

SNARK Lower Bounds via Communication Complexity*

Rishabh Bhadauria^{†1}, Alexander R. Block², Prantar Ghosh³, and Justin Thaler⁴

¹rishabh.bhadauria@gmail.com

²University of Illinois at Chicago. alexander.r.block@gmail.com

³Tennessee Tech University. pghosh@tnitech.edu

⁴a16z crypto research and Georgetown University. justin.r.thaler@gmail.com

December 18, 2025

Abstract

We initiate the study of lower bounding the verification time of Succinct Non-interactive ARguments of Knowledge (SNARKs) built in the Polynomial Interactive Oracle Proof + Polynomial Commitment Scheme paradigm. The verification time of these SNARKs is generally dominated by the polynomial commitment scheme, and so we want to understand if polynomial commitment schemes admit lower bounds on the verification time. By recognizing that polynomial commitment schemes are also often built by applying cryptography to some information-theoretic core protocol, we seek to separate this core from the cryptography in a way that meaningfully captures the verification time required by the polynomial commitment scheme verifier.

We provide strong evidence that several polynomial commitment schemes have (nearly) *optimal* verifier times. Our evidence comes from connecting polynomial commitment schemes to certain information-theoretic protocols known as *communication protocols* from the field of communication complexity, a link which we believe to be of independent interest. Through this lens, we model the verifier work in the cryptographic protocols as information (i.e., number of bits) exchanged between parties in the communication protocols, allowing us to leverage lower bounds from communication complexity. These lower bounds give strong evidence that the verifier time in these polynomial commitment schemes must be at least the number of bits exchanged in the communication protocol.

We extract the communication protocol cores of three polynomial commitment schemes and lower bound the bits exchanged in these cores. The lower bounds we obtain match (up to poly-logarithmic factors) the best-known (asymptotic) verification times of the polynomial commitment schemes we examine in this work. Specifically, we show that for univariate/multilinear polynomials of size $N = 2^n$:

- the communication core of Hyrax PCS (Wahby et al., S&P 2016) requires $\Omega(\sqrt{N})$ bits to be exchanged;
- the communication core of Bulletproofs PCS (Bootle et al., EUROCRYPT 2016; Bünz et al., S&P 2018) requires $\Omega(N)$ bits to be exchanged; and
- the communication core of Dory PCS (Lee, TCC 2021) requires $\Omega(\log(N))$ bits to be exchanged.

Our results strongly suggest a negative answer to a longstanding open question on whether the Bulletproofs verifier can run in sublinear time.

*© IACR 2026. This article is the full version of the article submitted by the authors to the IACR on 25 September, 2025, available at https://doi.org/10.1007/978-3-032-12287-2_15.

[†]Work partially done while at Georgetown University.

Contents

1	Introduction	2
1.1	Our Contributions	2
1.1.1	Evidence vs. Provable Optimality of Verification Times	3
1.1.2	Polynomial Commitment Schemes from any Communication Protocol?	3
1.1.3	Verifiers with PCP Proof String Access?	4
1.2	Technical Overview	4
1.2.1	Communication Complexity and Merlin-Arthur Communication	4
1.2.2	Review: Polynomial Commitment Schemes	5
1.2.3	Hyrax: Review and Information-Theoretic Core	5
1.2.4	Bulletproofs: Review and Information-Theoretic Core	8
1.2.5	Dory: Review and Information-Theoretic Core	11
1.3	Related Work	14
2	Preliminaries	15
3	Communication Complexity	18
3.1	Online Merlin-Arthur Communication Protocols	18
3.1.1	OMA Lower Bounds	19
3.2	Simultaneous Messaging with a Helper Communication Protocols	19
3.3	SMH Model Lower Bounds	21
4	Hyrax: Information-Theoretic Core and Evidence of Optimal Verifier Time	22
4.1	Information Theoretic Core of Hyrax-VMV	23
4.1.1	Hyrax-VMV Lower Bound	24
4.2	Evidence that Hyrax-PCS has Optimal Verifier Time	26
4.3	Hyrax-VMV as an SMH Protocol	27
5	Bulletproofs: Information-Theoretic Core and Evidence of Optimal Verifier Time	27
5.1	Bulletproofs is Sum-Check (Revisited)	28
5.1.1	Bulletproofs as a Standard Sum-Check Protocol	28
5.2	Information-Theoretic Core of Bulletproofs	29
5.2.1	Bulletproofs as a SMH Protocol	30
5.3	Evidence that Bulletproofs-PCS has Optimal Verifier Time	30
5.4	Why the OMA Model does not Suffice for Bulletproofs	31
5.5	Rounds vs. Communication Tradeoffs for Bulletproofs	33
6	Dory: Information-Theoretic Core and Evidence of Optimal Verifier Time	34
6.1	Proto-Dory is (a variant of) Sum-Check	35
6.1.1	Proto-Dory as a SMH Communication Protocol	37
6.2	Full Dory as Sum-Check	40
6.2.1	Full Dory as a SMH Protocol	40
6.3	Evidence that Dory-PCS has Optimal Verifier Time	43
7	Conclusions and Open Problems	45

1 Introduction

Succinct Non-interactive ARguments of Knowledge (SNARKs) allow a computationally powerful prover to convince a computationally weak verifier of the truth of some (e.g., **NP**) statement via a (non-interactive) proof of sublinear size. It is well-known that modern efficient SNARK constructions can be divided into two components: an information-theoretically secure component (typically a polynomial interactive oracle proof (PIOP) [BFS20]), and a cryptographic protocol (typically a polynomial commitment scheme (PCS) [KZG10]). These two interactive components are combined to obtain a succinct interactive argument, which is then rendered non-interactive via the Fiat-Shamir transformation [FS86].

Examining SNARKs built from this combination of PIOPs and PCSs, what is less recognized is that the PCS itself can typically be divided into two of the same components: an information-theoretic one and a cryptographic one. Notable exceptions (i.e., ones which put this display at the forefront of their design) include PCSs based on (or inspired by) the FRI protocol [BBH⁺18, ZCF24, ACF⁺24], and those based on Ligero/Brakedown [AHI⁺23, BFH⁺20, BBH⁺22, GLS⁺23, BFK⁺24]. With these notable exceptions, the vast majority of other PCSs in the literature (e.g., discrete-log based schemes [BCC⁺16, BBB⁺18, WTS⁺18, BHR⁺20], pairing-based schemes [KZG10, XZZ⁺19, Lee21, BHV⁺23], groups of unknown order based schemes [BFS20, BHR⁺21], and lattice-based schemes [ACL⁺22, BCF⁺22, AFL⁺24, CMN⁺24, NS24, HSS24, FMN24]) do not highlight or examine the PCS from this same perspective.

1.1 Our Contributions

In this work, we attempt to elucidate this viewpoint by isolating the information-theoretic cores of three well-studied polynomial commitment schemes. This perspective and goal bring two benefits. First is expository: understanding the information-theoretic core of these protocols gives us insights into what is “really happening” in these commitment schemes; i.e., what is the cryptography guaranteeing? Second, this view allows us to provide evidence that certain commitment schemes are *optimal*, in terms of *verification time*. Specifically, we can show *unconditionally* that (certain components of) the information-theoretic cores of these commitment schemes cannot be improved (up to constant factors), so building “better” commitment schemes to circumvent our results will have to deviate significantly from the current design paradigms.

We examine three polynomial commitment schemes through this lens: Hyrax [WTS⁺18], Bulletproofs [BCC⁺16, BBB⁺18], and Dory [Lee21]. We separate the information-theoretic cores from the cryptography of these protocols, allowing us to understand what the cryptography is actually achieving; i.e., which properties are needed in the underlying cores that the cryptography guarantees. Moreover, by isolating these cores, we obtain strong evidence that the verification time of all the aforementioned schemes is (nearly) optimal. We do this by translating the verifier work in the polynomial commitment schemes to the *number of bits exchanged* between parties in the information-theoretic cores, then obtain *lower bounds* on this number of bits.

The information-theoretic cores we extract from these polynomial commitment schemes are information-theoretic protocols known as *communication protocols* from the field of *communication complexity* (see Section 1.2.1). Communication protocols are information-theoretic protocols with computationally unbounded parties that seek to understand how many bits are required to be exchanged between parties in order to compute some discrete function (with inputs spread across parties). That is, the complexity measure of a communication protocol is explicitly this number of exchanged bits.

Our results frame Hyrax, Bulletproofs, and Dory as communication protocols by stripping away the cryptography and understanding what guarantees the cryptography provides, then translating those guarantees into a suitable communication protocol. This framing allows us to leverage known and new lower bounds for communication protocols; in particular, we show that the communication cores of Hyrax, Bulletproofs,

and Dory require a certain number of bits to be exchanged between parties (otherwise soundness of the core is violated). The lower bounds we obtain nearly match the verification time of the specified schemes, giving strong evidence that the verification times of Hyrax, Bulletproofs, and Dory are *(nearly) optimal*. In other words, to achieve more efficient verification time in these schemes, our results suggest that significant modifications would be needed: any modification that can be captured using the same communication protocols we obtain for Hyrax, Bulletproofs, and Dory is subject to our lower bounds. Informally, we show the following results.

Theorem 1.1 (Informal; see [Corollaries 4.2, 5.2](#) and [6.2](#)). *For $N = 2^n$ and any n -variate multilinear or degree- N univariate polynomial over a finite field \mathbb{F} ,*

1. *The communication protocol core of Hyrax-PCS requires $\Omega(\sqrt{N})$ bits of communication;*
2. *The communication protocol core of Bulletproofs-PCS requires $\Omega(N)$ bits of communication;*
3. *The communication protocol core of Dory-PCS requires $\Omega(\log(N))$ bits of communication.*

Note that each scheme under consideration operates using group-based assumptions, with Hyrax and Bulletproofs utilizing the discrete-logarithm assumption, and Dory additionally employing pairing-friendly groups. In light of this, the verification costs of Hyrax, Bulletproofs, and Dory nearly match the lower bounds of [Theorem 1.1](#), up to polylog $|\mathbb{G}|$ factors, assuming single \mathbb{G}/\mathbb{F} operations (including pairings) require polylog $|\mathbb{G}|$ bit-operations and $|\mathbb{G}| \geq |\mathbb{F}|$ for group \mathbb{G} with scalar field \mathbb{F} .

Notably, (2) in [Theorem 1.1](#) gives evidence towards a resolution of a long-standing open question on whether the Bulletproofs verifier is inherently linear-time. Our results show that linear communication is inherent to the underlying information-theoretic core, and any attempt to improve the Bulletproofs verifier time will require significant deviation from this core; Dory is one such deviation from Bulletproofs, which we highlight in [Section 1.2.5](#). In particular, our results suggest that prior work on obtaining a sublinear-time Bulletproofs verifier (e.g., [\[Lu22, Lu24\]](#), both of which have been withdrawn) was doomed to fail. Indeed, [\[Gab24\]](#) gave an attack against the work presented in [\[Lu24\]](#), and this attack matches (in spirit) the proof of our communication lower bound.

1.1.1 Evidence vs. Provable Optimality of Verification Times

We are very careful in how we state our results: we are giving (what we believe to be) *strong evidence* that the aforementioned protocols have optimal verifier time. Our evidence is strong in the sense that the communication protocols we extract from Hyrax, Bulletproofs, and Dory *nearly match* the communication lower bounds, and these lower bounds hold for *any communication protocol* which computes the same discrete functions as the communication cores of Hyrax, Bulletproofs, and Dory. However, this is not a formal proof, and any such formal proof would require proving a statement of the form “for every polynomial commitment scheme satisfying a set of properties P , there exists a communication protocol which satisfies a set of properties P' ” (here, P' is a function of P , the protocol, etc.), where the properties P, P' capture the efficiency, communication, etc., of these stated protocols. Such a statement seems difficult to even formally write down, let alone prove; nonetheless, we leave this as exciting future work to explore, as it would allow us to derive lower bounds on many more polynomial commitment schemes in a straightforward manner.

1.1.2 Polynomial Commitment Schemes from any Communication Protocol?

One may naturally wonder if we could reverse our results: namely, build a polynomial commitment scheme from any specified communication protocol, possibly subject to certain restrictions. Our results make

no claims of this form, and we strongly believe this to not be possible. It is quite likely one could cook up a communication protocol, even under some possible restrictions (e.g., like ones we introduce for our communication cores), that look nothing like a polynomial commitment scheme (as in, impossible to translate into a meaningful cryptographic protocol). This makes our results and evidence even stronger since communication complexity lower bounds are of the form “for a function f , *any* communication protocol computing f has a communication lower bound,” and the communication cores we derive from the aforementioned polynomial commitment schemes have (nearly) matching communication complexity.

1.1.3 Verifiers with PCP Proof String Access?

One natural question to ask is if our lower bounds can be circumvented by a polynomial commitment scheme verifier with query access to a PCP proof string, since in this case the actual work done by the verifier is much less than reading the proof string, which could potentially circumvent our lower bounds. However, we do not believe this would circumvent our results. This is because any of the homomorphic commitment schemes we model in this work (and for which our lower bounds apply to) can also be used to implement PCP-query access for cryptographic verifiers: e.g., commit to the (multilinear extension of) the PCP proof string, and now queries to the PCP proof string now become evaluations to the committed polynomial. In fact, any polynomial commitment scheme built via PCPs is likely to already have a cryptographic prover commit to the string (e.g., via a Merkle tree), and provide openings to the verifier queries.

1.2 Technical Overview

Proving [Theorem 1.1](#) relies on showing suitable lower bounds on the information-theoretic cores of the three polynomial commitment schemes we examine in this work. The first step is, thus, understanding *what* these information-theoretic cores are. A natural attempt might be to simply strip away the cryptography from these PCSs and obtain some information-theoretic interactive (oracle) proof; however, to the best of our knowledge, we do not know of any lower bounds on the verifier time of interactive (oracle) proof systems, so this path does not seem viable. The only result we are aware of that provide some evidence of lower bounds on the verifier time are due to Bangalore et al. [[BBH⁺22](#)], who show evidence of a $\text{poly}(\lambda)$ upper bound on the *space* required by certain verifiers for interactive arguments based on symmetric-key primitives (e.g., hash functions). They also provide evidence of a proof-size lower bound based on existing techniques, which translates to a lower bound on the verification time. In light of these limited results, we take a different approach and model the information-theoretic cores of these PCSs as communication protocols, allowing us to leverage tools from the rich field of communication complexity to obtain our lower bounds.

1.2.1 Communication Complexity and Merlin-Arthur Communication

The field of communication complexity, originally introduced by Yao [[Yao79](#)], studies the communication costs of computing discrete functions with inputs split between two or more parties. The most relevant communication model to us is (*online*) *Merlin-Arthur Communication* [[BFS86](#), [CCM⁺14](#)]. A communication protocol in this model is parameterized by a function $f: X \times Y \rightarrow Z$, two (randomized) parties Alice and Bob, and an omniscient but untrusted helper Merlin. Alice and Bob, given inputs $x \in X$ and $y \in Y$ respectively, wish to compute $f(x, y)$ and can settle for some (constant) probability of error. Merlin-Arthur communication asks: if Alice and Bob must communicate $R(f)$ bits to compute $f(x, y)$ on their own, can Merlin, who knows (x, y) , convince “Arthur” (defined as Alice and Bob together) that $f(x, y) = z$ while keeping the total communication between *all* parties $o(R(f))$?

Online Merlin-Arthur (OMA) communication modifies the above communication model and allows Bob and Merlin to interact arbitrarily, with the restriction that Bob’s messages do not depend on the input y , and restricts Alice to only send a *single message* to Bob after Bob and Merlin’s interaction (in particular, Alice cannot receive communication from either Bob or Merlin). Notably in these models, Alice and Bob are always honest, trusted parties, and all parties are computationally unbounded.

Why Communication Complexity? Intuitively, communication complexity seeks to measure the amount of *information that must flow* between parties in a protocol in order to compute a function f . As all parties within the protocol are computationally unbounded, the cost of such a protocol is strictly determined by the number of bits sent. Thus, if we can frame (the underlying interactive protocol of) SNARKs as communication protocols, it gives us a tangible connection to SNARK performance in two distinct ways: (1) the SNARK proof is explicitly a transfer of information from the SNARK prover to the SNARK verifier; and (2) the SNARK verifier work can be isolated and modeled as the number of bits the party “Arthur” needs to receive in order to verify this proof (i.e., verify the computation of some function).

Communication complexity, then, is a clean information-theoretic setting to formulate the information-theoretic cores of (certain) cryptographic protocols. As all parties are computationally unbounded, the “work” being done by cryptographic parties can be explicitly measured as the number of bits sent in the underlying communication protocol. This, then, allows us to utilize well-established (online) communication lower bounds for many functions, including inner product and set disjointness [BFS86, KS92, K1a03, AW09], matrix-vector products [KS93, SWY+21, CMT13], and (sparse) index [Ab196, CCG+14, GS24]. In this work, we will use OMA communication (and an extension that we introduce) to cleanly isolate the work being performed by parties in Hyrax [WTS+18], Bulletproofs [BCC+16, BBB+18], and Dory [Lee21], and establish lower bounds via known communication complexity lower bounds.

1.2.2 Review: Polynomial Commitment Schemes

We briefly give a high-level overview of polynomial commitment schemes, originally introduced by Kate, Zaverucha, and Goldberg [KZG10]. A polynomial commitment scheme, or a PCS in short, is a special cryptographic commitment scheme which first allows a committer/prover to commit to a polynomial f as some value C . Then, a receiver/verifier, given C , can request C be opened to a specific evaluation x of the underlying polynomial f via an evaluation protocol, which (informally) proves the statement “ $y = f(x)$ and $f(x)$ is consistent with C .” See Section 2 for the full definition.

All three polynomial commitment schemes we consider in this work utilize group-based hardness assumptions (specifically, discrete-log and pairings). For ease of presentation, we let \mathbb{G} be a finite cyclic group written in additive notation with scalar finite field \mathbb{F} . Throughout this section, we assume that $N = 2^n$. Moreover, we discuss all polynomial commitment schemes as *multilinear* PCSs; that is, as polynomial commitment schemes for multilinear polynomials. For any function $f: \{0, 1\}^n \rightarrow \mathbb{F}$, the *multilinear extension* of f , denoted as $\tilde{f}: \mathbb{F}^n \rightarrow \mathbb{F}$ is the unique polynomial such that $\deg(X_i) \leq 1$ for all $i \in [n]$ and $f(c) = \tilde{f}(c)$ for all $c \in \{0, 1\}^n$. Similarly, for any vector $u \in \mathbb{F}^N$, we let \tilde{u} denote the multilinear extension of u , where we equivalently view u as a function from $\{0, 1\}^n$ to \mathbb{F} in the natural way. We also naturally define extensions of vectors/functions over \mathbb{G} with scalar field \mathbb{F} : for $g \in \mathbb{G}^N$, $\tilde{g}: \mathbb{F}^n \rightarrow \mathbb{G}$ is the multilinear extension of g .

1.2.3 Hyrax: Review and Information-Theoretic Core

The Hyrax PCS [WTS+18] is built from a more general commitment scheme for verifying *vector-matrix-vector* products; we call this scheme *Hyrax-VMV*. In Hyrax-VMV, the prover is given a matrix $A \in \mathbb{F}^{a \times b}$ and

commits to this matrix by computing a Pedersen commitment [Ped91] per row of \mathbf{A} with random generators $\mathbf{g} \xleftarrow{\$} \mathbb{G}^b$; that is, the commitment is the vector $\mathbf{C} = \mathbf{A}\mathbf{g}^\top \in \mathbb{G}^a$. Then, for vectors $\mathbf{L} \in \mathbb{F}^a$ and $\mathbf{R} \in \mathbb{F}^b$ given by the verifier after receiving a commitment \mathbf{C} , Hyrax-VMV takes as input $(\mathbf{C}, \mathbf{L}, \mathbf{R})$ and is a proof of knowledge that the prover knows matrix \mathbf{A} such that $\mathbf{C} = \mathbf{A}\mathbf{g}^\top$ and $y = \mathbf{L}\mathbf{A}\mathbf{R}^\top$. The proof consists of a single message from the prover: a vector \mathbf{L}' (the purported value of $\mathbf{L}\mathbf{A}$) and a value y' (the purported value of $\mathbf{L}\mathbf{A}\mathbf{R}^\top$). Verification proceeds in 3 steps: (1) sample $\alpha \xleftarrow{\$} \mathbb{F}$; (2) check $\alpha \langle \mathbf{L}, \mathbf{C} \rangle \stackrel{?}{=} \alpha \langle \mathbf{L}', \mathbf{g} \rangle$, and (3) check $\alpha y' \stackrel{?}{=} \alpha \langle \mathbf{L}', \mathbf{R} \rangle$. Hyrax-VMV can be used as a PCS for any polynomial f such that evaluation of f at any point can be expressed as a vector-matrix-vector product; examples include multilinear polynomials and univariate polynomials (see Section 4.2 for full details).

Information-Theoretic Core of Hyrax. The above discussion leads us to frame the information-theoretic (IT) core of Hyrax as an online Merlin-Arthur (OMA) communication protocol for verifying (random) vector-matrix-vector products, which consists of two parts: Merlin convince “Arthur” that (1) $\mathbf{L}' = \mathbf{L}\mathbf{A}$, and (2) that $y = \langle \mathbf{L}', \mathbf{R} \rangle$. In particular, vector-matrix-vector products are verified by simultaneously verifying a matrix-vector product (the claim “ $\mathbf{L}' = \mathbf{L}\mathbf{A}$ ”) and an inner product (the claim “ $y = \langle \mathbf{L}\mathbf{A}, \mathbf{R} \rangle$ ”). Somehow, we need to account for the cryptographic commitment \mathbf{C} of Hyrax-VMV in the communication model. Moreover, since parties are computationally unbounded in the communication model, there is no such thing as a cryptographic commitment. Therefore, we first must understand what the commitment \mathbf{C} is actually contributing in Hyrax-VMV.

The commitment $\mathbf{C} = \mathbf{A}\mathbf{g}^\top$ is actually a random fingerprint of \mathbf{A} using public randomness \mathbf{g} , which cryptographic verifier uses to cross-check the proof. Traditionally, random fingerprinting of a matrix \mathbf{A} consists of computing $\mathbf{A}\mathbf{r}^\top$ for $\mathbf{r} \xleftarrow{\$} \mathbb{F}^b$ (see [Fre77]), where the fingerprint is computed and given by a *trusted* party. It is not difficult to see in Hyrax-VMV that if the (untrusted) prover were to commit to \mathbf{A} with a public random vector $\mathbf{r} \in \mathbb{F}^b$, then Hyrax-VMV no longer remains sound via the following attack. The prover can honestly send $\mathbf{C} = \mathbf{A}\mathbf{r}^\top$ and later convince the verifier that $\tilde{y} = \mathbf{L}\mathbf{A}\mathbf{R}^\top$ for some $\tilde{y} \neq y$ (the true value of $\mathbf{L}\mathbf{A}\mathbf{R}^\top$) by computing $\tilde{\mathbf{L}} \neq \mathbf{L}\mathbf{A}$ such that $\langle \mathbf{L}, \mathbf{C} \rangle = \langle \tilde{\mathbf{L}}, \mathbf{r} \rangle$ and $\langle \tilde{\mathbf{L}}, \mathbf{R} \rangle \neq \mathbf{L}\mathbf{A}\mathbf{R}^\top$.

However, in Hyrax-VMV, the discrete-log assumption prevents this attack and allows the prover to compute the fingerprint $\mathbf{A}\mathbf{r}^\top$ *without explicitly knowing \mathbf{r}* . The commitment $\mathbf{C} = \mathbf{A}\mathbf{g}^\top$ acts exactly like a random fingerprint of the matrix \mathbf{A} , where the randomness \mathbf{r} is hidden within \mathbf{g} . To see this, we can reimagine sampling the random generators \mathbf{g} as follows: (1) sample generator $g \in \mathbb{G}$; (2) sample random vector $\mathbf{r} \xleftarrow{\$} \mathbb{F}^b$; (3) define $\mathbf{g}_i = g \cdot \mathbf{r}_i$ for all i . Since g is a generator, each \mathbf{g}_i is a generator. Now, by hardness of discrete-log, the prover cannot uncover the hidden \mathbf{r} , which allows the prover to provide the fingerprint of \mathbf{A} without explicit access to the randomness, avoiding the above attack.

By this above discussion, we conclude that in the OMA protocol for Hyrax-VMV, Bob will need a random fingerprint of the matrix \mathbf{A} in order to cross-check Merlin’s claims. Moreover, the randomness used for this fingerprint must be hidden from Merlin. This is handled by having Alice and Bob sample shared private randomness \mathbf{r} and having Alice send the fingerprint $\mathbf{C} = \mathbf{A}\mathbf{r}^\top$ after Merlin’s interaction with Bob. Intuitively, Alice being honest and sending a correct fingerprint is capturing the cryptographic prover’s hands being tied by the discrete-log assumption and the knowledge soundness of Hyrax-VMV: if the prover tries to cheat, he either has to break the discrete-log assumption, or we can extract \mathbf{A} from an accepting transcript with a dishonest prover by the knowledge soundness property.

We now give the full OMA protocol for Hyrax-VMV in Figure 1 and discuss the protocol here. First, in Hyrax-VMV, let $\mathbf{A} \in \mathbb{F}^{a \times b}$ be the matrix that the prover has, as its private input, and let $\mathbf{L} \in \mathbb{F}^a$, $\mathbf{R} \in \mathbb{F}^b$ be the (public) inputs of the verifier. In the OMA protocol, Alice is given \mathbf{A} , Bob is given \mathbf{L}, \mathbf{R} , and Merlin receives both of these inputs. Before the start of Bob and Merlin’s interaction, Alice and Bob sample shared private




Figure 1: Hyrax-VMV as an OMA communication protocol. Dashed/red lines indicate private shared randomness that is independent of parties' inputs. The protocol executes steps (1) through (3) in order.

randomness $\mathbf{r} \xleftarrow{\$} \mathbb{F}^b$. Now, intuitively, the OMA protocol executes Hyrax-VMV in reverse: the protocol begins with Merlin sending $(y = \mathbf{L}\mathbf{A}\mathbf{R}^\top, \mathbf{L}' = \mathbf{L}\mathbf{A})$ to Bob. Then, Alice sends the fingerprint of her input $\mathbf{C} = \mathbf{A}\mathbf{r}^\top$ to Bob. Finally, Bob performs the same verification procedure as the Hyrax-VMV verifier.

How does this OMA protocol capture Hyrax-VMV? As mentioned previously, the cryptographic protocol is utilizing a random fingerprint, hidden within the discrete-log of \mathbf{g} , to force the prover to send the correct values; i.e., the soundness of the protocol relies on the randomness hidden within \mathbf{g} . In the described OMA protocol, Alice plays the role of the committer, and Merlin plays the role of the prover. Crucially, Merlin *does not know the vector \mathbf{r}* , as it is private randomness shared between Alice and Bob. This ties Merlin's hands: if a dishonest Merlin is trying to convince Bob that $\mathbf{L}\mathbf{A}\mathbf{R}^\top = \tilde{y}$ for some incorrect value \tilde{y} , then with high probability (over the choice of \mathbf{r}) Bob will reject the interaction. See Section 4 for full details.

Evidence for Optimality of the Hyrax-PCS Verifier via OMA Lower Bounds. With Hyrax-VMV framed as an OMA protocol (which is the most difficult part of our results), we can leverage the rich line of works lower bounding communication protocols to lower bound the amount of work needed by the Hyrax-VMV verifier. As stated previously, Hyrax-VMV and its OMA protocol simultaneously prove a matrix-vector product ($\mathbf{L}' = \mathbf{L}\mathbf{A}$) and an inner product ($y = \langle \mathbf{L}', \mathbf{R} \rangle$). In Figure 1, let hcost denote the number of bits Merlin sends to Bob, and let vcost denote the number of bits Alice sends to Bob (note here we are not counting the shared randomness in the costs). By prior results, we have the following lower bounds: (1) For $\mathbf{A} \in \mathbb{F}^{a \times b}$ and $\mathbf{x} \in \mathbb{F}^b$, any OMA protocol for verifying $f(\mathbf{A}, \mathbf{x}) = \mathbf{A}\mathbf{x}^\top$ requires $\text{hcost} \cdot \text{vcost} = \Omega(\min(a, b)^2)$ [CMT13] (see Theorem 3.4); (2) For $\mathbf{u}, \mathbf{v} \in \mathbb{F}^m$, any OMA protocol for verifying $f(\mathbf{u}, \mathbf{v}) = \langle \mathbf{u}, \mathbf{v} \rangle$ requires $\text{hcost} \cdot \text{vcost} = \Omega(m)$ [Kla03, CCM⁺14] (see Theorem 3.5).

Let the total cost of the OMA protocol be defined as $\text{hcost} + \text{vcost}$. Setting $m = b$, we minimize $\text{hcost} + \text{vcost}$ subject to $\text{hcost} \cdot \text{vcost} = \Omega(\max(b, \min(a, b)^2))$. With a bit of manipulation and calculus, we can minimize this as $\text{hcost} + \text{vcost} = \Omega(\sqrt{\max(b, \min(a, b)^2)}) = \Omega(\max(\sqrt{b}, \min(a, b)))$. This implies that the cryptographic verifier in Hyrax-VMV must perform $\Omega(\max(\sqrt{b}, \min(a, b)))$ bit-operations. Now, for Hyrax-PCS, let $N = 2^n$ be the instance size (i.e., a PCS for n -variate multilinear polynomials or degree $N - 1$ univariate polynomials). Then, Hyrax-PCS sets $a = b = \sqrt{N}$, which yields a lower bound of $\text{hcost} + \text{vcost} = \Omega(\sqrt{N})$. This gives us strong evidence that the Hyrax-PCS verifier must perform $\Omega(\sqrt{N})$ bit-operations. Notice that in Hyrax-VMV, the verifier computes 2 inner products over \mathbb{G} of length a and one inner product over \mathbb{F} of length b , leading to $O(\sqrt{N})$ \mathbb{G} -operations in Hyrax-PCS, strongly suggesting that the verifier is optimal up to polylog $|\mathbb{G}|$ factors; see Theorem 4.1 and Corollary 4.2 in Section 4 for full details.

1.2.4 Bulletproofs: Review and Information-Theoretic Core

The heart of the Bulletproofs polynomial commitment scheme is a proof of knowledge (PoK) of an opening to a Pedersen vector commitment. Given random generators $\mathbf{g} \xleftarrow{\$} \mathbb{G}^N$ as public parameters, the prover gives some commitment C and proves it knows a vector $\mathbf{w} \in \mathbb{F}^N$ such that $C = \langle \mathbf{g}, \mathbf{w} \rangle$. Bulletproofs utilizes a “split-and-fold” recursive technique to continuously reduce the claim $C = \langle \mathbf{g}, \mathbf{w} \rangle$ for length N vectors \mathbf{g}, \mathbf{w} to a new claim $C' = \langle \mathbf{g}', \mathbf{w}' \rangle$ for length $N/2$ vectors \mathbf{g}', \mathbf{w}' , until the final vectors are of length 1 and the prover simply sends the final “folded” witness. Here, \mathbf{g}' and \mathbf{w}' are a random linear combination of the left and right halves of the (respective) vectors \mathbf{g} and \mathbf{w} , where the randomness is provided by the verifier during the interaction. Bulletproofs readily extend to a PCS for any polynomial where polynomial evaluation can be expressed as an inner product (e.g., multilinear and univariate polynomials); see [Section 5.3](#) for full details.

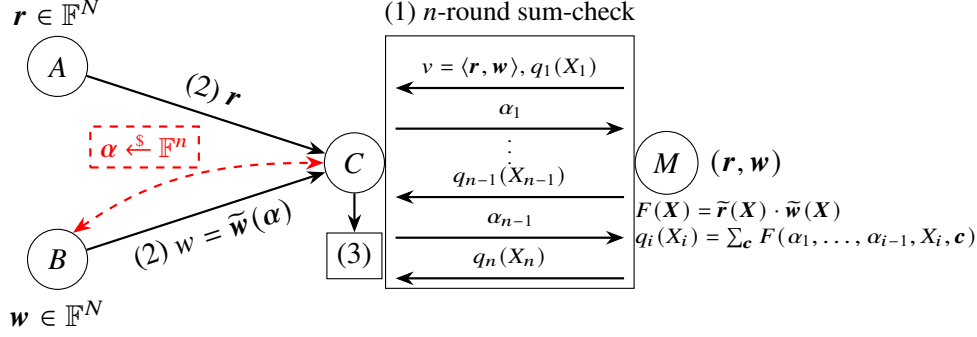
Bootle et al. [BCS21] originally observed that the Bulletproofs PoK is an instantiation of the sum-check protocol [LFK⁺92] for proving the inner product of two vectors \mathbf{w}, \mathbf{g} . The sum-check proves $C = \sum_{\mathbf{c} \in \{0,1\}^n} \tilde{\mathbf{g}}(\mathbf{c}) \tilde{\mathbf{w}}(\mathbf{c})$, where $\tilde{\mathbf{w}}$ and $\tilde{\mathbf{g}}$ are appropriately defined multilinear extensions of \mathbf{w} and \mathbf{g} , respectively.¹ The Bulletproofs sum-check allows us to do something impossible in an interactive proof: it allows the prover and verifier to run the sum-check with respect to a hidden vector \mathbf{r} . As with our discussion of Hyrax, this vector \mathbf{r} is hidden in the group generators \mathbf{g} . But because of hardness of discrete-log, \mathbf{r} is hidden, and due to group homomorphism, the prover and verifier can still operate the sum-check over the claim $v = \langle \mathbf{r}, \mathbf{w} \rangle$ without either party receiving access to \mathbf{r} directly.

As just mentioned, this is *impossible* in an information-theoretic setting: any information-theoretic prover needs the vector \mathbf{r} to perform the sum-check protocol. In Hyrax, the hidden random vector in \mathbf{g} was utilized as a random fingerprint that the prover had no control over, and when translating Hyrax to a communication protocol, we preserved this property: the randomness was hidden from Merlin. However, any interactive proof or communication protocol for proving $\langle \mathbf{w}, \mathbf{r} \rangle$ needs to give the prover (or Merlin) explicit access to \mathbf{r} , else they could never run the protocol in the first place!

This leads us to the following question: if the vector hidden in \mathbf{g} of Bulletproofs isn’t operating as a random fingerprint hidden from the prover, what is the “crypto” actually achieving here? To answer this question, we first recall what is needed to achieve soundness in the sum-check protocol. Consider a sum-check protocol for proving $C = \langle \mathbf{w}, \mathbf{r} \rangle$, which performs the sum-check over the polynomial $F(X) = \tilde{\mathbf{r}}(X) \cdot \tilde{\mathbf{w}}(X)$ (the multilinear extensions of \mathbf{r}, \mathbf{w}). Crucial to soundness of the sum-check protocol is that the verifier can evaluate $F(\alpha)$ at some random point α at the end of the protocol, which in our context means that the verifier needs to evaluate $\tilde{\mathbf{r}}(\alpha)$ and $\tilde{\mathbf{w}}(\alpha)$.

In Bulletproofs, the verifier can compute $\tilde{\mathbf{r}}(\alpha)$ implicitly by computing $\tilde{\mathbf{g}}(\alpha)$ (i.e., \mathbf{r} is hidden in the exponents of \mathbf{g}), but *cannot* compute $\tilde{\mathbf{w}}(\alpha)$, as \mathbf{w} is the prover’s private input. Instead, the prover sends a value w which is claimed to be $\tilde{\mathbf{w}}(\alpha)$. In the information-theoretic setting, such a sum-check protocol would be completely unsound, as the prover can always send some w^* to pass the final check of the sum-check protocol: the final check is of the form $C^* \stackrel{?}{=} \tilde{\mathbf{r}}(\alpha) \cdot w^*$, and the prover knows both C^* and $\tilde{\mathbf{r}}(\alpha)$ before sending w^* , so can trivially pass this check. However, in the cryptographic setting, the discrete-log assumption gives us the powerful knowledge soundness property: in order to send some w^* which would violate soundness of Bulletproofs, the cryptographic prover would need to break the discrete-log assumption, or one could extract the entire witness \mathbf{w} from the transcript of the protocol. Thus, the cryptographic prover is forced to send the correct value $w = \tilde{\mathbf{w}}(\alpha)$. This tells us that any information-theoretic core of Bulletproofs must enforce this invariant: the verifier must receive the correct and honest $\tilde{\mathbf{w}}(\alpha)$.

¹[BCS21] utilize a weighted sum-check protocol and non-standard definitions for multilinear extensions; the reason seems to be so that the resulting sum-check protocol is identical to the original Bulletproofs protocol.



- (3) C does the following.
- Set $v_1 = v$.
 - For each $i \in [n]$: (a) Check $v_i = q_i(0) + q_i(1)$ and $\deg(q_i) \leq 2$; (b) Set $v_{i+1} = q_i(\alpha_i)$.
 - Compute $r = \tilde{r}(\alpha)$. Check $v_{i+1} = r \cdot w$. Accept (i.e., output v_1) if and only if all checks pass.

Figure 2: BP-Sum-Check as an n -round SMH protocol for $N = 2^n$. The red dashed line indicates private randomness shared between the respective parties. The protocol executes (1), followed by (2), followed by (3).

Information-Theoretic Core of Bulletproofs. With this in mind, we turn to modeling Bulletproofs as a communication protocol. However, the OMA model, which we used for Hyrax, will not suffice (see the end of this section and Section 5.4 for discussion). Instead, we will need a new communication model to properly model Bulletproofs (and Dory in Section 1.2.5). This model, which we call the *simultaneous messaging with a helper* (SMH) model, is a generalization of the OMA model with four parties: Alice, Bob, Charlie, and Merlin. Alice has input a , Bob has input b , Merlin receives (a, b) , and the goal is to help Charlie compute $f(a, b)$. Notice here that Charlie has no input; moreover, Alice and Bob cannot speak to each other and can only send a single message to Charlie. Meanwhile, Merlin and Charlie can interact via an interactive proof. This model is readily captured by OMA model where “Bob” in the OMA protocol is simply a super player of Bob and Charlie (or Alice and Charlie) in the SMH protocol.

With our new model in hand, we can specify a SMH communication protocol for Bulletproofs. We give the protocol in Figure 2 and give an overview here. As with Hyrax, framing Bulletproofs as a communication protocol rearranges the execution of the protocol. To begin, for $N = 2^n$, Alice has input $r \in \mathbb{F}^N$ and is *deterministic* (we’ll come back to this later); Bob has input $w \in \mathbb{F}^N$ and additionally shares private randomness $\alpha \xleftarrow{\$} \mathbb{F}^n$ with Charlie; Merlin receives both (r, w) but does not receive α . Then, the communication protocol begins by having Charlie and Merlin execute the sum-check protocol to prove that $\langle r, w \rangle = v$ for some value $v \in \mathbb{F}$. At the end of the sum-check protocol, Alice sends Charlie the whole vector r and Bob sends Charlie $w = \tilde{w}(\alpha)$ (the multilinear extension of w evaluated at α). Finally, with all of this information, Charlie can verify whether Merlin was honest during the execution of the sum-check protocol.

Why does this capture Bulletproofs? As we discussed previously, Bulletproofs is the sum-check protocol where the verifier computes $\tilde{r}(\alpha)$ implicitly from the public randomness g , and *honestly* receives $w = \tilde{w}(\alpha)$ from the prover, as otherwise, either the discrete-log assumption is broken or knowledge soundness allows us to extract the entire witness w from the interaction. In our SMH protocol, we model the verifier receiving this correct $\tilde{w}(\alpha)$ by having the honest party Bob send it after Charlie and Merlin’s interaction. Key here is that the randomness α is hidden from Merlin before the start of the sum-check, otherwise the protocol is unsound. Moreover, Bob and Charlie need this shared randomness, otherwise if Charlie used some different randomness during the sum-check (say some β), then Bob would have no way of sending $\tilde{w}(\beta)$ to Charlie.

As with the Bulletproofs verifier, Charlie must also compute $\tilde{r}(\alpha)$ from Alice's input r . Here, this is trivial as Alice sends r to Charlie. In particular, we believe this is the correct modeling of Bulletproofs; i.e., Alice is deterministic and simply sends her entire input. We have two main reasons for this belief. First, in the actual prescribed Bulletproofs protocol, the verifier simply reads g from the public parameters and computes $\tilde{g}(\alpha)$ locally after running the sum-check protocol with the prover. Second, consider what it would mean if Alice, Bob, and Charlie all shared the random string α . In this case, Alice could send $\tilde{r}(\alpha)$ to Charlie at the end of the protocol yielding very succinct communication. However, how would this translate to the cryptographic setting? It is not clear; at best, this seems to be a variant of Bulletproofs with a trusted setup that includes $\tilde{g}(\alpha)$ with α known by the verifier then later revealed to the prover during the interactive protocol. Such a protocol is inherently different from Bulletproofs. Another possibility is to ask Merlin in the SMH protocol to send $\tilde{r}(\alpha)$ as the last message, but the information-theoretic dishonest Merlin can construct a purported value r^* to verify any the entire transcript, even if Bob sends the correct $\tilde{w}(\alpha)$. Moreover, such a Merlin sending $\tilde{r}(\alpha)$ does not match the prescribed Bulletproofs prover, leading to a mismatch between the cryptographic protocol and the information-theoretic core. See [Section 5](#) for full discussion.

All together, Charlie in our SMH protocol for Bulletproofs gets from Alice and Bob all the necessary information needed in order to verify the sum-check protocol executed with Merlin. Charlie receives r and $\tilde{w}(\alpha)$ from Alice and Bob, respectively, computes $\tilde{r}(\alpha)$, and uses this to verify the final claim of the sum-check protocol. This establishes the necessary conditions for the soundness of the sum-check protocol, yielding a valid SMH protocol for the inner product function.

Evidence for Optimality of the Bulletproofs-PCS Verifier via New SMH Lower Bounds. With our SMH protocol for Bulletproofs established, we now turn towards our goal of lower bounding the communication needed for parties in our protocol. However, the SMH model is a new communication model we introduce, so there are no prior known lower bounds we can leverage. Therefore, we prove new lower bounds on the communication complexity of computing inner product in the SMH model.

Theorem 1.2 (Informal; see [Theorem 3.10](#)). *Any SMH protocol for computing $f(x, y) = \langle x, y \rangle$ for vectors of length n such that Alice is deterministic requires Alice to send $\Omega(n)$ bits.*

Intuitively, the above proof follows since if Alice has input x and sends a $n - 1$ bit message m_A , then Merlin can find another vector $x' \neq x$ and $\langle x, y \rangle \neq \langle x', y \rangle$ (for Bob's input y) such that if Alice has input x' , her message would *also* be m_A , due to the Pigeonhole Principle and because Alice is deterministic. So Charlie would accept the same interaction with Merlin if Alice had input x or x' , violating soundness of the protocol.

Letting $hcost$ denote the number of bits exchanged between Charlie and Merlin, $bcost$ denote the number of bits Bob sends to Charlie, and $acost$ denote the number of bits Alice sends to Charlie. In our SMH protocol for Bulletproofs, we have specified that Alice is deterministic, thus in accordance with the above lower bound, we have $acost = \Omega(N)$, which implies that $hcost + acost + bcost = \Omega(N)$. This gives us strong evidence that the Bulletproofs verifier must perform $\Omega(N)$ bit-operations, otherwise we could translate the Bulletproofs verifier doing $o(N)$ work a communication protocol where Alice sends $o(N)$ bits but remains deterministic, violating our lower bound. Now, prescribed Bulletproofs verifier reads $O(N \log |\mathbb{G}|)$ bits (the vector g) from the public input, then computes $g(\alpha)$, which can be done in roughly $N \log |\mathbb{G}| / (\log(N) + \log \log |\mathbb{G}|)$ group operations via Pippenger's algorithm [[Pip80](#)], which gives roughly $O(N \text{polylog } |\mathbb{G}|)$ bit-operations. This strongly suggests that the Bulletproofs verifier is optimal up to $\text{polylog } |\mathbb{G}|$ factors; in particular, there is strong evidence to support that the Bulletproofs verifier is inherently *linear*.

More Evidence of Optimality: Framing an Attack on Prior Work in Our Model. As mentioned previously, two withdrawn papers [[Lu22](#), [Lu24](#)] attempted to obtain a sublinear Bulletproofs verifier. This

was done by attempting to augment Bulletproofs by changing how the “foldings” were performed, and by having the prover send some additional information to help the verifier efficiently compute the final folding of the public parameters \mathbf{g} in sublinear time. However, Gabizon [Gab24] showed an attack against the protocol where the prover could commit to some vector \mathbf{w} and open to some *other vector* $\mathbf{c} \neq \mathbf{w}$. Such an attack directly translates to the proof of our SMH lower bound: in the proof of [Theorem 1.2](#) (see [Theorem 3.10](#)), if Alice sends less than n bits when she is deterministic and her input \mathbf{x} is of length n , Merlin can find another $\mathbf{x}' \neq \mathbf{x}$ such that Alice’s message on input \mathbf{x} or \mathbf{x}' remains the same.

Why Does OMA Not Suffice for Bulletproofs? A natural question to ask is why did we introduce our new SMH communication model instead of framing Bulletproofs as an OMA model protocol, as we did with Hyrax. We give a high-level overview on why we believe the OMA model is not the correct model for capturing Bulletproofs; see [Section 5.4](#) for the full discussion.

In the OMA model, we only have 2 parties, Alice and Bob, and the helper Merlin. The first thing to think about in this model is which parties receive which inputs? In [Figure 2](#), there is the input \mathbf{w} , representing the cryptographic prover’s private witness, and there is \mathbf{r} , representing the “hidden” group exponents. Now, in an OMA protocol, Bob plays the role of the cryptographic verifier, and we want to translate the amount of bits communicated to Bob to a lower bound on the time the cryptographic verifier needs for verification.

We have two possible scenarios: (1) Alice has input \mathbf{r} and Bob has input \mathbf{w} ; (2) Alice has input \mathbf{w} and Bob has input \mathbf{r} . For (1), we do not believe this makes sense for modeling the cryptographic verifier, because it is, in essence, giving the verifier the prover’s private witness \mathbf{w} . In the cryptographic protocol, if the verifier has this witness, then there is no need for the verifier to even talk to the prover!

So, intuitively, that only leaves us with situation (2), where Alice has input \mathbf{w} and Bob has input \mathbf{r} . Here, we have the choice of letting Alice be deterministic or randomized. In the deterministic case, we actually run into the same lower bound as [Theorem 1.2](#), but in the OMA model (see [Corollary 3.11](#)): Alice will need to send her entire input to Bob. But, this is now just the trivial proof system where the prover sends his witness \mathbf{w} , so the verifier does not need to interact with the prover in the first place. Now suppose we let Alice and Bob share randomness α , Bob and Merlin run the sum-check protocol (i.e., Bulletproofs), and Alice sends $\tilde{\mathbf{w}}(\alpha)$ to Bob. This communication protocol sends $O(\log(N))$ field elements, which is much less than the linear lower bound we are trying to argue. Notably, Bob still needs to compute $\tilde{\mathbf{r}}(\alpha)$ in this case, but we can no longer capture the cost of this computation using communication complexity.

For these reasons, which are more deeply explored in [Section 5.4](#), we do not believe that the OMA model is the right model for capturing the verifier time of Bulletproofs.

1.2.5 Dory: Review and Information-Theoretic Core

The heart of Dory is nearly identical to Bulletproofs: it is a proof of knowledge of an opening to a Pedersen vector commitment, but now over a *pairing-friendly group* \mathbb{G} ;² such a commitment, due to Abe et al. [AFG⁺16, Gro09], is also known as an AFGHO commitment. Given random generator $\mathbf{g} \xleftarrow{\$} \mathbb{G}^N$ as public parameters, the prover gives some commitment C and proves it knows a vector $\mathbf{w} \in \mathbb{G}^N$ such that $C = \langle \mathbf{g}, \mathbf{w} \rangle_e := \sum_i e(\mathbf{g}_i, \mathbf{w}_i) \in \mathbb{G}_T$, where $e(\cdot, \cdot)$ is the pairing function of \mathbb{G} and \mathbb{G}_T is the target group.

Pairings are used to enable *logarithmic verification*, as pairings allow parties to commit to elements of \mathbb{F} and \mathbb{G} . In particular, pairings allow Dory to achieve what Bulletproofs could not: to force the cryptographic prover to send the correct value of $\tilde{\mathbf{g}}(\alpha)$ to the verifier, whereas in Bulletproofs, the verifier needed to compute

²For our purposes, it suffices to only consider symmetric pairings. See [Section 2](#) for a review of pairing-friendly groups and [Section 6](#) for a full discussion of Dory.

this value themselves. Dory achieves this as follows. The first round of Dory is identical to the first round of Bulletproofs: using verifier randomness, the claim $C = \langle \mathbf{g}, \mathbf{w} \rangle_e$ about vectors of size N is reduced to a new claim $C' = \langle \mathbf{g}', \mathbf{w}' \rangle_e$ about vectors of size $N/2$, where \mathbf{g}' and \mathbf{w}' are random linear combinations of the left and right halves of \mathbf{g}, \mathbf{w} . Since we are operating over a pairing-friendly group \mathbb{G} , we can force the prover to *commit to \mathbf{g}'* using some new (public) commitment key $\mathbf{\Gamma} \in \mathbb{G}^{N/2}$, resulting in a new claim $D = \langle \mathbf{g}', \mathbf{\Gamma} \rangle$. At the start of the second round of Dory, the prover sends this new claim D to the verifier, and then the prover and verifier run two instances of the Bulletproofs PoK protocol in parallel for C' and D . This process continues every round, requiring a new public commitment key per round and resulting in i parallel invocations of the Bulletproofs PoK for inputs of size $N/2^{i-1}$ in round i .

At the end of the protocol, the verifier has $n + 1$ claims to verify for $N = 2^n$. Assuming that the intermediate commitment keys (e.g., $\mathbf{\Gamma}$) are part of the public parameters, logarithmic verification is achievable with (linear-time) *preprocessing*. Let $\mathbf{g} = \mathbf{\Gamma}^{(1)}$ and let $\mathbf{\Gamma}^{(i)} \in \mathbb{G}^{N/2^{i-1}}$ be the commitment key for round i of the protocol. Then, preprocessing of the form $(\Delta_{i,L}, \Delta_{i,R}) = (\langle \mathbf{\Gamma}_L^{(i)}, \mathbf{\Gamma}^{(i+1)} \rangle_e, \langle \mathbf{\Gamma}_R^{(i)}, \mathbf{\Gamma}^{(i+1)} \rangle_e)$ gives the verifier enough information to perform only a *logarithmic* amount of work to verify the interaction, where the subscripts L, R respectively denote the left and right halves of the vectors.

The keen-eyed readers may notice that the above protocol actually has $O(\log^2(N))$ proof size/verification time; indeed, the above protocol is a simpler variant of Dory which we call *Proto-Dory*.³ Obtaining the full Dory protocol with logarithmic verification from Proto-Dory essentially amounts to taking a random linear combination of the new claim introduced in round i with the claim from the previous round; we refer the reader to [Section 6](#) for our treatment of the full Dory protocol and to [\[Tha22\]](#) for an excellent overview of both these protocols as this overview, while sufficient for understanding Dory at a high-level, is oversimplified and removes many of the subtle (and necessary) details for the actual protocol. However, for the purposes of the overview, Proto-Dory suffices for understanding the information-theoretic core of Dory. We also remark that, like Bulletproofs, (Proto-)Dory readily yields a PCS for any polynomials that can express polynomial evaluation as an inner product (e.g., multilinear and univariate); see [Section 6.3](#) for full details.

IT Core of Proto-Dory. As with Bulletproofs, we model Proto-Dory as a communication protocol in our new SMH model. From the above discussion, Proto-Dory is many invocations of Bulletproofs in parallel over pairing-friendly groups. In particular, in round i of the protocol, the prover and verifier are running i parallel Bulletproofs invocations over vectors of size $N/2^{i-1}$. This observation allows us to make a new connection: Proto-Dory is *also* an instantiation of the sum-check protocol (or, rather, many parallel invocations of the sum-check). To the best of our knowledge, we are the first to make this connection and observation explicit; moreover, in [Section 6](#), we show how the full Dory protocol is also an instantiation of (a slight variant) of the standard sum-check protocol.⁴

This yields a clear way to model Proto-Dory as a communication protocol, which we present in [Figure 3](#) and review here. Proto-Dory looks nearly identical to Bulletproofs ([Figure 2](#)), with a crucial difference that Alice is now *randomized*. First, Bob and Charlie both share a private random string α ; this is the sum-check randomness used by Charlie and enables Bob to send $\tilde{\mathbf{w}}(\alpha)$ at the end of the protocol. Now, the randomization of Alice is capturing the preprocessing of Proto-Dory that both the prover needs to run the sum-check, and the verifier needs to verify the sum-check. The vectors $\gamma^{(i+1)}$ are taking the role of the commitment keys in Proto-Dory, whereas the message sent by Alice to Charlie represents the preprocessing (expressed as polynomials) performed before the start of the cryptographic protocol (during the setup phase).

Why does this capture Proto-Dory? We first focus on Bob and Charlie’s specification. (Proto-)Dory is

³Proto-Dory actually appears in [\[Lee21\]](#) as a build-up to the “full” Dory protocol.

⁴In this light, it is possible to apply the results of [\[BCS21\]](#) to Dory; however, we leave this as future work to explore.

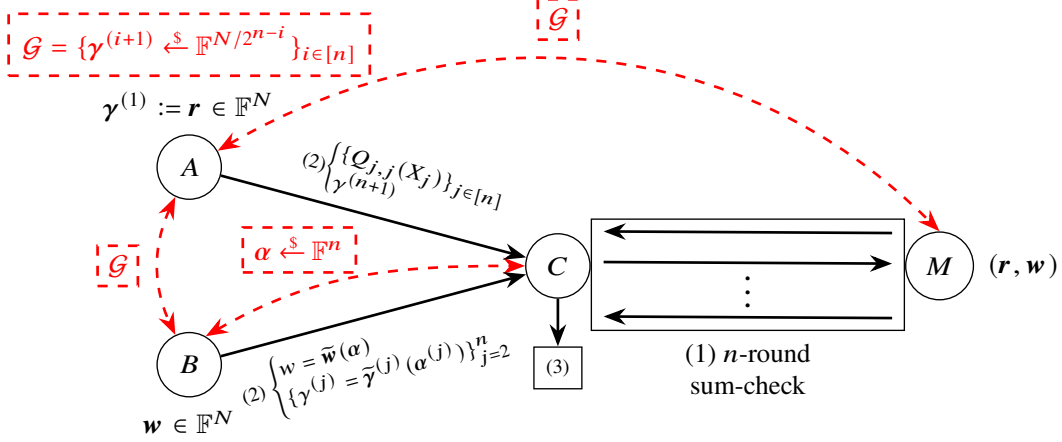
also a proof of knowledge under the discrete-log assumption in pairing-friendly groups. At the end of the actual protocol (see [Algorithm 6.1](#)), as with Bulletproofs the prover provides the value w claimed to be $\tilde{w}(\alpha)$, which the verifier needs in order to verify the sum-check protocol since there is no oracle access to \tilde{w} and the verifier cannot evaluate this polynomial at random points. By the proof of knowledge and discrete-log assumption (in pairing-friendly groups), the cryptographic prover is forced to send the correct w . In the communication setting, this correct value is enforced by having Bob send it; to do so, Bob and Charlie need to share the randomness ahead of time. Second, the interaction between Charlie and Merlin is exactly the interaction that occurs in Proto-Dory: during round i of the sum-check, Merlin and Charlie are running i invocations of the sum-check protocol in parallel. In [Figure 3](#), the polynomial S_i is the “main” sum-check polynomial for verifying the inner product $\langle \mathbf{r}, \mathbf{w} \rangle$, whereas the $(i - 1)$ polynomials $Q_{j,i}$ for $j \in [i - 1]$ are the parallel sum-check invocations. Third, Alice’s message to Charlie is exactly the preprocessing performed during the setup phase, just expressed as polynomials, which Charlie needs to verify the sum-check claims from the interaction with Merlin.

Now, why do Alice, Bob, and Merlin share random $\gamma^{(i)}$ for all $i \in [n]$? In the actual (Proto-)Dory protocol (see [Algorithms 6.1](#) and [6.2](#)), these vectors are actually recursively defined as the left halves of the previous vector: i.e., $\gamma^{(i+1)} = \gamma_L^{(i)}$ for $i \in [n]$, where the subscript L denotes the left half of the vector. First, if we let Alice share these vectors in this way, then Alice is now *deterministic*, and we run into the lower bound of [Theorem 1.2](#). Moreover, under such a setting, we are able to demonstrate an attack that Merlin can perform to violate the soundness of the protocol (in line with the proof of [Theorem 3.10](#), see [Remark 6.5](#) for details). Finally, this violates the very model we are introducing, where we enforce that shared randomness be independent of the parties’ inputs. So these vectors need to be random in order for the SMH protocol to be sound and in order for Merlin to execute the parallel sum-check with Charlie. Another natural question, as with Bulletproofs, is why do Alice, Bob, and Charlie not simply share α ? As mentioned in [Section 1.2.4](#), it is not clear how this translates to the cryptographic setting, as it would require the setup of Dory to have the vector α hidden and only known to the verifier before running the proof of knowledge with the prover, leading to some trusted setup that later becomes public.

The final discrepancy to discuss is Merlin’s final message to Charlie and Bob’s new messages. First, we additionally require Merlin to include in his final message a degree-1 polynomial \hat{g} , claimed to be the polynomial $\tilde{\gamma}^{(1)}(\alpha_1, \dots, \alpha_{n-1}, X_n)$, the multilinear extension of Alice’s input. Second, we additionally require Bob to send $\gamma^{(j)} = \tilde{\gamma}^{(j)}(\alpha_j, \dots, \alpha_n)$ for $j \in \{2, \dots, n\}$. Intuitively, Bob’s messages are again capturing the proof of knowledge property of the protocol: these values, in the cryptographic protocol, are forced to be sent honestly by the prover, unless the discrete-log assumption is broken, or we can extract the prover’s full witness from the protocol transcript. However, Bob can send every value honestly except for $\gamma^{(1)} = \tilde{\gamma}^{(1)}(\alpha)$, as this value depends solely on Alice’s input. If Bob could compute this value, Bob would need Alice’s entire input and make the problem trivial: Bob just sends $\langle \mathbf{w}, \gamma^{(1)} \rangle$ to Charlie.

Therefore, we delegate the responsibility of providing $\gamma^{(1)}$ to Merlin in the SMH protocol. We do this by forcing Merlin to send the polynomial $\hat{g}(X_n)$: crucially, Merlin sends this before ever knowing the value α_n , so (by soundness of the sum-check protocol), the only way for Merlin to pass the final verification checks performed by Charlie if \hat{g} is dishonest is for Charlie to sample some bad randomness. This happens with probability at most $1/|\mathbb{F}|$, so soundness remains in the protocol, effectively forcing Merlin to send the correct value, otherwise be caught with high probability.

Evidence for Optimality of the Dory-PCS Verifier. While Proto-Dory is a $(\log(N) + 1)$ -round protocol with $O(\log^2(N))$ verification time, the full Dory protocol is $(\log(N) + 1)$ -rounds with $O(\log(N))$ verification time (ignoring polylog $|\mathbb{G}|$ factors). In both cases, when extending Proto-Dory and Dory to a polynomial



- (0) B and C sample shared private randomness $\alpha \in \mathbb{F}^n$. A , B , and M sample shared private randomness $\gamma^{(i+1)} \in \mathbb{F}^{N/2^{n-i}}$ for $i \in [n]$.
- (1) M and C engage in the n -round sum-check protocol over the polynomial $F(X) = \tilde{r}(X)\tilde{w}(X)$. For $i \in [n]$:
- For all $j \in [i-1]$, M defines $S_i(X_i) = \sum_{\mathbf{c}} F(\alpha_1, \dots, \alpha_{i-1}, X_i, \mathbf{c})$, $P_{i-1}(X_{[i-1,n]}) = \tilde{\gamma}^{(i-1)}(X_{[i-1,n]})\tilde{\gamma}^{(i)}(X_{[i,n]})$, and $Q_{j,i}(X_i) = \sum_{\mathbf{c}} P_j(\alpha_j, \dots, \alpha_{i-1}, X_i, \mathbf{c})$.
 - M sends $S_i(X_i)$ and $\{Q_{j,i}(X_i)\}_{j \in [i-1]}$. If $i = 1$, then send $v = \langle r, w \rangle$. If $i = n$, then send $\hat{g}(X_n) = \tilde{\gamma}^{(1)}(\alpha_1, \dots, \alpha_{n-1}, X_n)$.
 - C sends α_i .
- (2) A computes $Q_{j,j}(X_j) = \sum_{\mathbf{c} \in \{0,1\}^{n-j}} \tilde{\gamma}^{(j)}(X_j, \mathbf{c}) \cdot \tilde{\gamma}^{(j+1)}(\mathbf{c})$ for all $j \in [n]$ and sends $(\{Q_{j,j}(X_j)\}_{j \in [n]}, \gamma^{(n+1)})$. B computes and sends: $\gamma^{(j)} = \tilde{\gamma}^{(j)}(\alpha^{(j)})$ for $j \in [2, n]$ and $\alpha^{(j)} := (\alpha_j, \dots, \alpha_n)$, and $w = \tilde{w}(\alpha)$.
- (3) C does the following checks.
- Set $v_1 = v$. For each $i \in [n]$:
 - Check if $v_i \stackrel{?}{=} S_i(0) + S_i(1)$; reject if not.
 - For each $j \in [i-1]$, check $D_{j,i} \stackrel{?}{=} Q_{j,i}(0) + Q_{j,i}(1)$ and $\deg(Q_{j,i}) \leq 2$.
 - Set $v_{j+1} = q_j(\alpha_j)$ and $D_{j,i+1} = Q_{j,i}(\alpha_i)$ for $j \in [i]$
 - Check that $\deg(\hat{g}) \leq 1$ and compute $\gamma^{(1)} = \hat{g}(\alpha_n)$.
 - Check if $v^{(n+1)} \stackrel{?}{=} w \cdot \gamma^{(1)}$ and $D_{j,n+1} \stackrel{?}{=} \gamma^{(j)} \cdot \gamma^{(j+1)}$ for all $j \in [n]$.
 - Accept (i.e., output v_1) if and only if all checks pass.

Figure 3: Proto-Dory (Algorithm 6.1) as a n -round SMH protocol for $N = 2^n$. The dashed/red lines indicate private randomness shared before C and M 's interaction. The protocol executes (1) through (3) in order.

commitment scheme, these costs are only a constant factor larger; in particular, extending to a polynomial commitment scheme just requires a single addition of another parallel sum-check instance in the first round of the protocol. This translates to an n -round SMH protocol nearly identical to Figure 3 (and Figure 7) with some minor modifications (see Section 6.3 for full details). In some sense, our lower bound is now trivially shown: we specify an n -round SMH protocol for computing inner product, and by definition, Charlie and Merlin need to exchange at least 1 bit per round, leading to $\Omega(n)$ communication. This gives us the $\Omega(\log(N))$ lower bound (for $N = 2^n$) for the SMH protocol representing (Proto-)Dory, whereas the verifier time in Dory is $O(\log(N) \text{ polylog } |\mathbb{G}|)$ bit-operations, giving strong evidence that Dory has optimal verifier time.

1.3 Related Work

As we mentioned previously, to the best of our knowledge, we are the first to make the connection between polynomial commitment schemes and communication complexity. Here, we discuss some related communication models and other seemingly related work.

Communication Complexity. The first and most popular communication model is the 2-player model by Yao [Yao79] as described earlier. For a comprehensive discussion on many other communication models, see

the book by Kushivilitz and Nisan [KN97]. We briefly describe the models that are most relevant to our work.

Chandra, Furst, and Lipton [CFL83] extended the 2-player communication model to the k -player or *multi-party* model, where k players communicate with each other to compute a function of their private inputs. Babai, Kimmel, and Lokam [BKL95, BGK⁺03] introduced the *Simultaneous Message* (SM) model, which can be seen as a restriction of the multi-party model: the players cannot interact with each other but simultaneously send messages to a “referee”, who doesn’t directly see any of the inputs but generates the output from the players’ messages. We note that our SMH model should be seen as an extension of the SM model rather than the general multi-player model: SMH extends the 2-player SM model with a referee by introducing a helper in the standard Merlin-Arthur sense.

We now discuss some other “prover-enhanced” communication models studied in the literature. The *non-deterministic model* [KN97] is an analog of **NP**, where the prover sends a deterministic help message to the verifier (the pair Alice and Bob), who then communicate with each other using a deterministic protocol to compute a function of their inputs. Other variants of Merlin-Arthur communication (where randomization is allowed) include the setting where Merlin first sends a help message to Arthur (represented collectively by two or more players), which we can think of being written on a blackboard shared between the players, who then interact with each other (possibly based on Merlin’s message) in unrestricted number of rounds to evaluate the function [Kla03, CCM⁺14]. Natural generalizations to allow multiple rounds of interaction between Merlin and Arthur have been studied: Chakrabarti et al. [CCM⁺19] defined the **OMA**[k] model that we use in this paper and also study the **OIP**[k] model, the online communication analog of **IP**[k] similar to **OMA**[k], except **OIP**[k] allows Bob and Merlin to interact based on Bob’s input.

Other Work. A (seemingly) related work by Boyle et al. [BJS⁺25] introduces the concept of simultaneous-message and succinct secure computation. This concept seems similar to the simultaneous message model [BKL95], as it involves 3 parties, Alice, Bob, and Charlie, where Alice has some large input X , Bob has a small input y , and Charlie wishes to compute $f(X, y)$ while minimizing the bits exchanged between parties. However, this concept is introduced in the context of secure computation, where Alice and Bob wish to keep their inputs private. Moreover, Alice and Bob in this setting are allowed to share messages with each other prior to sending messages to Charlie, which is different from the SM model discussed previously and is therefore not directly related to our work or communication complexity. Another work by Cormode et al. [CDG⁺24] studies streaming zero-knowledge proofs that use streaming algebraic commitment protocols, which may be adapted to certain communication models using standard streaming-to-communication simulations. However, the motivation behind their work is in designing zero-knowledge streaming protocols for specific problems, which is entirely different from ours.

2 Preliminaries

For $n \in \mathbb{Z}^+$, we let $[n] := \{1, 2, \dots, n\}$. For a set S , we let $x \xleftarrow{\$} S$ denote the process of sampling an element of S uniformly at random. We let \mathbb{F}_q denote a finite field of size $q = p^a$ for some prime p and $a \in \mathbb{Z}^+$. When clear from context, we simply write \mathbb{F} . We let \mathbb{G} denote a finite cyclic group of prime order. For ease of presentation, we express group operations as additive rather than multiplicative. We let lower case bold denote vectors of field and/or group elements; e.g., $\mathbf{a} \in \mathbb{F}^n$ or $\mathbf{b} \in \mathbb{G}^n$. For any vector \mathbf{a} of length n , we let \mathbf{a}_i denote the i th element of \mathbf{a} . For $\mathbf{a}, \mathbf{b} \in \mathbb{F}^n$, we let $\langle \mathbf{a}, \mathbf{b} \rangle := \sum_i \mathbf{a}_i \cdot \mathbf{b}_i \in \mathbb{F}$ denote the inner product of \mathbf{a} and \mathbf{b} . Moreover, for group \mathbb{G} with scalar field \mathbb{F} , given $\mathbf{g} \in \mathbb{G}^n$ and $\mathbf{a} \in \mathbb{F}^n$, we let $\langle \mathbf{g}, \mathbf{a} \rangle = \langle \mathbf{a}, \mathbf{g} \rangle \in \mathbb{G}$ denote the inner product of \mathbf{g} and \mathbf{a} .

Multilinear Polynomials and Extensions. Let $f: \{0, 1\}^n \rightarrow \mathbb{F}$ be a function. Then the *multilinear extension* of f , denoted as $\tilde{f}: \mathbb{F}^n \rightarrow \mathbb{F}$ is the unique polynomial such that $\tilde{f}(\mathbf{c}) = f(\mathbf{c})$ for every $\mathbf{c} \in \{0, 1\}^n$, and the individual degree of \tilde{f} is at most 1. Such a polynomial can be uniquely defined as

$$\tilde{f}(X_1, \dots, X_n) := \sum_{\mathbf{c} \in \{0, 1\}^n} f(\mathbf{c}) \cdot \text{eq}(\mathbf{X}, \mathbf{c}), \text{ where} \quad (1)$$

$$\text{eq}(\mathbf{X}, \mathbf{c}) := \prod_{i=1}^n X_i c_i + (1 - X_i)(1 - c_i). \quad (2)$$

Here, for binary inputs $\mathbf{b}, \mathbf{c} \in \{0, 1\}^n$, $\text{eq}(\mathbf{b}, \mathbf{c}) = 1$ if and only if $\mathbf{b} = \mathbf{c}$, otherwise $\text{eq}(\mathbf{b}, \mathbf{c}) = 0$. We abuse notation and similarly let $\tilde{\mathbf{u}}$ denote the multilinear extension of a vector $\mathbf{u} \in \mathbb{F}^n$, where we naturally interpret \mathbf{u} as a function $\mathbf{u}: \{0, 1\}^n \rightarrow \mathbb{F}$ via the mapping $i \mapsto u_{i+1}$ for $i \in \{0, 1\}^n$.

We also utilize functions $g: \{0, 1\}^n \rightarrow \mathbb{G}$ (cf. [BCS21]) and analogously define multilinear extensions of g as above. This extension is denoted as $\tilde{g}: \mathbb{F}^n \rightarrow \mathbb{G}$, where $\text{eq}: \mathbb{F}^n \times \mathbb{F}^n \rightarrow \mathbb{F}$ is defined as in Eq. (2) and \mathbb{F} is the scalar field of \mathbb{G} . Now, given two polynomials $F(\mathbf{X}) \in \mathbb{F}[\mathbf{X}]$ and $G(\mathbf{X}) \in \mathbb{G}[\mathbf{X}]$, $H(\mathbf{X}) = F(\mathbf{X}) \cdot G(\mathbf{X}) \in \mathbb{G}[\mathbf{X}]$ is also a polynomial. We similarly naturally interpret $\mathbf{g} \in \mathbb{G}^N$ as a function $\mathbf{g}: \{0, 1\}^n \rightarrow \mathbb{G}$ and let $\tilde{\mathbf{g}}: \mathbb{F}^n \rightarrow \mathbb{G}$ denote the multilinear extension of \mathbf{g} .

Bilinear Pairings and Pairing Polynomials. Let $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ be pairing-friendly additive groups of prime order p , with scalar field \mathbb{F} , generators g_1, g_2, g_T , and bilinear pairing function $e: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$. For $\mathbf{a} \in \mathbb{G}_1^n, \mathbf{b} \in \mathbb{G}_2^n$, we let $\langle \mathbf{a}, \mathbf{b} \rangle_e := e(\mathbf{a}, \mathbf{b}) = \sum_i e(a_i, b_i) \in \mathbb{G}_T$. When clear from context, we drop the subscript e above. Similarly, for $a \in \mathbb{G}_1, b \in \mathbb{G}_2$, we let $a \cdot b := e(a, b)$.

We also define the notion of polynomials over pairing friendly groups. A degree- d , n -variate polynomial $f_1(\mathbf{X}) \in \mathbb{G}_1[\mathbf{X}]$ is a function $f_1: \mathbb{F}^n \rightarrow \mathbb{G}_1^n$ with degree at most d in each variable. Analogously, $f_2(\mathbf{X}) \in \mathbb{G}_2[\mathbf{X}]$ is a function $f_2: \mathbb{F}^n \rightarrow \mathbb{G}_2$ where each variable has degree at most d . Notably, $h(\mathbf{X}) = f_1(\mathbf{X}) \cdot f_2(\mathbf{X}) \in \mathbb{G}_T[\mathbf{X}]$ is a map $h: \mathbb{F}^n \rightarrow \mathbb{G}_T$. Intuitively, the coefficients of h are obtained by “pairing” the coefficients of f_1 and f_2 (i.e., replacing \cdot with e as above). Finally, we define the multilinear extension of functions $h_1: \{0, 1\}^n \rightarrow \mathbb{G}_1, h_2: \{0, 1\}^n \rightarrow \mathbb{G}_2$, and $h_3: \{0, 1\}^n \rightarrow \mathbb{G}_T$ analogously as with the prior section, and denote these functions by $\tilde{h}_1, \tilde{h}_2, \tilde{h}_3$, respectively. As above, we again naturally interpret $\mathbf{h}_1 \in \mathbb{G}_1^N$ and $\mathbf{h}_2 \in \mathbb{G}_2^N$ as functions $\mathbf{h}_1: \{0, 1\}^n \rightarrow \mathbb{G}_1, \mathbf{h}_2: \{0, 1\}^n \rightarrow \mathbb{G}_2$, and let $\tilde{\mathbf{h}}_1: \mathbb{F}^n \rightarrow \mathbb{G}_1, \tilde{\mathbf{h}}_2: \mathbb{F}^n \rightarrow \mathbb{G}_2$ denote the multilinear extensions of these vectors.

Interactive Protocols. An interactive protocol $\Pi = (P, V)$ consists of two players (i.e., interactive algorithms), a *prover* P and a *verifier* V . A protocol proceeds in *rounds*, where P and V each send a single message to each other in some order. In this work, we only consider protocols where during each round, the prover speaks first, followed by the verifier, and in the final round, only the prover sends a message. We say such a protocol is *k-rounds* (or has $(2k - 1)$ -moves) if P and V interact as described above for k -rounds. In particular, the prover sends k messages and the verifier sends $(k - 1)$ messages.

Let $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ be a relation (e.g., an **NP**-relation). We say that an interactive protocol Π is an *interactive proof* for R if for all $(x, w) \in R$, we have $\langle P(x, w), V(x) \rangle = 1$ with probability 1 (*completeness*), and if for any $x \notin L_R$ and any computationally unbounded prover P^* , the probability $\langle P^*(x), V(x) \rangle = 1$ is negligible (in the length of x) (*soundness*). Here, $\langle P(x, w), V(x) \rangle$ denotes the interaction between P and V according to protocol Π , where P is given (x, w) as input and V is given x as input, and L_R denotes the language of the relation R (i.e., $L_R = \{x: \exists w \text{ s.t. } (x, w) \in R\}$). We say that Π is an *interactive argument* if soundness holds with respect to any probabilistic polynomial time algorithm P^* . Finally, we say that Π is a *non-interactive argument* if $k = 1$. For completeness, we give the formal definition below.

Definition 2.1 (Interactive Proofs/Arguments). An interactive proof (resp., argument) for a language \mathcal{L} with relation $\mathcal{R}(\mathcal{L})$ is a tuple of algorithms (Gen, P, V) such that the following properties are satisfied:

- **Gen** is a non-interactive algorithm that on input 1^λ for security parameter $\lambda \in \mathbb{N}$ outputs public parameters pp .
- **Completeness**: For any $(x, w) \in \mathcal{R}(\mathcal{L})$ and $\text{pp} \leftarrow \text{Gen}(1^\lambda)$, we have $\Pr[\langle P(w), V \rangle(\text{pp}, x) = 1] = 1$, where the probability is taken over the generation of pp and the random coins of P and V , $\langle P, V \rangle(\text{pp}, x)$ denotes the verifier output after P and V interact with common input (pp, x) , and P additionally receives w as input.
- **ϵ -Soundness**: For any $x \notin \mathcal{L}$, $\text{pp} \leftarrow \text{Gen}(1^\lambda)$, and any unbounded (resp., polynomial-time) interactive algorithm P^* , we have $\Pr[\langle P^*, V \rangle(\text{pp}, x) = 1] \leq \epsilon(|x|, \lambda)$.
- **Efficiency**: The verifier runs in time $\text{poly}(|x|, \lambda)$. If (P, V) is an argument, then P runs in time $\text{poly}(|x|, |w|, \lambda)$.

The protocol (Gen, P, V) is said to be public coin if each message sent by V are independent uniform random strings of some bounded length and the output of V does not depend on any secret state.

Polynomial Commitment Schemes. A Polynomial Commitment Scheme (PCS) [KZG10] is a tuple of three algorithms Setup , Commit , Open and a PPT protocol Eval between two parties: a prover (or committer) P and a verifier V . The algorithm Setup , on input 1^λ for security parameter $\lambda \in \mathbb{N}$, outputs some set of public parameter pp , which all other algorithms take as input. In the Commit algorithm, P generates a commitment C to a polynomial $p(\cdot)$ (usually univariate or multilinear). In the Open algorithm, a party verifies that a given commitment C corresponds to a polynomial $p(\cdot)$, using the opening or randomness r during the commitment phase. In the Eval protocol, P and V engage in an interactive proof where P attempts to convince V that a claimed evaluation y is correct for a given input x , i.e., that $y = p(x)$. The verifier V outputs 1 (accepts) if the proof is successful, and 0 otherwise. Such an interactive proof often has the desirable *knowledge soundness* property. While we do not build polynomial commitment schemes in this work, for completeness we provide the formal definition of a polynomial commitment scheme below.

Definition 2.2 (Polynomial Commitment). A polynomial commitment scheme allows one party to commit to a polynomial. Subsequently, it enables them to convince another party that the polynomial, when evaluated at a particular point, results in a specific value. It consists of the following algorithms:

- $\text{pp} \leftarrow \text{PC.setup}(1^\lambda, \mathcal{F})$: Given the security parameter and a family of polynomials, output the public parameters pp .
- $\text{com} \leftarrow \text{PC.commit}(\phi, \text{pp})$: On input a polynomial ϕ and the public parameter pp , output the commitment C .
- $(y, \pi) \leftarrow \text{PC.open}(\phi, \text{pp}, x)$: On input the polynomial, the public parameters, and the point to be evaluated at, output $y = \phi(x)$, and the proof.
- $\{0, 1\} \leftarrow \text{PC.verify}(\text{pp}, \text{com}, x, y, \pi)$: Verify the proof against the claim $\phi(x) = y$, output 1 if verification passes, otherwise 0.

A polynomial commitment scheme satisfies the following properties:

- **Completeness**. For any polynomial $\phi \in \mathcal{F}$, and any evaluation point x . The following probability holds:

$$\Pr \left[\text{PC.verify}(\text{pp}, \text{com}, x, y, \pi) = 1 \mid \begin{array}{l} \text{pp} \leftarrow \text{PC.setup}(1^\lambda, \mathcal{F}), \\ \text{com} \leftarrow \text{PC.commit}(\phi, \text{pp}), \\ (y, \pi) \leftarrow \text{PC.open}(\phi, \text{pp}, x), \end{array} \right] = 1$$

- *Knowledge soundness.* For any PPT adversary \mathbf{A} , there exists an expected PPT extractor \mathbf{E} that on given oracle access to \mathbf{A} , outputs the polynomial such that

$$\Pr \left[\begin{array}{l} \phi' \leftarrow \mathbf{E}^{\mathbf{A}}(\text{pp}, \text{com}, x, y, \pi), \\ \text{PC.commit}(\phi') = \text{com}, \\ \phi'(x) \neq y \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \text{PC.setup}(1^\lambda, \mathcal{F}), \\ (\text{com}, y, x, \pi) \leftarrow \mathbf{A}(\text{pp}), \\ \text{PC.verify}(\text{pp}, \text{com}, x, y, \pi) = 1 \end{array} \right] \leq \text{negl}(\lambda)$$

3 Communication Complexity

3.1 Online Merlin-Arthur Communication Protocols

Online Merlin-Arthur (OMA) protocols [BFS86, CCM⁺14] are a class of communication protocols which extend classical Merlin-Arthur protocols to simulate k -round streaming interactive proof (SIP) protocols. OMA protocols consist of a fixed, publicly known function f , two parties, Alice (A) and Bob (B) along with a helper party Merlin (M). Alice and Bob each have their own private inputs x_A, x_B , respectively, and may toss private random coins. Moreover, Alice and Bob may share random coins which are hidden from Merlin. Bob and Merlin may exchange k messages to help Bob compute $f(x_A, x_B)$, after which Alice may send a single message m_A to Bob.

Definition 3.1 (Online Merlin-Arthur Protocols). *Let $f: X \times Y \rightarrow Z$ be an arbitrary function. We say that a protocol Π for computing f is an $\mathbf{OMA}[k]$ protocol if Π consists of two parties A and B and a helper M such that for all $x \in X$ and $y \in Y$, the following hold:*

1. *A receives x as input, B receives y as input, and M receives (x, y) as input;*
2. *M and B participate in a $(2k - 1)$ -move protocol, where B's messages do not depend on the input y and all parties send at least 1 bit per round, resulting in transcript τ ;*
3. *A sends message m_A to B after this protocol; and*
4. *B outputs some $z \leftarrow B(\tau, m_A, y)$.*

Moreover, all parties may be randomized and sample both private and shared randomness prior to Step (1) above, where shared randomness denotes random strings that are agreed upon between two or more parties.

See Figure 4 for a pictorial overview of the $\mathbf{OMA}[k]$ -model. Next, we define the security and cost of $\mathbf{OMA}[k]$ -model protocols.

Definition 3.2 (Secure $\mathbf{OMA}[k]$ Protocol). *We say that an $\mathbf{OMA}[k]$ protocol Π for a function f is (γ, ϵ) -secure if for all $x \in X$ and $y \in Y$, the following hold.*

1. *(Completeness) For all honest helpers M , it holds that $\Pr_{z \leftarrow \Pi(x, y)}[z = f(x, y)] \geq 1 - \gamma$.*
2. *(Soundness) For any dishonest helper M^* , it holds that $\Pr_{z \leftarrow \Pi(x, y; M^*)}[z \neq f(x, y)] \leq \epsilon$.*

Here, the probability is taken over any private or shared randomness sampled by parties. For brevity, we say that a $(1/3, 1/3)$ -secure $\mathbf{OMA}[k]$ protocol is simply a secure $\mathbf{OMA}[k]$ protocol.

Definition 3.3 ($\mathbf{OMA}[k]$ Protocol Cost). *Let Π be an $\mathbf{OMA}[k]$ protocol for a function f . We define the verification cost as $\text{vcost}(\Pi) := \max\{1, \log |m_A|\}$ and the help cost as $\text{hcost}(\Pi) := \max\{1, \log |\tau|\}$. Then, we define the cost of Π as $C(\Pi) := \text{vcost}(\Pi) + \text{hcost}(\Pi)$. Finally, we let $\mathbf{C}_{\mathbf{OMA}[k]}(f) := \min_{\Pi \in \mathbf{OMA}[k](f)} \{C(\Pi)\}$ denote the $\mathbf{OMA}[k]$ -cost of f , where $\mathbf{OMA}[k](f)$ denotes the set of all secure $\mathbf{OMA}[k]$ protocols for f .*




Figure 4: Graphical representation of a **OMA**[k]-model protocol for function $f: X \times Y \rightarrow Z$. Parties may or may not be randomized. If randomized, parties may sample private coins, or shared random coins. Red dashed lines indicate shared randomness between parties. τ represents the transcript of the $(2k - 1)$ -move protocol executed between B and M . Parties B and M never share randomness prior to the protocol execution.

3.1.1 OMA Lower Bounds

Here, we state two important lower bounds for the OMA model that we use in our work (see [Section 4](#)).

Theorem 3.4 ([CMT13]). *Given a (dense) matrix $\mathbf{A} \in \mathbb{F}^{a \times b}$ and a vector $\mathbf{x} \in \mathbb{F}^b$ for any finite field \mathbb{F} , any secure **OMA**[1] protocol Π for $f(\mathbf{A}, \mathbf{x}) = \mathbf{A} \cdot \mathbf{x}^\top$ requires $\text{hcost}(\Pi) \cdot \text{vcost}(\Pi) = \Omega(\min(a, b)^2)$.*

Theorem 3.5 ([Kla03, CCM⁺19]). *Let $k \in \mathbb{N}$ be a constant. Given two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{F}^m$ for any finite field \mathbb{F} , any secure **OMA**[k] protocol Π for $f(\mathbf{u}, \mathbf{v}) = \langle \mathbf{u}, \mathbf{v} \rangle$ requires $\text{hcost}(\Pi)^k \cdot \text{vcost}(\Pi) = \Omega(m)$.*

3.2 Simultaneous Messaging with a Helper Communication Protocols

In this work, we introduce a new communication model called *Simultaneous Messaging with a Helper*, which we believe to be of independent interest. This model is a natural extension of the *Simultaneous Message* (SM) model [Yao79, BKL95, BK97]. In the SM model, there is a fixed and publicly known function f . Two parties Alice (A) and Bob (B) each have a private input x_A, x_B , respectively, and send messages m_A, m_B to a third party, the referee Charlie (C). The goal is for Charlie to correctly output $f(x_A, x_B)$ given m_A, m_B , while minimizing the quantity $|m_A| + |m_B|$. Note that in this model, all parties are computationally unbounded, may or may not be randomized, and may or may not be allowed to share randomness with other parties.

We naturally extend the SM model to the *Simultaneous Messaging with a Helper* (SMH) model. Inspired by the OMA model, we add a fourth party to the SM model, a helper Merlin (M), who receives both Alice and Bob's inputs and is allowed to communicate with Charlie for k rounds (or participate in a $(2k - 1)$ -move protocol) to help Charlie compute $f(x_A, x_B)$ correctly.

Definition 3.6 (Simultaneous Messaging with a Helper). *Let $f: X \times Y \rightarrow Z$ be an arbitrary function. We say that a protocol Π for computing f is a **SMH**[k] protocol if Π consists of two parties A and B , a referee C , and a helper M , such that for all $x \in X$ and $y \in Y$, the following hold:*

1. *A receives x as input, B receives y as input, and M receives (x, y) as input;*
2. *M and C participate in a $(2k - 1)$ -move protocol, where all parties send at least 1 bit per round, resulting in transcript τ ;*
3. *After this interaction, A and B send messages m_A, m_B , respectively, to C; and*




Figure 5: Graphical representation of a **SMH**[k]-model protocol for function $f : X \times Y \rightarrow Z$. Parties may or may not be randomized. If randomized, parties may sample private coins, or shared random coins. Red dashed lines indicate shared randomness between parties. τ represents the transcript of the $(2k - 1)$ -move protocol executed between C and M . Parties C and M never share randomness prior to the protocol execution.

4. C outputs $z \leftarrow C(\tau, m_A, m_B)$.

Moreover, all parties may be randomized and may sample both private and shared randomness prior to Step (1) above, and C 's output depends on any privately sampled randomness or randomness shared between C and other parties. Here, by shared randomness, we mean that two (or more) parties agree on common random strings prior to Step (1) above.

See Figure 5 for a pictorial overview of the **SMH**[k]-model. Next we define the security and cost of **SMH**[k]-model protocols.

Definition 3.7 (Secure **SMH**[k] Protocol). We say that a **SMH**[k] protocol Π for a function f is (γ, ϵ) -secure if for all $x \in X$ and $y \in Y$, the following hold.

1. (Completeness) There exists an honest M such that $\Pr_{z \leftarrow \Pi(x, y)}[z = f(x, y)] \geq 1 - \gamma$.
2. (Soundness) For any dishonest helper M^* , it holds that $\Pr_{z \leftarrow \Pi(x, y; M^*)}[z \neq f(x, y)] \leq \epsilon$.

Here, the probability is taken over any private or shared randomness sampled by the parties. For brevity, we say that a $(1/3, 1/3)$ -secure **SMH**[k] protocol is simply a secure **SMH**[k] protocol.

Definition 3.8 (**SMH**[k] Protocol Cost). Let Π be a **SMH**[k] protocol for a function f . We define $\text{acost}(\Pi) := \max\{1, \log |m_A|\}$, $\text{bcost}(\Pi) := \max\{1, \log |m_B|\}$, and $\text{hcost}(\Pi) = \max\{1, \log |\tau|\}$ respectively as Alice's cost, Bob's cost, and the help cost of Π . Then, we let $C(\Pi) := \text{acost}(\Pi) + \text{bcost}(\Pi) + \text{hcost}(\Pi)$ denote the cost of Π . Moreover, we let $C_{\text{SMH}[k]}(f) := \min_{\Pi \in \text{SMH}[k](f)} \{C(\Pi)\}$ the **SMH**[k]-cost of f , where $\text{SMH}[k](f)$ denotes the set of all secure **SMH**[k] protocols for f .

Note that any **SMH**[k] protocol can be simulated by an **OMA**[k] protocol.

Lemma 3.9. Let Π be a (γ, ϵ) -secure **SMH**[k] protocol for a function f . Then there exists a (γ, ϵ) -secure **OMA**[k] protocol Π' for computing f .

Proof. Let Π be such an **SMH** $[k]$ protocol. Recall that by definition, this means for function $f: X \times Y \rightarrow Z$, the protocol Π has parties A, B, C , and M . Party A gets input $x \in X$, party B gets input $y \in Y$, and party M gets (x, y) as input. Parties M and C interact over k -rounds of communication, parties A and B may or may not share randomness with each other, with C , or with M .

We show how to simulate Π via an **OMA** $[k]$ protocol. Let Π' be a protocol in the **OMA** model with parties A', B' , and M' . These three parties will simulate the protocol Π as follows.

- Party B' will play the role of a “super party.” In particular, B' will simulate both parties B and C in its head. Any messages and/or randomness B sends to C in protocol Π is simulated in the head of B' .
- Party M' will interact with party B' exactly as party M interacts with party C in Π . Moreover, if B or C shares randomness with M in Π , B' shares this same randomness with M' .
- Party A' will exactly simulate party A by treating M' as M and by treating B' as both B and C .

Clearly, if Π had completeness error γ and soundness error ϵ , then Π' also has the same completeness and soundness errors. \square

3.3 SMH Model Lower Bounds

We introduce the **SMH** $[k]$ model in order to appropriately model cryptographic protocols in an information-theoretic setting where we can lower bound verification costs via communication complexity. Thus, we need to establish lower bounds on the communication complexity of certain functions in the **SMH** $[k]$ model. Key to our results is computing the inner product between two vectors in the **SMH** $[k]$ model, subject to certain restrictions on parties in the protocol (we will highlight why these restrictions are needed in [Section 5](#)). We prove the following theorem.

Theorem 3.10. *Let Π be any **SMH** $[k]$ protocol for $f(x, y) = \langle x, y \rangle$ and $x, y \in \mathbb{F}^n$ for any finite field \mathbb{F} , subject to the following constraints:*

- *Alice is deterministic;*
- *Bob, Charlie, and Merlin may sample private random coins; and*
- *Bob and Charlie can share randomness (hidden from Merlin).*

Then, $C(\Pi) = \Omega(n)$. In particular, Alice must send $\Omega(n)$ bits to Charlie.

Proof. We show the result for $\mathbb{F} := \{0, 1\}$. Since one can always embed the binary field into large finite fields, the lower bound applies for any finite field. Let x denote Alice’s input and y denote Bob’s input and suppose that Alice is deterministic. Then for any $x \neq x' \in \{0, 1\}^n$, there exists $y \in \{0, 1\}^n$ such that $\langle x, y \rangle \neq \langle x', y \rangle$, which can be constructed as follows. Suppose x and x' disagree at bit i ; i.e., $x_i \neq x'_i$. Without loss of generality, assume that $x_i = 1$ and $x'_i = 0$. Then define y as $y_i = 1$ and $y_j = 0$ for $j \neq i$. Clearly, $\langle x, y \rangle = 1$, while $\langle x', y \rangle = 0$.

Now suppose that Alice sends at most $(n - 1)$ bits to Charlie in protocol Π . Since there are 2^n possible strings x that Alice can receive as input and there are only 2^{n-1} possible $(n - 1)$ -bit messages m_A that Alice can send to Charlie, by the Pigeonhole Principle there must exist a pair of strings $x \neq x' \in \{0, 1\}^n$ such that $m_A(x) = m_A(x')$. That is, Alice sends the same message m_A on inputs x and x' . Without loss of generality, suppose that $x_i = 1$ and $x'_i = 0$ for some index $i \in [n]$. Now suppose Bob has input $y \in \{0, 1\}^n$ such that $\langle x, y \rangle = 1$ and $\langle x', y \rangle = 0$. By our previous argument above, such an input y must exist. Since Π is an **SMH** $[k]$ protocol, we know that there exists a Charlie-Merlin transcript τ such that

$$\Pr_{\Pi}[C(\tau, m_A(x), m_B(y)) = 1] \geq 2/3,$$

where the probability is taken over all random coins sampled by parties in the protocol. However, we know that $m_A(\mathbf{x}) = m_A(\mathbf{x}')$, so it must also hold that

$$\Pr_{\Pi}[C(\tau, m_A(\mathbf{x}'), m_B(\mathbf{y})) = 1] \geq 2/3.$$

This violates the soundness of Π since $\langle \mathbf{x}', \mathbf{y} \rangle = 0$, but Charlie outputs 1 with probability at least $2/3$. Hence, Alice must send at least n -bits of communication to Charlie, which implies $C(\Pi) = \Omega(n)$. \square

As a direct corollary, the above lower bound also holds for **OMA** $[k]$ protocols under the same restrictions.

Corollary 3.11. *Let Π be a secure **OMA** $[k]$ protocol for $f(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle$, where $\mathbf{x}, \mathbf{y} \in \mathbb{F}^n$, such that party Alice is deterministic and parties Bob and Merlin can use private randomness. Then, $C(\Pi) = \Omega(n)$.*

Proof. The proof is identical to that of [Theorem 3.10](#) with deterministic Alice and randomized Bob simulating parties B and C in the SMH protocol. \square

4 Hyrax: Information-Theoretic Core and Evidence of Optimal Verifier Time

In this section, we provide strong evidence that the Hyrax polynomial commitment scheme, which we denote by Hyrax-PCS, has optimal verifier time. We do this by first extracting the generalized scheme for proving knowledge of vector-matrix-vector product, which we call *Hyrax-VMV*. Then, we extract the information-theoretic core of Hyrax-VMV as a communication protocol; specifically, an **OMA** protocol. We then show that any **OMA** protocol computing the same function as Hyrax-VMV has certain lower bounds on the number of bits exchanged between parties, giving evidence that the Hyrax-VMV verifier must perform at least the same number of bit-operations. Finally, we instantiate Hyrax-VMV to obtain Hyrax-PCS and show that the communication lower bounds of the **OMA** protocol nearly match the verification costs of Hyrax-PCS.

To begin, we present the formal Hyrax-VMV protocol in [Algorithm 4.1](#) and give an overview of the protocol here. Hyrax-VMV allows a prover/committer to commit to a matrix $\mathbf{A} \in \mathbb{F}^{a \times b}$ and later prove that $\mathbf{y} = \mathbf{L}\mathbf{A}\mathbf{R}^\top$ for vectors $\mathbf{L} \in \mathbb{F}^a$ and $\mathbf{R} \in \mathbb{F}^b$ supplied by the verifier/receiver. For a group of prime order \mathbb{G} with scalar field \mathbb{F} , the public parameters of the scheme are randomly sampled group generators $\mathbf{g} := (g_1, \dots, g_b) \in \mathbb{G}^b$. To commit to the matrix \mathbf{A} , the prover computes vector $\mathbf{C} = \mathbf{A} \cdot \mathbf{g}^\top \in \mathbb{G}^a$; or, equivalently, $\mathbf{C}_i = \langle \mathbf{A}_{i,*}, \mathbf{g} \rangle$ for $i \in [a]$, where $\mathbf{A}_{i,*}$ denotes the i^{th} row of \mathbf{A} .

After receiving a commitment \mathbf{C} , the verifier sends vectors \mathbf{L}, \mathbf{R} .⁵ Given these vectors, the proof provided by the prover is a pair $(\mathbf{L}', \mathbf{y}) \in \mathbb{F}^a \times \mathbb{F}$, the purported values of $\mathbf{L}\mathbf{A}$ and $\mathbf{L}\mathbf{A}\mathbf{R}^\top$, respectively. Verification proceeds by using \mathbf{C} to check consistency between \mathbf{L} and \mathbf{L}' , and by using \mathbf{R} to check consistency between \mathbf{y} and \mathbf{L}' . The verifier samples $\alpha \xleftarrow{\$} \mathbb{F}$ and checks: (1) $\alpha \langle \mathbf{L}, \mathbf{C} \rangle \stackrel{?}{=} \alpha \langle \mathbf{L}', \mathbf{g} \rangle$; and (2) $\alpha \mathbf{y} \stackrel{?}{=} \alpha \langle \mathbf{L}', \mathbf{R} \rangle$.

Our goal is to extract the information-theoretic core of Hyrax-VMV as an **OMA** protocol and prove lower bounds on the communication needed in this protocol to compute the same function as [Algorithm 4.1](#). We work towards this goal and proving the following theorem (the proof of which can be found in [Section 4.1.1](#)).

Theorem 4.1. *Let $a, b \in \mathbb{N}$, let $\mathbf{L} \in \mathbb{F}^a$, $\mathbf{R} \in \mathbb{F}^b$, and $\mathbf{A} \in \mathbb{F}^{a \times b}$ be inputs to Hyrax-VMV. Then, there exists an **OMA** $[1]$ protocol Π which verifies the same function as Hyrax-VMV with inputs $\mathbf{L}, \mathbf{R}, \mathbf{A}$ such that $C(\Pi) = \Omega(\max(\sqrt{b}, \min(a, b)))$.*

As a direct corollary, we obtain strong evidence that Hyrax-PCS has (nearly) optimal verifier time.

⁵Note that the scheme in [\[WTS⁺18\]](#) is *not sound* if the vectors are given to the prover *before* he commits to his input.

Algorithm 4.1: Hyrax-VMV

- Public Parameters:** $a, b \in \mathbb{N}$, \mathbb{G} , $\mathbf{g} \in \mathbb{G}^b$.
Private Input: $\mathbf{A} \in \mathbb{F}^{a \times b}$.
- 1 **Commit Phase:** ($\mathbf{A} \in \mathbb{F}^{a \times b}$, $\mathbf{g} \in \mathbb{G}^b$)
 └ **Output:** $\mathbf{C} = \mathbf{A} \cdot \mathbf{g}^\top \in \mathbb{G}^a$.
 - 2 **Proof of Knowledge Phase:** ($\mathbf{C} \in \mathbb{G}^a$, $\mathbf{L} \in \mathbb{F}^a$, $\mathbf{R} \in \mathbb{F}^b$, $\mathbf{g} \in \mathbb{G}^b$; $\mathbf{A} \in \mathbb{F}^{a \times b}$).
 └ **Output:** Accept or reject.
 - 3 Prover sends $y = \mathbf{L} \cdot \mathbf{A} \cdot \mathbf{R}^\top \in \mathbb{F}$ and $\mathbf{L}' = \mathbf{L} \cdot \mathbf{A} \in \mathbb{F}^b$ to verifier.
 - 4 Verifier samples $\alpha \xleftarrow{\$} \mathbb{F}$ and checks if $\alpha \langle \mathbf{L}, \mathbf{C} \rangle \stackrel{?}{=} \alpha \langle \mathbf{L}', \mathbf{g} \rangle$ and $\alpha y \stackrel{?}{=} \alpha \langle \mathbf{L}', \mathbf{R} \rangle$.
 - 5 └ Verifier outputs accept if and only if both checks pass.
-

Corollary 4.2. *Let $N = 2^n$ and \mathbb{G} be a group with scalar field \mathbb{F} .*

- *For inputs $f: \{0, 1\}^n \rightarrow \mathbb{F}$, $\mathbf{u} \in \mathbb{F}^n$, and $\mathbf{g} \in \mathbb{G}^{\sqrt{N}}$ to Hyrax-PCS for multilinear polynomials, there exists an **OMA**[1] protocol Π which verifies the same function as Hyrax-PCS with inputs $f, \mathbf{u}, \mathbf{g}$ such that $C(\Pi) = \Omega(\sqrt{N})$.*
- *For inputs $f \in \mathbb{F}^{<N}[X]$, $u \in \mathbb{F}$, and $\mathbf{g} \in \mathbb{G}^{\sqrt{N}}$ to Hyrax-PCS for univariate polynomials, there exists an **OMA**[1] protocol Π which verifies the same function as Hyrax-PCS with inputs f, u, \mathbf{g} such that $C(\Pi) = \Omega(\sqrt{N})$.*

Moreover, the prescribed Hyrax-PCS verifier (for both multilinear and univariate polynomials) performs at most $O(\sqrt{N} \text{polylog}(|\mathbb{G}|))$ bit-operations, assuming a single \mathbb{G}/\mathbb{F} operation requires $\text{polylog}(|\mathbb{G}|)$ bit-operations.

Corollary 4.2 follows directly from Theorem 4.1 since Hyrax-PCS is an instantiation of Hyrax-VMV with $a = b = \sqrt{N}$ for some $N = 2^n$, structured vectors \mathbf{L}, \mathbf{R} , and structured matrix \mathbf{A} , and by noting that the proof size of Hyrax-VMV is $b + 1$ elements of \mathbb{F} and $2a$ elements of \mathbb{G} . Here, we make the implicit assumption that $|\mathbb{G}| \geq |\mathbb{F}|$. The full proof of Corollary 4.2 is given in Section 4.2.

We build up to the proof of Theorem 4.1 by first extracting the information-theoretic core of Hyrax-VMV as an OMA protocol. Given this information-theoretic core, we apply prior known lower bounds to the problems of vector-matrix-vector products in the OMA model. Then, we show how these lower bounds directly translate to a lower bound on the verifier work of Hyrax-VMV, completing the proof. Finally, we show how to instantiate Algorithm 4.1 to obtain Hyrax-PCS, directly giving Corollary 4.2.

4.1 Information Theoretic Core of Hyrax-VMV

Hyrax-VMV is a proof of knowledge for proving vector-matrix-vector products, so the information-theoretic core of Hyrax-VMV must also be a protocol for verifying vector-matrix-vector products. Crucial here is first understanding (a) how Hyrax-VMV verifies vector-matrix-vector products, and (b) what the cryptography (in this case, the discrete-log assumption) is guaranteeing in the protocol. Understanding (1) gives us a direct mapping to a potential communication protocol for Hyrax-VMV, and understanding (2) is crucial to specifying the protocol, in particular, what guarantees the cryptography gives in Hyrax-VMV must also be guaranteed in the underlying information-theoretic protocol.

Towards (1), examining Algorithm 4.1 shows us that in order to prove vector-matrix-vector products, Hyrax-VMV proves/verifies in parallel 2 claims: a vector-matrix product claim (i.e., $\mathbf{L}' = \mathbf{L}\mathbf{A}$) and an inner product claim ($y = \langle \mathbf{L}', \mathbf{R} \rangle$). This gives us a straightforward way of specifying Algorithm 4.1 as an OMA protocol; see Figure 1 in Section 1.2.3. At a high-level, Figure 1 captures Algorithm 4.1 as follows.

- Party M plays the role of the prover in the **Eval Phase** of [Algorithm 4.1](#).
- Party A plays the role of the committer in the **Commit Phase** of [Algorithm 4.1](#).
- Party B plays the role of the verifier of [Algorithm 4.1](#).

There are two key differences between [Figure 1](#) and [Algorithm 4.1](#). In [Figure 1](#):

1. Party A generates the “commitment” $C = \mathbf{A}\mathbf{r}^\top$ using private shared randomness $\mathbf{r} \in \mathbb{F}^b$;
2. Party A sends the “commitment” C after parties M and B execute the **Eval Phase** of the protocol.

First, item (2) is here due to the **OMA** model, which requires the interaction between parties B and M to be *independent* of the message sent by A and of B ’s input. Since B does not send anything to M , this second restriction is already satisfied. The first restriction is enforced by having A send their message after all interaction between B and M has been executed. For item (1), the vector $\mathbf{r} \xleftarrow{\$} \mathbb{F}^b$ is playing the role of the vector $\mathbf{g} \in \mathbb{G}^b$ of [Algorithm 4.1](#). Recall that \mathbf{g} is a randomly sampled vector of generators for \mathbb{G} . Notice that such a vector can be equivalently sampled as follows: (i) sample a single generator $g \in \mathbb{G}$; sample random $\mathbf{r} \xleftarrow{\$} \mathbb{F}^b$ (ii); and (iii) set generator $\mathbf{g}_i = \mathbf{r}_i \cdot g$. Since g is a generator, \mathbf{g}_i is also a generator for all i .

In [Algorithm 4.1](#), if the prover *knew* the vector $\mathbf{r} \in \mathbb{F}^b$, then it can easily break the binding of the commitment scheme. However, in the context of [Algorithm 4.1](#), the prover does not know \mathbf{r} due to the hardness of the discrete-log problem in \mathbb{G} . Thus, this vector \mathbf{r} is hidden from the prover in [Algorithm 4.1](#).

In the **OMA** model, all parties are computationally unbounded. If party M knew the vector \mathbf{r} that parties A and B share, then party M can trivially break the soundness of the protocol and convince B that $y' = \mathbf{L}\mathbf{A}\mathbf{r}^\top$ for some $y' \neq y$, where y is the real value of $\mathbf{L}\mathbf{A}\mathbf{r}^\top$. So the vector \mathbf{r} is hidden from party M by specifying it as the private randomness shared between A and B (before the interaction between M and B). In particular, we have answered (b) from the start of this section: the cryptography is forcing the prover in Hyrax-VMV to provide a correct random fingerprint of the matrix \mathbf{A} using the vector \mathbf{g} , otherwise the hardness of discrete-log is broken, or we can extract the prover’s true input \mathbf{A} due to the proof of knowledge property of Hyrax-VMV. Thus, in our **OMA** protocol for Hyrax-VMV, this invariant must also be enforced; this is handled by having party A (who is honest) send the random fingerprint of \mathbf{A} using randomness that is hidden from party M .

4.1.1 Hyrax-VMV Lower Bound

With the information-theoretic core of [Algorithm 4.1](#) established in [Figure 1](#), we now give a lower bound on the communication cost of [Figure 1](#), which we then use to obtain [Theorem 4.1](#).

Proof of [Theorem 4.1](#). As in [Definition 3.3](#), we let hcost denote the number of bits exchanged between M and B , and let vcost denote the number of bits A sends to B . Our goal is to minimize the quantity $\text{hcost} + \text{vcost}$. We do this by analyzing $\text{hcost} \cdot \text{vcost}$. However, we first establish that [Figure 1](#) is a valid (i.e., secure according to [Definition 3.2](#)) **OMA** protocol via the following claim.

Claim 4.3. *Let $a, b \in \mathbb{N}$ and let $f(\mathbf{A}, \mathbf{L}, \mathbf{R}) := \mathbf{L}\mathbf{A}\mathbf{R}^\top$ for $\mathbf{L} \in \mathbb{F}^a$, $\mathbf{R} \in \mathbb{F}^b$, and $\mathbf{A} \in \mathbb{F}^{a \times b}$ for any finite field \mathbb{F} . Then, [Figure 1](#) is a $(0, 2/|\mathbb{F}|)$ -secure **OMA**[1] protocol for computing f .*

Perfect completeness is immediate. For soundness, suppose that M wishes to convince B that $y^* = \mathbf{L}\mathbf{A}\mathbf{r}^\top$ for some $y^* \neq y := \mathbf{L}\mathbf{A}\mathbf{r}^\top$. Thus, M must send some \mathbf{L}^* such that (y^*, \mathbf{L}^*) pass the checks of B . If party B samples $\alpha = 0$, then all of their checks pass; this happens with exactly $1/|\mathbb{F}|$ probability. Suppose this is not the case. Then, $\mathbf{r} \xleftarrow{\$} \mathbb{F}^b$ is uniformly sampled and hidden from M , so the value $\langle \mathbf{L}^*, \alpha \rangle$ is uniformly distributed in \mathbb{F} . So the probability $\langle \mathbf{L}, \mathbf{C} \rangle = \langle \mathbf{L}^*, \mathbf{r} \rangle$ is exactly $1/|\mathbb{F}|$, which is required to pass B ’s check for $\alpha \neq 0$. If this check does pass, then B checks if $\alpha y^* \stackrel{?}{=} \alpha \langle \mathbf{L}^*, \mathbf{R} \rangle$. Note that M can guarantee this check passes

Algorithm 4.2: Hyrax-PCS

Public Parameters: $N = 2^n$, \mathbb{G} , $\mathbf{g} \in \mathbb{G}^{\sqrt{N}}$.
Private Input: $f: \{0, 1\}^n \rightarrow \mathbb{F}$ for multilinear PCS, $f \in \mathbb{F}^{<N}[X]$ for univariate PCS.

1 **Commit Phase:** ($f \in \mathbb{F}^N$, $\mathbf{g} \in \mathbb{G}^{\sqrt{N}}$)
Output: $\mathbf{C} \in \mathbb{G}^N$
2 Define matrix \mathbf{F} as follows:

- If multilinear PCS, $\mathbf{F}[i, j] = f((i-1) + \sqrt{N}(j-1))$ (i.e., evaluations of f) for $i, j \in [\sqrt{N}]$.
- If univariate PCS, $\mathbf{F}[i, j] = f_{(i-1) + \sqrt{N}(j-1)}$ (i.e., coefficients of f) for $i, j \in [\sqrt{N}]$.

Output $\mathbf{C} = \mathbf{F} \cdot \mathbf{g}^\top$.

3 **Eval Phase:** ($\mathbf{C} \in \mathbb{G}^{\sqrt{N}}$, $\mathbf{u} \in \mathbb{F}^n$ (for multilinear PCS) or $r \in \mathbb{F}$ (for univariate PCS), $\mathbf{g} \in \mathbb{G}^{\sqrt{N}}$; f)
Output: Verifier accepts or rejects.
4 If the PCS is multilinear, the prover and verifier compute \mathbf{L}, \mathbf{R} as

$$\mathbf{L} = \left(\chi^+(\mathbf{u}, 0), \chi^+(\mathbf{u}, 1), \dots, \chi^+(\mathbf{u}, \sqrt{N}-1) \right) \quad \mathbf{R} = \left(\chi^-(\mathbf{u}, 0), \chi^-(\mathbf{u}, 1), \dots, \chi^-(\mathbf{u}, \sqrt{N}-1) \right),$$

5 where $\chi(X, Y) := X_i Y_i + (1 - X_i)(1 - Y_i)$ and

$$\chi^+(\mathbf{u}, \mathbf{c}) := \prod_{k=1}^{n/2} \chi(u_k, c_k) \quad \chi^-(\mathbf{u}, \mathbf{c}) := \prod_{k=1}^{n/2} \chi(u_{k+n/2}, c_k).$$

6 If the PCS is univariate, the prover and verifier compute \mathbf{L}, \mathbf{R} as :

$$\mathbf{L} = \left(1, r^{\sqrt{N}}, r^{2\sqrt{N}}, \dots, r^{(\sqrt{N}-1)\sqrt{N}} \right) \quad \mathbf{R} = \left(1, r, r^2, \dots, r^{\sqrt{N}-1} \right).$$

7 The prover additionally computes $\mathbf{F} \in \mathbb{F}^{\sqrt{N} \times \sqrt{N}}$ as in the commit phase.
8 The prover computes and sends $\mathbf{y} = \mathbf{L} \cdot \mathbf{F} \cdot \mathbf{R}^\top \in \mathbb{F}$ and $\mathbf{L}' = \mathbf{L} \cdot \mathbf{F} \in \mathbb{F}^{\sqrt{N}}$.
9 The verifier samples $\alpha \xleftarrow{\$} \mathbb{F}$ and checks the following two equations:

$$\alpha \langle \mathbf{L}, \mathbf{C} \rangle \stackrel{?}{=} \alpha \langle \mathbf{L}', \mathbf{g} \rangle \quad \alpha \mathbf{y} \stackrel{?}{=} \alpha \langle \mathbf{L}', \mathbf{R} \rangle.$$

10 The verifier accepts if and only if both checks pass.

by setting \mathbf{L}^* such that $\mathbf{y}^* = \mathbf{L}^* \mathbf{R}^\top$, so in this case the check passes if the above check also passes. Otherwise, this check fails for $\alpha \neq 0$. So the total soundness error is upper bounded by $2/|\mathbb{F}|$, establishing [Claim 4.3](#).

We now turn towards minimizing $\text{hcost} + \text{vcost}$. Notice that [Figure 1](#) is an **OMA** protocol for simultaneously verifying one matrix-vector product and $(\mathbf{L}' = \mathbf{L}\mathbf{A})$, and one inner product ($\mathbf{y} = \langle \mathbf{L}', \mathbf{R} \rangle$). In our protocol, B verifies the matrix-vector product $\mathbf{L}\mathbf{A}$ and the inner product $\langle \mathbf{L}', \mathbf{R} \rangle$, where $\mathbf{L} \in \mathbb{F}^a$ and $\mathbf{L}', \mathbf{R} \in \mathbb{F}^b$. To verify the matrix-vector product, [Theorem 3.4](#) tells us that $\text{vcost} \cdot \text{hcost} = \Omega(\min(a, b)^2)$, whereas to verify the inner product, [Theorem 3.5](#) tells us that $\text{vcost} \cdot \text{hcost} = \Omega(b)$. Combining these bounds, we have a lower bound of $\text{hcost} \cdot \text{vcost} = \Omega(\max(b, \min(a, b)^2))$.

Let $\ell = \max(b, \min(a, b)^2)$. Now, we minimize $\text{hcost} + \text{vcost}$ subject to the constraint that $\text{hcost} \cdot \text{vcost} = \Omega(\ell)$. This constraint implies $\text{vcost} \geq \ell / \text{hcost}$, which implies $\text{hcost} + \text{vcost} \geq \text{hcost} + \ell / \text{hcost}$. Taking the derivative of the right-hand side with respect to hcost , we see that the minimum value is obtained at $\text{hcost} = \sqrt{\ell}$, and thus $\text{hcost} + \text{vcost}$ is minimized at $\Omega(\sqrt{\ell}) = \Omega(\max(\sqrt{b}, \min(a, b)))$. This implies that party B must receive at least $\Omega(\max(\sqrt{b}, \min(a, b)))$ bits in any **OMA** protocol which simultaneously verifies a matrix-vector product and an inner product. \square

4.2 Evidence that Hyrax-PCS has Optimal Verifier Time

With the Hyrax-PCS presented in [Algorithm 4.2](#), we now show [Corollary 4.2](#), giving us strong evidence that Hyrax-PCS (for both multilinear and univariate polynomials) has optimal verifier time (up to polylog $|\mathbb{G}|$ factors). In particular, this entails showing how to build Hyrax-PCS from Hyrax-VMV, then applying our lower bound in conjunction with analyzing the work done by the verifier in Hyrax-PCS.

Proof of Corollary 4.2. Hyrax-PCS is simply an instantiation of Hyrax-VMV with more structured matrices and vectors. We begin with Hyrax-PCS for multilinear polynomials; the construction for univariate polynomials is similar and given at the end of the proof. In the multilinear PCS setting, the prover has a function $f: \{0, 1\}^n \rightarrow \mathbb{F}$ such that n is even. Setting $N = 2^n$, the prover then defines a $\sqrt{N} \times \sqrt{N}$ matrix \mathbf{F} representing f as $\mathbf{F}[i, j] := f((i-1) + \sqrt{N} \cdot (j-1))$ for $i, j \in [\sqrt{N}]$. Here, we interpret $(i-1) + \sqrt{N}(j-1)$ as an n -bit string in the natural way. So to commit to (the multilinear extension of) f , the prover outputs $\mathbf{C} \leftarrow \text{Hyrax-VMV.Commit}(\mathbf{F}, \mathbf{g})$.

Now for the evaluation proofs, let $\mathbf{u} \in \mathbb{F}^n$ be the evaluation point specified by the verifier after being provided some commitment \mathbf{C} . Let $\chi(X, Y) := X_i Y_i + (1 - X_i)(1 - Y_i)$ (e.g., see [Eq. \(2\)](#)). Then, for any $\mathbf{c} \in \{0, 1\}^{n/2}$, the verifier defines

$$\chi^+(\mathbf{u}, \mathbf{c}) := \prod_{k=1}^{n/2} \chi(u_k, c_k) \quad \chi^-(\mathbf{u}, \mathbf{c}) := \prod_{k=1}^{n/2} \chi(u_{k+n/2}, c_k).$$

Finally, vectors \mathbf{L}, \mathbf{R} are defined as $\mathbf{L} = (\chi^+(\mathbf{u}, i))_{i=0}^{\sqrt{N}-1}$ and $\mathbf{R} = (\chi^-(\mathbf{u}, i))_{i=0}^{\sqrt{N}-1}$, where we interpret i as a vector in $\{0, 1\}^{n/2}$ in the natural way. The prover and verifier then run $\text{Hyrax-VMV.Eval}(\mathbf{C}, \mathbf{L}, \mathbf{R}, \mathbf{g}; \mathbf{F})$ with the above definitions. From Hyrax [\[WTS⁺18\]](#), it can be shown that for \mathbf{L}, \mathbf{R} , and \mathbf{A} defined this way, we have $\mathbf{L}\mathbf{A}\mathbf{R}^\top = \tilde{f}(\mathbf{u})$, where \tilde{f} is the multilinear extension of f (defined in [Eq. \(1\)](#)).


Now, translating Hyrax-PCS to the **OMA** model is nearly identical to Hyrax-VMV. The only changes are: (1) we specify $a = b = \sqrt{N}$ for $N = 2^n$ and even n ; and (2) how the vectors \mathbf{L}, \mathbf{R} are sampled by the verifier. These two changes are reflected by the following minor modifications to [Figure 1](#).

1. Merlin's input is (\mathbf{F}, \mathbf{u}) , Bob's input is $\mathbf{u} \in \mathbb{F}^n$, and Alice's input is $\mathbf{F} \in \mathbb{F}^{\sqrt{N} \times \sqrt{N}}$. Notice that $f: \{0, 1\}^n \rightarrow \mathbb{F}$ uniquely defines \mathbf{F} and vice versa.
2. Merlin, during step (1) of the protocol, will appropriately define \mathbf{L} and \mathbf{R} as in Hyrax-PCS (described above).
3. Bob, during step (3) of the protocol, will also appropriately define \mathbf{L}, \mathbf{R} as in Hyrax-PCS.

In particular, the messages sent by both Alice and Merlin to Bob remain the same; they just now all have length $\Theta(\sqrt{N})$. Moreover, this is clearly just an instance of [Figure 1](#) with $a = b = \sqrt{N}$, and has the same soundness error as [Figure 1](#) ([Claim 4.3](#)) since each entry of \mathbf{u} is uniformly random.

By [Theorem 4.1](#), this implies a lower bound of $\Omega(\sqrt{N})$ for the **OMA** protocol representing Hyrax-PCS. Notice that in Hyrax-PCS, the verifier computes 3 inner products of size \sqrt{N} ; two are inner products between a \mathbb{G} and a \mathbb{F} vector, and one is an inner product between two \mathbb{F} vectors. This implies at most $O(\sqrt{N} \text{polylog } |\mathbb{G}|)$ bit-operations for these computations, assuming \mathbb{G}/\mathbb{F} operations can be done in polylog $|\mathbb{G}|$ bit-operations.

For univariate polynomials, the above argument is identical; we simply need to specify how the prover constructs the matrix \mathbf{A} from a univariate polynomial and how the verifier constructs \mathbf{L}, \mathbf{R} from an evaluation point. Let $p \in \mathbb{F}[X]$ be a univariate polynomial such that $\deg(p) < N$; write $p(X) = p_0 + p_1 X + \dots + p_{N-1} X^{N-1}$. Then, the prover defines a $\sqrt{N} \times \sqrt{N}$ matrix \mathbf{P} representing p as $\mathbf{P}[i, j] = p_{(i-1) + \sqrt{N}(j-1)}$ for $i, j \in [\sqrt{N}]$. So to commit to p , the prover outputs $\mathbf{C} \leftarrow \text{Hyrax-VMV.Commit}(\mathbf{P}, \mathbf{g})$.



(3) C does the following.

- Sample $\alpha \xleftarrow{\$} \mathbb{F}$.
- Check $\alpha \langle L, C \rangle \stackrel{?}{=} \alpha \langle L', r \rangle$ and $\alpha y \stackrel{?}{=} \alpha \langle L', R \rangle$.
- Output accept (i.e., output y) if and only if all checks pass.

Figure 6: Hyrax-VMV as a **SMH**[1] protocol. The red dashed line indicates private randomness shared between the respective parties. The protocol executes (1), followed by (2), followed by (3).

For the evaluation proofs, let $r \in \mathbb{F}$ be the evaluation point specified by the verifier after receiving some commitment C . The verifier defines $R = (1, r, r^2, \dots, r^{\sqrt{N}-1})$ and $L = (1, r^{\sqrt{N}}, r^{2\sqrt{N}}, \dots, r^{(\sqrt{N}-1)\sqrt{N}})$. Then, the prover and verifier run $\text{Hyrax-VMV.Eval}(C, L, R, g; \mathbf{P})$ with the above specifications. One can readily verify that $p(r) = LPR^T$ under these definitions of L, R, \mathbf{P} . The rest of the proof follows identically as in the case of multilinear Hyrax-PCS, except that the final **OMA** protocol for the univariate Hyrax-PCS has soundness error at most $n/|\mathbb{F}|$ due to using powers of r instead of uniformly random vectors. \square

4.3 Hyrax-VMV as an SMH Protocol

In this section, we reformulate Hyrax-VMV as a **SMH**[1] protocol, mostly for completeness and to showcase that the protocol is essentially identical to Figure 1. However, we believe this to be a useful exercise to motivate our **SMH** model, showcasing that the **OMA** model is not special for Hyrax.

The protocol, presented in Figure 6, is nearly identical to Figure 1. The only change now is that we have parties A, B , and C , along with helper M . We now give vectors L, R as input to B . Moreover, the protocol is simple: A and C share private random string $r \in \mathbb{F}^b$, M sends the same messages to C as in the **OMA** protocol, and now we simply have B send the vectors L, R . In other words, B is deterministic; moreover, C simultaneously verifies both a vector-matrix product and an inner product. The inner product claim has size b , so Theorem 3.10 implies that B must send $\Omega(b)$ bits to C . For $b = a = \sqrt{N}$, this matches the **OMA** lower bound we obtain for Hyrax-PCS.

5 Bulletproofs: Information-Theoretic Core and Evidence of Optimal Verifier Time

We now turn towards our second main result: giving strong evidence that the Bulletproofs' verification time is *inherently linear*. That is, the underlying information theoretic core of Bulletproofs requires a linear amount

of communication. We prove the following theorem in [Section 5.2.1](#).

Theorem 5.1. *Let \mathbb{G} be a prime-order group with scalar field \mathbb{F} , and let $N = 2^n$. Then, for private prover input $\mathbf{w} \in \mathbb{F}^N$ and public parameters $\mathbf{g} \in \mathbb{G}^N$ for the Bulletproofs proof of knowledge for opening a Pedersen vector commitment over \mathbb{G} , there exists a **SMH**[n] protocol Π which verifies the same function as Bulletproofs with inputs \mathbf{w}, \mathbf{g} such that $C(\Pi) = \Omega(N)$.*

As a direct corollary, we obtain strong evidence that Bulletproofs-PCS has (nearly) optimal verifier time.

Corollary 5.2. *Let $N = 2^n$, \mathbb{G} be a finite group with scalar field \mathbb{F} .*

- *For inputs $f: \{0, 1\}^n \rightarrow \mathbb{F}$, $\mathbf{u} \in \mathbb{F}^n$, and $\mathbf{g} \in \mathbb{G}^N$ to Bulletproofs-PCS for multilinear polynomials, there exists a **SMH**[n] protocol Π which verifies the same function as Bulletproofs-PCS with inputs $f, \mathbf{u}, \mathbf{g}$ such that $C(\Pi) = \Omega(N)$.*
- *For inputs $f \in \mathbb{F}^{<N}[X]$, $u \in \mathbb{F}$, and $\mathbf{g} \in \mathbb{G}^N$ to Bulletproofs-PCS for univariate polynomials, there exists a **SMH**[n] protocol Π which verifies the same function as Bulletproofs-PCS for inputs f, u, \mathbf{g} such that $C(\Pi) = \Omega(N)$.*

Moreover, the prescribed Bulletproofs-PCS verifier (for both multilinear and univariate polynomials) performs at most $O(N \text{ polylog } |\mathbb{G}|)$ bit-operations, assuming a single \mathbb{G}/\mathbb{F} operation requires $\text{polylog } |\mathbb{G}|$ bit-operations.

We prove [Corollary 5.2](#) in [Section 5.3](#). The rest of this section is dedicated to showing this lower bound. First, we cast Bulletproofs as an instantiation of the sum-check protocol. Then, we capture the information-theoretic core of Bulletproofs as a communication protocol. Finally, we show a linear communication lower bound in this information-theoretic core, giving strong evidence that the Bulletproofs verifier is optimal.

5.1 Bulletproofs is Sum-Check (Revisited)

Let $\mathbf{g} \in \mathbb{G}^N$ be (random) group generators and $\mathbf{w} \in \mathbb{F}^N$, where \mathbb{F} is the scalar field of \mathbb{G} , and $N = 2^n$. At a high-level, Bulletproofs is a proof of knowledge of an opening of the Pedersen commitment $C = \langle \mathbf{g}, \mathbf{w} \rangle$; i.e., the prover “knows” a vector \mathbf{w} (hidden from the verifier) that is within the Pedersen commitment C under public parameters \mathbf{g} . It utilizes a “split-and-fold” recursive technique to continuously reduce the claim $C = \langle \mathbf{g}, \mathbf{w} \rangle$ for length N vectors \mathbf{g}, \mathbf{w} to a new claim $C' = \langle \mathbf{g}', \mathbf{w}' \rangle$ for length $N/2$ vectors \mathbf{g}', \mathbf{w}' , until the final vectors are of length 1 and the prover simply sends the final “folded” witness. Here, \mathbf{g}' and \mathbf{w}' are a random linear combination of the left and right halves of the (respective) vectors \mathbf{g} and \mathbf{w} , where the randomness is given by the verifier in the protocol.

It was originally observed by Bootle et al. [[BCS21](#)] that the Bulletproofs PoK is simply an instantiation of the sum-check protocol [[LFK⁺92](#)] for proving the inner product of two vectors \mathbf{w}, \mathbf{g} . The sum-check proves $C = \sum_{\mathbf{c} \in \{0,1\}^n} \tilde{\mathbf{g}}(\mathbf{c}) \cdot \tilde{\mathbf{w}}(\mathbf{c})$, where $\tilde{\mathbf{w}}$ and $\tilde{\mathbf{g}}$ are appropriately defined multilinear extensions of \mathbf{w} and \mathbf{g} , respectively. [[BCS21](#)] give a different definition of $\tilde{\mathbf{w}}$ and $\tilde{\mathbf{g}}$ than what is “standard” (e.g., [[LFK⁺92](#)]). However, we are interested in presenting Bulletproofs as a sum-check protocol using the “standard” definition of multilinear extensions; i.e., [Eqs. \(1\) and \(2\)](#).

5.1.1 Bulletproofs as a Standard Sum-Check Protocol

Let $\mathbf{g} \in \mathbb{G}^N$ be the public parameters of the Bulletproofs protocol and let $\mathbf{w} \in \mathbb{F}^N$ be the private input (the witness) of the Bulletproofs prover. Now, the Bulletproofs protocol precisely corresponds to a sum-check over the polynomial $F(X) = \tilde{\mathbf{g}}(X) \cdot \tilde{\mathbf{w}}(X)$. The initial commitment to the vector \mathbf{w} under generators \mathbf{g} is

Algorithm 5.1: BP-Sum-Check

Public Params: $N = 2^n$, \mathbb{G} with scalar field \mathbb{F} , $\mathbf{g} \in \mathbb{G}^N$.
Private Input: $\mathbf{w} \in \mathbb{F}^N$.
1 **Commit Phase:** ($\mathbf{g} \in \mathbb{G}^N, \mathbf{w} \in \mathbb{F}^N$)
 Output: $C = \langle \mathbf{g}, \mathbf{w} \rangle \in \mathbb{G}$
2 **Proof of Knowledge Phase:** ($C_1 \in \mathbb{G}, \mathbf{g} \in \mathbb{G}^N; \mathbf{w} \in \mathbb{F}^N$)
 Output: Accept or reject.
3 The prover defines n -variate polynomial $F(X) := \tilde{\mathbf{g}}(X) \cdot \tilde{\mathbf{w}}(X) \in \mathbb{G}[X]$.
4 **for** $i \in [n]$ **do**
5 The prover defines and sends polynomial $q_i(X_i): \mathbb{F} \rightarrow \mathbb{G}$ as

$$q_i(X_i) := \sum_{\mathbf{c} \in \{0,1\}^{n-i}} F(\alpha_1, \dots, \alpha_{i-1}, X_i, \mathbf{c}).$$

6 The verifier checks $C_i \stackrel{?}{=} q_i(0) + q_i(1)$ and $\deg(q_i) \leq 2$; reject if not.
7 The verifier samples and sends $\alpha_i \xleftarrow{\$} \mathbb{F}$, and defines $C_{i+1} := q_i(\alpha_i)$.
8 The prover sends $w = \tilde{\mathbf{w}}(\alpha)$ for $\alpha \in \mathbb{F}^n$.
9 The verifier computes $g = \tilde{\mathbf{g}}(\alpha)$. The verifier checks $C_{n+1} \stackrel{?}{=} g \cdot w$ and rejects if not equal;
 otherwise the verifier outputs accept.

$C = \sum_{\mathbf{c}} F(\mathbf{c}) = \langle \mathbf{g}, \mathbf{w} \rangle$, which follows by the definition of multilinear extensions: $F(\mathbf{c}) = \tilde{\mathbf{g}}(\mathbf{c}) \cdot \tilde{\mathbf{w}}(\mathbf{c}) = \mathbf{g}_{\mathbf{c}} \cdot \mathbf{w}_{\mathbf{c}}$, where we index the vectors \mathbf{g} and \mathbf{w} using the binary vector $\mathbf{c} \in \{0,1\}^n$ in the natural way. Then, proving knowledge of \mathbf{w} such that $C = \langle \mathbf{g}, \mathbf{w} \rangle$ amounts to running a sum-check protocol using the polynomial F defined above. We present this formal sum-check protocol in [Algorithm 5.1](#).

Remark 5.3 (Sum-Check with Knowledge Soundness). Bootle et al. [BCS21] explore which properties needed by instances of the sum-check protocol in order to imply knowledge soundness. We do not explore this further in this work, and we make no claims on the connection between soundness in our communication protocols and (knowledge) soundness in our cryptographic protocols.

Remark 5.4 (Optimizing the Proof Size and Verifier Checks). During each round of the sum-check, the prover sends 3 group elements to specify the polynomial q_i . This can be optimized using a standard sum-check optimization: (1) the prover sends $S_i = \{q_i(0), q_i(2)\}$; (2) the verifier sets $q_i(1) := C_i - q_i(0)$ and interpolates $\{q_i(0), q_i(1), q_i(2)\}$ to obtain $q_i(X_i)$; and (3) the verifier sets $C_{i+1} = q_i(\alpha_i)$. This saves one group element per round of the sum-check, and the verifier does not have to check consistency every round. Furthermore, this protocol has the same soundness guarantees as the standard sum-check, and, in fact, sends the same number of group elements per round as the original Bulletproofs protocol.

5.2 Information-Theoretic Core of Bulletproofs

Given Bulletproofs is a standard sum-check protocol, we now turn to distilling the information-theoretic core of Bulletproofs as a communication protocol. Indeed, one motivation for recapturing Bulletproofs as a (standard) sum-check is due to the prevalence of sum-check in interactive proofs, SNARK design, and communication complexity.

In [Section 3](#), we discussed two flavors of communication protocols: the **OMA** model and our new **SMH** model. To capture the IT core of Bulletproofs, we use our new **SMH** model; that is, we'll present

an analogue of [Algorithm 5.1](#) as a **SMH** protocol. Note that while it is possible to capture some form of the information-theoretic core of Bulletproofs as a **OMA** model protocol, we do not believe this framing meaningfully captured the properties of Bulletproofs (see [Section 5.4](#)).

5.2.1 Bulletproofs as a SMH Protocol

Recall that a **SMH** $[k]$ communication protocol ([Section 3.2](#)) involves 4 parties: A , B , C , and M , where A and B communicate “one-way” with C , and C and M execute k -round (or $(2k - 1)$ -move) interactive proof. We present the formal **SMH** protocol for [Algorithm 5.1](#) in [Figure 2](#). Roughly speaking, the 4 parties in our protocol operate as follows.

- Party A receives input $\mathbf{r} \in \mathbb{F}^N$, representing the scalars of the group generators. Moreover, A will be *deterministic*. Intuitively, A represents the public parameters/setup of Bulletproofs.
- Party B receives input $\mathbf{w} \in \mathbb{F}^N$, representing the input to the prover of the cryptographic protocol.
- Party M receives both (\mathbf{r}, \mathbf{w}) as input and plays the role of the interactive prover from **Eval Phase** of [Algorithm 5.1](#).
- Party C will play the role of the cryptographic verifier. It will run a n -round sum-check protocol with M , where M tries to convince C of the claim $v = \langle \mathbf{r}, \mathbf{w} \rangle$. Moreover, C and B will share private randomness $\alpha \xleftarrow{\$} \mathbb{F}^n$ for $N = 2^n$, which is used in the sum-check.
- At the end of this sum-check, A sends their entire vector \mathbf{r} to C . B sends the value $w = \tilde{\mathbf{w}}(\alpha)$ to C .
- C then performs all the appropriate sum-check verification checks, outputting accept or reject.

Notably, in this information theoretic setting, (1) the vector \mathbf{r} , which intuitively represents the “hidden” vector in the group generators, is given in the clear to both C and M ; and (2) C is guaranteed the correct evaluation $\tilde{\mathbf{w}}(\alpha)$ from B . (2) captures exactly what happens in the cryptographic setting: in Bulletproofs, the prover sends the value w claimed to be $\tilde{\mathbf{w}}(\alpha)$. Now, by knowledge soundness and hardness of the discrete-log assumption, w will be the correct value, otherwise either the prover has broken the discrete-log assumption, or we can extract \mathbf{w} from the transcript of the interaction. For (1), it would be impossible for C and M to execute the sum-check protocol without M knowing the vector \mathbf{r} . Similarly, C could not compute $\tilde{\mathbf{r}}(\alpha)$ without the vector \mathbf{r} (moreover, A cannot send $\tilde{\mathbf{r}}(\alpha)$ since A does not receive α).

With our **SMH** $[n]$ protocol ([Figure 2](#)) specified for $n = \log(N)$, we now prove [Theorem 5.1](#).

Proof of Theorem 5.1. We begin by establishing the security of [Figure 2](#); namely, it is a valid **SMH** protocol.

Claim 5.5. *Figure 2 is a $(0, 2n/|\mathbb{F}|)$ -secure **SMH** $[n]$ protocol for N -sized inner products.*

This claim immediately follows from the completeness and soundness of the sum-check protocol. In particular, C always receives the correct $w = \tilde{\mathbf{w}}(\alpha)$, and will always be able to compute the correct $r = \tilde{\mathbf{r}}(\alpha)$, which gives C access to the correct value $F(\alpha)$ at the end of the sum-check. Since F has n variables, each with degree at most 2, the soundness error of the sum-check is at most $2n/|\mathbb{F}|$.

With the claim established, [Figure 2](#) is a valid **SMH** $[n]$ protocol for verifying the inner product of A ’s and B ’s inputs. By construction, party A is deterministic, so [Theorem 3.10](#) tells us that for soundness to hold, party A must send $\Omega(N)$ bits of communication, where N is the length of A ’s input. \square

5.3 Evidence that Bulletproofs-PCS has Optimal Verifier Time

Bulletproofs is readily adaptable into a polynomial commitment scheme for any polynomials which can express evaluation as an inner product; in particular, multilinear and univariate polynomials satisfy this

property. In fact, to obtain a PCS from Bulletproofs, we first give a more general protocol for proving an opening to a Pedersen commitment, and that the inner product of this opening and some public vector is some value. This allows us to prove [Corollary 5.2](#).

Proof of [Corollary 5.2](#). We modify [Algorithm 5.1](#) to obtain a new scheme which satisfies the above property; we present this new algorithm in [Algorithm 5.2](#). As with [Algorithm 5.1](#), we let $\mathbf{w} \in \mathbb{F}^N$ be the prover's input and $\mathbf{g} \in \mathbb{G}^N$ be the public parameters of the scheme. The commitment phase remains the same: the prover provides $C = \langle \mathbf{g}, \mathbf{w} \rangle$. Now, we modify the proof of knowledge phase as follows. First, there is an additional *public input*: $\mathbf{y} \in \mathbb{F}^N$, provided by the verifier, and $v \in \mathbb{F}$, provided by the prover. Now, for the proof of knowledge, the prover now proves it knows \mathbf{w} such that $C = \langle \mathbf{g}, \mathbf{w} \rangle$ and $v = \langle \mathbf{y}, \mathbf{w} \rangle \in \mathbb{F}$. This is actually simple to handle: another instantiation of the sum-check protocol is added in parallel, one for the polynomial $H(X) := \tilde{\mathbf{y}}(X) \cdot \tilde{\mathbf{w}}(X)$.

Translating [Algorithm 5.2](#) to an **SMH** protocol, we augment [Figure 2](#) by additionally giving party A the input \mathbf{y} and by having party M send the claimed value $v = \langle \mathbf{y}, \mathbf{w} \rangle$ as part of the first message during the sum-check protocol and run an additional sum-check in parallel for this claim. Moreover, we also specify the function f computed by the **SMH** $[n]$ protocol to be $f(\mathbf{w}, \mathbf{y}, \mathbf{r}) = (\langle \mathbf{w}, \mathbf{y} \rangle, \langle \mathbf{w}, \mathbf{r} \rangle)$. Then, party A also sends \mathbf{y} along with \mathbf{r} as her message, and party C additionally augments their checks in line with the checks performed by the verifier in [Algorithm 5.2](#). Clearly, our lower bound still remains for this protocol as now the protocol verifies two inner products in parallel.

To obtain a polynomial commitment scheme, we impose structure on the vectors \mathbf{w} and \mathbf{y} such that if p is the polynomial in question, then $p(x) = \langle \mathbf{w}, \mathbf{y} \rangle$. We first show how to instantiate [Algorithm 5.2](#) to obtain a multilinear polynomial commitment scheme. Let $p: \{0, 1\}^n \rightarrow \mathbb{F}$ be a function and \tilde{p} be its multilinear extension. Then, set $\mathbf{w} = (p(0), p(1), \dots, p(2^n - 1))$, where we interpret $i \in \{0, 1, \dots, 2^n - 1\}$ as an n -bit string in the natural way. Simply put, \mathbf{w} is the list of evaluations of p (and thus \tilde{p}) on the Boolean hypercube.

Now, if we wish to evaluate \tilde{p} at some point $\mathbf{x} \in \mathbb{F}^n$, we can express this as an inner product between the above \mathbf{w} and $\mathbf{y} \in \mathbb{F}^N$ defined as $y_i := \text{eq}(\mathbf{x}, i - 1)$, where eq is defined in [Eq. \(2\)](#), and we again interpret $i - 1$ as an n -bit string in the natural way. Clearly, by [Eq. \(1\)](#), we have that $\tilde{p}(\mathbf{x}) = \langle \mathbf{w}, \mathbf{y} \rangle$ for our defined \mathbf{w} and \mathbf{y} . Therefore, [Algorithm 5.2](#) for \mathbf{w} and \mathbf{y} defined this way for the prover's input p and the verifier's challenge \mathbf{x} yields a polynomial commitment scheme for multilinear polynomials.

Now, if p is a univariate polynomial of degree at most $N - 1$, we write $\mathbf{w} = (p_0, p_1, \dots, p_{N-1})$, where p_i is the i^{th} coefficient of p . For any evaluation point x , we write $\mathbf{y} = (1, x, x^2, \dots, x^{N-1})$. Now clearly, we have $p(x) = \langle \mathbf{w}, \mathbf{y} \rangle$. Doing this transformation in [Algorithm 5.2](#) readily yields a polynomial commitment scheme for univariate polynomials. \square

5.4 Why the OMA Model does not Suffice for Bulletproofs

While it is possible to frame Bulletproofs as a **OMA** protocol (e.g., one can apply [Lemma 3.9](#) to [Figure 2](#)), we give our reasoning for why we do not believe this is the right model for capturing Bulletproofs. As with the **SMH** protocol for Bulletproofs, we can reimagine Bulletproofs as a **OMA** $[n]$ protocol with parties Alice, Bob, and Merlin such that Merlin and Bob will interact, followed by Alice sending a single message to Bob. In this **OMA** protocol, Bob and Merlin execute the same n -round sum-check as in [Figure 2](#), and Bob will perform all the same verification checks as C .

The main question with framing Bulletproofs as an **OMA** protocol is the following: which party gets which input? First, suppose that Alice is given the input \mathbf{w} and Bob is given the input \mathbf{r} . If Alice is deterministic and does not share randomness with Bob, then we obtain our same lower bound as in [Figure 2](#)

Algorithm 5.2: BP-Sum-Check-2

Public Params: $N = 2^n$, \mathbb{G} with scalar field \mathbb{F} , $\mathbf{g} \in \mathbb{G}^N$.
Private Input: $\mathbf{w} \in \mathbb{F}^N$.

- 1 **Commit Phase:** ($\mathbf{g} \in \mathbb{G}^N, \mathbf{w} \in \mathbb{F}^N$)
 Output: $C = \langle \mathbf{g}, \mathbf{w} \rangle \in \mathbb{G}$
- 2 **Proof of Knowledge:** ($C_1 \in \mathbb{G}, \mathbf{g} \in \mathbb{G}^N, \mathbf{y} \in \mathbb{F}^N, v_1 \in \mathbb{F}; \mathbf{w} \in \mathbb{F}^N$)
 Output: Accept or reject.
- 3 The prover defines n -variate polynomials $F(X) := \tilde{\mathbf{g}}(X) \cdot \tilde{\mathbf{w}}(X) \in \mathbb{G}[X]$ and
 $H(X) := \tilde{\mathbf{y}}(X) \cdot \tilde{\mathbf{w}}(X) \in \mathbb{F}[X]$.
- 4 **for** $i \in [n]$ **do**
- 5 The prover defines and sends polynomials $q_i(X_i) : \mathbb{F} \rightarrow \mathbb{G}$ and $p_i(X_i) : \mathbb{F} \rightarrow \mathbb{F}$ as
$$q_i(X_i) := \sum_{\mathbf{c} \in \{0,1\}^{n-i}} F(\alpha_1, \dots, \alpha_{i-1}, X_i, \mathbf{c}); \quad p_i(X_i) := \sum_{\mathbf{c} \in \{0,1\}^{n-i}} H(\alpha_1, \dots, \alpha_{i-1}, X_i, \mathbf{c}).$$
- 6 The verifier checks $C_i \stackrel{?}{=} q_i(0) + q_i(1)$, $v_i \stackrel{?}{=} p_i(0) + p_i(1)$, $\deg(q_i) \leq 2$, and $\deg(p_i) \leq 2$;
 reject if not.
- 7 The verifier samples and sends $\alpha_i \xleftarrow{\$} \mathbb{F}$, and defines $C_{i+1} := q_i(\alpha_i)$ and $v_{i+1} = q_{i+1}(\alpha_i)$.
- 8 The prover sends $\mathbf{w} = \tilde{\mathbf{w}}(\alpha)$ for $\alpha \in \mathbb{F}^n$.
- 9 The verifier computes $\mathbf{g} = \tilde{\mathbf{g}}(\alpha)$ and $\mathbf{y} = \tilde{\mathbf{y}}(\alpha)$. The verifier checks $C_{n+1} \stackrel{?}{=} \mathbf{g} \cdot \mathbf{w}$ and $v_{n+1} \stackrel{?}{=} \mathbf{y} \cdot \mathbf{w}$
 and rejects if not equal; otherwise the verifier outputs accept.

due to [Corollary 3.11](#): Alice must send the entire vector \mathbf{w} to Bob. In this case, if Alice needs to do this, then Bob does not need to interact with Merlin at all, as Bob can simply compute $\langle \mathbf{w}, \mathbf{r} \rangle$ from Alice's message. Moreover, this does not make sense if we try to map this behavior back to the cryptographic protocol: it essentially says that if the prover sends their input \mathbf{w} to the verifier, then the verifier can compute the inner product. This is just the trivial proof system for computing inner product, and is not what Bulletproofs is doing as a cryptographic protocol.

Now, if Alice is randomized with input \mathbf{w} and shares randomness $\alpha \in \mathbb{F}^n$ with Bob, then Bob and Merlin can execute the sum-check protocol and Alice's message is $\tilde{\mathbf{w}}(\alpha)$. Then, Bob computes $\tilde{\mathbf{r}}(\alpha)$ and performs the sum-check verification. In this case, this is identical to what the verifier of Bulletproofs does, where the verifier computes $\tilde{\mathbf{g}}(\alpha)$ to implicitly compute $\tilde{\mathbf{r}}(\alpha)$. Notice that this **OMA** protocol communicates $O(\log(N))$ field elements, but clearly the cryptographic verifier in this case performs $O(N)$ group operations! So we have lost our connection between the communication complexity and the verifier time in the cryptographic protocol. In particular, the work done by Bob to compute $\tilde{\mathbf{r}}(\alpha)$ is not captured in the **OMA** model. However, we believe the **OMA** model may be a viable path to obtain lower bounds on *proof sizes* of interactive arguments; we leave this as future work for exploration.

Next, consider switching the inputs of Alice and Bob: let Alice's input be \mathbf{r} and Bob's input be \mathbf{w} . In this case, again if Alice is deterministic, by [Corollary 3.11](#), she must send $\Omega(N)$ bits, and in the protocol, she would send her whole vector \mathbf{r} as her message, and Bob can compute $\langle \mathbf{r}, \mathbf{w} \rangle$ without interacting with Merlin. If Alice shares private randomness α with Bob, then Bob and Merlin can execute the sum-check protocol and have Alice send $\tilde{\mathbf{r}}(\alpha)$ as her message, allowing Bob to compute the sum-check verification.

However, both of these scenarios does not make sense in the context of the cryptographic protocol. If Bob is playing the role of the cryptographic verifier, then his input being \mathbf{w} represents the cryptographic

verifier having the prover's private input w already. So the verifier does not need to check anything or talk to any other parties! Thus, we do not believe that the **OMA** model is sufficient to capture the verifier time in Bulletproofs (or in Dory).

5.5 Rounds vs. Communication Tradeoffs for Bulletproofs

Bulletproofs as originally given in [BCC⁺16, BBB⁺18] has $O(\log(N))$ rounds, where $N = 2^n$ is the length of the prover's input, allowing the proof size to be succinct and also of size $O(\log(N))$. As shown by [BCS21] and restated by us in this section, Bulletproofs is an instantiation of the sum-check protocol. This allows us to trade off rounds versus proof size. In this section, we generalize the Bulletproofs sum-check to allow for this trade off. Specifically, for constant r , we show how to construct an r -round Bulletproofs sum-check with proof size $O(r \cdot N^{1/r})$.

An important observation is that the number of rounds in a sum-check protocol corresponds to the number of variables in the polynomial over which the sum-check operates. To construct an r -round sum-check protocol, the underlying polynomial must be defined over r variables. Given that we have N points to interpolate, we require a multivariate polynomial of degree d such that it satisfies the condition $(d+1)^r \geq N$ to ensure all the points are interpolated. Hence, we can derive $d = N^{1/r} - 1$. Notably, in the standard Bulletproof (and sum-check) protocol for proving the inner product between two vectors, $r = \log N$ and $d = 1$ which satisfies the condition $(d+1)^r \geq N$.

To generalize to any r -round protocol, we modify the extension of vectors and polynomials from multilinear to r -variate, d -degree multivariate polynomials, for degree bound $d = N^{1/r} - 1$. We then apply the sum-check protocol over this multivariate polynomial. This generalization results in an r -round Bulletproofs/sum-check protocol, presented in Algorithm 5.3, where all modifications are highlighted with a gray box. The key change lies in representing the vectors w, g, y using a multivariate polynomial extension. This representation can be constructed in two primary ways. The first method involves direct interpolation: we construct the polynomial $\tilde{w}(X_1, \dots, X_r)$ by interpolating over the evaluation points corresponding to the vector entries:

$$\tilde{w}(i_1, \dots, i_r) = w[i_1 \cdot (d+1)^{r-1} + i_2 \cdot (d+1)^{r-2} \dots + i_{r-1}].$$

For all $i_1, \dots, i_r \in [0, d]$, the index mapping is defined as $\ell = i_1 \cdot (d+1)^{r-1} + i_2 \cdot (d+1)^{r-2} \dots + i_{r-1}$, which converts the multi-index (i_1, \dots, i_r) in a $(d+1)$ -ary number system into a unique index in the standard decimal system (i.e., $\ell \in [N]$). This mapping ensures a one-to-one correspondence between the vector entries and the multivariate evaluation domain. Using the same approach, we construct the multivariate interpolating polynomials $\tilde{g}(X_1, \dots, X_r)$ and $\tilde{y}(X_1, \dots, X_r)$ corresponding to the vectors g and y , respectively.

The second method for constructing the multivariate polynomial is to define it as a multivariate extension, analogous to the multilinear extension described in Eq. (1) and Eq. (2). In this approach, the polynomial $\tilde{w}(\cdot)$ can be represented as:

$$\begin{aligned} \tilde{w}(X_1, \dots, X_r) &:= \sum_{c \in [0, d]^r} w_{\ell(c)} \cdot \text{eq}(X, c), \text{ where } \ell(c) = c_1(d+1)^{r-1} + \dots + c_d \text{ and} \\ \text{eq}(X, c) &:= \prod_{i=1}^r \prod_{j=0, j \neq c_i}^d \frac{X_i - j}{c_i - j} \end{aligned}$$

Other minor modifications include requiring the verifier to compute d evaluations of the polynomials $q_i(\cdot)$ and $p_i(\cdot)$, and to verify that their degrees do not exceed $2d$.

Translating Algorithm 5.3 into a **SMH** protocol follows the same methodology used to convert Algorithm 5.2 into its **SMH** counterpart. The primary differences are how the multivariate extensions $\tilde{w}, \tilde{g}, \tilde{y}$ are

Algorithm 5.3: BP-Sum-Check-Round-vs-Comm

- Public Params:** $r, d = N^{1/r} - 1$, \mathbb{G} with scalar field \mathbb{F} , $\mathbf{g} \in \mathbb{G}^N$.
- Private Input:** $\mathbf{w} \in \mathbb{F}^N$.
- 1 **Commit Phase:** ($\mathbf{g} \in \mathbb{G}^N, \mathbf{w} \in \mathbb{F}^N$)
 | **Output:** $C = \langle \mathbf{g}, \mathbf{w} \rangle \in \mathbb{G}$
 - 2 **Proof of Knowledge:** ($C_1 \in \mathbb{G}, \mathbf{g} \in \mathbb{G}^N, \mathbf{y} \in \mathbb{F}^N, v_1 \in \mathbb{F}; \mathbf{w} \in \mathbb{F}^N$)
 | **Output:** Accept or reject.
 - 3 The prover defines r -variate degree- d polynomials $\tilde{\mathbf{w}}, \tilde{\mathbf{y}} \in \mathbb{F}[X]$ and $\tilde{\mathbf{g}} \in \mathbb{G}[X]$ and compute
 | $F(X) := \tilde{\mathbf{g}}(X) \cdot \tilde{\mathbf{w}}(X) \in \mathbb{G}[X]$ and $H(X) := \tilde{\mathbf{y}}(X) \cdot \tilde{\mathbf{w}}(X) \in \mathbb{F}[X]$.
 - 4 **for** $i \in [r]$ **do**
 - 5 The prover defines and sends polynomials $q_i(X_i) : \mathbb{F} \rightarrow \mathbb{G}$ and $p_i(X_i) : \mathbb{F} \rightarrow \mathbb{F}$ as

$$q_i(X_i) := \sum_{\mathbf{c} \in [0, d]^{r-i}} F(\alpha_1, \dots, \alpha_{i-1}, X_i, \mathbf{c}) \quad p_i(X_i) := \sum_{\mathbf{c} \in [0, d]^{r-i}} H(\alpha_1, \dots, \alpha_{i-1}, X_i, \mathbf{c}).$$
 - 6 The verifier checks $C_i \stackrel{?}{=} \sum_{a \in [0, d]} q_i(a)$, $v_i \stackrel{?}{=} \sum_{a \in [0, d]} p_i(a)$, $\deg(q_i) \leq 2d$, and $\deg(p_i) \leq 2d$;
 reject if not.
 - 7 The verifier samples and sends $\alpha_i \xleftarrow{\$} \mathbb{F}$, and defines $C_{i+1} := q_i(\alpha_i)$ and $v_{i+1} = p_i(\alpha_i)$.
 - 8 The prover sends $w = \tilde{\mathbf{w}}(\alpha)$ for $\alpha \in \mathbb{F}^n$.
 - 9 The verifier computes $g = \tilde{\mathbf{g}}(\alpha)$ and $y = \tilde{\mathbf{y}}(\alpha)$. The verifier checks $C_{n+1} \stackrel{?}{=} g \cdot w$ and $v_{n+1} \stackrel{?}{=} y \cdot w$
 and rejects if not equal; otherwise the verifier outputs accept.
-

defined, as well as the checks performed by party C (similar to the changes made in the protocol). Additionally, the number of messages sent by party M and the number of interaction rounds between party M and party C are modified. However, these changes do not impact the overall complexity analysis of the verifier or Party C . Overall, these changes give a **SMH** $[r]$ for computing the inner product function, where M sends $O(r \cdot N^{1/r})$ messages to C .

6 Dory: Information-Theoretic Core and Evidence of Optimal Verifier Time

We now turn our attention to Dory [Lee21]. Essentially, Dory is Bulletproofs in the setting of pairing-friendly groups, but with *logarithmic verification*. This succinct verification crucially relies on the pairing setting, where parties can commit to scalars in \mathbb{F} as well as elements of the underlying group(s) \mathbb{G} . The seemingly simple addition of this fact allows for succinct verification. We seek to extract the information-theoretic core out of Dory to understand what is happening at a deeper level. We prove the following theorem in Section 6.2.1.

Theorem 6.1. *Let \mathbb{G}, \mathbb{G}_T be prime-order, pairing friendly groups with pairing function $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$, with scalar field \mathbb{F} , and let $N = 2^n$. Then, for private prover input $\mathbf{w} \in \mathbb{G}^N$ and public parameters $\mathbf{g} \in \mathbb{G}^N$ for the Dory proof of knowledge for opening a AFGHO vector commitment over \mathbb{G} , there exists a **SMH** $[n]$ protocol Π which verifies the same function as Dory with inputs \mathbf{w}, \mathbf{g} such that $C(\Pi) = \Omega(\log(N))$.*

As a direct corollary, we obtain strong evidence that Dory-PCS has (nearly) optimal verifier time.

Corollary 6.2. *Let \mathbb{G}, \mathbb{G}_T be prime-order, pairing friendly groups with pairing function $e: \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$, with scalar field \mathbb{F} , and let $N = 2^n$.*

- *For inputs $f: \{0, 1\}^n \rightarrow \mathbb{F}$, $\mathbf{u} \in \mathbb{F}^n$, and $\mathbf{g} \in \mathbb{G}^N$ to Dory-PCS for multilinear polynomials, there exists a **SMH**[n] protocol Π which verifies the same function as Dory-PCS with inputs $f, \mathbf{u}, \mathbf{g}$ such that $C(\Pi) = \Omega(\log(N))$.*
- *For inputs $f \in \mathbb{F}^{<N}[X]$, $u \in \mathbb{F}$, and $\mathbf{g} \in \mathbb{G}^N$ to Dory-PCS for univariate polynomials, there exists a **SMH**[n] protocol Π which verifies the same function as Dory-PCS for inputs f, u, \mathbf{g} such that $C(\Pi) = \Omega(\log(N))$.*

Moreover, the prescribed Dory-PCS verifier (for both multilinear and univariate polynomials) performs at most $O(\log(N) \text{ polylog } |\mathbb{G}|)$ bit-operations, assuming a single \mathbb{G}/\mathbb{F} operation (including pairings) requires $\text{polylog } |\mathbb{G}|$ bit-operations.

We proceed as follows. First, we examine a simpler version of Dory we call Proto-Dory with $O(\log^2(N))$ verification time (when operating over vectors of length N).⁶ We then frame Proto-Dory as another sum-check protocol, as we did with Bulletproofs. From this sum-check, we extract the information theoretic core as a communication protocol in our **SMH** model. Finally, we frame the full Dory protocol (with $O(\log(N))$ verification) as a sum-check, and give its communication analogue in the **SMH** model. To the best of our knowledge, we are the first to formalize the connection between Dory and the sum-check protocol.

Remark 6.3. For simplicity, we state (Proto-)Dory in the context of *symmetric* pairings (i.e., $\mathbb{G}_1 = \mathbb{G}_2$), even though the symmetric external Diffie-Hellman assumption does not hold in this case. We do this for two reasons: (1) it greatly simplifies the presentation; and (2) the information-theoretic core is agnostic to the use of (a)symmetric pairings (since there will be no pairings in this model anyway).

Similar to Bulletproofs, (Proto-)Dory is a proof of knowledge for the statement $C = \langle \mathbf{g}, \mathbf{w} \rangle \in \mathbb{G}_T$, where $\mathbf{g}, \mathbf{w} \in \mathbb{G}^N$ for pairing friendly group \mathbb{G} , and where $\langle \mathbf{g}, \mathbf{w} \rangle := \langle \mathbf{g}, \mathbf{w} \rangle_e$ for pairing function e defined in Section 2. In other words, it is (essentially) the same proof of knowledge statement being proven in Bulletproofs: given generator $h \in \mathbb{G}$ and vector $\mathbf{u} \in \mathbb{F}^N$, we can commit to \mathbf{u} by setting $\mathbf{w}_i = h \cdot \mathbf{u}_i$ for all i and computing $C = \langle \mathbf{g}, \mathbf{w} \rangle_e$. At first, Dory proceeds identically as Bulletproofs: it proves the statement $C = \langle \mathbf{g}, \mathbf{w} \rangle$ via a sum-check over the polynomial $F(X) = \tilde{\mathbf{g}}(X) \tilde{\mathbf{w}}(X)$ for appropriately defined multilinear extensions (see Section 2). However, recall in Bulletproofs that the verifier at the end of the protocol must compute the value $\tilde{\mathbf{g}}(\alpha)$, which lead to our linear lower bound. In Dory, we take explicit advantage of the pairing setting: Dory instead *forces* the prover to commit to $\tilde{\mathbf{g}}(\alpha)$ and prove that the commitment is valid. This is done via another set of random generators $\mathbf{\Gamma}$, which will be commitment keys used to commit to \mathbf{g} (and its correct foldings) during every round of the sum-check. For a more thorough treatment on the insights of Dory, see [Tha22, Chapter 15].

6.1 Proto-Dory is (a variant of) Sum-Check

We first present a simpler version of Dory, which we call Proto-Dory, with $O(\log^2(N))$ verification for vectors of length $N = 2^n$. The public parameters of Proto-Dory are group generators $\mathbf{g} \in \mathbb{G}^N$ and $\mathbf{\Gamma} \in \mathbb{G}^{N/2}$. To aid in the sum-check protocol, Proto-Dory uses pairings to precompute some preprocessing polynomials that the verifier will use throughout the sum-check protocol.

Consider the first round of the Proto-Dory sum-check, which is identical to the first round of the Bulletproofs sum-check. The prover is trying to convince the verifier it knows \mathbf{w} such that $C = \langle \mathbf{g}, \mathbf{w} \rangle$. So

⁶Proto-Dory is presented in [Lee21] as a warm-up protocol to the $O(\log(N))$ verification time Dory.

the prover and verifier engage in the sum-check protocol over the polynomial $F(X) = \tilde{g}(X) \cdot \tilde{w}(X)$. Let $S_1(X_1)$ be the sum-check polynomial sent by the prover this round. After the first round of the sum-check, the verifier has the claim $C_2 = S_1(\alpha_1)$, which is equivalent to the claim (if the prover is honest) $C_2 = \langle g', w' \rangle$ for $g' = (1 - \alpha_1)g_L + \alpha_1 g_R$ and $w' = (1 - \alpha_1)w_L + \alpha_1 w_R$.

Intuitively, to overcome the limitations of the discrete-log setting of Bulletproofs, Dory takes advantage of pairings and gives the verifier the ability to compute the commitment $D_1 = \langle g', \Gamma \rangle$ *without explicitly computing g'* . Once the verifier has this commitment, then during the next round of the protocol, the prover and verifier engage in *two parallel sum-check protocols*, one for the claim $C_2 = \sum_c F(\alpha_1, c)$ and another for the claim $D = \langle g', \Gamma \rangle = \sum_c \tilde{g}(\alpha_1, c) \cdot \tilde{\Gamma}(c)$.

How does Dory ensure the verifier can compute the claim D_1 above without reading the entire g or Γ ? This is handled during a linear-time preprocessing phase. Notice that the verifier can compute D_1 for any α_1 sampled if given the polynomial $Q(X_1) = \sum_c \tilde{g}(X_1, c) \cdot \tilde{\Gamma}(c)$, where $\tilde{g}: \mathbb{G}^n \rightarrow \mathbb{G}_T$ and $\tilde{\Gamma}: \mathbb{G}^{n-1} \rightarrow \mathbb{G}_T$. Notice further that this polynomial Q has degree at most 1, so the verifier can compute D by reading two group elements (i.e., the description of Q) from the preprocessing.

Proto-Dory's preprocessing enables the verifier to repeat this process each round of the sum-check. In particular, in round i of the sum-check, the prover is trying to prove i parallel sum-check claims. The first claim is still for proving the original inner product claim $C = \langle g, w \rangle$. The second claim is for proving the first folding $D_1 = \langle g', \Gamma \rangle$. The third claim now makes sure that the prover folds Γ to Γ' under challenge α_2 correctly, using some other commitment key $\Delta \in \mathbb{G}^{N/4}$; that is, proving the claim $D_2 = \langle \Gamma', \Delta \rangle$. This process is repeated during every round of the sum-check, leading to the $O(\log^2(N))$ proof size.

The formal protocol for Proto-Dory as a sum-check is presented in [Algorithm 6.1](#). First, given public $g \in \mathbb{G}^N$ and $\Gamma \in \mathbb{G}^{N/2}$, the linear-time preprocessing phase constructs n degree-1 polynomials $Q_{j,j}(X_j)$ as follows. Set $\Gamma^{(1)} = g$, $\Gamma^{(2)} = \Gamma$, and for $i \in \{3, \dots, n\}$, define $\Gamma^{(i)} = \Gamma_L^{(i-1)} \in \mathbb{G}^{N/2^{i-1}}$. Finally, set $\Gamma^{(n+1)} = \Gamma_L^{(n)} \in \mathbb{G}$. Here, the subscript L represents the left half of the vector. Then we define $Q_{j,j}(X_j) = \sum_c \tilde{\Gamma}^{(j)}(X_j, c) \cdot \tilde{\Gamma}^{(j+1)}(c)$ for all $j \in [n]$. The preprocessing outputs all $Q_{j,j}$ along with $\Gamma^{(n+1)}$.⁷

At the start of round i of the sum-check, the prover (in parallel) proves i sum-check claims. First, it computes $S_i(X_i) = \sum_c F(\alpha_1, \dots, \alpha_{i-1}, X_i, c)$ for verifier challenges $\alpha_j \xleftarrow{\$} \mathbb{F}$ for $j \in [i-1]$; this corresponds to the sum-check for the original claim $C = \langle g, w \rangle$. Then, for each $j \in [i-1]$, the prover additionally computes polynomials $Q_{j,i}(X_i) = \sum_c P_j(\alpha_j, \dots, \alpha_{i-1}, X_i, c)$, where $P_j(X_{[j:n]}) = \tilde{\Gamma}^{(j)}(X_{[j:n]}) \tilde{\Gamma}^{(j+1)}(X_{[j+1:n]})$. Here, $X_{[j:n]}$ denotes the variables X_j, \dots, X_n . Intuitively, each polynomial $Q_{j,i}$ corresponds to the sum-check claim $D_j = \langle (1 - \alpha_j)\Gamma^{(j)} + \alpha_j \Gamma^{(j)}, \Gamma^{(j+1)} \rangle = \sum_c \tilde{\Gamma}^{(j)}(\alpha_j, c) \cdot \tilde{\Gamma}^{(j+1)}(c)$. At the end of round i , the verifier updates its i sum-check claims as $C_{i+1} = S_i(X_i)$ and $D_{j,i+1} = Q_{j,i}(\alpha_i)$ for $j \in [i-1]$. Additionally, the verifier now adds a new sum-check claim $D_{i,i+1} = Q_{i,i}(\alpha_i)$ by using the preprocessing polynomial $Q_{i,i}$, increasing the number of sum-check claims from i to $i+1$ for the next round.

In the final round of the protocol, the prover sends the claimed evaluations $w = \tilde{w}(\alpha)$ and $\Gamma^{(j)} = \tilde{\Gamma}^{(j)}(\alpha_j, \dots, \alpha_n)$ for $j \in [n]$. Using these values and the preprocessing value $\Gamma^{(n+1)}$, the verifier checks consistency of the values sent by the prover with the set of claims it has been tracking C_{n+1} and $D_{j,n+1}$ for $j \in [n]$. Intuitively, because Dory is a proof of knowledge, the prover is forced to send the correct values unless it can break the discrete-log assumption, otherwise (1) the verifier will reject (with high probability), or (2) we can extract the prover's witness from the protocol transcript.

⁷Notice that $Q_{j,j}(0) = \langle \Gamma_L^{(j)}, \Gamma^{(j+1)} \rangle$ and $Q_{j,j}(1) = \langle \Gamma_R^{(j)}, \Gamma^{(j+1)} \rangle$. That is, $Q_{j,j}$ contains the commitments to the left and right halves of $\Gamma^{(j)}$ under the commitment key $\Gamma^{(j+1)}$, which is exactly the output of the Dory preprocessing [\[Lee21\]](#).

Algorithm 6.1: Proto-Dory-Sum-Check

Public Parameters: $N = 2^n$, pairing-friendly $(\mathbb{G}, \mathbb{G}_T, \mathbb{F})$, $\mathbf{g} \in \mathbb{G}^N$, $\mathbf{\Gamma} \in \mathbb{G}^{N/2}$.
Private Input: $\mathbf{w} \in \mathbb{G}^N$.

- 1 **Preprocessing** ($\mathbf{g} \in \mathbb{G}^N$, $\mathbf{\Gamma} \in \mathbb{G}^{N/2}$)
 Output: $\mathbf{\Gamma}^{(n+1)} \in \mathbb{G}$ and $\{Q_{j,j}(X_j)\}_{j \in [n]}$.
 2 Define $\mathbf{\Gamma}^{(1)} := \mathbf{g}$, $\mathbf{\Gamma}^{(2)} := \mathbf{\Gamma}$ and set $\mathbf{\Gamma}^{(i)} := \mathbf{\Gamma}_L^{(i-1)} \in \mathbb{G}^{N/2^{i-1}}$ for $i \in \{3, \dots, n\}$. Set $\mathbf{\Gamma}^{(n+1)} := \mathbf{\Gamma}_L^{(n)} \in \mathbb{G}$.
 3 For $j \in [n]$, define polynomials

$$Q_{j,j}(X_j) = \sum_{\mathbf{c} \in \{0,1\}^{n-j}} \tilde{\mathbf{\Gamma}}^{(j)}(X_j, \mathbf{c}) \cdot \tilde{\mathbf{\Gamma}}^{(j+1)}(\mathbf{c}) \in \mathbb{G}_T[X_j].$$
- 4 Output $\mathbf{\Gamma}^{(n+1)}$ and $\{Q_{j,j}(X_j)\}_{j \in [n]}$.
- 5 **Commit Phase** ($\mathbf{w}, \mathbf{g} \in \mathbb{G}^N$)
 Output: $C = \langle \mathbf{g}, \mathbf{w} \rangle \in \mathbb{G}_T$
- 6 **Proof of Knowledge** ($C \in \mathbb{G}_T$, $\mathbf{g} \in \mathbb{G}^N$, $\mathbf{\Gamma} \in \mathbb{G}^{N/2}$, $\{Q_{j,j}(X_j) \in \mathbb{G}_T[X_j]\}_{j \in [n]}$, $\mathbf{\Gamma}^{(n+1)} \in \mathbb{G}$; $\mathbf{w} \in \mathbb{G}^N$)
 Output: Accept or reject.
 7 The prover defines $P_j(X_{[j:n]}) = \tilde{\mathbf{\Gamma}}^{(j)}(X_{[j:n]}) \cdot \tilde{\mathbf{\Gamma}}^{(j+1)}(X_{[j+1:n]}) \in \mathbb{G}_T[X_{[j:n]}]$ for $j \in [i-1]$ and
 $S(X) = \tilde{\mathbf{g}}(X) \cdot \tilde{\mathbf{w}}(X) \in \mathbb{G}_T[X]$.
 8 The verifier defines $C_1 := C$.
 9 **for** $i \in [n]$ **do**
 10 The prover defines polynomials

$$S_i(X_i) := \sum_{\mathbf{c} \in \{0,1\}^{n-i}} S(\alpha_1, \dots, \alpha_{i-1}, X_i, \mathbf{c})$$

$$Q_{j,i}(X_i) := \sum_{\mathbf{c} \in \{0,1\}^{n-j-i-1}} P_j(\alpha_j, \dots, \alpha_{i-1}, X_i, \mathbf{c}) \quad \forall j \in [i-1].$$
- 11 The prover sends $S_i(X_i)$ and $\{Q_{j,i}(X_i)\}_{j \in [i-1]}$.
- 12 The verifier checks $C_i \stackrel{?}{=} S_i(0) + S_i(1)$, $\{Q_{j,i}(0) + Q_{j,i}(1) \stackrel{?}{=} D_{j,i}\}_{j \in [i-1]}$, and that all polynomials have degree at most 2; reject if not.
- 13 The verifier samples and sends $\alpha_i \xleftarrow{\$} \mathbb{F}$.
- 14 The verifier defines $C_{i+1} = S_i(\alpha_i)$ and $\{D_{j,i+1} = Q_{j,i}(\alpha_i)\}_{j \in [i]}$.
- 15 The prover sends $\mathbf{w} = \tilde{\mathbf{w}}(\alpha)$ and $\{\mathbf{\Gamma}^{(j)} = \tilde{\mathbf{\Gamma}}^{(j)}(\alpha_j, \dots, \alpha_n)\}_{j \in [n]}$.
- 16 The verifier checks $C_{n+1} \stackrel{?}{=} \mathbf{w} \cdot \mathbf{\Gamma}^{(1)}$ and $\{D_{j,n+1} \stackrel{?}{=} \mathbf{\Gamma}^{(j)} \cdot \mathbf{\Gamma}^{(j+1)}\}_{j \in [n]}$, rejecting if not all are true.

6.1.1 Proto-Dory as a SMH Communication Protocol

As with Bulletproofs, the benefit of framing Proto-Dory as a sum-check protocol is that we can cleanly capture the information-theoretic core using our new SMH communication model. We present the formal **SMH**[n] protocol for Proto-Dory in Figure 3. This protocol is nearly identical to Figure 2, with two notable exceptions. First, we have M send an additional message $\hat{R}(X_n)$ in round $i = n$, and second is that A is now *randomized*.

In more detail, party A of Figure 3 has the same input as in our Bulletproofs protocol (Figure 2): $\mathbf{r} \in \mathbb{F}^N$ representing the scalar vector “hidden” within $\mathbf{g} \in \mathbb{G}^N$. Now, we allow A , B , and M to sample shared randomness $\mathbf{\gamma}^{(i+1)} \xleftarrow{\$} \mathbb{F}^{N/2^{n-i}}$ for $i \in [n]$, intuitively representing the “hidden” scalars within vectors $\mathbf{\Gamma}^{(i+1)} \in \mathbb{G}^{N/2^{n-i}}$. There is one major difference when translating Algorithm 6.1 to Figure 3 regarding how the commitment keys $\mathbf{\Gamma}^{(i)}$ are sampled and how the private shared randomness $\mathbf{\gamma}^{(i)}$ are sampled. In Figure 3, the vectors $\mathbf{\gamma}^{(i)}$ for $i \in [2, n+1]$ are sampled uniformly at random, whereas in Algorithm 6.1 only the first vector $\mathbf{\Gamma} \in \mathbb{G}^{N/2}$ was random (and the remaining were the left halves). In the actual specification of Dory [Lee21], this vector is chosen as $\mathbf{\Gamma}^{(2)} = \mathbf{g}_L$, with each subsequent commitment key being set as $\mathbf{\Gamma}^{(j+1)} = \mathbf{\Gamma}_L^{(j)}$, with the subscript L denoting the left half of the vector. Note that in the actual definitions of our models (Definitions 3.1 and 3.6), parties must sample shared randomness *before* receiving any inputs to the protocol. However, suppose we do allow parties to sample shared randomness after receiving their inputs.

When selecting vectors $\mathbf{\Gamma}^{(i)}$ in Algorithm 6.1, they do not need to be sampled randomly, as the hardness

of discrete-log, intuitively, lets us treat each vector $\Gamma^{(i)}$ as random, even though they are structured. This, however, is *not possible* in the information theoretic setting. In particular, one can demonstrate an attack in the information theoretic setting when using such structured vectors; see [Remark 6.5](#) for more discussion. Notably, if we mimic how Dory chooses the intermediate keys, where they are simply a function of the initial input \mathbf{g} , then this would translate to a SMH protocol with a *deterministic* party A , and thus we would run into the lower bound of [Theorem 3.10](#). The aforementioned attack lines up with the proof of [Theorem 3.10](#): party M can run the protocol with respect to some $\mathbf{r}' \neq \mathbf{r}$ but party A 's messages would send the same messages for both \mathbf{r} and \mathbf{r}' as their input. However, if we select the commitment key $\gamma^{(2)}$ uniformly and independently at random from $\gamma^{(1)} = \mathbf{r}$ in the communication protocol, then we can set $\gamma^{(i+1)} = \gamma_L^{(i)}$ for $i \geq 2$ and retain the soundness of the protocol; for ease of presentation, we choose all $\gamma^{(i+1)}$ uniformly at random.

Now, why do we allow A to share this same randomness with B ? Allowing A and B to share this same randomness enables B to send the correct values $\gamma^{(i+1)} = \tilde{\gamma}^{(i+1)}(\alpha_{i+1}, \dots, \alpha_n)$ for $i \in [n-1]$ at the end of the protocol, along with the correct $w = \tilde{w}(\alpha)$. Intuitively, this is capturing the proof of knowledge property of Dory, which forces the cryptographic prover to send these correct values, otherwise we can extract the prover's witness \mathbf{w} from the transcript of the interaction, or the prover has broken the discrete-log assumption.

The final subtle difference between [Algorithm 6.1](#) and [Figure 3](#) is the extra message $\hat{R}(X_n)$ sent by M in round n of the sum-check. We enforce this to again, intuitively, try to capture the cryptography in Proto-Dory. In [Algorithm 6.1](#), the discrete-log assumption and proof of knowledge property force the cryptographic prover to correctly send $\gamma^{(i)} = \tilde{\gamma}^{(i+1)}(\alpha_i, \dots, \alpha_n)$ for $i \in [n]$. By giving party B the random vectors $\gamma^{(j)}$, we allow party B to emulate this behavior for all but the value $\gamma^{(1)}$. Under the specification of [Figure 3](#), $\gamma^{(1)} = \mathbf{r}$, which is A 's input. Notice that if A could share their input with B , then there is no need to interact with M and no need to share any random strings with any party: B can simply send $\langle \mathbf{w}, \mathbf{r} \rangle$ directly. The randomness allows B to aid C in the computation of $\langle \mathbf{w}, \mathbf{r} \rangle$ even though B does not have access to A 's input \mathbf{r} .

However, to compute all the necessary checks for the sum-check protocol, C is missing one final piece of information: the value $\gamma^{(1)} = \tilde{\mathbf{r}}(\alpha)$. We give C access to this value through M sending $\hat{R}(X_n)$, which is claimed to be the polynomial $\tilde{\mathbf{r}}(\alpha_1, \dots, \alpha_{n-1}, X_n)$. This is to simplify the analysis; we could also add one more round of interaction where M receives α_n and sends $\hat{\gamma}^{(n+1)}$ which is claimed to be $\tilde{\mathbf{r}}(\alpha)$. In this case, M could try to cheat by computing $\hat{\gamma}^{(n+1)} = v^{(n+1)}/w$ (M can compute w now that it has received α_n), but we explicitly disallow this because in the cryptographic setting, this would translate to the cryptographic prover finding a non-trivial discrete-log relation. So, intuitively, forcing M to send $\hat{R}(X_n)$ is trying to also enforce that the cryptographic prover would be forced to send $\tilde{\mathbf{r}}(\alpha)$ honestly.

Note also that C actually doesn't need to receive this value from M . Instead, C could recover $\gamma^{(1)}$ through computing $\gamma^{(1)} := v^{(n+1)}/w$, then running all the remaining checks with respect to the honestly received $\gamma^{(j)}$ for $j \geq 2$. But then C is doing something impossible in the cryptographic setting, since computing this $\gamma^{(1)}$ is actually happening in the exponents of the cryptographic scheme, and this would require C to find a non-trivial discrete-log relation. So we explicitly disallow this and have C compute this value from M 's extra message in the last round of the sum-check protocol.

With the discussion on the discrepancies between [Algorithm 6.1](#) and [Figure 3](#) complete, we now show that the protocol of [Figure 3](#) is a valid SMH protocol.

Claim 6.4. *Let $N = 2^n$ and let $f(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle$ for any $\mathbf{x}, \mathbf{y} \in \mathbb{F}^N$ and any finite field \mathbb{F} . Then, [Figure 3](#) is a $(0, O(n^2)/|\mathbb{F}|)$ -secure SMH[n] protocol for f .*

Proof. Perfect completeness follows from the sum-check protocol. Similarly, soundness follows almost directly from the sum-check protocol. During round i of the sum-check, C is verifying i claims, and constructing claim $i+1$ using random challenge α_i . Each of these claims is about a degree at most 2

polynomial, so the total soundness error contribution is at most $2i/|\mathbb{F}|$. Taking a union bound over all rounds $i \in [n]$ gives us a $2n^2/|\mathbb{F}|$ upper bound.

However, in the last round of the sum-check, M sends $\hat{R}(X_n)$, which is claimed to be $\tilde{r}(\alpha_1, \dots, \alpha_{n-1}, X_n)$. If M has been cheating so far and has lucked out with all the sum-check verification checks for $i \in [n]$, the final checks that need to pass are $\hat{v}^{(n+1)} \stackrel{?}{=} w \cdot \hat{R}(\alpha_n)$ and $\hat{D}_{j,n+1} \stackrel{?}{=} \gamma^{(j)} \cdot \gamma^{(j+1)}$ for $j \in [n]$. Here, we use “ $\hat{\cdot}$ ” to denote potentially tampered values sent by M . Notice that since \hat{R} is sent without knowing α_n , M never learns the value of $w = \tilde{w}(\alpha)$. So the probability the first check passes if \hat{R} is not honestly sent is $1/|\mathbb{F}|$. Notice also that the values $\hat{D}_{j,n+1}$ are defined by polynomials sent by M without knowing the value α_n , while $\gamma^{(j+1)}$ is honestly sent by B for $j \in [n]$. So the probability each of these checks pass is also $1/|\mathbb{F}|$. This gives a final soundness upper bound of $O(n^2/|\mathbb{F}|)$. \square

Remark 6.5. We demonstrate how to break the soundness of the above protocol if we choose $\gamma^{(1)} = \mathbf{r}$ and $\gamma^{(i+1)} = \gamma_L^{(i)}$ for all $i \in [n]$. Consider $N = 4$ and $n = 2$. Let $\mathbf{r} = \gamma^{(1)} = (r_{00}, r_{01}, r_{10}, r_{11})$. Then, we have $\gamma^{(2)} = (r_{00}, r_{01})$ and $\gamma^{(3)} = r_{00}$.

We construct another vector $\mathbf{s} \neq \mathbf{r}$ such that party A 's messages to C are identical when A 's input is \mathbf{s} and \mathbf{r} . Let $\mathbf{s} = (s_{00}, s_{01}, s_{10}, s_{11})$. First, note that A sends two polynomials $Q_{1,1}, Q_{2,2}$, and the value $\gamma^{(3)}$. Since we want A to send the same messages for both input \mathbf{s}, \mathbf{r} , we set $s_{00} = r_{00} = \gamma^{(3)}$. Now, the polynomial $Q_{2,2}(X_2)$ can be expressed as $Q_{2,2}(X_2) = (1 - X_2)\langle \gamma_L^{(2)}, \gamma^{(3)} \rangle + X_2\langle \gamma_R^{(2)}, \gamma^{(3)} \rangle = (1 - X_2)r_{00}^2 + X_2r_{01}r_{00}$. Since we have set $s_{00} = r_{00}$, in order for this message to be identical for input \mathbf{s} , we must set $s_{01} = r_{01}$.

Finally, we demonstrate how to set $s_{10} \neq r_{10}$ or $s_{11} \neq r_{11}$ such that the message $Q_{1,1}(X_1)$ remains the same. From the definition of $Q_{1,1}(X_1)$, it can be shown that for A 's input \mathbf{r} :

$$\begin{aligned} Q_{1,1}(X_1) &= (1 - X_1)\langle \gamma_L^{(1)}, \gamma^{(2)} \rangle + X_1\langle \gamma_R^{(1)}, \gamma^{(2)} \rangle \\ &= (1 - X_1)\langle \gamma^{(2)}, \gamma^{(2)} \rangle + X_1\langle \gamma_R^{(1)}, \gamma^{(2)} \rangle \\ &= (1 - X_1)\langle (r_{00}, r_{01}), (r_{00}, r_{01}) \rangle + X_1\langle (r_{10}, r_{11}), (r_{00}, r_{01}) \rangle \\ &= (1 - X_1)(r_{00}^2 + r_{01}^2) + X_1(r_{10}r_{00} + r_{11}r_{01}) \end{aligned}$$

Now, if A 's input is \mathbf{s} such that $s_{00} = r_{00}$ and $s_{01} = r_{01}$, we can write $Q_{1,1}(X_1)$ as

$$Q_{1,1}(X_1) = (1 - X_1)(r_{00}^2 + r_{01}^2) + X_1(s_{10}r_{00} + s_{11}r_{01}).$$

In order for A to send the same message for both \mathbf{r} and \mathbf{s} as input, the above equations tell us that the following must hold:

$$r_{10}r_{00} + r_{11}r_{01} = s_{10}r_{00} + s_{11}r_{01}.$$

Letting s_{10} be any value such that $s_{10} \neq r_{10}$, we can rewrite the above (assuming $r_{01} \neq 0$) as

$$s_{11} = \frac{r_{00}(r_{10} - s_{10}) + r_{11}r_{01}}{r_{01}}.$$

So party A sends the same $Q_{1,1}, Q_{2,2}$, and $\gamma^{(3)}$ to party C with inputs $\mathbf{r} \neq \mathbf{s}$, breaking soundness of the protocol (in fact, demonstrating the exact attack derived in the proof of [Theorem 3.10](#)).

6.2 Full Dory as Sum-Check

Building upon Proto-Dory, the full Dory protocol compresses the proofs from $O(\log^2(N))$ size to $O(\log(N))$ size. While the full details of this compression are (somewhat) messy, the intuition behind the compression is quite simple. Recall that in round i of the Proto-Dory sum-check, the prover and verifier were running i parallel instances of the sum-check protocol in order to prove the claims $\langle \mathbf{g}, \mathbf{w} \rangle$, $\langle (1 - \alpha_1)\mathbf{g}_L + \alpha_1\mathbf{g}_R, \mathbf{\Gamma}^{(2)} \rangle$, $\langle (1 - \alpha_2)\mathbf{\Gamma}_L^{(2)} + \alpha_2\mathbf{\Gamma}_R^{(2)}, \mathbf{\Gamma}^{(3)} \rangle, \dots, \langle (1 - \alpha_{i-1})\mathbf{\Gamma}_L^{(i-1)} + \alpha_{i-1}\mathbf{\Gamma}_R^{(i-1)}, \mathbf{\Gamma}^{(i)} \rangle$. This essentially lead to the $O(\log^2(N))$ sized proofs. The key insight to the full Dory protocol is that we can still force the prover to track *all* of these claims, but only using a *constant number* of claims (specifically 3) per round. We outline the intuition behind this insight.

After round 1 of Proto-Dory, we went from one claim to two claims. Rather than have the verifier track these claims, Dory actually has the verifier specify an additional random value β_1 and force the prover to “fold” together all the claims via a random linear combination (through β_1). Formally, the full Dory sum-check is simultaneously proving 3 claims. Let $\mathbf{w}, \mathbf{g}, \mathbf{\Gamma}^{(1)} \in \mathbb{G}^N$. Then Dory is a sum-check protocol for simultaneously proving the claims $C_1 = \sum_{\mathbf{c}} \tilde{\mathbf{w}}(\mathbf{c})\tilde{\mathbf{g}}(\mathbf{c})$, $C_2 = \sum_{\mathbf{c}} \tilde{\mathbf{w}}(\mathbf{c})\tilde{\mathbf{\Gamma}}^{(1)}(\mathbf{c})$, and $C_3 = \tilde{\mathbf{g}}(\mathbf{c})\tilde{\mathbf{\Gamma}}^{(1)}(\mathbf{c})$. Let $\mathbf{\Gamma}^{(2)} \in \mathbb{G}^{N/2}$.

The Dory sum-check will reduce the above claims to three new claims of the form $C'_1 = \sum_{\mathbf{c}} \tilde{\mathbf{w}}'(\mathbf{c})\tilde{\mathbf{g}}'(\mathbf{c})$, $C'_2 = \sum_{\mathbf{c}} \tilde{\mathbf{w}}'(\mathbf{c})\tilde{\mathbf{\Gamma}}^{(2)}(\mathbf{c})$, and $C'_3 = \tilde{\mathbf{g}}'(\mathbf{c})\tilde{\mathbf{\Gamma}}^{(2)}(\mathbf{c})$, where for verifier challenges α_1, β_1 , we have $\mathbf{w}' = (1 - \alpha_1)(\mathbf{w} + \beta_1\mathbf{\Gamma}^{(1)})_L + \alpha_1(\mathbf{w} + \beta_1\mathbf{\Gamma}^{(1)})_R$ and $\mathbf{g}' = (1 - \alpha_1)(\mathbf{g} + \beta_1\mathbf{\Gamma}^{(1)})_L + \alpha_1(\mathbf{g} + \beta_1\mathbf{\Gamma}^{(1)})_R$. As polynomials, notice that

$$\tilde{\mathbf{w}}'(X_{[2:n]}) = \tilde{\mathbf{w}}(\alpha_1, X_{[2:n]}) + \beta_1\tilde{\mathbf{\Gamma}}^{(1)}(\alpha_1, X_{[2:n]}) \quad \tilde{\mathbf{g}}'(X_{[2:n]}) = \tilde{\mathbf{g}}(\alpha_1, X_{[2:n]}) + \beta_1\tilde{\mathbf{\Gamma}}^{(1)}(\alpha_1, X_{[2:n]}).$$

So Dory has forced the prover to track the original polynomials $\tilde{\mathbf{w}}, \tilde{\mathbf{g}}$, and the preprocessing polynomial $\tilde{\mathbf{\Gamma}}^{(1)}$. This continues recursively each round: after round 2, the prover is forced to track $\tilde{\mathbf{w}}, \tilde{\mathbf{g}}, \tilde{\mathbf{\Gamma}}^{(1)}$, and $\tilde{\mathbf{\Gamma}}^{(2)}$. Meanwhile, the verifier only needs to track a constant number of claims, and to update these claims, the verifier uses the preprocessing polynomials, along with two additional auxiliary polynomials given by the prover each round of the protocol. We give the full Dory protocol as a sum-check protocol in [Algorithm 6.2](#).

6.2.1 Full Dory as a SMH Protocol

With the specification of full Dory as a sum-check protocol, we now move extract the information-theoretic core of Dory as a **SMH** protocol, which will be similar to [Figure 3](#). The full protocol is presented in [Figure 7](#), and we give an overview of the protocol here. Intuitively, A is capturing the public setup of Dory, B is capturing the (most) of the information that the cryptographic prover sends at the end of the protocol, and M is capturing a potentially dishonest prover interacting with the cryptographic verifier.

Before the start of the protocol, B and C share private randomness $\alpha, \beta \xleftarrow{\$} \mathbb{F}^n$, and A, B , and M share private randomness $\gamma^{(i)} \xleftarrow{\$} \mathbb{F}^{N/2^{n+1-i}}$ for $i \in [n+1]$. Now, C and M engage in a sum-check protocol that is identical to the proof of knowledge of [Algorithm 6.2](#), except that in round $i = n$, M additionally sends a degree at most 1 polynomial $\hat{R}(X_n)$, claimed to be the polynomial $\tilde{r}(\alpha_1, \dots, \alpha_{n-1}, X_n)$. Next, A computes preprocessing polynomials Q_j, \hat{Q}_j identically to the preprocessing in [Algorithm 6.2](#), except using the sampled vectors $\gamma^{(j)}$. A sends these polynomials and the value $\gamma^{(n+1)}$ to C . B computes and sends $w = \tilde{\mathbf{w}}(\alpha)$ and $\gamma^{(j)} = \tilde{\gamma}^{(j)}(\alpha_j, \dots, \alpha_n)$ for $j \in [2, n]$. Finally, C performs all the verification checks of [Algorithm 6.2](#), plus one more check with the polynomial $\hat{R}(X_n)$: C checks that this polynomial has degree at most 1. Moreover, C sets $W^{(n+1)} = w + \sum_{j \in [n]} \beta_j \gamma^{(j)}$ and $R^{(n+1)} = \hat{R}(\alpha_n) + \sum_{j \in [n]} \beta_j \gamma^{(j)}$. Note that with standard algebraic manipulation, if M is honest, it can be seen that $W^{(n+1)}$ and $R^{(n+1)}$ computed in this way are identical to the honest prover messages $W^{(n+1)}, G^{(n+1)}$ from [Algorithm 6.2](#).

Algorithm 6.2: Full-Dory-Sum-Check (Preprocessing and Commit)

Public Parameters: Instance size $N = 2^n$, prime-order pairing-friendly group \mathbb{G} with scalar field \mathbb{F} , and random $\mathbf{g}, \mathbf{\Gamma} \in \mathbb{G}^N$.
Private Input: Prover has private input $\mathbf{w} \in \mathbb{G}^N$.

- 1 **Preprocessing**

Input: Group elements $\mathbf{g}, \mathbf{\Gamma}$ from public parameters.
Output: $\{(Q_j(X_j), \hat{Q}_j(X_j))\}_{j \in [n]}$ of degree at most 1 and $\Gamma^{(n+1)} \in \mathbb{G}$.

2 Define $\Gamma^{(1)} := \mathbf{\Gamma}$ and set $\Gamma^{(i)} := \Gamma_L^{(i-1)} \in \mathbb{G}^{N/2^{i-1}}$ for $i \in \{2, \dots, n\}$. Set $\Gamma^{(n+1)} := \Gamma_L^{(n)}$.

3 For $j \in [n]$, define polynomials $P_j(X_j, \dots, X_n) := \tilde{\Gamma}^{(j)}(X_j, \dots, X_n) \cdot \tilde{\Gamma}^{(j)}(X_j, \dots, X_n)$,
 $\hat{P}_j(X_j, \dots, X_n) := \tilde{\Gamma}^{(j)}(X_j, \dots, X_n) \cdot \tilde{\Gamma}^{(j+1)}(X_{j+1}, \dots, X_n)$, and

$$Q_j(X_j) = \sum_{\mathbf{c} \in \{0,1\}^{n-j}} P_j(X_j, \mathbf{c}) \qquad \hat{Q}_j(X_j) = \sum_{\mathbf{c} \in \{0,1\}^{n-j}} \hat{P}_j(X_j, \mathbf{c}).$$

4 Output $\Gamma^{(n+1)}$ and $\{Q_j(X_j), \hat{Q}_j(X_j)\}_{j \in [n]}$.
- 5 **Commit Phase** ($\mathbf{w}, \mathbf{g}, \mathbf{\Gamma} \in \mathbb{G}^N$)

Output: Commitments $C_1 = \langle \mathbf{g}, \mathbf{w} \rangle, C_2 = \langle \mathbf{w}, \mathbf{\Gamma} \rangle, C_3 = \langle \mathbf{g}, \mathbf{\Gamma} \rangle \in \mathbb{G}_T$
- 6 **Proof of Knowledge**

Input: Commitments $C_1, C_2, C_3 \in \mathbb{G}_T$, public parameters $\mathbf{g}, \mathbf{\Gamma} \in \mathbb{G}^N$, public preprocessing $(\Gamma^{(n+1)}, \{(Q_j(X_j), \hat{Q}_j(X_j))\}_{j \in [n]})$; private input $\mathbf{w} \in \mathbb{G}$.
Output: Verifier accepts or rejects.

7 The verifier defines $C_j^{(1)} := C_j$ for $j \in [3]$. The prover defines $J^{(n+1)} = \Gamma^{(n+1)}$ and

$$W^{(1)}(\mathbf{X}) = \tilde{\mathbf{w}}(\mathbf{X}), \qquad G^{(1)}(\mathbf{X}_{[1,n]}) = \tilde{\mathbf{g}}(\mathbf{X}), \qquad \{J^{(k)}(\mathbf{X}_{[k,n]}) = \tilde{\Gamma}^{(k)}(\mathbf{X}_{[k,n]})\}_{k \in [n]}.$$

8 **for** $i \in [n]$ **do**

9 The prover defines and sends polynomials

$$\begin{aligned} S_1^{(i)}(X_i) &= \sum_{\mathbf{c} \in \{0,1\}^{n-i}} W^{(i)}(X_i, \mathbf{c}) \cdot G^{(i)}(X_i, \mathbf{c}) & S_2^{(i)}(X_i) &= \sum_{\mathbf{c} \in \{0,1\}^{n-i}} W^{(i)}(X_i, \mathbf{c}) \cdot J^{(i)}(X_i, \mathbf{c}) \\ S_3^{(i)}(X_i) &= \sum_{\mathbf{c} \in \{0,1\}^{n-i}} G^{(i)}(X_i, \mathbf{c}) \cdot J^{(i)}(X_i, \mathbf{c}) & S_4^{(i)}(X_i) &= \sum_{\mathbf{c} \in \{0,1\}^{n-i}} W^{(i)}(X_i, \mathbf{c}) \cdot J^{(i+1)}(\mathbf{c}) \\ S_5^{(i)}(X_i) &= \sum_{\mathbf{c} \in \{0,1\}^{n-i}} G^{(i)}(X_i, \mathbf{c}) \cdot J^{(i+1)}(\mathbf{c}). \end{aligned}$$

10 The verifier checks $C_j^{(i)} \stackrel{?}{=} S_j^{(i)}(0) + S_j^{(i)}(1)$ for $j \in [3]$ and $\deg(S_j^{(i)}) \leq 2$ for $j \in [5]$.

11 The verifier samples and sends $\alpha_i, \beta_i \xleftarrow{\$} \mathbb{F}$.

12 The verifier defines

$$\begin{aligned} C_1^{(i+1)} &= S_1^{(i)}(\alpha_i) + \beta_i (S_2^{(i)}(\alpha_i) + S_3^{(i)}(\alpha_i)) + \beta_i^2 Q_i(\alpha_i) & C_2^{(i+1)} &= S_4^{(i)}(\alpha_i) + \beta_i \hat{Q}_i(\alpha_i) \\ C_3^{(i+1)} &= S_5^{(i)}(\alpha_i) + \beta_i \hat{Q}_i(\alpha_i) \end{aligned}$$

13 The prover defines

$$\begin{aligned} W^{(i+1)}(\mathbf{X}_{[i+1,n]}) &= W^{(i)}(\alpha_i, \mathbf{X}_{[i+1,n]}) + \beta_i J^{(i)}(\alpha_i, \mathbf{X}_{[i+1,n]}) \\ G^{(i+1)}(\mathbf{X}_{[i+1,n]}) &= G^{(i)}(\alpha_i, \mathbf{X}_{[i+1,n]}) + \beta_i J^{(i)}(\alpha_i, \mathbf{X}_{[i+1,n]}). \end{aligned}$$

14 The prover sends $w^{(n+1)} = W^{(n+1)}$ and $g^{(n+1)} = G^{(n+1)}$ to the verifier.

15 The verifier checks $C_1^{(n+1)} \stackrel{?}{=} w^{(n+1)} \cdot g^{(n+1)}, C_2^{(n+1)} \stackrel{?}{=} w^{(n+1)} \cdot \Gamma^{(n+1)}$, and $C_3^{(n+1)} \stackrel{?}{=} g^{(n+1)} \cdot \Gamma^{(n+1)}$; reject if not all hold.

Given the above high-level overview of Figure 7, we now prove Theorem 6.1.

Proof of Theorem 6.1. We begin by showing that Figure 7 is a valid SMH protocol.

Claim 6.6. Let $N = 2^n$ and let $f(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle$ for any $\mathbf{x}, \mathbf{y} \in \mathbb{F}^N$ and any finite field. Then, Figure 7 is a

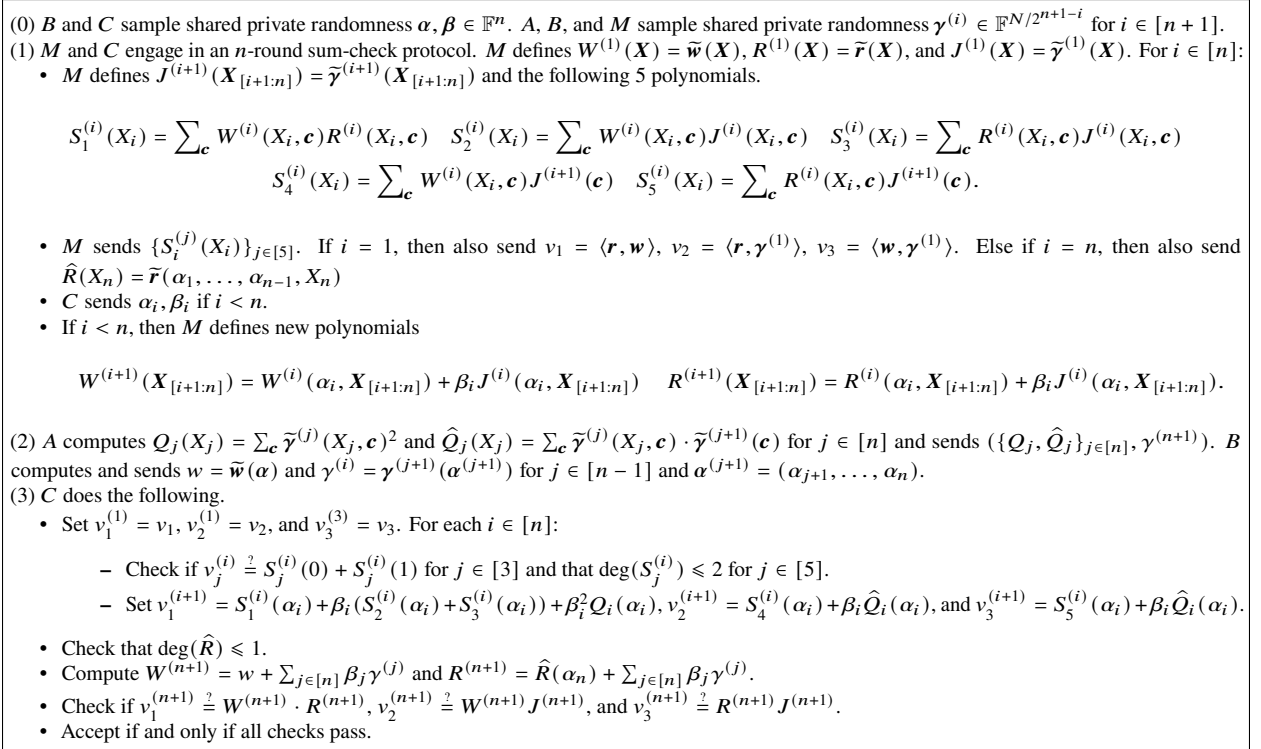



Figure 7: Dory-Sum-Check as a **SMH**[n] Protocol for $N = 2^n$. The red dashed line indicates private randomness shared between the respective parties. The protocol executes (1), followed by (2), followed by (3).

(0, $O(n/|\mathbb{F}|)$)-secure **SMH**[n] protocol for f .

Completeness immediately follows from completeness of the sum-check protocol, along with the observation that taking a random linear combination of two multilinear extensions (i.e., how M computes $W^{(i+1)}$ and $R^{(i+1)}$) is itself a multilinear extension (so all polynomials in the protocol have degree at most 2).

Soundness also immediately follows from the sum-check protocol. During round i of the sum-check, since C is taking a random β_i combination of polynomials, the soundness error incurred is proportional to the largest power of β_i taken in the combination. Notably, C updates $v_1^{(i+1)}$ using β_i^2 , and the other two values

with only β_i . Thus, the total contribution to the soundness error of this round is $4/|\mathbb{F}|$. Additionally, C could sample a bad α_i as a root of one of the polynomials it is checking consistency with. Since all the polynomials have degree at most 2, this again contributes at most $2/|\mathbb{F}|$ to the soundness error per polynomial. So we have an upper bound of $14/|\mathbb{F}|$ for the soundness error in round i . Now, in the final round of the protocol, when M sends the additional message $\hat{R}(X_n)$, M sends this without knowing the final challenge α_n , so it only passes C 's check using $\hat{R}(\alpha_n)$ with probability at most $1/|\mathbb{F}|$. A union bound over all rounds i plus the soundness error of this final check gives us the final bound of $O(n/|\mathbb{F}|)$, establishing the claim. With the claim established, the rest of the proof is trivial. The specified protocol has n rounds, and by definition, the amount of information exchanged between parties is $\Omega(n) = \Omega(\log(N))$ bits. \square

6.3 Evidence that Dory-PCS has Optimal Verifier Time

As with Bulletproofs, (Proto-)Dory is readily adaptable into a polynomial commitment scheme for any polynomials which can express evaluation as an inner product (again, multilinear and univariate polynomials satisfy this property). We present this modification below as we prove [Corollary 6.2](#).

Proof of Corollary 6.2. We modify [Algorithm 6.2](#) to obtain a new scheme satisfying the above property, but for brevity specify it here. The modification we present is first for univariate polynomials; the case of multilinear polynomials is similar.

First, suppose that the prover has private input a univariate polynomial $p \in \mathbb{F}^{<N}[X]$. To commit to p , let $\mathbf{p} = (p_0, \dots, p_{N-1}) \in \mathbb{F}^N$ denote the coefficients of p . We additionally add a new generator $h \in \mathbb{G}$ to the public parameters. Then, the prover first computes a vector $\mathbf{w} = (h \cdot p_0, \dots, h \cdot p_{N-1}) \in \mathbb{G}$ and commits to p via $C_1 = \langle \mathbf{w}, \mathbf{g} \rangle$ for public parameters \mathbf{g} , identical to [Algorithm 6.2](#) (note it also computes $C_2 = \langle \mathbf{w}, \mathbf{\Gamma} \rangle$ as well; C_3 remains the same). Now, at the start of the proof of knowledge phase of [Algorithm 6.2](#), the verifier additionally specifies an evaluation point $r \in \mathbb{F}$. The prover and verifier then run the proof of knowledge of [Algorithm 6.2](#) with the prover using \mathbf{w} defined above, and additionally adding another sum-check polynomial $F(X) := \tilde{\mathbf{p}}(X) \cdot \tilde{\mathbf{r}}(X)$, where $\mathbf{r} := (1, r, r^2, r^3, \dots, r^{N-1})$, and with the prover additionally sending $y \in \mathbb{F}$ in the first round, claimed to be the value $\langle \mathbf{p}, \mathbf{r} \rangle = p(r)$.

In the final round of the proof of knowledge (round $n + 1$), the prover additionally sends a value $q \in \mathbb{F}$, claimed to be $\tilde{\mathbf{p}}(\alpha)$, and a value $\sigma^{(n+1)} \in \mathbb{G}$, claimed to be $\sum_i \beta_i J^{(i)}(\alpha_i, \dots, \alpha_n)$. The verifier additionally checks that $(h \cdot q) + \sigma^{(n+1)} = w^{(n+1)}$, where $w^{(n+1)} \in \mathbb{G}$ is claimed to be $\tilde{\mathbf{w}}(\alpha) + \sum_i \beta_i \tilde{\mathbf{\Gamma}}^{(i)}(\alpha_i, \dots, \alpha_n)$ sent by the prover (i.e., the final folded sum-check polynomial $W^{(n+1)}$). Now, let $y^{(n+1)}$ be the final value obtained from the prover and verifier running the sum-check on the polynomial $F(X)$, claimed to be $F(\alpha)$. In order to verify this check as well, the verifier needs to compute $\tilde{\mathbf{r}}(\alpha)$ and check $q \cdot \tilde{\mathbf{r}}(\alpha) = y^{(n+1)}$.

The verifier takes advantage of the tensor structure of \mathbf{r} to efficiently compute $\tilde{\mathbf{r}}(\alpha)$ in only logarithmic time. To see this, first consider the value $\sum_{\mathbf{c}} \tilde{\mathbf{r}}(\alpha_1, \mathbf{c})$. Using the definition of multilinear extensions, we see that this value is computed as

$$\sum_{\mathbf{c} \in \{0,1\}^{n-1}} \tilde{\mathbf{r}}(\alpha_1, \mathbf{c}) = \sum_{\mathbf{c} \in \{0,1\}^{n-1}} \sum_{\mathbf{b} \in \{0,1\}^n} \mathbf{r}_{\mathbf{b}} \cdot \text{eq}((\alpha_1, \mathbf{c}), \mathbf{b}) = \sum_{\mathbf{c} \in \{0,1\}^{n-1}} \mathbf{r}_{0\mathbf{c}}(1 - \alpha_1) + \mathbf{r}_{1\mathbf{c}}\alpha_1.$$

By definition of \mathbf{r} , $0\mathbf{c}$ indexes the left half of \mathbf{r} and $1\mathbf{c}$ indexes the right half of \mathbf{r} . The left half is given by $\mathbf{r}_L = (1, r, r^2, \dots, r^{N/2-1})$, and the right half is given by $\mathbf{r}_R = (r^{N/2}, r^{N/2+1}, \dots, r^{N-1})$. Notice that

$\mathbf{r}_R = r^{N/2} \cdot \mathbf{r}_L$. So we can rewrite the summation as

$$\begin{aligned} \sum_{\mathbf{c} \in \{0,1\}^{n-1}} \mathbf{r}_{0\mathbf{c}}(1 - \alpha_1) + \mathbf{r}_{1\mathbf{c}}\alpha_1 &= \sum_{\mathbf{c} \in \{0,1\}^{n-1}} \mathbf{r}_{0\mathbf{c}}(1 - \alpha_1) + r^{N/2} \cdot \mathbf{r}_{0\mathbf{c}} \cdot \alpha_1 \\ &= [(1 - \alpha_1) + \alpha_1 r^{N/2}] \sum_{\mathbf{c} \in \{0,1\}^{n-1}} \mathbf{r}_{0\mathbf{c}}. \end{aligned}$$

We can continue this expansion and manipulation for each round i with respect to $\sum_{\mathbf{c}} \tilde{\mathbf{r}}(\alpha_1, \dots, \alpha_i, \mathbf{c})$, which yields the following equality:

$$\tilde{\mathbf{r}}(\alpha) = \prod_{i=1}^n (1 - \alpha_i) + \alpha_i \cdot r^{N/2^i}.$$

This equation tells us that the verifier only needs the powers $r, r^2, r^4, \dots, r^{N/2}$ for $N = 2^n$ in order to compute $\tilde{\mathbf{r}}(\alpha)$, which by repeated squaring are computable in $O(\log(N))$ field operations, so $\tilde{\mathbf{r}}(\alpha)$ is computable in $O(\log(N))$ field operations, culminating in the final verification check by the verifier $y^{(n+1)} \stackrel{?}{=} q \cdot \tilde{\mathbf{r}}(\alpha)$.

Translating this modification of [Algorithm 6.2](#) to a **SMH** $[n]$ protocol, we additionally give party A the input $r \in \mathbb{F}$. Then, C and M execute the same protocol as [Figure 7](#), adding in the minor modifications above. The only additional change is that A now additionally sends C the vector $(r, r^2, r^4, \dots, r^{N/2}) \in \mathbb{F}^{\log(N)}$ so C can perform the additional checks above. The protocol, being $\log(N)$ rounds, still has a communication lower-bound of $\Omega(\log(N))$, which is (nearly) matched by the communication sent in the **SMH** $[n]$ protocol and the work of the verifier in [Algorithm 6.2](#) (up to polylog $|\mathbb{G}|$ factors).

For multilinear polynomials, the modifications to [Algorithm 6.2](#) are nearly identical to the above for univariate polynomials. Now, given $p: \{0, 1\}^n \rightarrow \mathbb{F}$ as a multilinear polynomial, the prover commits to the vector $\mathbf{p} = (p(0), p(1), \dots, p(N-1))$ (the vector of evaluations) by computing $\mathbf{w} = (h \cdot p(0), h \cdot p(1), \dots, h \cdot p(N-1))$ as above. The only other modification we need to address is how to efficiently compute $\tilde{\mathbf{r}}(\alpha)$ for a vector $\mathbf{r} \in \mathbb{F}^N$ such that $\langle \mathbf{p}, \mathbf{r} \rangle = p(\mathbf{u})$ for $\mathbf{u} \in \mathbb{F}^n$. First, using [Eq. \(1\)](#) (and work we already did in [Section 5.3](#)), the vector \mathbf{r} is defined as $\mathbf{r}_i = \text{eq}(\mathbf{u}, i-1)$ for $i \in [N]$.

We just need to argue that \mathbf{r} also has sufficient tensor structure to allow the verifier to compute $\tilde{\mathbf{r}}(\alpha)$ in logarithmic time. As before, first consider $\sum_{\mathbf{c}} \tilde{\mathbf{r}}(\alpha_1, \mathbf{c})$, which can be written as

$$\sum_{\mathbf{c} \in \{0,1\}^{n-1}} \tilde{\mathbf{r}}(\alpha_1, \mathbf{c}) = \sum_{\mathbf{c} \in \{0,1\}^{n-1}} \mathbf{r}_{0\mathbf{c}}(1 - \alpha_1) + \mathbf{r}_{1\mathbf{c}}\alpha_1.$$

Let $\mathbf{u}^{(1)} := \mathbf{u}$ and let $\mathbf{u}^{(2)} := (\mathbf{u}_2, \dots, \mathbf{u}_n)$. By definition of \mathbf{r} , we see that

$$\begin{aligned} \mathbf{r}_{0\mathbf{c}} &= \text{eq}(\mathbf{u}^{(1)}, 0\mathbf{c}) = (1 - \mathbf{u}_1) \cdot \prod_{j=1}^n (1 - \mathbf{u}_j)(1 - \mathbf{c}_j) + \mathbf{u}_j \mathbf{c}_j = (1 - \mathbf{u}_1) \cdot \text{eq}(\mathbf{u}^{(2)}, \mathbf{c}) \\ \mathbf{r}_{1\mathbf{c}} &= \text{eq}(\mathbf{u}^{(1)}, 1\mathbf{c}) = \mathbf{u}_1 \cdot \prod_{j=1}^n (1 - \mathbf{u}_j)(1 - \mathbf{c}_j) + \mathbf{u}_j \mathbf{c}_j = \mathbf{u}_1 \cdot \text{eq}(\mathbf{u}^{(2)}, \mathbf{c}). \end{aligned}$$

Thus, we can simplify the summation as

$$\begin{aligned} \sum_{\mathbf{c} \in \{0,1\}^{n-1}} \mathbf{r}_{0\mathbf{c}}(1 - \alpha_1) + \mathbf{r}_{1\mathbf{c}}\alpha_1 &= \sum_{\mathbf{c} \in \{0,1\}^{n-1}} (1 - \mathbf{u}_1) \text{eq}(\mathbf{u}^{(2)}, \mathbf{c})(1 - \alpha_1) + \mathbf{u}_1 \text{eq}(\mathbf{u}^{(2)}, \mathbf{c})\alpha_1 \\ &= [(1 - \alpha_1)(1 - \mathbf{u}_1) + \alpha_1 \mathbf{u}_1] \sum_{\mathbf{c} \in \{0,1\}^{n-1}} \text{eq}(\mathbf{u}^{(2)}, \mathbf{c}). \end{aligned}$$

As before, continuing to expand the summation for $\tilde{r}(\alpha_1, \dots, \alpha_i, \mathbf{c})$, we can write $\tilde{r}(\alpha)$ as:

$$\tilde{r}(\alpha) = \prod_{i=1}^n (1 - \alpha_i)(1 - \mathbf{u}_i) + \alpha_i \mathbf{u}_i.$$

Clearly, this computation can be done in $O(\log(N)) = O(n)$ field operations. Finally, in the **SMH**[n] protocol, party A receives the additional input $\mathbf{u} \in \mathbb{F}^n$ and simply sends \mathbf{u} to C at the end of the protocol, and C computes $\tilde{r}(\alpha)$ as specified above, yielding the same upper and lower bounds as in the univariate case. \square

7 Conclusions and Open Problems

In this work, we initiate the study of the connection between SNARK verification times and the communication complexity of the underlying information-theoretic cores of the SNARKs. Our work examines the polynomial commitment schemes of Hyrax, Bulletproofs, and Dory, extracts their information-theoretic cores as communication protocols, and lower bounds the number of bits that must be exchanged between parties in these protocols. This connection gives us strong evidence that these polynomial commitment schemes all have optimal verifier time. Moreover, our results unfortunately support evidence that the Bulletproofs verifier is inherently linear-time, possibly settling a long-standing open question.

Natural future directions in this line of work are extending the results in this paper to other polynomial commitment schemes. For example, based on groups of unknown order [BFS20, BHR⁺21], based on the FRI protocol [BBH⁺18, ZCF24, ACF⁺24], based on Ligero/Brakedown [AHI⁺23, BFH⁺20, BBH⁺22, GLS⁺23, BFK⁺24], and the KZG scheme [KZG10]. There is also potential to extend our results to the verification costs of *lookup arguments* [PT12, PT13, GW20, STW24, AST24]. Finally, it is natural to question whether the connections we made in this paper can extend to the Groth16 SNARK [Gro16]. We believe these are all exciting and promising future directions to explore.

Another direction is towards understanding the connections we began to make in this paper. Specifically, can we strengthen the connections between polynomial commitment schemes and communication protocols, and get provable guarantees to actually prove optimality of the verifier of polynomial commitments (and potentially other cryptographic protocols) using communication complexity? To the best of our knowledge, we are the first to explore such a connection and the first to consider verifier time lower bounds, so there is immense potential in this line of work to gain a deeper understanding of polynomial commitments and SNARKs at a deeper level.

We also believe that our **SMH** model is of independent interest in the field of communication complexity. Though we can simulate any **SMH** model protocol as an **OMA** model protocol, it is natural to wonder if we can obtain different (better or worse) lower bounds on computing functions in the **SMH** model versus what is known for the **OMA** model.

References

- [Ab196] Farid M. Ablayev. Lower bounds for one-way probabilistic communication complexity and their application to space complexity. *Theor. Comput. Sci.*, 157(2):139–159, 1996. doi: [10.1016/0304-3975\(95\)00157-3](https://doi.org/10.1016/0304-3975(95)00157-3). URL: [https://doi.org/10.1016/0304-3975\(95\)00157-3](https://doi.org/10.1016/0304-3975(95)00157-3).
- [ACF⁺24] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. STIR: reed-solomon proximity testing with fewer queries. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part X*, volume 14929 of *Lecture Notes in Computer Science*, pages 380–413. Springer, 2024. doi: [10.1007/978-3-031-68403-6_12](https://doi.org/10.1007/978-3-031-68403-6_12). URL: https://doi.org/10.1007/978-3-031-68403-6_12.
- [ACL⁺22] Martin R. Albrecht, Valerio Cini, Russell W. F. Lai, Giulio Malavolta, and Sri Aravinda Krishnan Thyagarajan. Lattice-based snarks: publicly verifiable, preprocessing, and recursively composable - (extended abstract). In *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 102–132. Springer, 2022. URL: https://doi.org/10.1007/978-3-031-15979-4_4.
- [AFG⁺16] Masayuki Abe, Georg Fuchsbaauer, Jens Groth, Kristiyan Haralambiev, and Miyako Ohkubo. Structure-preserving signatures and commitments to group elements. *J. Cryptol.*, 29(2):363–421, 2016. doi: [10.1007/S00145-014-9196-7](https://doi.org/10.1007/S00145-014-9196-7). URL: <https://doi.org/10.1007/s00145-014-9196-7>.
- [AFL⁺24] Martin R. Albrecht, Giacomo Fenzi, Oleksandra Lapiha, and Ngoc Khanh Nguyen. SLAP: succinct lattice-based polynomial commitments from standard assumptions. In *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part VI*, volume 14656 of *Lecture Notes in Computer Science*, pages 90–119. Springer, 2024. URL: https://doi.org/10.1007/978-3-031-58754-2_4.
- [AHI⁺23] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramaniam. Liger: lightweight sublinear arguments without a trusted setup. *Des. Codes Cryptogr.*, 91(11):3379–3424, 2023. URL: <https://doi.org/10.1007/s10623-023-01222-8>.
- [AST24] Arasu Arun, Srinath T. V. Setty, and Justin Thaler. Jolt: snarks for virtual machines via lookups. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part VI*, volume 14656 of *Lecture Notes in Computer Science*, pages 3–33. Springer, 2024. doi: [10.1007/978-3-031-58751-1_1](https://doi.org/10.1007/978-3-031-58751-1_1). URL: https://doi.org/10.1007/978-3-031-58751-1_1.
- [AW09] Scott Aaronson and Avi Wigderson. Algebrization: A new barrier in complexity theory. *ACM Trans. Comput. Theory*, 1(1):2:1–2:54, 2009. doi: [10.1145/1490270.1490272](https://doi.org/10.1145/1490270.1490272). URL: <https://doi.org/10.1145/1490270.1490272>.
- [BBB⁺18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Gregory Maxwell. Bulletproofs: short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*, pages 315–334. IEEE Computer Society, 2018. doi: [10.1109/SP.2018.00020](https://doi.org/10.1109/SP.2018.00020). URL: <https://doi.org/10.1109/SP.2018.00020>.

- [BBH⁺18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast reed-solomon interactive oracle proofs of proximity. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, Prague, Czech Republic, July 9-13, 2018*, volume 107 of *LIPIcs*, 14:1–14:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. DOI: [10.4230/LIPICS.ICALP.2018.14](https://doi.org/10.4230/LIPICS.ICALP.2018.14). URL: <https://doi.org/10.4230/LIPICS.ICALP.2018.14>.
- [BBH⁺22] Laasya Bangalore, Rishabh Bhadauria, Carmit Hazay, and Muthuramakrishnan Venkitasubramaniam. On black-box constructions of time and space efficient sublinear arguments from symmetric-key primitives. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography - 20th International Conference, TCC 2022, Chicago, IL, USA, November 7-10, 2022, Proceedings, Part I*, volume 13747 of *Lecture Notes in Computer Science*, pages 417–446. Springer, 2022. DOI: [10.1007/978-3-031-22318-1_15](https://doi.org/10.1007/978-3-031-22318-1_15). URL: https://doi.org/10.1007/978-3-031-22318-1_15.
- [BCC⁺16] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 327–357. Springer, 2016. DOI: [10.1007/978-3-662-49896-5_12](https://doi.org/10.1007/978-3-662-49896-5_12). URL: https://doi.org/10.1007/978-3-662-49896-5_12.
- [BCF⁺22] David Balbás, Dario Catalano, Dario Fiore, and Russell W. F. Lai. Functional commitments for circuits from falsifiable assumptions. *IACR Cryptol. ePrint Arch.*:1365, 2022. URL: <https://eprint.iacr.org/2022/1365>.
- [BCS21] Jonathan Bootle, Alessandro Chiesa, and Katerina Sotiraki. Sumcheck arguments and their applications. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 742–773. Springer, 2021. DOI: [10.1007/978-3-030-84242-0_26](https://doi.org/10.1007/978-3-030-84242-0_26). URL: https://doi.org/10.1007/978-3-030-84242-0_26.
- [BFH⁺20] Rishabh Bhadauria, Zhiyong Fang, Carmit Hazay, Muthuramakrishnan Venkitasubramaniam, Tiancheng Xie, and Yupeng Zhang. Liger++: A new optimized sublinear IOP. In *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*, pages 2025–2038. ACM, 2020. URL: <https://doi.org/10.1145/3372297.3417893>.
- [BFK⁺24] Alexander R. Block, Zhiyong Fang, Jonathan Katz, Justin Thaler, Hendrik Waldner, and Yupeng Zhang. Field-agnostic snarks from expand-accumulate codes. In *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part X*, volume 14929 of *Lecture Notes in Computer Science*, pages 276–307. Springer, 2024. URL: https://doi.org/10.1007/978-3-031-68403-6_9.
- [BFS20] Benedikt Bünz, Ben Fisch, and Alan Szeponiec. Transparent snarks from DARK compilers. In *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 677–706. Springer, 2020. URL: https://doi.org/10.1007/978-3-030-45721-1_24.

- [BFS86] László Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory (preliminary version). In *27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986*, pages 337–347. IEEE Computer Society, 1986. doi: [10.1109/SFCS.1986.15](https://doi.org/10.1109/SFCS.1986.15). URL: <https://doi.org/10.1109/SFCS.1986.15>.
- [BGK⁺03] László Babai, Anna Gál, Peter G. Kimmel, and Satyanarayana V. Lokam. Communication complexity of simultaneous messages. *SIAM J. Comput.*, 33(1):137–166, 2003. doi: [10.1137/S0097539700375944](https://doi.org/10.1137/S0097539700375944). URL: <https://doi.org/10.1137/S0097539700375944>.
- [BHR⁺20] Alexander R. Block, Justin Holmgren, Alon Rosen, Ron D. Rothblum, and Pratik Soni. Public-coin zero-knowledge arguments with (almost) minimal time and space overheads. In *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part II*, volume 12551 of *Lecture Notes in Computer Science*, pages 168–197. Springer, 2020. URL: https://doi.org/10.1007/978-3-030-64378-2_7.
- [BHR⁺21] Alexander R. Block, Justin Holmgren, Alon Rosen, Ron D. Rothblum, and Pratik Soni. Time- and space-efficient arguments from groups of unknown order. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part IV*, volume 12828 of *Lecture Notes in Computer Science*, pages 123–152. Springer, 2021. doi: [10.1007/978-3-030-84259-8_5](https://doi.org/10.1007/978-3-030-84259-8_5). URL: https://doi.org/10.1007/978-3-030-84259-8_5.
- [BHV⁺23] Rishabh Bhaduria, Carmit Hazay, Muthuramakrishnan Venkitasubramaniam, Wenxuan Wu, and Yupeng Zhang. Private polynomial commitments and applications to MPC. In *Public-Key Cryptography - PKC 2023 - 26th IACR International Conference on Practice and Theory of Public-Key Cryptography, Atlanta, GA, USA, May 7-10, 2023, Proceedings, Part II*, volume 13941 of *Lecture Notes in Computer Science*, pages 127–158. Springer, 2023. URL: https://doi.org/10.1007/978-3-031-31371-4_5.
- [BJS⁺25] Elette Boyle, Abhishek Jain, Sacha Servan-Schreiber, and Akshayaram Srinivasan. Simultaneous-message and succinct secure computation. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part V*, volume 15605 of *Lecture Notes in Computer Science*, pages 207–239. Springer, 2025. doi: [10.1007/978-3-031-91092-0_8](https://doi.org/10.1007/978-3-031-91092-0_8). URL: https://doi.org/10.1007/978-3-031-91092-0_8.
- [BK97] László Babai and Peter G. Kimmel. Randomized simultaneous messages: solution of a problem of yao in communication complexity. In *Proceedings of the Twelfth Annual IEEE Conference on Computational Complexity, Ulm, Germany, June 24-27, 1997*, pages 239–246. IEEE Computer Society, 1997. doi: [10.1109/CCC.1997.612319](https://doi.org/10.1109/CCC.1997.612319). URL: <https://doi.org/10.1109/CCC.1997.612319>.
- [BKL95] László Babai, Peter G. Kimmel, and Satyanarayana V. Lokam. Simultaneous messages vs. communication. In Ernst W. Mayr and Claude Puech, editors, *STACS 95, 12th Annual Symposium on Theoretical Aspects of Computer Science, Munich, Germany, March 2-4, 1995, Proceedings*, volume 900 of *Lecture Notes in Computer Science*, pages 361–372. Springer, 1995. doi: [10.1007/3-540-59042-0_88](https://doi.org/10.1007/3-540-59042-0_88). URL: https://doi.org/10.1007/3-540-59042-0_88.

- [CCG⁺14] Amit Chakrabarti, Graham Cormode, Navin Goyal, and Justin Thaler. Annotations for sparse data streams. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 687–706. SIAM, 2014. DOI: [10.1137/1.9781611973402.52](https://doi.org/10.1137/1.9781611973402.52). URL: <https://doi.org/10.1137/1.9781611973402.52>.
- [CCM⁺14] Amit Chakrabarti, Graham Cormode, Andrew McGregor, and Justin Thaler. Annotations in data streams. *ACM Trans. Algorithms*, 11(1):7:1–7:30, 2014. DOI: [10.1145/2636924](https://doi.org/10.1145/2636924). URL: <https://doi.org/10.1145/2636924>.
- [CCM⁺19] Amit Chakrabarti, Graham Cormode, Andrew McGregor, Justin Thaler, and Suresh Venkatasubramanian. Verifiable stream computation and arthur-merlin communication. *SIAM J. Comput.*, 48(4):1265–1299, 2019. DOI: [10.1137/17M112289X](https://doi.org/10.1137/17M112289X). URL: <https://doi.org/10.1137/17M112289X>.
- [CDG⁺24] Graham Cormode, Marcel Dall’Agnol, Tom Gur, and Chris Hickey. Streaming zero-knowledge proofs. In Rahul Santhanam, editor, *39th Computational Complexity Conference, CCC 2024, Ann Arbor, MI, USA, July 22-25, 2024*, volume 300 of *LIPICs*, 2:1–2:66. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. DOI: [10.4230/LIPICs.CCC.2024.2](https://doi.org/10.4230/LIPICs.CCC.2024.2). URL: <https://doi.org/10.4230/LIPICs.CCC.2024.2>.
- [CFL83] Ashok K. Chandra, Merrick L. Furst, and Richard J. Lipton. Multi-party protocols. In David S. Johnson, Ronald Fagin, Michael L. Fredman, David Harel, Richard M. Karp, Nancy A. Lynch, Christos H. Papadimitriou, Ronald L. Rivest, Walter L. Ruzzo, and Joel I. Seiferas, editors, *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 94–99. ACM, 1983. DOI: [10.1145/800061.808737](https://doi.org/10.1145/800061.808737). URL: <https://doi.org/10.1145/800061.808737>.
- [CMN⁺24] Valerio Cini, Giulio Malavolta, Ngoc Khanh Nguyen, and Hoeteck Wee. Polynomial commitments from lattices: post-quantum security, fast verification and transparent setup. In *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part X*, volume 14929 of *Lecture Notes in Computer Science*, pages 207–242. Springer, 2024. URL: https://doi.org/10.1007/978-3-031-68403-6_7.
- [CMT13] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Streaming graph computations with a helpful advisor. *Algorithmica*, 65(2):409–442, 2013. DOI: [10.1007/s00453-011-9598-Y](https://doi.org/10.1007/s00453-011-9598-Y). URL: <https://doi.org/10.1007/s00453-011-9598-y>.
- [FMN24] Giacomo Fenzi, Hossein Moghaddas, and Ngoc Khanh Nguyen. Lattice-based polynomial commitments: towards asymptotic and concrete efficiency. *J. Cryptol.*, 37(3):31, 2024. URL: <https://doi.org/10.1007/s00145-024-09511-8>.
- [Fre77] Rusins Freivalds. Probabilistic machines can use less running time. In Bruce Gilchrist, editor, *Information Processing, Proceedings of the 7th IFIP Congress 1977, Toronto, Canada, August 8-12, 1977*, pages 839–842. North-Holland, 1977.
- [FS86] Amos Fiat and Adi Shamir. How to prove yourself: practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology - CRYPTO ’86, Santa Barbara, California, USA, 1986, Proceedings*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, 1986. DOI: [10.1007/3-540-47721-7_12](https://doi.org/10.1007/3-540-47721-7_12). URL: https://doi.org/10.1007/3-540-47721-7_12.

- [Gab24] Ariel Gabizon. https://x.com/rel_zeta_tech/status/1794842766708535544, May 2024. See also https://x.com/rel_zeta_tech/status/1794780789789376774.
- [GLS⁺23] Alexander Golovnev, Jonathan Lee, Srinath T. V. Setty, Justin Thaler, and Riad S. Wahby. Brakedown: linear-time and field-agnostic snarks for R1CS. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part II*, volume 14082 of *Lecture Notes in Computer Science*, pages 193–226. Springer, 2023. DOI: [10.1007/978-3-031-38545-2_7](https://doi.org/10.1007/978-3-031-38545-2_7). URL: https://doi.org/10.1007/978-3-031-38545-2_7.
- [Gro09] Jens Groth. Homomorphic trapdoor commitments to group elements. *IACR Cryptol. ePrint Arch.*:7, 2009. URL: <http://eprint.iacr.org/2009/007>.
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 305–326. Springer, 2016. DOI: [10.1007/978-3-662-49896-5_11](https://doi.org/10.1007/978-3-662-49896-5_11). URL: https://doi.org/10.1007/978-3-662-49896-5_11.
- [GS24] Prantar Ghosh and Vihan Shah. New lower bounds in merlin-arthur communication and graph streaming verification. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference, ITCS 2024, Berkeley, CA, USA, January 30 - February 2, 2024*, volume 287 of *LIPICs*, 53:1–53:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. DOI: [10.4230/LIPICs.ITCS.2024.53](https://doi.org/10.4230/LIPICs.ITCS.2024.53). URL: <https://doi.org/10.4230/LIPICs.ITCS.2024.53>.
- [GW20] Ariel Gabizon and Zachary J. Williamson. Plookup: A simplified polynomial protocol for lookup tables. *IACR Cryptol. ePrint Arch.*:315, 2020. URL: <https://eprint.iacr.org/2020/315>.
- [HSS24] Intak Hwang, Jinyeong Seo, and Yongsoo Song. Concretely efficient lattice-based polynomial commitment from standard assumptions. In *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part X*, volume 14929 of *Lecture Notes in Computer Science*, pages 414–448. Springer, 2024. URL: https://doi.org/10.1007/978-3-031-68403-6_13.
- [Kla03] Hartmut Klauck. Rectangle size bounds and threshold covers in communication complexity. In *18th Annual IEEE Conference on Computational Complexity (Complexity 2003), 7-10 July 2003, Aarhus, Denmark*, pages 118–134. IEEE Computer Society, 2003. DOI: [10.1109/CCC.2003.1214415](https://doi.org/10.1109/CCC.2003.1214415). URL: <https://doi.org/10.1109/CCC.2003.1214415>.
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997. ISBN: 978-0-521-56067-2.
- [KS92] Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. *SIAM J. Discret. Math.*, 5(4):545–557, 1992. DOI: [10.1137/0405044](https://doi.org/10.1137/0405044). URL: <https://doi.org/10.1137/0405044>.
- [KS93] Tracy Kimbrel and Rakesh K. Sinha. A probabilistic algorithm for verifying matrix products using $O(n^2)$ time and $\log_2 n + O(1)$ random bits. *Inf. Process. Lett.*, 45(2):107–110, 1993. DOI: [10.1016/0020-0190\(93\)90224-W](https://doi.org/10.1016/0020-0190(93)90224-W). URL: [https://doi.org/10.1016/0020-0190\(93\)90224-W](https://doi.org/10.1016/0020-0190(93)90224-W).

- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *Lecture Notes in Computer Science*, pages 177–194. Springer, 2010. DOI: [10.1007/978-3-642-17373-8_11](https://doi.org/10.1007/978-3-642-17373-8_11). URL: https://doi.org/10.1007/978-3-642-17373-8_11.
- [Lee21] Jonathan Lee. Dory: efficient, transparent arguments for generalised inner products and polynomial commitments. In Kobbi Nissim and Brent Waters, editors, *Theory of Cryptography - 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8-11, 2021, Proceedings, Part II*, volume 13043 of *Lecture Notes in Computer Science*, pages 1–34. Springer, 2021. DOI: [10.1007/978-3-030-90453-1_1](https://doi.org/10.1007/978-3-030-90453-1_1). URL: https://doi.org/10.1007/978-3-030-90453-1_1.
- [LFK⁺92] Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992. DOI: [10.1145/146585.146605](https://doi.org/10.1145/146585.146605). URL: <https://doi.org/10.1145/146585.146605>.
- [Lu22] Frank Y. C. Lu. Kevlar: transparent, efficient, polynomial commitment scheme with logarithmic verification and communication costs on efficient groups. *IACR Cryptol. ePrint Arch.*:702, 2022. URL: <https://eprint.iacr.org/2022/702>. Withdrawn.
- [Lu24] Frank Y. C. Lu. Smartbean: transparent, concretely efficient, polynomial commitment scheme with logarithmic verification and communication costs that runs on any group. *IACR Cryptol. ePrint Arch.*:785, 2024. URL: <https://eprint.iacr.org/2024/785>. Withdrawn.
- [NS24] Ngoc Khanh Nguyen and Gregor Seiler. Greyhound: fast polynomial commitments from lattices. In *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part X*, volume 14929 of *Lecture Notes in Computer Science*, pages 243–275. Springer, 2024. URL: https://doi.org/10.1007/978-3-031-68403-6_8.
- [Ped91] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *Advances in Cryptology - CRYPTO '91, 11th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1991, Proceedings*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140. Springer, 1991. DOI: [10.1007/3-540-46766-1_9](https://doi.org/10.1007/3-540-46766-1_9). URL: https://doi.org/10.1007/3-540-46766-1_9.
- [Pip80] Nicholas Pippenger. On the evaluation of powers and monomials. *SIAM J. Comput.*, 9(2):230–250, 1980. DOI: [10.1137/0209022](https://doi.org/10.1137/0209022). URL: <https://doi.org/10.1137/0209022>.
- [PT12] Mihai Pătraşcu and Mikkel Thorup. The power of simple tabulation hashing. *J. ACM*, 59(3):14:1–14:50, 2012. DOI: [10.1145/2220357.2220361](https://doi.org/10.1145/2220357.2220361). URL: <https://doi.org/10.1145/2220357.2220361>.
- [PT13] Mihai Pătraşcu and Mikkel Thorup. Twisted tabulation hashing. In Sanjeev Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 209–228. SIAM, 2013. DOI: [10.1137/1.9781611973105.16](https://doi.org/10.1137/1.9781611973105.16). URL: <https://doi.org/10.1137/1.9781611973105.16>.

- [STW24] Srinath T. V. Setty, Justin Thaler, and Riad S. Wahby. Unlocking the lookup singularity with lasso. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part VI*, volume 14656 of *Lecture Notes in Computer Science*, pages 180–209. Springer, 2024. doi: [10.1007/978-3-031-58751-1_7](https://doi.org/10.1007/978-3-031-58751-1_7). URL: https://doi.org/10.1007/978-3-031-58751-1_7.
- [SWY⁺21] Xiaoming Sun, David P. Woodruff, Guang Yang, and Jialin Zhang. Querying a matrix through matrix-vector products. *ACM Trans. Algorithms*, 17(4):31:1–31:19, 2021. doi: [10.1145/3470566](https://doi.org/10.1145/3470566). URL: <https://doi.org/10.1145/3470566>.
- [Tha22] Justin Thaler. Proofs, arguments, and zero-knowledge. *Found. Trends Priv. Secur.*, 4(2-4):117–660, 2022. doi: [10.1561/33000000030](https://doi.org/10.1561/33000000030). URL: <https://doi.org/10.1561/33000000030>.
- [WTS⁺18] Riad S. Wahby, Ioanna Tzialla, Abhi Shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*, pages 926–943. IEEE Computer Society, 2018. doi: [10.1109/SP.2018.00060](https://doi.org/10.1109/SP.2018.00060). URL: <https://doi.org/10.1109/SP.2018.00060>.
- [XZZ⁺19] Tiancheng Xie, Jiaheng Zhang, Yupeng Zhang, Charalampos Papamanthou, and Dawn Song. Libra: succinct zero-knowledge proofs with optimal prover computation. In *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 733–764. Springer, 2019. URL: https://doi.org/10.1007/978-3-030-26954-8_24.
- [Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In Michael J. Fischer, Richard A. DeMillo, Nancy A. Lynch, Walter A. Burkhard, and Alfred V. Aho, editors, *Proceedings of the 11th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA*, pages 209–213. ACM, 1979. doi: [10.1145/800135.804414](https://doi.org/10.1145/800135.804414). URL: <https://doi.org/10.1145/800135.804414>.
- [ZCF24] Hadas Zeilberger, Binyi Chen, and Ben Fisch. Basefold: efficient field-agnostic polynomial commitment schemes from foldable codes. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part X*, volume 14929 of *Lecture Notes in Computer Science*, pages 138–169. Springer, 2024. doi: [10.1007/978-3-031-68403-6_5](https://doi.org/10.1007/978-3-031-68403-6_5). URL: https://doi.org/10.1007/978-3-031-68403-6_5.