

# Polynomial IOPs for Linear Algebra Relations

Alan Szepieniec  
alan@nervos.org  
Nervos Foundation

Yuncong Zhang  
shjdzhangyuncong@sjtu.edu.cn  
Shanghai Jiao Tong University

**Abstract.** This paper proposes new Polynomial IOPs for arithmetic circuits. They rely on the monomial coefficient basis to represent the matrices and vectors arising from the arithmetic constraint satisfaction system, and build on new protocols for establishing the correct computation of linear algebra relations such as matrix-vector products and Hadamard products. Our protocols give rise to concrete proof systems with succinct verification when compiled down with a cryptographic compiler whose role is abstracted away in this paper. Depending only on the compiler, the resulting SNARKs are either transparent or rely on a trusted setup.

**Keywords:** SNARK · Polynomial IOP · Zero-Knowledge · Succinct Verification

## 1 Introduction

Succinct Non-Interactive Arguments of Knowledge (SNARKs) enable a resource-constrained verifier to cryptographically verify the authentic computations of an untrusted prover. The technology is particularly well-suited to the cryptocurrency setting, where participants are typically anonymous, untrusted, and where the success of the network depends on the capability of lightweight nodes to verify the network’s consensus (however that is defined). In this setting, there is a large monetary incentive for malicious behavior.

Despite the flurry of rapid related and unrelated developments by diverse parties, two trends are emerging as good practices in this domain.

1. *Functional separation in the compilation pipeline.* The compilation process for general purpose zero-knowledge proofs is separated into multiple steps with clear boundaries. At the input of this pipeline is a computation, represented either as program source code or as a circuit. A technique called arithmetization turns this computation into a constraint system involving native operations over a finite field. The next step transforms this constraint system into an abstract proof system between two parties, prover  $P$  and verifier  $V$ , that are interactive Turing machines with access to unrealistic or unrealizable resources such as PCP oracles. These abstract proof systems in this step are called *interactive oracle proofs (IOPs)* and typically achieve statistical or even perfect security. In the last step, the *cryptographic compilation*, the unrealistic resources are replaced by cryptographic approximations that achieve the same functionality at the expense of introducing computational hardness assumptions for security.

2. *Polynomial IOP formalism.* The abstract information-theoretical proof system in the step before cryptographic compilation could in principle rely on a variety of unrealistic resources, and build a sound proof system from their mathematical properties. However, for the purpose of establishing soundness, the Schwartz-Zippel lemma is an indispensable tool. The strategy is to reduce the satisfaction of arithmetic constraints arising from the constraint system to series of identities of *low-degree polynomials*. By evaluating these polynomials in random points, their equality is tested probabilistically. If the left and right hand sides of an equation represent identical polynomials, they are identical everywhere, but if they are unequal they are different *almost* everywhere. The Schwartz-Zippel lemma provides an exact concrete quantification of the security lost due to this probabilistic approximation. A *Polynomial IOP* is the abstract proof system tailored to this strategy. In this formalism, the prover sends low degree polynomials to the verifier, and rather than reading the entire list of coefficients, the verifier queries these polynomials in a given point through an oracle interface. The cryptographic compiler uses a *polynomial commitment scheme* to simulate this unrealistic resource.

These trends are visible in the rise of universal SNARKs with universal and updatable structured reference strings (SRS's) such as Sonic [11], PLONK [8], and Marlin [7]. The common idea here is to use the cryptographic pairing-based mathematics only to realize *polynomial commitment scheme*, typically the KZG scheme [10]. Since the SRS is used only for the KZG scheme, it is independent of the preceding abstract proof system and the circuit it encodes; this independence is precisely what enables updates to the SRS and its adaptation to any circuit. PLONK and Marlin independently formalize this abstraction and introduce the terms *Polynomial Protocol* and *Algebraic Holographic Proof (AHP)*, respectively. This paper adopts the terminology of Bünz *et al.* [6], who introduce a new polynomial commitment scheme (and hence a cryptographic compiler) based on groups of unknown order and in the process explore the landscape of protocols it can apply to.

These trends are *also* visible in the rise of IOPs based on Reed-Solomon codes [1,4,3]. The underlying abstract protocols here are not explicitly Polynomial IOPs. However, their common feature is the reliance on Reed-Solomon codewords as the proof oracles. Since Reed-Solomon codewords are obtained by evaluating polynomials in a domain of points whose cardinality is larger than the polynomials' degree, these proof oracles uniquely identify the originating low-degree polynomials. As a result, a Reed-Solomon IOP is a Polynomial IOP in disguise.

Despite the spontaneous convergence onto Polynomial IOPs as a useful formalism, there seems to be little agreement about the optimal interface between Polynomial IOPs and the arithmetic constraint systems that they realize. Arithmetic constraint systems typically express constraints using matrix algebra: in terms of vectors, and matrix multiplication, but also *Hadamard products*, which is a fancy word for the element-wise products of pairs of equal-length vectors. The

set of operations that Polynomial IOPs natively offer are somewhat different. As a result, how the Polynomial IOP represents the objects in the arithmetic constraint system and how it simulates the equations that constrain them, are the key questions in the design process of Polynomial IOPs. The various answers to these questions are what set the various Polynomial IOPs for arithmetic circuits apart.

- Marlin and Aurora represent the objects of the arithmetic constraint system as the Reed-Solomon codewords of polynomials. Standard techniques establish the correct computation of a Hadamard product of such codewords. The computation of a linear transform applied to such a codeword is reduced to checking the sum of a related codeword.
- PLONK represents the vector of wire values as the values of a polynomial in a domain of points. A permutation argument establishes the assignment of wires to gates and the standard techniques for Reed-Solomon codewords establish the consistency of inputs and outputs to addition and multiplication gates.
- Sonic represents the vectors of left, right, and output wires of a series of multiplication gates as the coefficient vectors of three polynomials. The consistency of these multiplication gates, and of a linear transform, is established by checking several properties of bivariate polynomials. The paper furthermore explains under which conditions these bivariate polynomials can be simulated with univariate ones.

*Contributions.* In this paper we propose a new Polynomial IOP for arithmetic circuits, called Claymore<sup>1</sup>. Succinct verification is achieved with a trusted pre-processing phase. When compiled with any polynomial commitment scheme, the result is a concrete zk-SNARK with universal updatable structured reference string, or transparent setup, depending only on the nature of the polynomial commitment scheme.

The arithmetic constraint system chosen to represent the arithmetic circuit is the Hadamard Product Relation (HPR), in which the witness consists of three vectors representing the left, right, and output wires of a list of multiplication gates. We note that Sonic realizes a similar constraint satisfaction relation by reducing both the multiplication and linear constraints into one large equation. In Claymore, the multiplication gate consistency and linear consistency are achieved in two separate steps, both of which rely on a collection of subprotocols for linear algebra relations that we develop along the way. The separate steps are later merged as an explicit optimization.

Like Sonic but unlike Marlin and PLONK, Claymore opts for the *monomial coefficient basis* to represent the vectors of the arithmetic constraint system. DenseClaymore represents the linear transform as a dense matrix in the monomial coefficient basis and this choice results in the smallest number of polynomials in the transcript across all Polynomial IOPs. The price to pay for this brevity is the  $O(n^2)$  scaling of the polynomials’ degree, where  $n$  is the size of the circuit.

---

<sup>1</sup> A type of Scottish sword.

A question that naturally arises when using this basis, is whether it is also equipped to deal with the sparse linear transformations that typically come from long-winded computations. We answer this question positively by providing methods for dealing with sparse linear algebra relations, culminating in a sparse variant of Claymore. This variant concretely outperforms Marlin and Sonic in terms of the number of polynomials in the transcript. While this number is smaller still for PLONK, one notes that PLONK does not support arbitrary fan-in for linear constraints, whereas Marlin and Claymore (both variants) do.

Additionally, we compare the new and existing Polynomial IOPs both abstractly and concretely. In the abstract comparison we determine how the key performance-driving parameters of the Polynomial IOP evolve as a function of the circuit size. In the concrete comparison, we transform the various Polynomial IOPs with three different cryptographic compilers into concrete zk-SNARKs in order to compare the size of the resulting proofs. In this comparison, all proofs establish the integrity of the same benchmark computation.

*Motivation and applications.* The motivation for this work is chiefly theoretical. We study the interface between arithmetic circuits and Polynomial IOPs in isolation of other constraints and demands. As a result of this focus, our protocol is arguably simpler than other protocols that achieve nominally the same thing. Complexity is the friend of mistakes, and our protocol may therefore be the preferred option for this reason even in circumstances where it is inferior in terms of performance.

The dense variant of Claymore performs extremely well for shallow arithmetic circuits, such as the verification circuits of lattice-based and MQ-based signature schemes, which typically involve operations on large matrices and vectors over a small finite field. As a result, a DenseClaymore-SNARK is an outstanding candidate for achieving post-quantum signature aggregation, or signature schemes with various fancy properties that zero-knowledge proofs enable.

Using SNARKs in combination with other cryptographic tools points to a useful property that SNARKs frequently lack — they typically require a finite field with a particular structure, such as a large multiplicative subgroup of smooth order. Marlin and PLONK have this property, while Sonic is defined for arbitrary fields. Using the SNARK in combination with a different cryptosystem that requires an incompatible field, requires the SNARK to simulate the cryptosystem’s field operations using the arithmetic constraint system of the SNARK. In contrast, Claymore (like Sonic) induces no such costly simulation overhead as it works for any large enough finite field.

The protocols proposed here promote *simplicity through modularity*. However, we observe that once the basic protocol has been composed, there are available optimizations that improve its characteristics at the cost of violating the boundaries between modules. This observation highlights the utility of separating *design* from *optimization* considerations. Note that it is only the optimized SparseClaymore protocol that outperforms Marlin in the target metric, number of polynomials. The unoptimized version is inferior in all respects. Furthermore, the proof of zero-knowledge relies on a batching-related optimization that ap-

plies to both variants of Claymore; without this optimization the proof is tricky and complex. Lastly, the optimizations stand on their own, and can possibly improve other Polynomial IOPs beyond Claymore.

## 2 Preliminaries

### 2.1 Indexed Relations

Owing to their convenience, we use *indexed relations* [7]. An indexed relation is a set  $\mathcal{R}$  of tuples  $(\mathbf{i}, \mathbf{x}, \mathbf{w})$ , whose three components are called the *index*, *instance*, and *witness*, respectively. The separation between index and instance captures the intuition that some properties of concrete proofs for  $\mathcal{R}$  should be computable from  $\mathbf{i}$  even before  $\mathbf{x}$  is known. For instance,  $\mathbf{i}$  can be the description of an arithmetic circuit,  $\mathbf{x}$  the values of the output wires, and  $\mathbf{w}$  an assignment of values to all wires that makes the all gates consistent. The projection  $\{(\mathbf{i}, \mathbf{x}) \mid (\mathbf{i}, \mathbf{x}, \mathbf{w}) \in \mathcal{R}\}$  of triples in  $\mathcal{R}$  onto the first two components is the *indexed language corresponding to  $\mathcal{R}$*  and is denoted by  $\mathcal{L}(\mathcal{R})$ .

### 2.2 Constraint Systems

A constraint system is a representation of a computation in terms of equations with unknown variables. When there is an assignment to the unknown variables that satisfies all equations, we say the constraint system is *satisfiable*, and this assignment is the *witness*. The *index* determines all fixed constants in the equations, and the *instance* determines known variables that can vary independently of the index but are ultimately known by all parties involved.

The following constraint system is adapted from Bootle *et al.* [5].

**Definition 1 (Hadamard Product Relation (HPR)).** *Let  $\mathbb{F}$  be a finite field. A triple  $(\mathbf{i}, \mathbf{x}, \mathbf{w})$  where  $\mathbf{i} = (m, n, M)$  with  $m, n \in \mathbb{N}$ , and  $M \in \mathbb{F}^{m \times (1+3n)}$ , where  $\mathbf{x} = \mathbf{x} \in \mathbb{F}^m$ , and where  $\mathbf{w} = (\mathbf{w}_1, \mathbf{w}_r, \mathbf{w}_o) \in \mathbb{F}^n \times \mathbb{F}^n \times \mathbb{F}^n$ ; satisfies the Hadamard Product Relation iff both*

$$\mathbf{x} = M \begin{pmatrix} 1 \\ \frac{\mathbf{w}_1}{\mathbf{w}_r} \\ \frac{\mathbf{w}_r}{\mathbf{w}_o} \end{pmatrix} \quad (1)$$

and

$$\mathbf{w}_1 \circ \mathbf{w}_r = \mathbf{w}_o \quad , \quad (2)$$

where  $\circ$  denotes the Hadamard (i.e., entry-wise) product; and in this case we write  $(\mathbf{i}, \mathbf{x}, \mathbf{w}) \in \mathcal{R}_{\text{HPR}}$ .

### 2.3 Interactive Proof Systems

**Definition 2 (Interactive Proof System).** Let  $\mathcal{R}$  be an indexed relation with corresponding relation language  $\mathcal{L}(\mathcal{R})$ . An interactive proof system is a pair  $(P, V)$  of stateful interactive Turing machines such that: the input to  $P$  is  $(i, x, w)$ , the input to  $V$  is  $(i, x)$ ;  $P$  and  $V$  exchange  $r = r(|i|)$  messages in total; and in the last step of the protocol  $V$  outputs a single bit  $b \in \{\top, \perp\}$ . The system satisfies two more properties:

- Completeness —  $V$  accepts members of  $\mathcal{L}(\mathcal{R})$ :  $(i, x) \in \mathcal{L}(\mathcal{R}) \Rightarrow b = \top$ .
- Soundness (with soundness error  $\sigma$ ) —  $V$  rejects non-members of  $\mathcal{L}(\mathcal{R})$  except with probability at most  $\sigma$  taken over the all random coins involved:  $\Pr[(i, x) \notin \mathcal{L}(\mathcal{R}) \Rightarrow b = \perp] \geq 1 - \sigma$ .

Soundness becomes a moot point when for the given index  $i$  every instance  $x$  has a matching witness  $w$  such that  $(i, x, w) \in \mathcal{R}$ . In this case a stronger notion called *knowledge soundness* [2] is preferred, which informally requires that any adversary that successfully convinces the verifier can be made to leak a witness by an extractor machine that has the same interface as the verifier but can additionally reset the adversary to an earlier point in time without forgetting the observed transcripts. In our context, all witnesses are encoded into oracles, and the prover displays knowledge of them simply by providing the oracles to the verifier. As a result, at our level of abstraction, knowledge soundness follows automatically from soundness. When the oracles are simulated by a concrete cryptographic tool, knowledge soundness becomes an important consideration that is not automatically satisfied. However, this cryptographic instantiation is beyond the scope of this paper.

A proof system is zero-knowledge [9] if, informally, an authentic transcript could have been produced by an adversary who is ignorant of the witness. More formally, the distribution of authentic transcripts must be samplable with public information only.

**Definition 3 (Honest-Verifier Zero-Knowledge).** Let  $\mathcal{R}$  be an indexed relation and let  $(P, V)$  be a proof system for  $\mathcal{R}$ . Let  $tr \leftarrow \langle P(i, x, w), V(i, x) \rangle$  denote the assignment to the variable  $tr$  of the transcript arising from the interaction between  $P$  with input  $(i, x, w)$  and  $V$  with input  $(i, x)$ . The proof system  $(P, V)$  is honest-verifier zero-knowledge if there exists a polynomial-time Turing machine  $S$  such that the distribution  $\mathcal{D}_0$  of authentic transcripts  $tr \leftarrow \langle P(i, x, w), V(i, x) \rangle$ , is identical to the distribution  $\mathcal{D}_1$  of simulated transcripts  $tr \leftarrow S(i, x)$ . When  $\mathcal{D}_0$  and  $\mathcal{D}_1$  are distinct, we consider the statistical distance and use the term Statistical Honest-Verifier Zero-Knowledge.

### 2.4 Polynomial IOP

Informally, a Polynomial IOP is an abstract proof system, where the prover sends polynomials and the verifier, instead of reading the polynomials in their entirety, is allowed to query the polynomial as oracles in select points.

**Definition 4 (Polynomial IOP).** Let  $\mathcal{R}$  be an indexed relation with corresponding indexed language  $\mathcal{L}(\mathcal{R})$ ,  $\mathbb{F}$  some finite field, and  $d \in \mathbb{N}$  a degree bound. A Polynomial IOP for  $\mathcal{R}$  with degree bound  $d$  is a pair of interactive machines  $(P, V)$ , satisfying the following description.

- $(P, V)$  is an interactive proof for  $\mathcal{L}(\mathcal{R})$  with  $r$  rounds, and with soundness error  $\sigma$ .
- $P$  sends polynomials  $f_i(X) \in \mathbb{F}[X]$  of degree at most  $d$  to  $V$ .
- $V$  is an oracle machine with access to a list of oracles, which contains one oracle for each polynomial it has received from the prover.
- When an oracle associated with a polynomial  $f_i(X)$  is queried on a point  $z_j \in \mathbb{F}$ , the oracle responds with the value  $f_i(z_j)$ .
- $V$  sends challenges  $\alpha_k \in \mathbb{F}$  to  $P$ .
- $V$  is public coin.

Definition 4 stipulates one global degree bound  $d$  for all polynomials. In Appendix A<sup>2</sup> we offer this alternative definition that stipulates individual degree bounds  $d_i$  for each polynomial  $f_i(X)$ . The same appendix presents a transformation between definitions to establish their equivalence. This transformation does lose some generality: queries in  $z_j = 0$  are not allowed, the global-bound protocol has one polynomial more, and the soundness error increases by at most  $\frac{2p+d-1}{|\mathbb{F}|-1}$ , where  $p$  is the original number of polynomials. However, these restrictions are not significant for typical applications of Polynomial IOPs, where the field  $\mathbb{F}$  is large. Therefore, without too much loss of generality, we may assume for the sake of a simpler presentation that the polynomials come with individual degree bounds.

With a minor extension, Polynomial IOPs can appropriately capture preprocessing. This extension introduces third machine, the *indexer*  $I$ . As its name suggests,  $I$  reads only  $i$ , and it outputs a list of polynomials to which  $V$  has oracle access.

**Definition 5 (Polynomial IOP with Preprocessing).** Let  $\mathcal{R}$  be an indexed relation with corresponding language  $\mathcal{L}(\mathcal{R})$ . A Polynomial IOP with Preprocessing is a tuple of interactive machines  $(I, P, V)$  such that  $(P, V)$  is a Polynomial IOP for  $\mathcal{L}(\mathcal{R})$  and such that

- $I$  takes  $i$  for input and outputs a list of polynomials of degree at most  $d$ ;
- $V$  has oracle access to these polynomials in addition to the polynomials it receives from  $P$ .

Some of the Polynomial IOPs in this paper are designed for modular composition. As a result,  $V$  does not begin with an empty list of polynomial oracles. In order to define the relations that these Polynomial IOPs realize, we denote by  $[f_i(X)]$  a polynomial  $f_i(X)$  that was sent to  $V$  by  $I$  or  $P$  at some earlier stage and to which  $V$  has oracle access.

<sup>2</sup> This and other appendices are available in the full version of this paper [13].

### 3 Dense Linear Algebra Relations

#### 3.1 Inner Product

Bünz *et al.* [6] are the first to sketch a Polynomial IOP that realizes an inner product relation between two vectors. It relies on the fact that the inner product of the coefficient vectors of  $f_a(X)$  and  $f_b(X)$  is the middle coefficient of  $f_a(X) \cdot X^d \cdot f_b(X^{-1})$ , assuming that both  $f_a(X)$  and  $f_b(X)$  are of degree  $d$ . To verify the middle coefficient is indeed the claimed inner product  $c$ ,  $\mathsf{V}$  needs two polynomials: the left half  $l(X)$  and the right half  $r(X)$ , both of degree  $d-1$ . Then the identity  $f_a(X) \cdot X^d \cdot f_b(X^{-1}) = l(X) + X^d \cdot c + X^{d+1} \cdot r(X)$  cannot hold in more than  $2d$  points unless  $c$  is the correct inner product.

Our variant of this protocol achieves the same result with the same number of queries but with one polynomial oracle less. This trade-off induces a doubling of the polynomial's degree and an increase-by-one in the number of distinct evaluation points. To see how this is achieved, observe that the coefficient on  $X^d$  of the polynomial  $f_a(X) \cdot X^d \cdot f_b(X^{-1}) - c \cdot X^d$  is zero. The same is true for  $h(X) = \bar{h}(X) \cdot \gamma^d - \bar{h}(\gamma X)$  for *any*  $\bar{h}(X)$  and any  $\gamma$  with a large enough multiplicative order. If  $f_a(X) \cdot X^d \cdot f_b(X^{-1}) = \sum_{i=0}^{2d} c_i X^i$  (with  $c_d = c$ ), then  $\mathsf{P}$  can obtain  $\bar{h}(X)$  by setting its  $i$ th coefficient to  $c_i/(\gamma^d - \gamma^i)$  when  $i \neq d$ , or uniformly at random when  $i = d$ . The verifier  $\mathsf{V}$  tests that the coefficient of  $X^d$  is indeed zero by sampling the left and right hand sides of the polynomial identity

$$\bar{h}(X) \cdot \gamma^d - \bar{h}(\gamma \cdot X) = f_a(X) \cdot X^d \cdot f_b(X^{-1}) - c \cdot X^d \quad (3)$$

in a uniformly random point  $z \xleftarrow{\$} \mathbb{F} \setminus \{0\}$ . The multiplicative order of  $\gamma$  must be larger than  $2d$  for this  $\bar{h}(X)$  to exist; for simplicity set  $\gamma$  to the smallest element that generates  $\mathbb{F} \setminus \{0\}$ ,  $\times$ .

Formally, the relation realized by inner product protocols is

$$\mathcal{R}_{\text{ip}} = \left\{ (\mathbf{i}, \mathbf{x}, \mathbf{w}) \left| \begin{array}{l} \mathbf{i} = d \\ \mathbf{x} = ([f_a(X)], [f_b(X)], c) \\ \mathbf{w} = (f_a(X), f_b(X)) \\ f_a(X) = \sum_{i=0}^d a_i X^i \\ f_b(X) = \sum_{i=0}^d b_i X^i \\ c = \sum_{i=0}^d a_i b_i \end{array} \right. \right\}. \quad (4)$$

**Theorem 1 (Security of InnerProduct).** *Protocol InnerProduct of Protocol 1 is a Polynomial IOP for  $\mathcal{L}(\mathcal{R}_{\text{ip}})$  with completeness and soundness with soundness error  $\sigma = \frac{2d}{|\mathbb{F}|-1}$ .*

*Proof.* The protocol revolves around the polynomial identity of Equation (3). The verifier tests this identity by sampling left and right hand sides in a random point  $z$ . Since this is an identity whenever  $c = \mathbf{a}^\top \mathbf{b}$ , completeness follows. For soundness, observe that when  $c \neq \mathbf{a}^\top \mathbf{b}$  then the coefficient of  $X^d$  on the right hand side of (3) is nonzero whereas the matching coefficient of the left hand side

<p><b>description:</b> decides <math>\mathcal{L}(\mathcal{R}_{\text{ip}})</math></p> <p><b>inputs:</b> <math>\mathbf{i} : \mathbf{d}</math>  <math>\mathbf{x} : ([f_{\mathbf{a}}(X)], [f_{\mathbf{b}}(X)], c)</math>  <math>\mathbf{w} : (f_{\mathbf{a}}(X), f_{\mathbf{b}}(X))</math></p> <p><b>begin</b></p> <ul style="list-style-type: none"> <li>P computes <math>f_c(X) = \sum_{i=0}^{2\mathbf{d}} c_i X^i \leftarrow f_{\mathbf{a}}(X) \cdot X^{\mathbf{d}} \cdot f_{\mathbf{b}}(X^{-1})</math></li> <li>P computes <math>\bar{h}(X) = \sum_{i=0}^{2\mathbf{d}} \bar{h}_i X^i</math> with <math>\bar{h}_i \leftarrow \frac{c_i}{\gamma^{\mathbf{d}-\gamma^i}}</math> for all <math>i \neq \mathbf{d}</math> and <math>\bar{h}_{\mathbf{d}} \xleftarrow{\\$} \mathbb{F}</math></li> <li>P sends <math>\bar{h}(X)</math> of degree at most <math>2\mathbf{d}</math> to V</li> <li>V samples <math>z \xleftarrow{\\$} \mathbb{F} \setminus \{0\}</math> and queries <math>([f_{\mathbf{a}}(X)], [f_{\mathbf{b}}(X)], [\bar{h}(X)], [\bar{h}(X)])</math> in <math>(z, z^{-1}, z, \gamma \cdot z)</math></li> <li>V receives <math>y_a = f_{\mathbf{a}}(z)</math>, <math>y_b = f_{\mathbf{b}}(z^{-1})</math>, <math>y_h = \bar{h}(z)</math>, and <math>y_h^* = \bar{h}(\gamma \cdot z)</math></li> <li>V tests <math>y_h \cdot \gamma^{\mathbf{d}} - y_h^* \stackrel{?}{=} y_a \cdot y_b \cdot z^{\mathbf{d}} - c \cdot z^{\mathbf{d}}</math></li> </ul>
--

**Protocol 1: InnerProduct**

is zero. There are at most  $2\mathbf{d}$  points  $z$  where left and right hand sides are equal, since both hands are bounded by this degree. By the Schwartz-Zippel lemma, the probability of a false accept is  $\sigma = \frac{2\mathbf{d}}{|\mathbb{F}|-1}$ .  $\square$

### 3.2 Batched Inner Product

We can batch multiple invocations of protocol `InnerProduct` into a single protocol that requires the prover to send only one polynomial oracle. Formally, the relation is given by

$$\mathcal{R}_{\text{bip}} = \left\{ (\mathbf{i}, \mathbf{x}, \mathbf{w}) \left| \begin{array}{l} \mathbf{i} = (m, \mathbf{d}) \\ \mathbf{x} = \{([f_{\mathbf{a}_i}(X)], [f_{\mathbf{b}_i}(X)], c_i)\}_{i=1}^m \\ \mathbf{w} = \{(f_{\mathbf{a}_i}(X), f_{\mathbf{b}_i}(X))\}_{i=1}^m \\ \forall i \in \{0, \dots, m-1\} \cdot f_{\mathbf{a}_i}(X) = \sum_{j=0}^{\mathbf{d}} a_{ij} X^j \\ \forall i \in \{0, \dots, m-1\} \cdot f_{\mathbf{b}_i}(X) = \sum_{j=0}^{\mathbf{d}} b_{ij} X^j \\ \forall i \in \{0, \dots, m-1\} \cdot c_i = \sum_{j=0}^{\mathbf{d}} a_{ij} b_{ij} \end{array} \right. \right\}. \quad (5)$$

**Theorem 2 (Security of BatchedInnerProduct).** *Protocol `BatchedInnerProduct` of Protocol 2 is a Polynomial IOP for  $\mathcal{L}(\mathcal{R}_{\text{bip}})$  with completeness and soundness with soundness error  $\sigma = \frac{2\mathbf{d}+m-1}{|\mathbb{F}|-1}$ .*

*Proof.* Let  $\sum_{j=0}^{2\mathbf{d}} c_{i,j} X^j = f_{\mathbf{a}_i}(X) \cdot f_{\mathbf{b}_i}(X^{-1}) \cdot X^{\mathbf{d}}$  for all  $i \in \{1, \dots, m\}$ . Furthermore, let  $\bar{H}(X, Y) = \sum_{i=1}^m \sum_{j=0}^{2\mathbf{d}} h_{i,j} X^j Y^{i-1}$  with  $h_{i,j} = \frac{c_{i,j}}{\gamma^{\mathbf{d}-\gamma^j}}$  for  $j \neq \mathbf{d}$ , arbitrary  $h_{i,\mathbf{d}}$  for  $i > 0$ , and  $h_{0,\mathbf{d}}$  such that  $\bar{h}(X) = \bar{H}(X, \alpha)$ .

The protocol revolves around the bivariate polynomial identity

$$\bar{H}(X, Y) \cdot \gamma^{\mathbf{d}} - \bar{H}(\gamma \cdot X, Y) = X^{\mathbf{d}} \cdot \sum_{i=1}^m (f_{\mathbf{a}_i}(X) \cdot f_{\mathbf{b}_i}(X^{-1}) - c_i) \cdot Y^{i-1}. \quad (6)$$

<p><b>description:</b> decides <math>\mathcal{L}(\mathcal{R}_{\text{bip}})</math></p> <p><b>inputs:</b> <math>\mathbf{i} : (m, \mathbf{d})</math>  <math>\mathbf{x} : \{([f_{\mathbf{a}_i}(X)], [f_{\mathbf{b}_i}(X)], c_i)\}_{i=1}^m</math>  <math>\mathbf{w} : \{(f_{\mathbf{a}_i}(X), f_{\mathbf{b}_i}(X))\}_{i=1}^m</math></p> <p><b>begin</b></p> <p>    P computes <math>f_{c_i}(X) \leftarrow f_{\mathbf{a}_i}(X) \cdot f_{\mathbf{b}_i}(X^{-1}) \cdot X^{\mathbf{d}}</math> for <math>i</math> from 1 to <math>m</math></p> <p>    V samples <math>\alpha \xleftarrow{\\$} \mathbb{F} \setminus \{0\}</math> and sends <math>\alpha</math> to P</p> <p>    P computes <math>f_c(X) \leftarrow \sum_{i=1}^m f_{c_i}(X) \cdot \alpha^{i-1}</math></p> <p>    P computes <math>\bar{h}(X) = \sum_{j=0}^{2\mathbf{d}} \bar{h}_j X^j</math> with <math>\bar{h}_j \leftarrow \frac{c_j}{\gamma^{\mathbf{d}-\gamma^j}}</math> for all <math>j \neq \mathbf{d}</math> and <math>\bar{h}_{\mathbf{d}} \xleftarrow{\\$} \mathbb{F}</math></p> <p>    P sends <math>\bar{h}(X)</math> of degree at most <math>2\mathbf{d}</math> to V</p> <p>    V samples <math>z \xleftarrow{\\$} \mathbb{F} \setminus \{0\}</math> and queries  <math>(\{([f_{\mathbf{a}_i}(X)], [f_{\mathbf{b}_i}(X)])\}_{i=1}^m, [\bar{h}(X)], [\bar{h}(X)])</math> in <math>(\{z, z^{-1}\}_{i=1}^m, z, \gamma \cdot z)</math></p> <p>    V receives <math>y_{a,i} = f_{\mathbf{a}_i}(z)</math>, <math>y_{b,i} = f_{\mathbf{b}_i}(z^{-1})</math> for <math>i</math> from 1 to <math>m</math>, and <math>y_h = \bar{h}(z)</math>,  <math>y_h^* = \bar{h}(\gamma \cdot z)</math></p> <p>    V tests <math>y_h \cdot \gamma^{\mathbf{d}} - y_h^* \stackrel{?}{=} z^{\mathbf{d}} \cdot \sum_{i=1}^m (y_{a,i} \cdot y_{b,i} - c_i) \cdot \alpha^{i-1}</math></p>
---

**Protocol 2: BatchedInnerProduct**

The verifier tests this identity by sampling left and right hand sides in a random point  $(z, \alpha)$ . Since this is an identity whenever  $c_i = \mathbf{a}_i^T \mathbf{b}_i$  for all  $i$  from 1 to  $m$ , completeness follows. For soundness, consider when for some  $i$ ,  $c_i \neq \mathbf{a}_i^T \mathbf{b}_i$ . Then left and right hand sides of (6) are unequal. There are at most  $2\mathbf{d} + m - 1$  points  $(z, \alpha)$  in which left and right hand sides are equal, since both hands are bounded by this total degree. By the (two-dimensional) Schwartz-Zippel lemma, the probability of a false accept is  $\sigma = \frac{2\mathbf{d}+m-1}{|\mathbb{F}|-1}$ .  $\square$

This inner product protocol and its batched version are convenient for zero-knowledge. The verifier V makes two queries to the polynomial  $\bar{h}(X)$ : one in  $z$  and one in  $\gamma \cdot z$ . Since  $\bar{h}(X)$  has one uniformly random coefficient, one of the responses is uniformly random. The other one is such that the tested equality is true. So the honest verifier learns no new information from  $\bar{h}(X)$ .

### 3.3 Modular Reduction

We start with a protocol that will be used as a subprotocol in the sequel. This protocol establishes that one polynomial,  $r(X)$ , is the remainder after division of a second polynomial  $f(X)$ , by a third,  $d(X)$ . This third polynomial is assumed to be known, but the protocol can be naturally amended to allow V only oracle access to  $[d(X)]$ . Formally, the relation is given by

$$\mathcal{R}_{\text{reduce}} = \left\{ (\mathbf{i}, \mathbf{x}, \mathbf{w}) \left| \begin{array}{l} \mathbf{i} = (\mathbf{d}_f, \mathbf{d}_r) \\ \mathbf{x} = ([f(X)], [r(X)], d(X)) \\ \mathbf{w} = (f(X), r(X)) \\ \exists q(X) \in \mathbb{F}[X] \cdot f(X) = q(X) \cdot d(X) + r(X) \\ \deg(f) \leq \mathbf{d}_f \\ \deg(r) \leq \mathbf{d}_r \end{array} \right. \right\}. \quad (7)$$

<p><b>description:</b> decides <math>\mathcal{L}(\mathcal{R}_{\text{reduce}})</math></p> <p><b>inputs:</b> <math>\mathbf{i} : (\mathbf{d}_f, \mathbf{d}_r)</math>  <math>\mathbf{x} : ([f(X)], [r(X)], d(X))</math>  <math>\mathbf{w} : (f(X), r(X))</math></p> <p><b>begin</b></p> <ul style="list-style-type: none"> <li>┌ P computes <math>q</math> such that <math>f(X) = q(X) \cdot d(X) + r(X)</math></li> <li>┌ P sends <math>q(X)</math> of degree at most <math>\mathbf{d}_f - \text{deg}(d)</math> to <math>\mathbf{V}</math></li> <li>┌ <math>\mathbf{V}</math> samples <math>z \xleftarrow{\\$} \mathbb{F} \setminus \{0\}</math> and queries <math>[f(X)], [q(X)],</math> and <math>[r(X)]</math> in <math>z</math></li> <li>┌ <math>\mathbf{V}</math> receives <math>y_f = f(z), y_q = q(z),</math> and <math>y_r = r(z)</math></li> <li>┌ <math>\mathbf{V}</math> tests <math>y_f \stackrel{?}{=} y_q \cdot d(z) + y_r</math></li> </ul>
---

**Protocol 3: ModReduce**

**Theorem 3 (Security of ModReduce).** *Protocol ModReduce of Protocol 3 is a Polynomial IOP for  $\mathcal{L}(\mathcal{R}_{\text{reduce}})$  with completeness and soundness with soundness error  $\sigma = \mathbf{d}_f / |\mathbb{F}|$ .*

*Proof.* Completeness follows from construction: dividing  $f(X)$  by  $d(X)$  gives quotient  $q(X)$  and remainder  $r(X)$ . Therefore,  $f(X) = q(X) \cdot d(X) + r(X)$  is an identity of polynomials and guaranteed to hold everywhere including in the point  $z$ .

For soundness, observe that when  $r(X) \not\equiv f(X) \pmod{d(X)}$  then  $d(X)$  does not divide  $f(X) - r(X)$ . As a result,  $f(X) \neq q(X) \cdot d(X) + r(X)$  is an inequality of polynomials with degree  $\text{deg}(d) + \text{deg}(q) = \mathbf{d}_f$ . Due to the Schwartz-Zippel lemma, the left and right hand sides can evaluate to the same value in at most  $\mathbf{d}_f$  choices for  $z$ . The probability of  $\mathbf{V}$  accepting when  $r(X) \not\equiv f(X) \pmod{d(X)}$  is therefore  $\sigma = \mathbf{d}_f / |\mathbb{F}|$ .

What is left to argue is that  $\mathbf{P}$  fails to convince  $\mathbf{V}$  when the congruence  $r(X) \equiv f(X) \pmod{d(X)}$  holds, but  $r(X)$  is not equal to the remainder after division of  $f(X)$  by  $d(X)$ . The representatives of the congruence class of  $r(X)$  are apart by polynomials of degree at least  $\text{deg}(d)$ , there is only one representative of degree at most  $\mathbf{d}_r < \text{deg}(d)$ . The index value  $\mathbf{d}_r$  therefore already constrains  $r(X)$  to a unique polynomial.  $\square$

**3.4 Matrix-Vector Product**

The next protocol involves two polynomials that represent vectors in the monomial coefficient basis. It establishes that the one vector is the result of applying a linear transformation to the other. This linear transformation itself can be known and computed explicitly by the verifier. However, for succinct verifiers it is more appealing to encode this matrix into a polynomial oracle. Depending on the context, either the protocol's preprocessing phase produces this oracle, or another external protocol does.

Specifically, let  $\mathbf{a} \in \mathbb{F}^n$  and  $\mathbf{b} \in \mathbb{F}^m$  and  $M \in \mathbb{F}^{m \times n}$  with the element in row  $i$  and column  $j$  (both indices starting at zero) indexed as  $M_{[i,j]}$ . These objects are represented as polynomials with  $a_{[i]}$  being  $i$ th element of  $\mathbf{a}$  and simultaneously

the coefficient of the monomial  $X^i$  in  $f_{\mathbf{a}}(X)$ , and similarly for  $\mathbf{b}$ ,  $b_{[i]}$ , and  $f_{\mathbf{b}}(X)$ . When encoded into polynomial form, the matrix is encoded in row-first order, specifically  $f_M(X) = \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} M_{[i,j]} X^{in+j}$ . The protocol establishes that  $\mathbf{b} = M\mathbf{a}$ . Formally, the relation is given by

$$\mathcal{R}_{\text{mvp}} = \left\{ (\mathbf{i}, \mathbb{x}, \mathbb{w}) \left| \begin{array}{l} \mathbf{i} = (m, n, M) \\ \mathbb{x} = ([f_{\mathbf{a}}(X)], [f_{\mathbf{b}}(X)]) \\ \mathbb{w} = (f_{\mathbf{a}}(X), f_{\mathbf{b}}(X)) \\ f_{\mathbf{a}}(X) = \sum_{i=0}^{n-1} \mathbf{a}_{[i]} X^i \text{ for some } \mathbf{a} \in \mathbb{F}^n \\ f_{\mathbf{b}}(X) = \sum_{i=0}^{m-1} \mathbf{b}_{[i]} X^i \text{ for some } \mathbf{b} \in \mathbb{F}^m \\ \mathbf{b} = M\mathbf{a} \end{array} \right. \right\}. \quad (8)$$

**description:** decides  $\mathcal{L}(\mathcal{R}_{\text{mvp}})$   
**inputs:**  $\mathbf{i} : (m, n, M)$   
 $\mathbb{x} : ([f_{\mathbf{a}}(X)], [f_{\mathbf{b}}(X)])$   
 $\mathbb{w} : (f_{\mathbf{a}}(X), f_{\mathbf{b}}(X))$   
*// pre-processing*  
**begin**  
 $\lfloor$  I computes  $f_M(X) \leftarrow \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} M_{[i,j]} X^{in+j}$   
 $\lfloor$  I sends  $f_M(X)$  of degree at most  $mn - 1$  to P and V  
**begin**  
 $\lfloor$  V samples  $\alpha \xleftarrow{\$} \mathbb{F} \setminus \{0\}$  and sends  $\alpha$  to P  
P computes  $r(X) \leftarrow f_M(X) \bmod X^n - \alpha$   
P sends  $r(X)$  of degree at most  $n - 1$  to V  
P and V run **ModReduce** with  $\mathbf{i}^{(1)} = (mn - 1, n - 1)$ ,  
 $\mathbb{x}^{(1)} = ([f_M(X)], [r(X)], X^n - \alpha)$ , and  $\mathbb{w}^{(1)} = (f_M(X), r(X))$   
V queries  $[f_{\mathbf{b}}(X)]$  in  $\alpha$  and receives  $y_{\alpha\tau\mathbf{b}} = f_{\mathbf{b}}(\alpha)$   
P and V run **InnerProduct** with  $\mathbf{i}^{(2)} = n - 1$ ,  $\mathbb{x}^{(2)} = ([r(X)], [f_{\mathbf{a}}(X)], y_{\alpha\tau\mathbf{b}})$ ,  
 $\mathbb{w}^{(2)} = (r(X), f_{\mathbf{a}}(X))$   
 $\lfloor$

**Protocol 4:** DenseMVP

**Theorem 4 (Security of DenseMVP).** *Protocol DenseMVP of Protocol 4 is a Polynomial IOP for  $\mathcal{L}(\mathcal{R}_{\text{mvp}})$  with completeness and soundness with soundness error  $\sigma = \frac{mn+m+2n-4}{|\mathbb{F}|-1}$ .*

*Proof.* Let  $\boldsymbol{\alpha}^\top = (\alpha^0, \alpha^1, \dots)$  and  $\mathbf{r}^\top = \boldsymbol{\alpha}^\top M$ , and consider the equations

$$\mathbf{b} = M\mathbf{a} \quad (9)$$

$$\boldsymbol{\alpha}^\top \mathbf{b} = \boldsymbol{\alpha}^\top M\mathbf{a} \quad (10)$$

$$\sum_{i=0}^{m-1} \alpha^i b_{[i]} = \mathbf{r}^\top \mathbf{a} \quad (11)$$

$$f_{\mathbf{b}}(\alpha) = \sum_{i=0}^{n-1} r_{[i]} a_{[i]} \quad (12)$$

$$y_{\boldsymbol{\alpha}^\top \mathbf{b}} = \text{coeffs}(r(X)) \cdot \text{coeffs}(f_{\mathbf{a}}(X)) \quad (13)$$

$$(\mathbf{i}^{(2)}, \mathbb{x}^{(2)}) = (n-1, ([r(X)], [f_{\mathbf{a}}(X)], y_{\boldsymbol{\alpha}^\top \mathbf{b}})) \in \mathcal{L}(\mathcal{R}_{\text{ip}}) , \quad (14)$$

where  $\text{coeffs} : \mathbb{F}[X] \rightarrow \mathbb{F}^n$  is the function that returns the vector of coefficients of its argument. Observe that  $\text{coeffs}(r(X)) = \mathbf{r}$ , by substituting  $X^n$  by  $\alpha$  in the expression for  $f_M(X)$ :

$$\sum_{i=0}^{m-1} \sum_{j=0}^{n-1} M_{[i,j]} X^{in+j} \xrightarrow{X^n \mapsto \alpha} r(X) = \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} M_{[i,j]} \alpha^i X^j \quad (15)$$

$$= \sum_{j=0}^{n-1} \left( \sum_{i=0}^{m-1} M_{[i,j]} \alpha^i \right) X^j \quad (16)$$

$$= \sum_{j=0}^{n-1} r_{[j]} X^j . \quad (17)$$

Completeness follows from the implications (9)  $\Rightarrow$  (10)  $\Leftrightarrow$  (11)  $\Leftrightarrow$  (12)  $\Rightarrow$  (13)  $\Rightarrow$  (14).

For soundness, there are 3 events that can cause  $\mathbf{V}$  to accept despite  $\mathbf{b} \neq M\mathbf{a}$ :

1. (9)  $\not\Leftrightarrow$  (10). The probability of this event is at most  $\frac{m-1}{|\mathbb{F}|-1}$  due to the Schwartz-Zippel lemma.
2. (12)  $\not\Leftrightarrow$  (13) because  $r(X)$  is not the remainder of  $f_M(X)$  after division by  $X^n - \alpha$ . The probability of this event is at most  $\frac{mn-1}{|\mathbb{F}|}$ , the soundness error of ModReduce.
3. (13)  $\not\Leftrightarrow$  (14), because  $y_{\boldsymbol{\alpha}^\top \mathbf{b}}$  is not the inner product of the coefficient vectors of  $r(X)$  and  $f_{\mathbf{a}}(X)$ . The probability of this event is at most  $\frac{2n-2}{|\mathbb{F}|}$ , the soundness error of InnerProduct.

By the union bound, the soundness error of DenseMVP is bounded by  $\sigma = \frac{mn+m+2n-4}{|\mathbb{F}|-1}$ .  $\square$

Note that after unrolling, the verifier the DenseMVP protocol tests two polynomial identities. One arises from expanding ModReduce, and the other arises from InnerProduct. Both polynomial identities involve the polynomial  $r(X)$ , and

as a result it can be eliminated and the identities merged. We present the unrolled and optimized version in Appendix C.1.

To see that this merger has no effect on soundness, observe that the inequality  $lhs_1 \neq lhs_2$  implies  $r(X) \neq lhs_1$  or  $r(X) \neq lhs_2$ . The verifier therefore accepts this false instance with a probability bounded by the same soundness error as the unoptimized protocol. This optimization strategy translates more generally to (some) other Polynomial IOPs: to eliminate a polynomial that is common to two identities, move it to the right hand side and then equate both left hand sides.

### 3.5 Hadamard Product

The next protocol establishes that the Hadamard (or component-wise) product of two vectors is equal to a third. These vectors are represented as the coefficient vectors of polynomials  $f_a(X)$ ,  $f_b(X)$ , and  $f_c(X)$  such that  $\mathbf{c} = \mathbf{a} \circ \mathbf{b}$  and  $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{F}^{d+1}$ . The protocol relies on the fact that  $\mathbf{c} = \mathbf{a} \circ \mathbf{b}$  implies  $\alpha^\top (\mathbf{a} \circ \mathbf{b}) = \alpha^\top \mathbf{c}$  for all vectors  $\alpha$ . In other words, one can simply sample a random scalar  $\alpha \xleftarrow{\$} \mathbb{F} \setminus \{0\}$ , and check the inner product of  $\alpha \circ \mathbf{a}$  with  $\mathbf{b}$  against the inner product  $\alpha^\top \mathbf{c}$ . Note that the right hand side of this check amounts to  $f_c(\alpha)$  and the operands in the left hand side amount to the coefficient vectors of  $f_a(\alpha X)$  and  $f_b(X)$ , respectively. Formally, the relation is given by

$$\mathcal{R}_{\text{hadamard}} = \left\{ (\mathbf{i}, \mathbf{x}, \mathbf{w}) \left| \begin{array}{l} \mathbf{i} = \mathbf{d} \\ \mathbf{x} = ([f_a(X)], [f_b(X)], [f_c(X)]) \\ \mathbf{w} = (f_a(X), f_b(X), f_c(X)) \\ \forall i \in \{0, \dots, \mathbf{d}\}. a_i b_i = c_i \end{array} \right. \right\}. \quad (18)$$

<p><b>description:</b> decides <math>\mathcal{L}(\mathcal{R}_{\text{hadamard}})</math></p> <p><b>inputs:</b> <math>\mathbf{i}: \mathbf{d}</math>  <math>\mathbf{x}: [f_a(X)], [f_b(X)], [f_c(X)]</math>  <math>\mathbf{w}: f_a(X), f_b(X), f_c(X)</math></p> <p><b>begin</b></p> <div style="border-left: 1px solid black; border-right: 1px solid black; padding-left: 10px;"> <p>V samples <math>\alpha \xleftarrow{\\$} \mathbb{F} \setminus \{0\}</math> and sends <math>\alpha</math> to P</p> <p>P evaluates <math>y \leftarrow f_c(\alpha)</math></p> <p>V queries <math>[f_c(X)]</math> in <math>\alpha</math> and receives <math>y = f_c(\alpha)</math></p> <p>P and V run <code>InnerProduct</code> with <math>\mathbf{i}^{(1)} = \mathbf{d}</math>, <math>\mathbf{x}^{(1)} = ([f_a(\alpha X)], [f_b(X)], y)</math>,  <math>\mathbf{w}^{(1)} = (f_a(\alpha X), f_b(X))</math>, where V simulates <math>[f_a(\alpha X)]</math> using <math>[f_a(X)]</math> and the scalar <math>\alpha</math></p> </div>
---

**Protocol 5: Hadamard**

**Theorem 5 (Security of Hadamard).** *Protocol Hadamard of Protocol 5 is a Polynomial IOP for  $\mathcal{L}(\mathcal{R}_{\text{hadamard}})$  with completeness and soundness with soundness error  $\sigma = 3\mathbf{d}/(|\mathbb{F}| - 1)$ .*

*Proof.* Consider the following sequence of equations.

$$\mathbf{a} \circ \mathbf{b} = \mathbf{c} \tag{19}$$

$$\boldsymbol{\alpha}^\top \cdot (\mathbf{a} \circ \mathbf{b}) = \boldsymbol{\alpha}^\top \cdot \mathbf{c} \tag{20}$$

$$\sum_{i=0}^d (\alpha^i \mathbf{a}_{[i]}) \mathbf{b}_{[i]} = \sum_{i=0}^d \alpha^i \mathbf{c}_{[i]} \tag{21}$$

$$\text{coeffs}(f_{\mathbf{a}}(\alpha X)) \cdot \text{coeffs}(f_{\mathbf{b}}(X)) = f_{\mathbf{c}}(\alpha) \tag{22}$$

$$\text{coeffs}(f_{\mathbf{a}}(\alpha X)) \cdot \text{coeffs}(f_{\mathbf{b}}(X)) = y \tag{23}$$

$$(\mathbf{i}, \mathbf{x}) = (d, ([f_{\mathbf{a}}(\alpha X)], [f_{\mathbf{b}}(X)], y)) \in \mathcal{L}(\mathcal{R}_{\text{InnerProduct}}) \tag{24}$$

Completeness follows from the sequence of implications (19)  $\Rightarrow$  (20)  $\Leftrightarrow$  (21)  $\Leftrightarrow$  (22)  $\Leftrightarrow$  (23)  $\Rightarrow$  (24).

For soundness, consider when the reverse implications fail.

- (19)  $\not\Leftarrow$  (20). This event happens with probability at most  $d/(|\mathbb{F}| - 1)$  due to the Schwartz-Zippel lemma.
- (23)  $\not\Leftarrow$  (24). This event happens with probability at most  $2d/(|\mathbb{F}| - 1)$ , the soundness error of `InnerProduct`.

Therefore, the probability that  $\mathcal{V}$  accepts even though  $\mathbf{a} \circ \mathbf{b} \neq \mathbf{c}$  is bounded by  $\sigma = 3d/(|\mathbb{F}| - 1)$ .  $\square$

## 4 Sparse Linear Algebra Relations

The purpose of this section is to present an analogue of the `DenseMVP` Polynomial IOP but that works when the matrix  $M$  is represented sparsely, *i.e.*, as a list of nonzero coefficients and their coordinates. The full, formal presentation of this protocol is rather lengthy, and so we defer it to Appendix B. Here we present an intuitive, high-level overview with just enough detail so that the reader could reconstruct the deferred formal presentation.

### 4.1 High-Level Overview

Let  $M \in \mathbb{F}^{m \times n}$  be a matrix with only  $K$  nonzero elements. It can be represented by a triple of functions (`col`, `row`, `val`) via  $M = \sum_{k=0}^{K-1} \mathbf{e}_{\text{row}(k)} \mathbf{e}_{\text{col}(k)}^\top \cdot \text{val}(k)$ , where  $\mathbf{e}_i$  is the  $i$ th unit vector, where `row`, `col` :  $\mathbb{N} \rightarrow \mathbb{N}$  indicate the column and row of the  $k$ th element, and where `val` :  $\mathbb{N} \rightarrow \mathbb{F}$  indicates its value. We detail a protocol to establish that  $\mathbf{y} = M\mathbf{x}$ . We first explain the steps from a high level point of view.

**From MVP to bivariate polynomial evaluation.** A key component of the dense matrix-vector multiplication protocol is the evaluation of  $(f_M(X) \bmod X^n - \alpha)$  at the point  $z$ , where  $f_M(X)$  is the polynomial associated with the matrix  $M$ , i.e.,  $f_M(X) = \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} M_{[i,j]} X^{in+j}$ . This step can equivalently be interpreted as the evaluation of the bivariate polynomial  $f_M(X, Y) = \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} M_{[i,j]} X^i Y^j$  in the point  $(\alpha, z)$ . In other words, if we can achieve sparse bivariate polynomial evaluation, then we can achieve sparse matrix-vector products.

**From bivariate polynomials to univariate monomial vectors.** The reduction goes one step further: it is possible to achieve sparse bivariate polynomial evaluation given a procedure that establishes that the vector of coefficients of a dense polynomial is the same as the vector of monomials of a sparse univariate polynomial when evaluated in a given point. To see this, observe that a sparse bivariate polynomial  $f(X, Y) = \sum_{k=0}^{K-1} c_k X^{a_k} Y^{b_k}$  can be evaluated in a point  $(x, y)$  using the polynomials  $f_c(X) = \sum_{k=0}^{K-1} c_k X^k$ ,  $f_x(X) = \sum_{k=0}^{K-1} x^{a_k} X^k$ , and  $f_y(X) = \sum_{k=0}^{K-1} y^{b_k} X^k$ , simply by performing one Hadamard and one InnerProduct subprotocol. This reduction does introduce a problem, namely  $f_x(X)$  and  $f_y(X)$  cannot be known before  $\mathbf{V}$  supplies  $x$  and  $y$ . So how does  $\mathbf{P}$  commit to them, and how does  $\mathbf{V}$  verify that the received oracles match with the commitment?

**From univariate monomial vector to bit matrix.** Let's focus on  $f_x(X)$ , as  $f_y(X)$  proceeds analogously. This polynomial can be represented by a bit matrix  $B$ , which takes the value 1 in cells  $(a_k, k)$  and 0 elsewhere. Let  $H$  denote the largest such  $a_k$ , i.e.,  $H = \max_k a_k$ . Let furthermore  $\mathbf{x} = (x^0, x, x^2, \dots, x^{H-1})^\top$ , and  $\mathbf{z} = (z^0, z, z^2, \dots, z^{K-1})^\top$ . Then  $B$  represents the polynomial  $f_x(X)$  since  $f_x(z) = \mathbf{x}B\mathbf{z}$ .

**From bit matrix to Lagrange and Vandermonde matrices.** The idea is to decompose the matrix  $B \in \mathbb{F}^{H \times K}$  as the product of two matrices,  $L \in \mathbb{F}^{H \times H}$  and  $R \in \mathbb{F}^{H \times K}$ . Let  $\mathcal{H} \subset \mathbb{F}$  be a set of  $H$  distinct elements of  $\mathbb{F}$  and  $\varphi : \mathbb{N} \rightarrow \mathcal{H}$  any mapping from  $\{0, \dots, H-1\}$  to  $\mathcal{H}$ .  $L$  is the Lagrange matrix, whose  $h$ th row is the coefficient vector of  $\mathcal{L}_h(X)$ , which is the Lagrange polynomial taking the value 1 in  $\varphi(h)$  and 0 in all other points of  $\mathcal{H}$ . Symbolically:

$$\mathcal{L}_h(X) = \sum_{i=0}^{H-1} L_{[h,i]} X^i = \prod_{\substack{i=0 \\ i \neq h}}^{H-1} \frac{X - \varphi(i)}{\varphi(h) - \varphi(i)}. \quad (25)$$

$R$  is the Vandermonde matrix, whose rows are the (Hadamard) powers of  $(\varphi(a_k))_{k=0}^{K-1}$ . Specifically:

$$R = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ \varphi(a_0) & \varphi(a_1) & \cdots & \varphi(a_{K-1}) \\ \varphi(a_0)^2 & \varphi(a_1)^2 & \cdots & \varphi(a_{K-1})^2 \\ \vdots & \vdots & \cdots & \vdots \\ \varphi(a_0)^{H-1} & \varphi(a_1)^{H-1} & \cdots & \varphi(a_{K-1})^{H-1} \end{pmatrix}. \quad (26)$$

To verify that  $LR = B$ , observe that the inner product between  $L_{[h,:]}$  and  $R_{[:,k]}$  is equal to  $\mathcal{L}_h(\varphi(a_k))$ . When  $V$  provides  $(x, z)$  hoping to obtain  $\mathbf{x}^\top B \mathbf{z}$ ,  $P$  will respond with  $[\mathbf{x}^\top L]$  and  $[R\mathbf{z}]$  (in polynomial form), and both proceed to an `InnerProduct` protocol. We will refer to these vectors as the Lagrange and Vandermonde vectors, respectively. The next question is, how and against what does  $V$  verify them?

**Verifying the Lagrange vector.** After sending  $x$  and receiving the vector (encoded as a polynomial oracle)  $[\mathbf{x}^\top L]$ ,  $V$  sends  $\gamma$  to  $P$ , who responds with the vector  $[L\gamma]$ , where  $\gamma = (1, \gamma, \gamma^2, \dots, \gamma^{H-1})^\top$ . Let  $Z_{\mathcal{H}}(X)$  be the unique monic polynomial of degree  $H - 1$  that vanishes on  $\mathcal{H}$ . By repeating the equation

$$\mathcal{L}_h(\gamma) \cdot (\gamma - \varphi(h)) = \frac{Z_{\mathcal{H}}(\gamma)}{\prod_{\substack{i=0 \\ i \neq h}}^{H-1} (\varphi(h) - \varphi(i))} \quad (27)$$

for every  $h$ ,  $V$  can check  $L\gamma$  using a `Hadamard` subprotocol, assuming that  $P$  or  $I$  previously committed to oracles for  $f_\varphi(X) = \sum_{h=0}^{H-1} \varphi(h)X^h$  and  $f_{\mathcal{H}}(X) = \sum_{h=0}^{H-1} \left( \prod_{i \in \{0 \dots H-1\} \setminus \{h\}} (\varphi(h) - \varphi(i)) \right)^{-1} \cdot X^h$ . The next step is to query the oracles  $[\mathbf{x}L]$  and  $[L\gamma]$  in  $\gamma$  and  $x$ , respectively and verify that the responses match.

**Verifying the Vandermonde Vector.** A similar technique allows  $V$  to verify the Vandermonde vector. After sending  $z$  and receiving  $[R\mathbf{z}]$  back,  $V$  sends  $\delta$ , and  $P$  responds with  $[\delta^\top R]$ . Next,  $V$  checks that for every  $k \in \{0, \dots, K - 1\}$ ,

$$\left( \sum_{i=0}^{H-1} (\delta \cdot \varphi(a_k))^i \right) \cdot (\delta \varphi(a_k) - 1) = (\delta \varphi(a_k))^H - 1 \quad (28)$$

using another `Hadamard` protocol and the precommitted oracle  $f_a(X) = \sum_{k=0}^{K-1} \varphi(a_k)X^k$ . Lastly,  $V$  queries  $[R\mathbf{z}]$  in  $\delta$  to see if the response matches with  $[\delta R]$  when queried in  $z$ .

**Batching Lagrange and Vandermonde Vectors.** In order to establish the correct evaluation of the bivariate polynomial, the prover must establish the correct production of two univariate monomial vectors. A naïve implementation invokes the Lagrange vector and the Vandermonde vector procedure twice. However, it turns out to be possible to merge these two invocations, and save a total of 4 polynomials. We treat this optimization explicitly in Appendix C.

## 5 A Polynomial IOP for Arithmetic Circuits

### 5.1 The Protocol

The next protocol, Protocol 6 puts many of the previously developed tools together into a Polynomial IOP (with preprocessing) for arithmetic circuits as captured by the HPR. To differentiate our protocol from other similar ones, we name it Claymore.

<pre> <b>description:</b> realizes <math>\mathcal{R}_{\text{hpr}}</math> <b>inputs:</b> <math>\mathbf{i}: (m, n, M)</math> with <math>M \in \mathbb{F}^{m \times (3n+1)}</math>            <math>\mathbf{x}: \mathbf{x} \in \mathbb{F}^n</math>            <math>\mathbf{w}: (\mathbf{w}_1, \mathbf{w}_r, \mathbf{w}_o) \in \mathbb{F}^n \times \mathbb{F}^n \times \mathbb{F}^n</math>            // preprocessing <b>begin</b>     I runs MVP.I on <math>\mathbf{i}^{(1)} = (m, 3n + 1, M)</math>            // online <b>begin</b>     P computes <math>f_{wl} \leftarrow \sum_{j=0}^{n-1} \mathbf{w}_{1[j]} X^j</math>, <math>f_{wr} \leftarrow \sum_{j=0}^{n-1} \mathbf{w}_{r[j]} X^j</math>, and     <math>f_{wo} \leftarrow \sum_{j=0}^{n-1} \mathbf{w}_{o[j]} X^j</math>     P sends <math>f_{wl}(X)</math>, <math>f_{wr}(X)</math>, and <math>f_{wo}(X)</math>, all of degrees at most <math>n - 1</math>, to V     P computes <math>f_{1w}(X) \leftarrow 1 + X f_{wl}(X) + X^{n+1} f_{wr}(X) + X^{2n+1} f_{wo}(X)</math>     P computes <math>f_x(X)</math>, whose coefficient vectors correspond to     <math>\mathbf{x} = M (1   \mathbf{w}_1^\top   \mathbf{w}_r^\top   \mathbf{w}_o^\top)^\top</math>     P and V run MVP with <math>\mathbf{i}^{(1)} = (m, 3n + 1, M)</math>, <math>\mathbf{x}^{(1)} = ([f_{1w}(X)], [f_x(X)])</math>,     <math>\mathbf{w}^{(1)} = (f_{1w}(X), f_x(X))</math> where V simulates <math>[f_{1w}(X)]</math> using     <math>f_{1w}(X) = 1 + X f_{wl}(X) + X^{n+1} f_{wr}(X) + X^{2n+1} f_{wo}(X)</math>, <math>[f_{wl}(X)]</math>,     <math>[f_{wr}(X)]</math>, and <math>[f_{wo}(X)]</math>; and where V computes <math>[f_x(X)]</math> locally using     <math>\mathbf{x} = \mathbf{x}</math>     P and V run Hadamard with <math>\mathbf{i}^{(2)} = n - 1</math>,     <math>\mathbf{x}^{(2)} = ([f_{wl}(X)], [f_{wr}(X)], [f_{wo}(X)])</math>, <math>\mathbf{w}^{(2)} = (f_{wl}(X), f_{wr}(X), f_{wo}(X))</math> </pre>
--

**Protocol 6:** Claymore

**Theorem 6 (Security of Claymore).** *Protocol Claymore of Protocol 6 is a Polynomial IOP for  $\mathcal{R}_{\text{HPR}}$  with completeness and soundness error  $\sigma \leq \sigma_{\text{Hadamard}} + \sigma_{\text{MVP}}$ .*

*Proof.* Completeness follows from construction. Since the arguments are computed honestly, the subprotocols succeed and guarantee equalities (1) and (2), respectively.

Soundness. If the HPR instance is a false instance, then  $\mathbf{x} \neq M(1|\mathbf{w}_1^\top|\mathbf{w}_r^\top|\mathbf{w}_o^\top)^\top$  or  $\mathbf{w}_1 \circ \mathbf{w}_r \neq \mathbf{w}_o$ . As a result either the Hadamard protocol succeeds despite being run on a false instance, or the MVP protocol succeeds despite being run on a false instance. The probabilities of these events are respectively at most  $\sigma_{\text{Hadamard}}$  and at most  $\sigma_{\text{MVP}}$ .  $\square$

## 5.2 The Role of Preprocessing

The preprocessing phase can be omitted. In this case,  $V$  must compute  $f_M(X)$  locally. This task requires  $O(mn)$  work, or only  $O(K)$  if the matrix  $M$  has only  $K$  nonzero elements and is represented as such. When this phase is omitted, Claymore should be compared to the Polynomial IOP underlying Aurora [4].

When used with preprocessing, Claymore achieves fast verification. Specifically, the matrix  $M$  which determines the circuit being proved, is processed by the indexer. For long and drawn-out computations, this matrix is typically sparse and the SparseMVP is suitable. However, for short or shallow computations, DenseMVP is the better option. Depending on the choice of MVP protocol, the matching soundness error should be considered.

## 5.3 Optimizations

**Reuse  $\alpha$  Across Hadamard Protocols.** DenseClaymore has only one invocation of the Hadamard subprotocol, but the (partially unrolled) SparseClaymore has many more. It is worth reusing the same  $\alpha$  for all these invocations as this reduces the number of unique evaluation points.

First, observe that all invocations to Hadamard can be shuffled around until they can all be run simultaneously – none of the inputs to any of the Hadamard protocols depend on the outputs of any other. Second, we can concatenate all the Hadamard relations and prove one batched relation

$$\mathbf{a}_0 \|\mathbf{a}_1\| \cdots \|\mathbf{a}_{k-1}\| \circ \mathbf{b}_0 \|\mathbf{b}_1\| \cdots \|\mathbf{b}_{k-1}\| = \mathbf{c}_0 \|\mathbf{c}_1\| \cdots \|\mathbf{c}_{k-1}\| \quad (29)$$

instead of  $k$  individual relations separately. This batching comes with no soundness degradation.

The batched equation can be verified with  $k$  separate InnerProduct protocols that prove the same inner product relations as would be proved without batching – except that  $\alpha$  is now the same everywhere. So neither  $P$  nor any other observer can determine whether  $V$  is verifying Equation 29 or  $k$  separate equations.

**Batch the Inner Product Protocols.** The unrolled SparseClaymore protocol consists of 10 invocations of InnerProduct protocol, and the unrolled DenseClaymore protocol consists of 2. We can replace these InnerProduct protocols with the

**Batched InnerProduct** protocol presented in Protocol 2. To see that this replacement does not affect the soundness, note that the **InnerProduct** subprotocols do not involve any verifier randomness and we can safely postpone them to the end of the **Claymore** protocol. Next, we replace them with a **BatchedInnerProduct**, unifying the degrees by the maximal degree of these polynomials. The negligible soundness degradation of this modification is captured concretely by Lemma 2 of Appendix A.

**Batch the Sparse Vector Protocols.** We also present an alternative version of **SparseBiEval** by batching the two instances of **VandermondeVector** and the two **LagrangeVector** protocols. This optimization eliminates four polynomial oracles at the cost of doubling the polynomial degrees. We present the protocol details and security proofs in Appendix C.2.

**Concatenate Left and Right Wire Vectors.** Instead of sending three witness polynomials  $(f_{wl}(X), f_{wr}(X), f_{wo}(X))$ , the prover can get away with sending only two:  $(f_{wi}(X), f_{wo}(X))$  where  $f_{wi}(X) = f_{wl}(X) + X^n \cdot f_{wr}(X)$ . This concatenation is already implicit in the matrix-vector product subprotocol. The input to the Hadamard subprotocol should be  $\mathbf{x} = ([f_{wi}(X)], [X^n \cdot f_{wi}(X)], [X^n \cdot f_{wo}(X)])$ . The subprotocol then establishes that

$$\begin{pmatrix} \mathbf{w}_l \\ \mathbf{w}_r \\ \mathbf{0}_n \end{pmatrix} \circ \begin{pmatrix} \mathbf{0}_n \\ \mathbf{w}_l \\ \mathbf{w}_r \end{pmatrix} = \begin{pmatrix} \mathbf{0}_n \\ \mathbf{w}_o \\ \mathbf{0}_n \end{pmatrix}, \quad (30)$$

which is clearly equivalent to the original Hadamard relation. With this technique, the polynomials are of degree  $3n - 1$ , and so the soundness error is  $(9n - 3)/(|\mathbb{F}| - 1)$  instead of  $(3n - 3)/(|\mathbb{F}| - 1)$ .

This optimization also preserves zero knowledge. To see this, observe that any distinguisher  $D$  that uses  $f_{wi}(X)$  can be simulated with a distinguisher  $D'$  that uses  $f_{wl}(X)$  and  $f_{wr}(X)$ . As a result, the optimized protocol lacks zero knowledge only if the protocol before applying the optimization also lacks it.

## 6 Zero-Knowledge

The strategy for achieving zero-knowledge consists of appending  $3q$  coefficients to the initial wire vectors  $\mathbf{w}_l$ ,  $\mathbf{w}_r$ , and  $\mathbf{w}_o$  such that each new vector has  $q$  uniformly random coefficients and such that their Hadamard relation remains. The randomizers will make the witness polynomials  $q$ -wise independent, meaning that no distinguisher restricted to at most  $q$  queries will obtain any information about the witness.

It is tricky to define zero-knowledge the context of Polynomial IOPs. The distinguisher  $D$  can always query the received oracles in enough points to interpolate and then extract the witness. The notion is only meaningful when the

<p><b>description:</b> realizes <math>\mathcal{R}_{\text{HPR}}</math></p> <p><b>inputs:</b> <math>\mathbf{i}: (m, n, M)</math> with <math>M \in \mathbb{F}^{m \times n}</math></p> <p style="padding-left: 20px;"><math>\mathbf{x}: \mathbf{x} \in \mathbb{F}^n</math></p> <p style="padding-left: 20px;"><math>\mathbf{w}: (\mathbf{w}_1, \mathbf{w}_r, \mathbf{w}_o) \in \mathbb{F}^n \times \mathbb{F}^n \times \mathbb{F}^n</math></p> <p style="padding-left: 20px;">additional parameters: <math>q</math></p> <p><b>offline preprocessing:</b></p> <div style="border-left: 1px solid black; border-right: 1px solid black; padding: 0 10px;"> <p><math>\mathbf{I}</math> runs <code>Claymore.I</code> on <math>\mathbf{i}^{(1)} = (m, n + 3q, M' =</math>  <math>(M_{[1,0:(n+1)]}   0_{m \times 3q}   M_{[1:(n+1):(2n+1)]}   0_{m \times 3q}   M_{[1:(2n+1):(3n+1)]}   0_{m \times 3q}))</math></p> </div> <p><b>online phase:</b></p> <div style="border-left: 1px solid black; border-right: 1px solid black; padding: 0 10px;"> <p><i>// compute witness polynomial with randomizers</i></p> <p><math>\mathbf{P}</math> samples <math>\mathbf{r}_{[0:q]}^{(l)}, \mathbf{r}_{[q:2q]}^{(r)}, \mathbf{r}_{[2q:3q]}^{(o)} \xleftarrow{\\$} \mathbb{F}^q</math> and sets <math>\mathbf{r}_{[q:2q]}^{(l)} = \mathbf{r}_{[0:q]}^{(r)} = 0_{q \times 1}</math>,  <math>\mathbf{r}_{[0:2q]}^{(o)} = 0_{2q \times 1}</math>, <math>\mathbf{r}_{[2q:3q]}^{(l)} = 1_{q \times 1}</math>, and <math>\mathbf{r}_{[2q:3q]}^{(r)} = \mathbf{r}_{[4:6]}^{(o)}</math> // <math>\mathbf{r}^{(l)} \circ \mathbf{r}^{(r)} = \mathbf{r}^{(o)}</math></p> <p><math>\mathbf{P}</math> and <math>\mathbf{V}</math> run <code>Claymore</code> with <math>\mathbf{i}^{(1)} = (m, n + 3q, M')</math>, <math>\mathbf{x}^{(1)} = \mathbf{x}</math>,  <math>\mathbf{w}^{(1)} = ((\mathbf{w}_1^\top   \mathbf{r}^{(l)\top}), (\mathbf{w}_r^\top   \mathbf{r}^{(r)\top}), (\mathbf{w}_o^\top   \mathbf{r}^{(o)\top}))</math></p> </div>
--

**Protocol 7: ZKClaymore**

number of queries bounded by some parameter. We furthermore restrict the distinguisher's queries to be distributed identically to that of an honest verifier; this restriction therefore corresponds to *honest-verifier* zero knowledge. As a result, we are not concerned with finding a complete description of the polynomials that make up the transcript. Instead, we are only concerned with the verifier's view of the transcript. This view corresponds to the list of queries and responses to the various oracles.

**Theorem 7.** *When  $q \geq 2$ , the Polynomial IOP ZKClaymore of protocol 7 has statistical honest-verifier zero-knowledge if all the InnerProduct subprotocols are replaced by a single invocation of BatchedInnerProduct. Concretely, the statistical distance between the verifier's view of authentic transcript versus the verifier's view of simulated transcript is bounded by  $\frac{3}{|\mathbb{F}|-1}$ , which is negligible in the field size.*

*Proof.* We show how  $\mathbf{S}$  produces the verifier view for  $(\mathbf{i}, \mathbf{x})$  without knowledge of  $\mathbf{w}$ . In the process, we establish that this view is indistinguishable from that of an authentic protocol execution.

The protocol ZKClaymore consists of an invocation to Hadamard protocol and an invocation to either the dense or sparse variant of MVP. Note that both protocols DenseMVP and SparseMVP consists of:

1. a query to  $f_x(X)$  at uniformly random  $\alpha \xleftarrow{\$} \mathbb{F} \setminus \{0\}$ ;
2. a protocol invocation (ModReduce in DenseMVP, or SparseBiEval in SparseMVP) with inputs that are independent of  $\mathbf{w}_1, \mathbf{w}_r, \mathbf{w}_o$ ;
3. an invocation of InnerProduct on input  $f_{1w}(X)$  and another polynomial  $r(X)$  in DenseMVP or  $f_{\alpha^\top M}(X)$  in SparseMVP), denoted by  $f_t(X)$  hereafter, that is also independent of  $\mathbf{w}_1, \mathbf{w}_r, \mathbf{w}_o$ .

Since  $S$  knows  $M$  and  $\mathbf{x}$ ,  $S$  can compute all polynomials that do not depend on witnesses honestly, *i.e.*, as the honest  $P$  would. We therefore restrict attention to polynomials that depend on the witness.

What remains is to demonstrate how to sample the verifier view for `InnerProduct` on input  $f_{1w}(X)$  and  $f_t(X)$ , and for `Hadamard` on input  $f_{wl}(X)$ ,  $f_{wr}(X)$  and  $f_{wo}(X)$ . These two subprotocols contribute two polynomial pairs each to the `BatchedInnerProduct` protocol. It suffices to sample the verifier view for the `BatchedInnerProduct` protocol just for these two polynomial pairs, because the remaining pairs are independent of the witness.

This verifier view consists of several elements, namely:

1. Uniformly random  $z, \alpha^*$  (the symbol  $\alpha^*$  is used for batching the various inner product relations into one)
2.  $y_h = \bar{h}(z), y_h^* = \bar{h}(\gamma \cdot z)$
3. The verifier view contributed by the `InnerProduct` protocol in MVP:
  - (a)  $y_{l_1} = f_{wl}(z)$
  - (b)  $y_{r_1} = f_{wr}(z)$
  - (c)  $y_{o_1} = f_{wo}(z)$
  - (d)  $y_t = f_t(z^{-1})$  ( $S$  samples this one honestly)
4. The verifier view contributed by the top-level `Hadamard` protocol:
  - (a) Uniformly random  $\beta$
  - (b)  $y_{o_2} = f_{wo}(\beta)$
  - (c)  $y_{l_2} = f_{wl}(\beta z)$
  - (d)  $y_{r_2} = f_{wr}(z^{-1})$

In the verifier view of an honest run, the above values satisfy:

$$\begin{aligned} y_h \cdot \gamma^{3n+3q} - y_h^* = & (y_{1w} \cdot y_t - f_x(\alpha)) \cdot z^{3n+3q} \\ & + \alpha^* \cdot ((y_{l_2} \cdot y_{r_2} - y_{o_2}) \cdot z^{n+q-1}) \\ & + \alpha^{*2} \cdot (\dots) , \end{aligned} \quad (31)$$

where the ellipses omit terms that are independent of the witness and thus already known to  $S$ , and where  $y_{1w} = 1 + z^{-1}y_{l_1} + z^{-n-q-1}y_{r_1} + z^{-2n-2q-1}y_{o_1}$ .

$S$  samples uniformly random  $\alpha^*, z, \beta \xleftarrow{\$} \mathbb{F} \setminus \{0\}$  and computes  $y_t$  honestly.

Consider the matrices

$$Z_l = \begin{pmatrix} 1 & z & z^2 & \dots & z^{n+3q-1} \\ 1 & \beta z & (\beta z)^2 & \dots & (\beta z)^{n+3q-1} \end{pmatrix} \quad (32)$$

$$Z_r = \begin{pmatrix} 1 & z & z^2 & \dots & z^{n+3q-1} \\ 1 & z^{-1} & z^{-2} & \dots & z^{-n-3q+1} \end{pmatrix} \quad (33)$$

$$Z_o = \begin{pmatrix} 1 & z & z^2 & \dots & z^{n+3q-1} \\ 1 & \beta & \beta^2 & \dots & \beta^{n+3q-1} \end{pmatrix} \quad (34)$$

which satisfy  $(y_{l_1}, y_{l_2})^\top = Z_l (\mathbf{w}_l^\top | \mathbf{r}^{(l)\top})^\top$ ,  $(y_{r_1}, y_{r_2})^\top = Z_r (\mathbf{w}_r^\top | \mathbf{r}^{(r)\top})^\top$ , and  $(y_{o_1}, y_{o_2})^\top = Z_o (\mathbf{w}_o^\top | \mathbf{r}^{(o)\top})^\top$ . Capture the relation between polynomials' values

and randomizers into a single equation:

$$\begin{pmatrix} y_{l_1} \\ y_{l_2} \\ y_{r_1} \\ y_{r_2} \\ y_{o_1} \\ y_{o_2} \end{pmatrix} = \begin{pmatrix} z^n & z^{n+1} & 0 & 0 & 0 & 0 \\ (\beta z)^n & (\beta z)^{n+1} & 0 & 0 & 0 & 0 \\ 0 & 0 & z^{n+q} & z^{n+q+1} & z^{n+2q} & z^{n+2q+1} \\ 0 & 0 & z^{-n-q} & z^{-n-q-1} & z^{-n-2q} & z^{-n-2q-1} \\ 0 & 0 & 0 & 0 & z^{n+2q} & z^{n+2q+1} \\ 0 & 0 & 0 & 0 & \beta^{n+2q} & \beta^{n+2q+1} \end{pmatrix} \begin{pmatrix} \mathbf{r}_{[0]}^{(l)} \\ \mathbf{r}_{[1]}^{(l)} \\ \mathbf{r}_{[q]}^{(r)} \\ \mathbf{r}_{[q+1]}^{(r)} \\ \mathbf{r}_{[2q]}^{(o)} \\ \mathbf{r}_{[2q+1]}^{(o)} \end{pmatrix} + \mathbf{c} , \quad (35)$$

where  $\mathbf{c} \in \mathbb{F}^6$  is a constant vector independent of the randomizers.

Let  $Z_l^*, Z_r^*, Z_o^*$  be the  $2 \times 2$  submatrices of  $Z_l, Z_r, Z_o$  whose columns are multiplied by  $\mathbf{r}_{[0:2]}^{(l)}, \mathbf{r}_{[q:q+2]}^{(r)}$  and  $\mathbf{r}_{[2q:2q+2]}^{(o)}$ , respectively; or equivalently, the  $2 \times 2$  submatrices on the diagonal of Equation 35. If these  $2 \times 2$  submatrices are all invertible, then  $(y_{l_1}, y_{l_2}, y_{r_1}, y_{r_2}, y_{o_1}, y_{o_2})$  are uniform because the randomizers  $\mathbf{r}_{[0:q]}^{(l)}, \mathbf{r}_{[q:2q]}^{(r)}, \mathbf{r}_{[2q:3q]}^{(o)}$  are. So  $S$  samples  $(y_{l_1}, y_{l_2}, y_{r_1}, y_{r_2}, y_{o_1}, y_{o_2})$  uniformly at random from  $\mathbb{F}^6$ .

The  $2 \times 2$  submatrices of  $Z_l^*, Z_r^*, Z_o^*$  are not all invertible if  $z = 1$ , if  $\beta = 1$ , or if  $\beta = z$ . The probability of this event is  $3/(|\mathbb{F}| - 1)$ .

Since we can solve for one of  $y_h, y_h^*$  given the other, we only need to show that  $y_h$  is uniformly random over  $\mathbb{F}$ . This is where the convenient arbitrary coefficient  $\bar{h}_d$  of  $\bar{h}(X)$  comes into play. This coefficient is chosen uniformly at random, and so  $S$  samples  $y_h \xleftarrow{\$} \mathbb{F}$ . Lastly,  $S$  solves Eqn. (31) for  $y_h^*$ .

To complete the argument, except with a negligible failure probability corresponding to the  $2 \times 2$  submatrices  $Z_l^*, Z_r^*, Z_o^*$  being singular,  $S$  samples a verifier view from a distribution that is identical to the distribution of verifier views of an authentic protocol execution. The distinguishing advantage of any distinguisher  $D$  is bounded by the  $S$ 's failure probability, which is  $3/(|\mathbb{F}| - 1)$ . This number also bounds the statistical distance in distributions of the view of the verifier of authentic transcripts versus simulated transcripts.  $\square$

## 7 Comparison

### 7.1 Abstract Comparison

We compare both variants of *Claymore* to some other Polynomial IOPs from the literature, namely *Sonic*, *PLONK*, *Marlin*, and *Aurora*. Of these Polynomial IOPs, the first three give rise to SNARKs after cryptographic compilation. In contrast, *Aurora* gives rise to a proof system generating short proofs but whose verifier complexity is linear in the size of the witness. Importantly, *Claymore* is comparable to both types of proof system: with preprocessing, it gives rise to a SNARK; when preprocessing is omitted, the proofs remain short at the expense of linear verifier complexity.

Table 1 contains an overview of the comparison. It considers the following key performance indicators for Polynomial IOPs.

- The number of polynomials sent by  $I$  during the offline preprocessing phase. This number determines the size of the universal or structured reference strings. While this number contributes to the complexity of  $I$ , this complexity is generally speaking not a make or break factor.
- The number of polynomials sent by  $P$  during the online proving phase. This number contributes to the size of the proof and to the complexity of both  $P$  and  $V$ .
- The number of evaluations. This number contributes to the size of the proof, as well as indirectly to the complexity of  $P$  and  $V$ .
- The number of distinct points for evaluation. Some cryptographic compilers (*e.g.*, [6]) enable the merger of two polynomial evaluations provided that they are being evaluated in the same point. This number limits the number of times this optimization can be applied.
- The maximum degree of all polynomials. This number contributes to indexer and prover complexity in two ways. First, before cryptographic compilation,  $I$  and  $P$  operate on polynomials of this degree and their complexity is affected accordingly. The exception is if the polynomials are sparse, or otherwise exhibit a structure that enables fast computation. Second, some cryptographic compilers induce overheads that are superlinear in this degree.

**Table 1.** Comparison between **Claymore** and other Polynomial IOPs from the literature, with respect to key performance indicators.

	# polynomials offline / online	# evaluations	# distinct points	max. degree
Sonic [11]	$12M/3M + 7$	$11M + 3$	$9M + 2$	$O(n)$
PLONK [8]	8 / 6	7	2	$12(n + a)$
Marlin [7]	9 / 12	18	3	$6k + 6$
Aurora [4]	- / 7	8	2	$\max(m, n)$
DenseClaymore	1 / 4	10	6	$m(3n + 1) - 1$
SparseClaymore	8 / 10	30	10	$3K - 1$

For **Sonic**,  $n$  refers to the number of multiplication gates and the degree of the largest polynomial is  $7n$ . However, due to their technique for simulating bivariate polynomials, the addition gates have fan-in bounded by a parameter  $M$ , corresponding to having at most  $M$  nonzero elements in every row of the linear transform matrix. The same bound applies to the number of nonzero elements in every column. As a result of converting the original circuit into one with this property, a number of multiplication gates may have to be added, thus explaining the Landau notation.

For **PLONK**,  $n$  refers to the number of multiplication gates and  $a$  refers to the number of addition gates, all of which have fan-in 2. We note that there is a variant of **PLONK** with larger proofs and smaller prover time, which is not shown in the table.

**Aurora** does not have a preprocessing phase and as a result the verifier’s complexity is linear in the number of nonzero elements in the matrices  $A, B, C$  from the R1CS tuple. **Marlin** uses the same mechanics but uses preprocessing to shrink the verifier’s workload for the matrix multiplication; this technique requires 9 polynomials in the uniform or structured reference string (3 per matrix) and a few more in the online protocol. The parameter  $k$  denotes the largest number of nonzero elements of  $\{A, B, C\}$ .

For **Claymore** the linear transform is either represented densely as an  $m \times n$  matrix, or sparsely as a list of  $K$  nonzero coefficients in this matrix.

**Marlin**, **PLONK**, and **Aurora** work in the Reed-Solomon codeword basis and crucially rely on the structure of the field or of its multiplicative group. In contrast, **Sonic** and **Claymore** work for any field.

## 7.2 Concrete Comparison

To make the comparison more concrete, we compile the various Polynomial IOPs into concrete SNARKs with various compilers. In the following we use  $P$  to denote the number of polynomials,  $Q$  for the number of queries,  $U$  for the number of unique points,  $|\mathbb{F}|$  for the size of a field element, and  $|\mathbb{G}|$  for the size of a group element.

For the benchmark computation we choose the following: to prove the membership of an element in a set by verifying the Merkle tree authentication path in zero knowledge. The set holds 1024 elements and its Merkle root is known, as is the member element. The witness consists of the element’s position in the tree, and the authentication path. The Merkle tree is constructed using the zero-knowledge-proof-friendly Rescue-Prime hash function [12] with  $m = 3$  state elements, rate equal to  $r = 2$ , over a prime field with  $p > 2^{256}$  elements, with  $N = 18$  rounds, targeting a security level of  $\lambda = 128$  bits against collisions.

After arithmetization, this computation can be cast into one of three constraint systems for arithmetic circuits.

- The computation can be represented as a **PLONK**-relation, in which case there are  $a = 13021$  additions,  $m = 17101$  multiplications, and 31234 distinct wires in total.
- The computation can be represented as a Hadamard Product Relation (HPR) with  $n = 4631$  multiplications and  $m = 7571$  linear relations with arbitrary fan-in. After coercing the constraint system to one whose matrix has at most  $M = 3$  nonzero elements on every row and on every column, there are  $n = 23850$  multiplications and  $m = 26790$  linear relations.
- The computation can be represented as a Rank-1 Constraint Satisfaction System (R1CS) with 12202 rank-1 constraints and no more than  $k = 38694$  nonzero elements in  $\{A, B, C\}$ .

The various cryptographic compilers differ in how they simulate evaluation queries. A polynomial oracle  $[f(X)]$  is represented by a cryptographic commitment. When  $V$  queries it in  $z$  and obtains the response  $y$ , one option is to simply

run the polynomial commitment’s evaluation protocol to prove that  $f(z) = y$ . Another option is for  $V$  to use the commitment to  $f(X)$  to derive the commitment to  $\frac{f(X)-y}{X-z}$ , and for  $P$  to prove that this polynomial has an appropriately bounded degree. The point is that degree bound checks might be simpler to combine than evaluation queries, depending on the nature of the polynomial commitment.

- KZG polynomial commitments. A commitment to a polynomial  $f(X)$  is a single group element. When  $V$  queries its value in  $z$ ,  $P$  responds with its value  $y = f(z)$  along with another group element, a commitment to  $\frac{f(X)-y}{X-z}$ . Note that the zeroifier  $X-z$  must divide  $f(X)-f(z)$ . One pairing evaluation allows  $V$  to verify the correct relation between the received commitments. Evaluation queries in the same point  $z$  but to different polynomials can be batched using random weights supplied by  $V$ . However, the scheme does not support batching evaluations in distinct points, at least in terms of communication cost. To establish the proper degree bounds,  $P$  must supply a commitment to a *degree bound check polynomial*  $F(X) = \sum_i \omega^{2i} \cdot f_i(X) + \omega^{2i+1} \cdot X^{d-\delta(i)} \cdot f_i(X)$  where the sum ranges over all prior polynomials  $f_i(X)$ , where  $\delta(i)$  is its degree bound, and where  $\omega$  is a random challenge supplied by  $V$ . One more batch-evaluation is necessary to establish that the commitment to  $F(X)$  is well formed. So the total size of a Polynomial IOP compiled with this KZG-based compiler is

$$(P + 1) \times |\mathbb{G}| + (U + 1) \times (|\mathbb{F}| + |\mathbb{G}|) . \quad (36)$$

Using the BLS-384 pairing group, we can represent group elements with  $|\mathbb{G}| = 385$  bits and scalar field elements with  $|\mathbb{F}| = 256$  bits.

- DARK polynomial commitments. Evaluation queries to the same polynomial in distinct points can be batched, and evaluation queries to distinct polynomials in the same point can be batched. Evaluation queries to distinct polynomials in distinct points cannot be batched.<sup>3</sup> We choose to batch all queries into one evaluation proof for each polynomial, as opposed to batching all polynomials for each distinct point. Every evaluation protocol consists of at most  $\lfloor \log_2 d \rfloor$  rounds where  $d$  is the global degree bound. In every round except the last,  $P$  sends two group elements and at most  $U$  field elements, where  $U$  is the number of unique evaluation points. In the last round,  $P$  sends an integer of  $\lfloor \log_2 d \rfloor \times \lambda + |\mathbb{F}|$  bits. So the total size of a Polynomial

---

<sup>3</sup> There is a natural method for this type of batching that relies on dividing out the zeroifier. In response to queries  $z_1$  to  $[f(X)]$  and  $z_2$  to  $[g(X)]$ ,  $P$  supplies  $y_1 = f(z_1), y_2 = g(z_2)$  and commitments to the polynomials  $\frac{f(X)-y_1}{X-z_1}$  and  $\frac{g(X)-y_2}{X-z_2}$ . At this point,  $V$  can verify that the new commitments are correctly related to the old, and he can use all commitments to derive a new commitment to “zero” polynomial. Precisely speaking, this commitment is to an integer polynomial whose coefficients are multiples of  $p$ . Then *one* evaluation protocol suffices to establish that  $y_1 = f(z_1)$  and  $y_2 = g(z_2)$ . This batching method has not been formally analyzed and such an analysis is out of the scope of this paper.

IOP proof compiled with this DARK strategy is

$$P \times \lfloor \log_2 d \rfloor \times (2|\mathbb{G}| + U|\mathbb{F}|) + \lfloor \log_2 d \rfloor \times \lambda + |\mathbb{F}| . \quad (37)$$

Concretely we use 2000 bit class group elements, so  $|\mathbb{G}| = 2000$ . The field elements are integers modulo  $p$ , a  $2\lambda = 256$  bit prime, so  $|\mathbb{F}| = \lfloor \log_2 p \rfloor = 256$ .

- FRI polynomial commitments. A FRI commitment is a Merkle root of a Reed-Solomon codeword, obtained by evaluating the polynomial in a domain that is  $\rho$  times larger than its degree, where  $\rho$  is known as the expansion factor. The FRI protocol establishes the bounded degree by opening Merkle leafs. As a result, the technique for dividing out the zerofier applies even without supplying a new commitment. The Reed-Solomon codeword of the polynomial  $\frac{f(X)-y}{X-z}$  can be computed given the Reed-Solomon codeword of the polynomial  $f(X)$ . So every query generates one field element. The FRI protocol is run on the degree bound check polynomial  $F(X) = \sum_i \omega^{2i} \cdot f_i(X) + \omega^{2i+1} \cdot X^{d-\delta(i)} \cdot f_i(X)$  where the sum ranges over all polynomials after dividing out the zerofiers. In this expression,  $\delta(i)$  is its updated degree bound, *i.e.*, after dividing out the zerofiers, and  $\omega$  is a random challenge supplied by  $V$ . The evaluations of this degree bound check polynomial on the codeword domain can be computed from evaluations of the constituent polynomials. The FRI protocol consists of  $\lfloor \log_2 d \rfloor$  rounds where  $d$  is the global degree bound. In every round,  $P$  supplies a new Merkle root along with  $s$  field elements and as many authentication paths of length at most  $\lfloor \log_2 \rho \cdot d \rfloor$ . The protocol bounds the polynomial’s degree with soundness error conjecturally approximately  $\rho^{-s}$  when  $\mathbb{F}$  is large enough. So the total size of a Polynomial IOP compiled with FRI is less than

$$P \times 2 \cdot \lambda + Q \times |\mathbb{F}| + \lfloor \log_2 d \rfloor \times s \times (\lfloor \log_2 \rho \cdot d \rfloor \times 2 \cdot \lambda + |\mathbb{F}|) . \quad (38)$$

Concretely, we set  $\rho = 16$ ,  $s = 32$ , and use  $|\mathbb{F}| = 256$ . Note that the hash function with which the FRI Merkle trees are built, must have at least  $2 \cdot \lambda = 256$  bit outputs.

Table 2 summarizes the results. It shows the size of the SNARK in bytes obtained from the given Polynomial IOP and compiled with the given compiler. The computation whose integrity is proved, is the Merkle tree membership relation described above. The source code for reproducing these numbers is available at <https://github.com/aszepieniec/claymore-benchmark>.

## 8 Conclusion

The protocols proposed in this paper challenge the notion that the Reed-Solomon codeword basis is the appropriate basis for representing objects from the arithmetic constraint system in a Polynomial IOP. Instead, the monomial coefficient basis provides a natural and intuitive representation for these objects. In this basis, the native operations on polynomials are identifiable with the matrix operations in the arithmetic constraint system. Moreover, this basis does not impose

**Table 2.** Comparison of concrete SNARK size.

	KZG	DARK	FRI
Sonic $P = 16, Q = 36, U = 20, M = 3, d = 166950$	3 222	388 720	384 640
PLONK $P = 6, Q = 7, U = 2, d = 361464$	578	61 232	424 352
Marlin $P = 12, Q = 18, U = 3, d = 232170$	947	121 888	383 936
DenseClaymore $P = 4, Q = 10, U = 6, d = 105327751$	802	72 416	825 792
SparseClaymore $P = 10, Q = 30, U = 10, d = 158951$	1 411	139 704	384 256

any restrictions on the structure of the field. The resulting Polynomial IOP for arithmetic constraint systems outperforms similar constructions based on the Reed-Solomon codeword basis, at least as far as the number of polynomials in the transcript is concerned.

The modular approach followed in this paper admits a piece by piece presentation and analysis that benefits simplicity and accessibility. Nevertheless, some optimizations violate the boundaries implicit in this modular structure. These optimizations are of independent interest as they may also apply elsewhere; perhaps they have straightforward analogues in the Reed-Solomon codeword domain. Some of the more important optimizations are summarized as follows.

- In some cases it is possible to eliminate polynomials. In particular, when a polynomial is queried exactly twice and is involved in exactly two polynomial identities. By moving this polynomial to the left hand side, and equating the right hand sides, the polynomial identities are merged and the polynomial in question is eliminated, all without impacting soundness.
- After unrolling a Polynomial IOP, `InnerProduct` subprotocols appear in great numbers and many places. They can all be batched. In addition to saving polynomials, this batching facilitates a simpler and more direct proof of zero-knowledge that would not be possible otherwise.
- Batching may apply in more places still. For instance, the sparse MVP procedure benefits from two invocations of a subprotocol that establishes that a given vector is a Lagrange vector, and two more that another vector is a Vandermonde vector. The Lagrange and Vandermonde vector protocols can be merged in order to save polynomials. In fact, this merger extends to multivariate polynomials in more than two variables.
- Concatenating the vectors of left and right wires saves one polynomial, but only if the `Hadamard` subprotocol can be made to work with the concatenated vector. In particular, this adaptation requires shifting the protocol’s second and third arguments. Performing this shift in either basis is possible, but this shift highlights an interesting difference. In the Reed-Solomon codeword basis, the query to the polynomial oracle is multiplied by a constant factor; in

the monomial coefficient basis, however, it is the response that is multiplied by a factor that depends on the query.

One of the questions that led to the present line of research was to find the Polynomial IOP with the smallest possible number of polynomials in the transcript. In this respect, we report a mitigated success: **DenseClaymore** has only four polynomials, down 33% from runner-up **PLONK**; and when restricting to polynomials whose degrees grow linearly with the size of the circuit, then **SparseClaymore** achieves two polynomials less than **Marlin**, the only competitor that also admits arbitrary fan-in for linear gates.

However, the concrete comparison shows that we were optimizing the wrong metric. Contrary to our expectation, the number of polynomials in the transcript is *not* the dominant factor of proof size. Indeed, this comparison highlights the importance of balancing a Polynomial IOP's number of polynomials against its other parameters.

**ACKNOWLEDGEMENTS.** Both authors are supported by the Nervos Foundation. This work was supported in part by the National Key Research and Development Project 2020YFA0712300.

## References

1. Ames, S., Hazay, C., Ishai, Y., Venkatasubramanian, M.: Liger: Lightweight Sub-linear Arguments Without a Trusted Setup. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017. pp. 2087–2104. ACM (2017), <https://doi.org/10.1145/3133956.3134104>
2. Bellare, M., Goldreich, O.: On defining proofs of knowledge. In: Brickell, E.F. (ed.) CRYPTO 1992s. Lecture Notes in Computer Science, vol. 740, pp. 390–420. Springer (1992), [https://doi.org/10.1007/3-540-48071-4\\_28](https://doi.org/10.1007/3-540-48071-4_28)
3. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable zero knowledge with no trusted setup. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part III. Lecture Notes in Computer Science, vol. 11694, pp. 701–732. Springer (2019). [https://doi.org/10.1007/978-3-030-26954-8\\_23](https://doi.org/10.1007/978-3-030-26954-8_23)
4. Ben-Sasson, E., Chiesa, A., Riabzev, M., Spooner, N., Virza, M., Ward, N.P.: Aurora: Transparent Succinct Arguments for R1CS. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. Lecture Notes in Computer Science, vol. 11476, pp. 103–128. Springer (2019), [https://doi.org/10.1007/978-3-030-17653-2\\_4](https://doi.org/10.1007/978-3-030-17653-2_4)
5. Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient Zero-Knowledge Arguments for Arithmetic Circuits in the Discrete Log Setting. In: Fischlin, M., Coron, J. (eds.) EUROCRYPT 2016, Part II. Lecture Notes in Computer Science, vol. 9666, pp. 327–357. Springer (2016), [https://doi.org/10.1007/978-3-662-49896-5\\_12](https://doi.org/10.1007/978-3-662-49896-5_12)
6. Bünz, B., Fisch, B., Szepieniec, A.: Transparent SNARKs from DARK Compilers. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020 Part I. Lecture Notes in Computer Science, vol. 12105, pp. 677–706. Springer (2020), <https://eprint.iacr.org/2019/1229.pdf>

7. Chiesa, A., Hu, Y., Maller, M., Mishra, P., Vesely, N., Ward, N.P.: Marlin: Pre-processing zkSNARKs with universal and updatable SRS. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020 Part I. Lecture Notes in Computer Science, vol. 12105, pp. 738–768. Springer (2020), <https://eprint.iacr.org/2019/1047.pdf>
8. Gabizon, A., Williamson, Z.J., Ciobotaru, O.: PLONK: Permutations over Lagrange-bases for Oecumenical Noninteractive arguments of Knowledge. IACR Cryptology ePrint Archive **2019**, 953 (2019), <https://eprint.iacr.org/2019/953>
9. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof-systems (extended abstract). In: Sedgewick, R. (ed.) ACM STOC. pp. 291–304. ACM (1985). <https://doi.org/10.1145/22145.22178>
10. Kate, A., Zaverucha, G.M., Goldberg, I.: Constant-size commitments to polynomials and their applications. In: Abe, M. (ed.) ASIACRYPT 2010. Lecture Notes in Computer Science, vol. 6477, pp. 177–194. Springer (2010), [https://doi.org/10.1007/978-3-642-17373-8\\_11](https://doi.org/10.1007/978-3-642-17373-8_11)
11. Maller, M., Bowe, S., Kohlweiss, M., Meiklejohn, S.: Sonic: Zero-knowledge snarks from linear-size universal and updatable structured reference strings. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) ACM CCS 2019. pp. 2111–2128. ACM (2019), <https://eprint.iacr.org/2019/099.pdf>
12. Szepieniec, A., Ashur, T., Dhooghe, S.: Rescue-Prime: a standard specification (sok). IACR Cryptology ePrint Archive **2020**, 1143 (2020), <https://eprint.iacr.org/2020/1143>
13. Szepieniec, A., Zhang, Y.: Polynomial IOPs for Linear Algebra Relations. IACR Cryptol. ePrint Arch. p. 1022 (2020), <https://eprint.iacr.org/2020/1022>