cryptography Team. *aPlonK*, a library for zero-knowledge proofs and validity rollups, 2022. <https://gitlab.com/nomadic-labs/privacy-team/>.
- Ove22. Overeality Labs. Infrastructure for web3 interoperability, 2022. <https://overeality.io/home>.
- PFM<sup>+</sup>22. Luke Pearson, Joshua Fitzgerald, Héctor Masip, Marta Bellés-Muñoz, and Jose Luis Muñoz-Tapia. Plonkup: Reconciling plonk with plookup. Cryptology ePrint Archive, Report 2022/086, 2022. <https://ia.cr/2022/086>.
- PHGR13. Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013.

- PS00. David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, June 2000.
- Sch91. Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, January 1991.
- Sta21. StarkWare. ethstark documentation. Cryptology ePrint Archive, Paper 2021/582, 2021. <https://eprint.iacr.org/2021/582>.
- Sup21. Supranational. blst, a bls12-381 signature library focused on performance and security., 2021. <https://github.com/supranational/blst/tree/757aa0a90c03779f70d0ddab6bc84b40861bb4b>.
- Val08. Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In Ran Canetti, editor, *TCC 2008*, volume 4948 of *LNCS*, pages 1–18. Springer, Heidelberg, March 2008.
- Wil18. Zachary J. Williamson. Aztec network (white paper), 2018. Whitepaper. <https://github.com/AztecProtocol/AZTEC/blob/master/AZTEC.pdf>.
- ZXZS20. Jiaheng Zhang, Tiancheng Xie, Yupeng Zhang, and Dawn Song. Transparent polynomial delegation and its applications to zero knowledge proof. In *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18-21, 2020*, pages 859–876. IEEE, 2020.

## A Additional definitions

### A.1 $\mathcal{PlonK}$ permutation identities

Given a permutation  $\sigma : [3n] \rightarrow [3n]$ ,  $\eta_1 = 1 \in \mathbb{Z}_p$ ,  $\eta_2, \eta_3 \in \mathbb{Z}_p$  two quadratic non-residues (declared during the preprocessing phase of  $\mathcal{PlonK}$ ) such that  $\eta_3 \notin \eta_2 \mathcal{H}_n$ , the permutation polynomials  $S_{\sigma_1}, S_{\sigma_2}, S_{\sigma_3}$  are defined as

$$S_{\sigma_k} := \sum_{i=1}^n \eta_{q_{ki}} \omega^{r_{ki}} L_i(X), \quad \forall k \in [3] .$$

with  $q_{ki}n + r_{ki} = \sigma(kn + i)$  such that  $0 \leq r_{ki} < n$ .

The permutation identities are parametrized by two scalars  $\beta, \gamma \in \mathbb{Z}_p$  and is formed by the following polynomials:

$$\begin{aligned} \text{perm-ids}_{\sigma, \beta, \gamma}^\sigma(A(X), B(X), C(X), Z(X)) := \\ \{ P(X, \eta_2 X, \eta_3 X, X) - P(S_{\sigma_1}(X), S_{\sigma_2}(X), S_{\sigma_3}(X), \omega X), (Z(X)-1)L_1(X) \} , \end{aligned}$$

where  $P(Y_1, Y_2, Y_3, Y_4)$  is defined as  $(A(X) + \beta Y_1 + \gamma)(B(X) + \beta Y_2 + \gamma)(C(X) + \beta Y_3 + \gamma)Z(Y_4)$

An honest prover, who has built polynomials  $A, B$  and  $C$  from a trace witness  $\tilde{w}$  that respects the permutation constraints induced by  $\sigma$ , can construct a polynomial  $Z$  that satisfies the permutation identities as follows (define  $W = (A, B, C)$ ).

$$Z := L_1 + \sum_{i=2}^n L_i \prod_{j=1}^{i-1} \prod_{k=0}^2 \frac{(W_k(\omega^j) + \beta \eta_k \omega^j + \gamma)}{(W_k(\omega^j) + \beta S_{\sigma_k}(\omega^j) + \gamma)} .$$

### A.2 Assumptions

**Definition 5 ( $q$ -DLOG Assumption).** For  $q \in \mathbb{N}$ , the  $q$ -discrete logarithm assumption (relative to bilinear group generator  $\mathcal{G}$ ) states that for any PPT algorithm  $\mathcal{A}$ , the following probability is negligible in  $\lambda$ :

$$\Pr [(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_t, e) \leftarrow \mathcal{G}(1^\lambda); \mathbf{x} := (1, \dots, x^q); y \leftarrow \mathcal{A}([\mathbf{x}]_1, [\mathbf{x}]_2) : y = x] .$$

## B Building blocks for multi-polynomial commitments

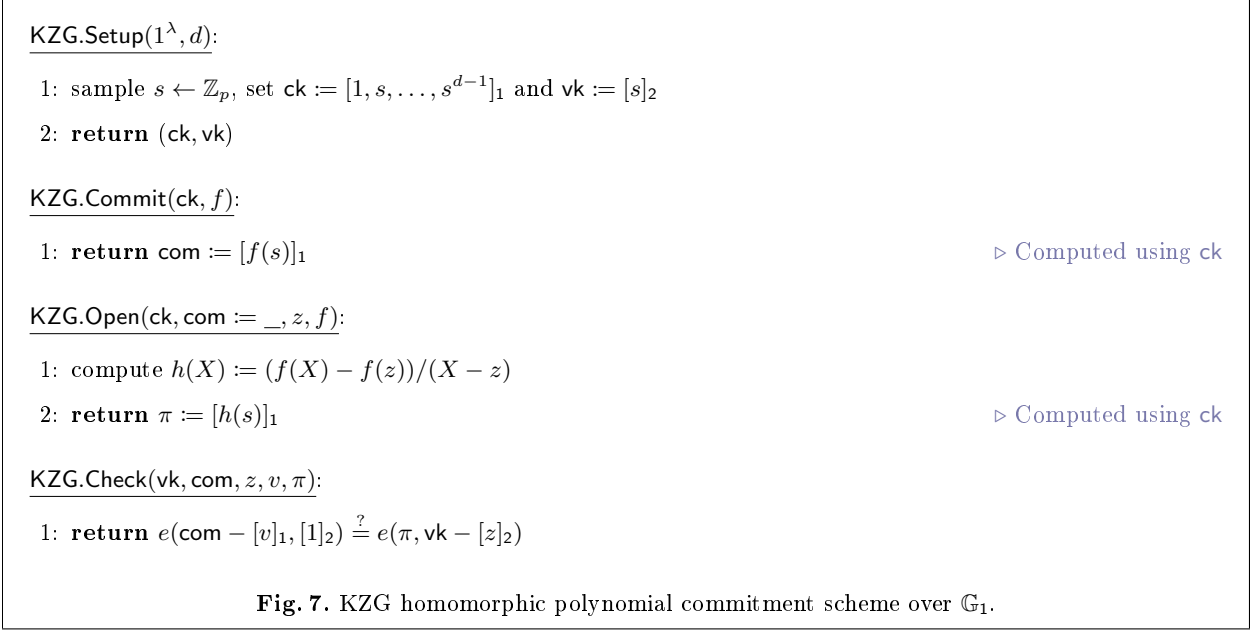
The multi-polynomial commitment scheme from Figure 2 relies on two main building blocks: (i) a homomorphic polynomial commitment scheme whose commitment space is  $\mathbb{G}_1$ , and (ii) an argument system for relation  $\text{PoK}\{\boldsymbol{\mu} : \langle \boldsymbol{\mu}, \mathbf{G} \rangle = C_G \wedge \langle \mathbf{r}, \boldsymbol{\mu} \rangle = P\}$ . In this section we provide a description of candidate instantiations for such building blocks. For the former, we choose the well-known KZG homomorphic polynomial commitment scheme [KZG10]. For the latter, we propose the protocol from Section B.2, a modified version of the inner-product argument [BBB+18]. Similar modifications of the inner-product argument have been proposed in the literature [BGH19b, BCL+21]. Here we present our own version, specialized for our use case, and prove its security for completeness.

### B.1 KZG polynomial commitment

We describe in Figure 7 the well-known KZG homomorphic polynomial commitment scheme [KZG10].

**Lemma 1.** *The polynomial commitment scheme from Figure 7 is complete, homomorphic, binding and knowledge sound in the algebraic group model under the  $q$ -DLOG assumption.*

*Proof.* *Completeness* and the *homomorphic property* can be checked by inspection. Breaking the *binding property* implies finding two different polynomials  $f, f'$  such that  $[f(s)]_1 = [f'(s)]_1$  which is equivalent to finding a non-trivial linear relation between the elements of  $\text{ck}$ , which is hard under the  $q$ -DLOG assumption. We refer to [GWC19, Section 3] for a proof of *knowledge soundness* in the algebraic group model.  $\square$



**Fig. 7.** KZG homomorphic polynomial commitment scheme over  $\mathbb{G}_1$ .

**Lemma 2.** *The polynomial commitment scheme from Figure 7 is inner-product binding (in the standard model) under the  $q$ -DLOG assumption.*

*Proof.* Let  $\mathcal{A}$  be an adversary against the inner-product binding property of the scheme from Figure 7, for some  $d, K \in \mathbb{N}$ . We build an adversary  $\mathcal{B}$  against the  $q$ -DLOG problem (for  $q = K - 1$ ).  $\mathcal{B}$  is given  $[\mathbf{x}]_1$  and  $[\mathbf{x}]_2$  for some  $x \leftarrow \mathbb{Z}_p$  and  $\mathbf{x} := (1, x, \dots, x^{K-1})$ .  $\mathcal{B}$  will now flip a fair coin:

- If the coin results in heads,  $\mathcal{B}$  will simulate  $\text{ck} := ([\mathbf{s}]_1, [\mathbf{s}]_2)$  by sampling  $s$  uniformly at random.  $\mathcal{B}$  will call  $\mathcal{A}$  on  $(\text{ck}, [\mathbf{x}]_1, [\mathbf{x}]_2)$ , receiving two distinct vectors of polynomials  $\mathbf{f}, \mathbf{f}' \in \mathbb{Z}_p[X]^k$  for some  $k \leq K$ , such that  $\sum_{i=1}^k x^{i-1} (f_i - f'_i)(s) = 0$ . If  $(f_i - f'_i)(s) \neq 0$  for some  $i$ , then  $\mathcal{B}$  can extract  $x$  by solving a non-trivial polynomial equation over  $\mathbb{Z}_p$ .  $\mathcal{B}$  will win the  $q$ -DLOG game in this branch with non-negligible probability unless  $(f_i - f'_i)(s) = 0$  with overwhelming probability for all  $i \in [k]$ .
- If the coin results in tails,  $\mathcal{B}$  will simulate  $\tau$  by sampling it uniformly at random. It will then instantiate ck with  $([\mathbf{x}]_1, [\mathbf{x}]_2)$  and call  $\mathcal{A}$  on  $(\text{ck}, [\boldsymbol{\tau}]_1, [\boldsymbol{\tau}]_2)$ , receiving two distinct vectors of polynomials  $\mathbf{f}, \mathbf{f}' \in \mathbb{Z}_p[X]^k$  for some  $k \leq K$ , such that  $\sum_{i=1}^k \tau^{i-1} (f_i - f'_i)(x) = 0$ . Assume  $\mathcal{B}$  is almost never successful in the previous branch, then  $(f_i - f'_i)(x) = 0$  with overwhelming probability for all  $i \in [k]$ . Since  $\mathbf{f} \neq \mathbf{f}'$  there must exist an index  $i$  such that  $(f_i - f'_i)(X)$  is not the zero polynomial. However,  $(f_i - f'_i)(x) = 0$ , thus  $\mathcal{B}$  can extract  $x$  by solving a non-trivial polynomial equation over  $\mathbb{Z}_p$ .

Algorithm  $\mathcal{B}$  must be successful in at least one of the branches with the same probability that  $\mathcal{A}$  is, so the total advantage of  $\mathcal{B}$  solving the  $q$ -DLOG problem is at least half the advantage of  $\mathcal{A}$  against the inner-product binding property.  $\square$

**Lemma 3.** *The polynomial commitment scheme from Figure 7 is inner-product extractable in the algebraic group model under the  $q$ -DLOG assumption.*

*Proof.* We proceed in the symbolic model, which immediately implies generic security. As explained above, symbolic security can be turned into a proof in the algebraic group model under the  $q$ -DLOG assumption.

Let  $\mathcal{A}$  be a symbolic adversary against the inner-product extractability of the scheme from Figure 7, for some  $d, K \in \mathbb{N}$ . On its first message,  $\mathcal{A}$  outputs polynomial  $G \in \mathbb{Z}_p[S, T]$  and scalars  $z \in \mathbb{Z}_p, \mathbf{v} \in \mathbb{Z}_p^k$ , for

some  $k \leq K$ .<sup>12</sup> After being given  $r \in \mathbb{Z}_p$ ,  $\mathcal{A}$  outputs polynomials  $\boldsymbol{\mu} \in \mathbb{Z}_p[S, T]^k$  and  $\pi \in \mathbb{Z}_p[S, T]$ .<sup>13</sup> We define an extractor  $\mathcal{E}$  that runs the first step of  $\mathcal{A}$ , parses the resulting polynomial  $G(S, T)$  as  $\sum_{j=1}^{2K-1} T^{j-1} f_j(S)$  and outputs  $\mathbf{f} := (f_1, \dots, f_k)$ . We need to show that if the following equations hold with non-negligible probability over the choice of  $r$ , and for  $\mathbf{T} := (1, T, \dots, T^{k-1})$ ,  $\mathbf{r} := (1, r, \dots, r^{k-1})$ ,

$$\langle \boldsymbol{\mu}(S, T), \mathbf{r} \rangle - \langle \mathbf{v}, \mathbf{r} \rangle = \pi(S, T)(S - z) \quad (5)$$

$$\langle \boldsymbol{\mu}(S, T), \mathbf{T} \rangle = G(S, T) \quad , \quad (6)$$

then the above extractor produces  $k$  polynomials  $\mathbf{f}$  satisfying  $\mathbf{f}(z) = \mathbf{v}$  and  $\langle \mathbf{f}, \mathbf{T} \rangle = G(S, T)$ .

First, given the structural conditions of footnote 13, polynomial  $\mu_i$ , for every  $i \in [k]$ , must be expressible as  $\mu_i := P_i(S) + Q_i(T)$  for some univariate polynomials  $P_i$  and  $Q_i$ . We can thus express equation (6) as  $\sum_{i=1}^k P_i(S)T^{i-1} + \sum_{i=1}^k Q_i(T)T^{i-1} = G(S, T)$ , which implies that  $f_j(S)$  is constant for every  $j > k$ . Now, by subtracting equations (6) and (5) and rearranging terms, we get:

$$G(S, T) = \langle \boldsymbol{\mu}(S, T), \mathbf{T} - \mathbf{r} \rangle + \pi(S, T)(S - z) + \langle \mathbf{v}, \mathbf{r} \rangle \quad . \quad (7)$$

Partially evaluating equation (7) on  $S = z$  leads to  $G(z, T) = \langle \boldsymbol{\mu}(z, T), \mathbf{T} - \mathbf{r} \rangle + \langle \mathbf{v}, \mathbf{r} \rangle$ , a polynomial equation on  $T$ , which evaluated on  $T = r$  gives  $G(z, r) = \langle \mathbf{v}, \mathbf{r} \rangle$ . Now, define  $h(X) := G(z, X) - \langle \mathbf{v}, \mathbf{X} \rangle$ , where  $\mathbf{X} := (1, X, \dots, X^{k-1})$  and observe that polynomial  $h$  is determined by the first message of  $\mathcal{A}$ , before  $r$  is chosen. Furthermore, if the last equality holds with non-negligible probability over the choice of  $r$ , then  $h(r) = 0$  for a non-negligible amount of values of  $r$  over  $\mathbb{Z}_p$ . Since  $\deg(h) < 2K$ , which is negligible,  $h$  must be the zero polynomial. Given that  $h$  can be expressed as  $h(X) = (\sum_{j=1}^{2K-1} X^{j-1} f_j(z)) - \sum_{j=1}^k X^{j-1} v_j$ , we can deduce that  $f_j(z) = v_j$  for every  $j \in [k]$  and  $f_j(z) = 0$  for all  $j > k$ . Furthermore,  $f_j$  is a constant polynomial for every  $j > k$ , so  $f_j(X)$  must be the zero polynomial for every  $j > k$ . Consequently,  $\langle \mathbf{f}, \mathbf{T} \rangle = G(S, T)$ , as desired.  $\square$

## B.2 Modified inner-product argument

**Lemma 4.** *The argument from Figure 3 has perfect completeness and computational witness-extended emulation for either extracting a non-trivial linear relation between group elements in  $\mathbf{G}, \mathbf{H}$  or extracting a valid witness for relation:*

$$\text{PoK}\{ \boldsymbol{\mu} \in \mathbb{G}_1^{2\kappa} : \langle \boldsymbol{\mu}, \mathbf{G} \rangle = C_G \wedge \langle \boldsymbol{\mu}, \mathbf{H} \rangle = C_H \wedge \langle \mathbf{r}, \boldsymbol{\mu} \rangle = P \} \quad .$$

*Proof.* We focus on the variant of Figure 3 which includes colored terms. A proof for the simpler version could be easily derived from this one. Correctness can be checked by inspection. For witness-extended emulation, we focus on the interactive version of the protocol from Figure 3 where every hash evaluation producing  $u_j$  is replaced by a round of interaction where the verifier samples  $u_j$  uniformly at random. We will see that there exists an efficient extractor  $\mathcal{E}$  that produces a witness from  $3^\kappa$  different valid transcripts. Since  $n = 2^\kappa$  is polynomial-size and  $3^\kappa$  equals  $n^{\log_2 3}$ , the result then follows from the General Forking Lemma from [BCC+16, BBB+18]. We proceed by induction on  $\kappa$ .

If  $\kappa = 0$  the extractor from a transcript simply returns  $\boldsymbol{\mu}^{(0)}$ , which is a valid witness if the transcript is valid (the transcript verification equations are those of the NP-relation when  $\kappa = 0$ ).

We now argue that if (for some  $\kappa > 0$ ) we have an extractor  $\mathcal{E}_{\kappa-1}$  for the protocol of size  $\kappa-1$ , we can build an extractor  $\mathcal{E}_\kappa$  for the protocol of size  $\kappa$ .  $\mathcal{E}_\kappa$  will run the prover until values  $(L_G^{(\kappa)}, L_H^{(\kappa)}, L_r^{(\kappa)}, R_G^{(\kappa)}, R_H^{(\kappa)}, R_r^{(\kappa)})$  have been fixed and sent to the verifier. Observe that the protocol after this first step is equivalent to a

<sup>12</sup>Since the bilinear pairing is available, polynomial  $G$  may contain crossed monomials  $S^i T^j$  for every  $0 \leq i < d$ , and  $0 \leq j < K$  or higher degree non-crossed monomials  $S^i$  for  $0 \leq i < 2d-1$ ,  $T^j$  for  $0 \leq j < 2K-1$ .

<sup>13</sup>In this case, since these polynomials correspond to elements of  $\mathbb{G}_1$ ,  $\pi$  and every  $\mu_i$  cannot contain crossed monomials that depend on both  $S$  and  $T$ .

protocol of size  $\kappa-1$  on a different public input:

$$\begin{aligned} \mathbf{G}' &:= u_\kappa^{-1} \mathbf{G}_L^{(\kappa)} + u_\kappa \mathbf{G}_R^{(\kappa)} & C'_G &:= C_G + u_\kappa^2 L_G^{(\kappa)} + u_\kappa^{-2} R_G^{(\kappa)} & r' &= r \\ \mathbf{H}' &:= u_\kappa^{-1} \mathbf{H}_L^{(\kappa)} + u_\kappa \mathbf{H}_R^{(\kappa)} & C'_H &:= C_H + u_\kappa^2 L_H^{(\kappa)} + u_\kappa^{-2} R_H^{(\kappa)} & P' &= (P + u_\kappa^2 L_r^{(\kappa)} + u_\kappa^{-2} R_r^{(\kappa)}) / (u_\kappa^{-1} + u_\kappa r^{2^{\kappa-1}}). \end{aligned} \quad (8)$$

This is possible because  $\mathbf{r}^{(\kappa-1)} := u_\kappa^{-1} \mathbf{r}_L^{(\kappa)} + u_\kappa \mathbf{r}_R^{(\kappa)}$  can be expressed as  $(u_\kappa^{-1} + u_\kappa r^{2^{\kappa-1}})(1, r, r^2, \dots, r^{2^{\kappa-1}-1})$ . Thus,  $\mathbf{r}^{(\kappa-1)}$  is again a vector containing the powers of a single element (actually the same  $r$ ), modulo multiplication by a constant. We got rid of the constant by dividing by it on the definition of  $P'$ .

At this point,  $\mathcal{E}_\kappa$  forks the execution of the prover three times, by providing three different challenges  $u_\kappa$ , say  $u_{\kappa,1}$ ,  $u_{\kappa,2}$  and  $u_{\kappa,3}$ . From this step, the protocol can be seen as an execution of the protocol of size  $\kappa-1$  on public inputs  $(2^{\kappa-1}, \mathbf{G}'_i, \mathbf{H}'_i, (C'_{G,i}, C'_{H,i}, r, P'_i))$ , for  $i = 1, 2, 3$ , defined as in equation (8) by using the corresponding challenge  $u_{\kappa,i}$ . We can leverage the extractor  $\mathcal{E}_{\kappa-1}$  to obtain, from  $3 \cdot 3^{\kappa-1} = 3^\kappa$  valid transcripts, witnesses  $\boldsymbol{\mu}_i \in \mathbb{G}_1^{2^{\kappa-1}}$  such that for all  $i \in \{1, 2, 3\}$ :

$$\langle \boldsymbol{\mu}_i, \mathbf{G}'_i \rangle = C'_{G,i} \quad \wedge \quad \langle \boldsymbol{\mu}_i, \mathbf{H}'_i \rangle = C'_{H,i} \quad \wedge \quad \sum_{j=0}^{2^{\kappa-1}} \mu_{i,j} r^j = P'_i .$$

Now, if  $u_{\kappa,i} \neq u_{\kappa,j}$  for  $i \neq j$ , which will occur with overwhelming probability, extractor  $\mathcal{E}_\kappa$  can solve a linear system of equations and find  $\nu_1, \nu_2, \nu_3 \in \mathbb{Z}_p$  such that

$$\sum_{i=1}^3 u_{\kappa,i}^{-2} \nu_i = 0 \quad \sum_{i=1}^3 \nu_i = 1 \quad \sum_{i=1}^3 u_{\kappa,i}^2 \nu_i = 0 .$$

Extractor  $\mathcal{E}_\kappa$  concludes by defining  $\boldsymbol{\mu} \in \mathbb{G}_1^{2^\kappa}$  as:

$$\boldsymbol{\mu} := \sum_{i=1}^3 (\nu_i u_{\kappa,i}^{-1} \boldsymbol{\mu}_i, \nu_i u_{\kappa,i} \boldsymbol{\mu}_i) .$$

Observe that this represents a valid witness for the relation of Lemma 4:

$$\begin{aligned} \langle \boldsymbol{\mu}, \mathbf{G} \rangle &= \sum_{i=1}^3 \nu_i u_{\kappa,i}^{-1} \langle \boldsymbol{\mu}_i, \mathbf{G}_L \rangle + \sum_{i=1}^3 \nu_i u_{\kappa,i} \langle \boldsymbol{\mu}_i, \mathbf{G}_R \rangle = \sum_{i=1}^3 \nu_i \langle \boldsymbol{\mu}_i, u_{\kappa,i}^{-1} \mathbf{G}_L + u_{\kappa,i} \mathbf{G}_R \rangle \\ &= \sum_{i=1}^3 \nu_i \langle \boldsymbol{\mu}_i, \mathbf{G}'_i \rangle = \sum_{i=1}^3 \nu_i C'_{G,i} = \sum_{i=1}^3 \nu_i (C_G + u_{\kappa,i}^2 L_G^{(\kappa)} + u_{\kappa,i}^{-2} R_G^{(\kappa)}) \\ &= (\sum_{i=1}^3 \nu_i) C_G + (\sum_{i=1}^3 \nu_i u_{\kappa,i}^2) L_G^{(\kappa)} + (\sum_{i=1}^3 \nu_i u_{\kappa,i}^{-2}) R_G^{(\kappa)} \\ &= C_G . \end{aligned}$$

Similarly,  $\langle \boldsymbol{\mu}, \mathbf{H} \rangle = C_H$ . Finally,

$$\begin{aligned} \langle \mathbf{r}, \boldsymbol{\mu} \rangle &= \sum_{i=1}^3 \nu_i u_{\kappa,i}^{-1} \langle \mathbf{r}_L, \boldsymbol{\mu}_i \rangle + \sum_{i=1}^3 \nu_i u_{\kappa,i} \langle \mathbf{r}_R, \boldsymbol{\mu}_i \rangle = \sum_{i=1}^3 \nu_i \langle u_{\kappa,i}^{-1} \mathbf{r}_L + u_{\kappa,i} \mathbf{r}_R, \boldsymbol{\mu}_i \rangle \\ &= \sum_{i=1}^3 \nu_i (u_{\kappa,i}^{-1} + r^{2^{\kappa-1}} u_{\kappa,i}) \langle \mathbf{r}_L, \boldsymbol{\mu}_i \rangle = \sum_{i=1}^3 \nu_i (u_{\kappa,i}^{-1} + r^{2^{\kappa-1}} u_{\kappa,i}) P'_i \\ &= \sum_{i=1}^3 \nu_i (P + u_{\kappa,i}^2 L_r^{(\kappa)} + u_{\kappa,i}^{-2} R_r^{(\kappa)}) = (\sum_{i=1}^3 \nu_i) P + (\sum_{i=1}^3 \nu_i u_{\kappa,i}^2) L_r^{(\kappa)} + (\sum_{i=1}^3 \nu_i u_{\kappa,i}^{-2}) R_r^{(\kappa)} \\ &= P . \end{aligned}$$

□

## C Proofs of the main body

### C.1 Proof of Theorem 1

Theorem 1 establishes that the scheme from Figure 2 is a complete, binding and knowledge sound multi-polynomial commitment scheme. We first establish the following helper lemma, which we prove after the proof of Theorem 1.

**Lemma 5.** *Let  $\mathcal{A}$  be an algebraic (stateful) algorithm. If the  $q$ -DLOG assumption holds, then the following probability is negligible in  $\lambda$ , for any  $K \in \mathbb{N}$  and any polynomial  $g \in \mathbb{Z}_p[X]$  with  $\deg(g) \leq K$ :*

$$\Pr \left[ \begin{array}{l} \tau, \rho \leftarrow \mathbb{Z}_p \\ G \leftarrow \mathcal{A}([\tau]_1, [\tau]_2) : G \neq [g(\tau)]_t \wedge e([\tau]_1 - [\rho]_1, \pi) = G - g(\rho) \\ \pi \leftarrow \mathcal{A}(\rho) \end{array} \right],$$

where  $\tau := (1, \tau, \dots, \tau^{K-1})$ .

*Proof (of Theorem 1).* *Completeness* can be checked by inspection.

For the *binding property*, note that **Commit-Evals** is binding by definition. We will prove that **Commit-Polys** is binding if  $\Psi$  is inner-product binding. For that, let  $\mathcal{A}$  be an adversary against the binding property of the scheme from Figure 2 (for  $d, K \in \mathbb{N}$ ). We will build an adversary  $\mathcal{B}$  against the inner-product binding property of  $\Psi$  (for the same  $d, K$ ), who succeeds with the same probability  $\mathcal{A}$  does. On input  $(\text{ck}_\Psi, [\tau]_1, [\tau]_2)$ ,  $\mathcal{B}$  simulates the commitment key of the multi-polynomial commitment as  $\text{ck} := (\text{ck}_\Psi, [\tau]_2)$  and the verification key as  $\text{vk} := (\text{vk}_\Psi, [\tau]_1)$ .<sup>14</sup> Adversary  $\mathcal{B}$  sends  $\text{ck}$  to  $\mathcal{A}$ , who will produce  $\mathbf{f} \in \mathbb{Z}_p^{<d}[X]^k$  and  $\mathbf{f}' \in \mathbb{Z}_p^{<d}[X]^{k'}$ , with  $k, k' < K$ , such that  $\mathbf{f} \neq \mathbf{f}'$  and  $\text{Commit-Polys}(\text{ck}, \mathbf{f}) = \text{Commit-Polys}(\text{ck}, \mathbf{f}')$ . Adversary  $\mathcal{B}$  will simply output  $(\mathbf{f}, \mathbf{f}')$ , which must be a valid forgery to its own inner-product binding game, since:

$$(k, \langle \Psi.\text{Commit}(\text{ck}_\Psi, \mathbf{f}), \tau \rangle) = \text{Commit-Polys}(\text{ck}, \mathbf{f}) = \text{Commit-Polys}(\text{ck}, \mathbf{f}') = (k', \langle \Psi.\text{Commit}(\text{ck}_\Psi, \mathbf{f}'), \tau \rangle),$$

which implies that  $\langle \Psi.\text{Commit}(\text{ck}_\Psi, \mathbf{f}), \tau \rangle = \langle \Psi.\text{Commit}(\text{ck}_\Psi, \mathbf{f}'), \tau \rangle$  and that  $|\mathbf{f}| = |\mathbf{f}'|$ , as desired.

Finally, we show that the scheme is *knowledge sound*. Let  $\mathcal{A}$  be an algorithm that on input  $\text{ck}$  produces  $(\text{com}_\mathbf{f}, z, \text{com}_\mathbf{v}, \pi)$  s.t.  $\text{Check}(\text{vk}, \text{com}_\mathbf{f}, z, \text{com}_\mathbf{v}, \pi) = 1$  with non-negligible probability. We will define an extractor  $\mathcal{E}$  that runs in expected polynomial time and produces  $\mathbf{f}$  such that  $\text{com}_\mathbf{f} = \text{Commit-Polys}(\text{ck}, \mathbf{f})$  and  $\text{com}_\mathbf{v} = \text{Commit-Evals}(\mathbf{f}(z))$  with overwhelming probability conditioned on  $\text{Check}(\text{vk}, \text{com}_\mathbf{f}, z, \text{com}_\mathbf{v}, \pi) = 1$ .

For that, we will extend algorithm  $\mathcal{A}$  (against knowledge soundness for some  $d, K \in \mathbb{N}$ ) into an algorithm  $\mathcal{B}$  against the inner-product extractability of  $\Psi$  (for the same  $d, K \in \mathbb{N}$ ). Given  $(\text{ck}_\Psi, [\tau]_1, [\tau]_2)$ ,  $\mathcal{B}$  will prepare the multi-polynomial commitment key  $\text{ck}$  as described above and run  $\mathcal{A}$  on it, producing  $(\text{com}_\mathbf{f}, z, \text{com}_\mathbf{v}, \pi)$ . Parse  $\text{com}_\mathbf{f}$  as  $(k, G)$ .  $\mathcal{B}$  will then run the knowledge extractor of the proof of relation (3) contained in  $\pi$  to obtain a vector  $\mathbf{v}$  s.t.  $\text{Commit-Evals}(\mathbf{v}) = \text{com}_\mathbf{v}$ . After that,  $\mathcal{B}$  will output  $(G, z, \mathbf{v})$  as its first message of the inner-product extractability game. On receiving  $r$ ,  $\mathcal{B}$  will rewind  $\mathcal{A}$  and provide  $r$  as the output of **Hash** on  $(\text{com}_\mathbf{f}, z, \text{com}_\mathbf{v})$ .<sup>15</sup> By the Forking Lemma [PS00, BN06], this second execution of  $\mathcal{A}$  results in a tuple  $(\text{com}_\mathbf{f}, z, \text{com}_\mathbf{v}, \tilde{\pi})$  which also satisfies  $\text{Check}(\text{vk}, \text{com}_\mathbf{f}, z, \text{com}_\mathbf{v}, \tilde{\pi}) = 1$  with non-negligible probability. Parse  $\tilde{\pi}$  as  $(\hat{\mu}, \hat{v}, \pi_\Psi, \pi_v, \pi_{\text{PA}}, \pi_\tau)$ . Now observe that by virtue of Lemma 5, steps 5-6 from the **Check** algorithm from Figure 2 serve as a replacement for the skipped steps 6-7 from the inner-product argument from Figure 3. Therefore,  $\mathcal{B}$  can leverage the extractor of the inner-product argument to obtain  $\boldsymbol{\mu} \in \mathbb{G}_1^k$  such that  $\langle \boldsymbol{\mu}, \tau[:k] \rangle = G$  and  $\langle \mathbf{r}, \boldsymbol{\mu} \rangle = \hat{\mu}$ , where  $\mathbf{r} = (1, r, \dots, r^{k-1})$ . Algorithm  $\mathcal{B}$  will output  $(\boldsymbol{\mu}, \pi_\Psi)$ .

Observe that  $\mathcal{B}$  is a successful algorithm against the inner-product extractability of  $\Psi$  in the sense that  $G \in \mathbb{G}_t$ ,  $z \in \mathbb{Z}_p$ ,  $\mathbf{v} \in \mathbb{Z}_p^k$ ,  $\boldsymbol{\mu} \in \mathbb{G}_1^k$  for some  $k \leq K$  and with non-negligible probability  $\langle \boldsymbol{\mu}, \tau[:k] \rangle = G$  and  $\Psi.\text{Check}(\text{ck}_\Psi, \hat{\mu}, z, \hat{v}, \pi_\Psi) = 1$ , where  $\hat{\mu} = \langle \boldsymbol{\mu}, \mathbf{r} \rangle$  and  $\hat{v} = \langle \mathbf{v}, \mathbf{r} \rangle$ , as ensured by the proof of relation (3) contained in  $\tilde{\pi}$  (and given the binding property of **Commit-Evals**). Therefore, given the inner-product extractability of  $\Psi$ , there exists an extractor  $\mathcal{E}_\mathcal{B}$  that produces  $k$  polynomials  $\mathbf{f}$  s.t.  $\mathbf{f}(z) = \mathbf{v}$  and  $\text{com}_\mathbf{f} = (k, G) = (k, \langle \Psi.\text{Commit}(\text{ck}_\Psi, \mathbf{f}), \tau[:k] \rangle)$  which equals  $\text{Commit-Polys}(\text{ck}, \mathbf{f})$ . Consequently, given  $\mathcal{A}$ , we can define an extractor  $\mathcal{E}$  for the knowledge soundness game that builds  $\mathcal{B}$  from  $\mathcal{A}$  and replays  $\mathcal{E}_\mathcal{B}$ , as desired.  $\square$

## Proof of Lemma 5

For simplicity, we describe a proof in the generic group model. As observed before, it can be lifted to a proof in the algebraic group model under the  $q$ -DLOG assumption.

<sup>14</sup>Verification keys are implicitly contained in commitment keys, thus we need to show how to simulate them.

<sup>15</sup>We can assume w.l.o.g. that  $\mathcal{A}$  made such query since its winning probability would be negligible otherwise.

*Proof.* Algorithm  $\mathcal{A}$ , will output a polynomial  $G \in \mathbb{Z}_p[T]$  of degree bounded by  $2K$ . (Note that  $\mathcal{A}$  can obtain  $[\tau^i]_t$  for every  $i \in [0, 2K - 2]$  by using the pairing, and no more relevant elements.) Then, after receiving  $\rho \in \mathbb{Z}_p$ ,  $\mathcal{A}$  will output a second polynomial  $\pi \in \mathbb{Z}_p[T]$  of degree bounded by  $K$ .  $\mathcal{A}$  is successful iff  $G(T) \neq g(T)$  and  $(T - \rho)\pi(T) = G(T) - g(\rho)$ . Once  $G$  is fixed, the probability that it evaluates to the same value as  $g$  on a uniformly chosen  $\rho$  is negligible. (It can be upper-bounded by  $2K/p$ , given that  $G$  and  $g$  are distinct polynomials with degree bounded by  $2K$ .) This means that, with overwhelming probability over the choice of  $\rho$ , the right-hand side of the previous equality is a polynomial that does not evaluate to 0 on  $\rho$ . On the other hand, the left-hand side is a polynomial that evaluates to 0 on  $\rho$ , independently of  $\pi(T)$ . We conclude that both polynomials cannot be equal and thus, the adversary cannot succeed in the symbolic model, what implies that the adversary's success probability in the generic group model is negligible.  $\square$

### Theorem 1 without the inner-product binding property or inner-product extractability

The scheme from Figure 2 could be instantiated with a polynomial commitment scheme  $\Psi$  that is not inner-product binding nor inner-product extractable if it is modified as described in Remark 2. This is because function  $\text{Com} : \mathbb{G}_1^k \rightarrow \mathbb{G}_2^2$  defined as  $\text{Com}(\mu) := (\langle \mu, \tau[:k] \rangle, \langle \mu, \tilde{\tau}[:k] \rangle)$  is binding, as long as  $\tau$  and  $\tilde{\tau}$  are sampled independently and only given as powers in  $\mathbb{G}_1$  and  $\mathbb{G}_2$ .

The *binding* of the modified multi-polynomial commitment scheme follows directly from the previous fact and the (standard) binding property of  $\Psi$ , which together imply that **Commit-Polys** is binding. Also, note that **Commit-Evals** is binding by definition.

The *knowledge soundness* of the multi-polynomial commitment scheme could be proven as follows. Let  $\mathcal{A}$  be an algorithm that on input  $\text{ck}$  produces  $(\text{com}_f, z, \text{com}_v, \pi)$  satisfying  $\text{Check}(\text{vk}, \text{com}_f, z, \text{com}_v, \pi) = 1$  with non-negligible probability  $\delta$ . We define an extractor  $\mathcal{E}$  that runs in expected polynomial time and produces  $\mathbf{f}$  such that  $\text{com}_f = \text{Commit-Polys}(\text{ck}, \mathbf{f})$  and  $\text{com}_v = \text{Commit-Evals}(\mathbf{f}(z))$  with overwhelming probability. From an analog version of Lemma 5 we could show that the steps 5-6 from the **Check** algorithm from Figure 2 (after the proper modifications) serve as a replacement for the skipped steps 6-7 from the inner-product argument from Figure 3. Say  $\mathcal{A}$  has performed  $q$  queries to the random oracle **Hash** and assume without loss of generality that  $\mathcal{A}$  has queried the random oracle on  $(\text{com}_f, z, \text{com}_v)$ , obtaining  $r$ , the scalar involved in the inner-product argument. Parse  $\text{com}_f$  as  $(k, G, H)$ .  $\mathcal{E}$  can leverage the extractor of the inner-product argument to obtain  $\mu \in \mathbb{G}_1^k$  such that  $\langle \mu, \tau[:k] \rangle = G$  and  $\langle \mu, \tilde{\tau}[:k] \rangle = H$ , and  $\langle r, \mu \rangle = \hat{\mu}$ . Furthermore,  $\mathcal{E}$  can run the knowledge soundness extractor of  $\Psi$  to obtain a polynomial  $\hat{f}$  such that  $\hat{f}(z) = \hat{v}$  and  $\hat{\mu} = \Psi.\text{Commit}(\text{ck}_\Psi, \hat{f})$ . Additionally,  $\mathcal{E}$  can run the knowledge extractor of the given proof of relation (3) to obtain a vector  $\mathbf{v}$  such that  $\text{Commit-Evals}(\mathbf{v}) = \text{com}_v$  and  $\langle \mathbf{r}_i, \mathbf{v} \rangle = \hat{v}$ , for every  $i \in [k]$ , where  $\mathbf{r} = (1, r, \dots, r^{k-1})$ . If the total running time of the above extractors is  $t$ , by the generalized Forking Lemma [PS00, BN06],  $\mathcal{E}$  can run the above extraction  $k$  times, on  $k$  different  $r_i$ , obtaining  $k$  different tuples  $(\mu_i \in \mathbb{G}_1^k, \hat{f}_i \in \mathbb{Z}_p[X], \mathbf{v}_i \in \mathbb{Z}_p^k)$ , such that for every  $i \in [k]$ :

$$\langle \mu_i, \tau \rangle = \text{com}_f \quad \wedge \quad \langle \mathbf{r}_i, \mathbf{v}_i \rangle = \hat{f}_i(z) \quad \wedge \quad \text{Commit-Evals}(\mathbf{v}_i) = \text{com}_v \quad \wedge \quad \langle \mathbf{r}_i, \mu_i \rangle = \Psi.\text{Commit}(\text{ck}_\Psi, \hat{f}_i) .$$

The expected running time of such extraction is  $\mathcal{O}(kqt/\delta)$ , which is polynomial since  $\delta$  is non-negligible and  $\mathcal{A}$  runs in polynomial-time. Now, given that function **Com** is binding, it must be  $\mu_i = \mu_j$  for all  $i, j \in [k]$ . Similarly, since **Commit-Evals** is binding, we must have  $\mathbf{v}_i = \mathbf{v}_j$  for every  $i, j \in [k]$ . So  $\mathcal{E}$  ended up with vectors  $\mu \in \mathbb{G}_1^k$  and  $\mathbf{v} \in \mathbb{Z}_p^k$  such that for every  $i \in [k]$ :

$$\langle \mathbf{r}_i, \mathbf{v} \rangle = \hat{f}_i(z) \quad \wedge \quad \langle \mathbf{r}_i, \mu \rangle = \Psi.\text{Commit}(\text{ck}_\Psi, \hat{f}_i) ,$$

for  $k$  different  $\mathbf{r}_i$ . Let  $R$  be the Vandermonde matrix formed by vectors  $\mathbf{r}_i$  for every  $i \in [k]$ , which is invertible given that  $r_i \neq r_j$  for different  $i, j$ . And let  $\hat{\mathbf{f}} := (\hat{f}_1, \dots, \hat{f}_k)$ . We have:

$$R\mathbf{v} = \hat{\mathbf{f}}(z) \wedge R\mu = \Psi.\text{Commit}(\text{ck}_\Psi, \hat{\mathbf{f}}) \quad \text{or equivalently} \quad \mathbf{v} = R^{-1}\hat{\mathbf{f}}(z) \wedge \mu = R^{-1}\Psi.\text{Commit}(\text{ck}_\Psi, \hat{\mathbf{f}}) ,$$

where  $\Psi.\text{Commit}(\text{ck}_\Psi, \hat{\mathbf{f}})$  is a shorthand for the column vector  $(\Psi.\text{Commit}(\text{ck}_\Psi, \hat{f}_1), \dots, \Psi.\text{Commit}(\text{ck}_\Psi, \hat{f}_k))$ . Extractor  $\mathcal{E}$  will output  $\mathbf{f} = R^{-1}\hat{\mathbf{f}}$ , which satisfies  $\mathbf{f}(z) = \mathbf{v}$  and, by the homomorphic property of  $\Psi$ ,  $\mu = \Psi.\text{Commit}(\text{ck}_\Psi, \mathbf{f})$ , where  $\langle \mu, \tau[:k] \rangle = G$  and  $\langle \mathbf{r}, \mu \rangle = \hat{\mu}$ , so **Commit-Polys**( $\text{ck}, \mathbf{f}$ ) =  $\text{com}_f$  as desired.



## C.2 Proof of Theorem 2

Theorem 2 establishes that the scheme described in Figure 4 constitutes a *SNARK* for relation:

$$\text{PoK} \left\{ \{ \mathbf{w}_i \in \mathbb{Z}_p^{m-\ell} \}_{i \in [k]} : (\mathbf{x}_i, \mathbf{w}_i) \in \mathbb{Z}_p^m \text{ satisfies } \mathcal{C} \ \forall i \in [k] \right\},$$

if  $\Psi$  is a complete, binding and knowledge sound multi-polynomial commitment scheme and  $\text{Hash} : \{0, 1\}^* \rightarrow \mathbb{Z}_p$  is a random oracle.

We first establish the following helper lemma.

**Lemma 6.** *Let  $\mathbf{E}$  be a vector of  $m$  multi-variate linear polynomials on  $n$  variables over  $\mathbb{Z}_p$ . Let  $k \in \mathbb{N}$  and  $\mathbf{x}_i \in \mathbb{Z}_p^n$  for  $i \in [k]$ . If  $\Pr[\delta \leftarrow \mathbb{Z}_p : \mathbf{E}(\sum_{i \in [k]} \delta^{i-1} \mathbf{x}_i) = \mathbf{0}]$  is non-negligible, then  $\mathbf{E}(\mathbf{x}_i) = \mathbf{0}$  for all  $i \in [k]$ .*

*Proof (by contrapositive).* Assume there exist  $i^* \in [k]$  and  $j^* \in [m]$  such that  $E_{j^*}(\mathbf{x}_{i^*}) \neq 0$ . Define  $P(\Delta) := E_{j^*}(\sum_{i \in [k]} \Delta^{i-1} \mathbf{x}_i)$ . By linearity, it holds that  $P(\Delta) = \sum_{i \in [k]} \Delta^{i-1} E_{j^*}(\mathbf{x}_{i^*})$ , so  $P$  is a non-zero polynomial given that  $E_{j^*}(\mathbf{x}_{i^*})$  is different from 0. Thus,  $P$  will not vanish on a uniformly sampled  $\delta$  except with negligible probability.  $\square$

*Proof (of Theorem 2).* Correctness can be checked by inspection. To show knowledge soundness we need to build an extractor which can produce valid witnesses for all the atomic statements from the interaction with a successful *aPlonK* prover. We rely on the knowledge soundness extractor of *PlonK* [GWC19]. Note that the scheme from Figure 4 is essentially the original *PlonK* construction with the following modifications:

- (1) The Fiat-Shamir random challenges  $\alpha, \beta, \gamma, \xi$  are shared across all proofs.
- (2) The  $\mathbb{T}$  polynomial is common for all proofs.
- (3) We use a shared permutation argument, thus there is a common polynomial  $\mathbb{Z}$  for all proofs with respect to the same *PlonK* constraint system.
- (4) We use a multi-polynomial commitment scheme instead of a standard polynomial commitment scheme.
- (5) The polynomial evaluations on  $\xi$  (and  $\omega\xi$ ) are not included in the proof. Instead, we include a commitment to them, together with a meta-verification proof of the relation from Figure 5.

As explained in Section 4.1, the *PlonK* extractor can be easily modified to support modifications (1)-(2), by extracting a witness for each of the statements when fixing all others.

The shared permutation argument in modification (3) guarantees that a random linear combination of the wires of each proof (through powers of a uniformly sampled random value  $\delta$ ) meets the copy-satisfiability constraints of permutation  $\sigma$ . If that is the case with overwhelming probability over the choice of  $\delta$ , then by virtue of Lemma 6, each of the wire evaluations of every individual statement being proved also meets the copy-satisfiability constraints. Thus applying the standard *PlonK* on a proof that satisfies the shared permutation argument will result in an extracted set of witnesses that meets the copy-satisfiability constraints as desired.

Modification (4) is minimal from the extractor point of view. Note that *PlonK*'s extractor relies on the definition of "knowledge soundness in the algebraic group model" from [GWC19, Section 3], a unified security notion that captures extractability and the binding property at the same time. Such definition is the basis of *PlonK*'s extractor (then formalized through so-called polynomial protocols). Note that our *binding property* (on *Commit-Polys*) and *knowledge soundness* for multi-polynomial commitments (Section 3) together, imply such definition, which guarantees that the same extractor can be used after modification (4).

Finally, note that modification (5) can be addressed as follows. We can use the extractor of the meta-verification proof for relation  $\mathcal{R}_{n,k}((\alpha, \beta, \gamma, \delta, \xi, \nu_w, \nu_z, \nu_{\bar{z}}, \nu_t, \nu_{pp}, \{\mathbf{x}_j\}_{j \in [k]}), \cdot)$  (see Figure 5) to obtain evaluations  $(\{a_j, b_j, c_j\}_{j \in [k]}, z, \bar{z}, t, e_{qL}, e_{qR}, e_{qO}, e_{qM}, e_{qC}, e_{s1}, e_{s2}, e_{s3})$  that satisfy all identities and are the actual evaluation of the corresponding committed polynomials at  $\xi$  (and  $\omega\xi$  in the case of  $\bar{z}$ ), as ensured by the multi-polynomial commitment *binding property* (on *Commit-Evals*) and the *knowledge soundness*. Such extracted evaluations thus satisfy all the properties that are verified in a standard *PlonK* proof and can consequently be used by the standard *PlonK* extractor.  $\square$