



# Libra: Succinct Zero-Knowledge Proofs with Optimal Prover Computation

Tiacheng Xie<sup>1</sup>(✉), Jiaheng Zhang<sup>1</sup>(✉), Yupeng Zhang<sup>1,2</sup>,  
Charalampos Papamanthou<sup>3</sup>, and Dawn Song<sup>1</sup>

<sup>1</sup> University of California, Berkeley, USA  
{[tianc.x](mailto:tianc.x@berkeley.edu), [jiaheng\\_zhang](mailto:jiaheng_zhang@berkeley.edu), [dawnsong](mailto:dawnsong@berkeley.edu)}@berkeley.edu

<sup>2</sup> Texas A&M University, College Station, USA  
[zhangyp@tamu.edu](mailto:zhangyp@tamu.edu)

<sup>3</sup> University of Maryland, College Park, USA  
[cpap@umd.edu](mailto:cpap@umd.edu)

**Abstract.** We present **Libra**, the first zero-knowledge proof system that has both optimal prover time and succinct proof size/verification time. In particular, if  $C$  is the size of the circuit being proved (i) the prover time is  $O(C)$  irrespective of the circuit type; (ii) the proof size and verification time are both  $O(d \log C)$  for  $d$ -depth log-space uniform circuits (such as RAM programs). In addition **Libra** features an one-time trusted setup that depends only on the size of the input to the circuit and not on the circuit logic. Underlying **Libra** is a new linear-time algorithm for the prover of the interactive proof protocol by Goldwasser, Kalai and Rothblum (also known as GKR protocol), as well as an efficient approach to turn the GKR protocol to zero-knowledge using small masking polynomials. Not only does **Libra** have excellent asymptotics, but it is also efficient in practice. For example, our implementation shows that it takes 200s to generate a proof for constructing a SHA2-based Merkle tree root on 256 leaves, outperforming all existing zero-knowledge proof systems. Proof size and verification time of **Libra** are also competitive.

## 1 Introduction

Zero-knowledge proofs (ZKP) are cryptographic protocols between two parties, a *prover* and a *verifier*, in which the prover can convince the verifier about the validity of a statement without leaking any extra information beyond the fact that the statement is true. Since they were first introduced by Goldwasser et al. [31], ZKP protocols have evolved from pure theoretical constructs to practical implementations, achieving proof sizes of just hundreds of bytes and verification times of several milliseconds, regardless of the size of the statement being proved. Due to this successful transition to practice, ZKP protocols have found numerous applications not only in the traditional computation delegation setting but most importantly in providing privacy of transactions in deployed cryptocurrencies (e.g., Zcash [9]) as well as in other blockchain research projects (e.g., Hawk [37]).

Despite such progress in practical implementations, ZKP protocols are still notoriously hard to scale for large statements, due to a particularly high overhead on generating the proof. For most systems, this is primarily because the

prover has to perform a large number of cryptographic operations, such as exponentiation in an elliptic curve group. And to make things worse the asymptotic complexity of computing the proof is typically more than linear, e.g.,  $O(C \log C)$  or even  $O(C \log^2 C)$ , where  $C$  is the size of the statement.

Unfortunately, as of today we are yet to construct a ZKP system whose prover time is *optimal*, i.e., linear in the size of the statement  $C$  (this is irrespective of whether the ZKP system has per-statement trusted setup, one-time trusted setup or no trusted setup at all). The only notable exception is the recent work by Bünz et al. [16] that however suffers from linear verification time—for a detailed comparison see Table 1. Therefore designing ZKP systems that enjoy linear prover time as well as succinct<sup>1</sup> proof size and verification time is an open problem, whose resolution can have significant practical implications.

**Our Contributions.** In this paper we propose *Libra*, the first ZKP protocol with *linear prover time* and *succinct proof size and verification time* in the size of the arithmetic circuit representing the statement  $C$ , when the circuit is *log-space uniform*. *Libra* is based on the doubly efficient interactive proof protocol proposed by Goldwasser et al. in [30] (referred as GKR protocol in this paper), and the verifiable polynomial delegation scheme proposed by Zhang et al. in [50]. As such it comes with *one-time trusted* setup (and not per-statement trusted setup) that depends only on the size of the input (witness) to the statement that is being proved. Not only does *Libra* have excellent asymptotic performance but also its prover outperforms in practice all other ZKP systems while verification time and proof size are also very competitive—see Table 1. Our concrete contributions are:

- **GKR with linear prover time.** *Libra* features a new linear-time algorithm to generate a GKR proof. Our new algorithm does not require any pattern in the circuit and our result subsumes all existing improvements on the GKR prover assuming special circuit structures, such as regular circuits in [43], data-parallel circuits in [43, 46], circuits with different sub-copies in [51]. See related work for more details.
- **Adding zero-knowledge.** We propose an approach to turn *Libra* into zero-knowledge efficiently. In particular, we show a way to mask the responses of our linear-time prover with small random polynomials such that the zero-knowledge variant of the protocol introduces minimal overhead on the verification time compared to the original (unmasked) construction.
- **Implementation and evaluation.** We implement *Libra*. Our implementation takes an arithmetic circuit with various types of gates (fan-in 2 and degree  $\leq 2$ , such as  $+$ ,  $-$ ,  $\times$ , AND, XOR, etc.) and compiles it into a ZKP protocol. We conduct thorough comparisons to all existing ZKP systems (see Sect. 1.1). We plan to release our system as an open-source implementation.

## 1.1 Comparing to Other ZKP Systems

Table 1 shows a detailed comparison between *Libra* and existing ZKP systems. First of all, *Libra* is the best among all existing systems in terms of practical

<sup>1</sup> In ZKP literature, “succinct” is poly-logarithmic in the size of the statement  $C$ .

**Table 1.** Comparison of **Libra** to existing ZKP systems, where  $(\mathcal{G}, \mathcal{P}, \mathcal{V}, |\pi|)$  denote the trusted setup algorithm, the prover algorithm, the verification algorithm and the proof size respectively. Also,  $C$  is the size of the log-space uniform circuit with depth  $d$ , and  $n$  is the size of its input. The numbers are for a circuit computing the root of a Merkle tree with 256 leaves (511 instances of SHA256).

	libSNARK [13]	Ligero [6]	Bulletproofs [16]	Hyrax [48]	libSTARK [8]	Aurora [11]	Libra
$\mathcal{G}$	$O(C)$ per-statement trusted setup	no trusted setup					$O(n)$ one-time trusted setup
$\mathcal{P}$	$O(C \log C)$	$O(C \log C)$	$O(C)$	$O(C \log C)$	$O(C \log^2 C)$	$O(C \log C)$	$O(C)$
$\mathcal{V}$	$O(1)$	$O(C)$	$O(C)$	$O(\sqrt{n} + d \log C)$	$O(\log^2 C)$	$O(C)$	$O(d \log C)$
$ \pi $	$O(1)$	$O(\sqrt{C})$	$O(\log C)$	$O(\sqrt{n} + d \log C)$	$O(\log^2 C)$	$O(\log^2 C)$	$O(d \log C)$
$\mathcal{G}$	1027 s	NA					210 s
$\mathcal{P}$	360 s	400 s	13,000 s	1,041 s	2,022 s	3199 s	201 s
$\mathcal{V}$	0.002 s	4 s	900 s	9.9 s	0.044 s	15.2 s	0.71 s
$ \pi $	0.13 KB	1,500 KB	5.5 KB	185 KB	395 KB	174.3 KB	51 KB

prover time. In terms of asymptotics, **Libra** is the only system with linear prover time and succinct verification and proof size for log-space uniform circuits. The only other system with linear prover time is **Bulletproofs** [16] whose verification time is linear, *even for log-space uniform circuits*. In the practical front, **Bulletproofs** prover time and verification time are high, due to the large number of cryptographic operations required for every gate of the circuit.

The proof and verification of **Libra** are also competitive to other systems. In asymptotic terms, our proof size is only larger than **libSNARK** [13] and **Bulletproofs** [16], and our verification is slower than **libSNARK** [13] and **libSTARK** [8]. Compared to **Hyrax** [48], which is also based on similar techniques with our work, **Libra** improves the performance in all aspects (yet **Hyrax** does not have any trusted setup). One can refer to Sect. 5 for a detailed description of our experimental setting as well as a more detailed comparison.

Finally, among all systems, **libSNARK** [13] requires a trusted setup for every statement, and **Libra** requires an one-time trusted setup that depends on the input size.

**Log-Space Uniform Circuits.** Though the prover time in **Libra** is optimal for all circuits, the verification time is succinct only when the circuit is structured (log-space uniform with logarithmic depth). This is the best that can be achieved for all ZKP protocols without per-circuit setup, as the verifier must read the entire circuit, which takes linear time in the worst case. We always refer to log-space uniform circuits when we say our scheme is succinct in this paper, to differentiate from schemes with linear verification time on all circuits (irrespective of whether the circuits are log-space uniform or not). Schemes such as **libSTARK** [8], **zkVSQL** [49] and **Hyrax** [48] also have such property.

In practice, with the help of auxiliary input and circuit squashing, most computations can be expressed as log-space uniform circuits with low depth, such as matrix multiplication, image scaling and Merkle hash tree in Sect. 5. Asymptotically, as shown in [8, 13, 51], all random memory access (RAM) programs can be

validated by circuits that are log-space uniform with log-depth in the running time of the programs (but linear in the size of the programs) by RAM-to-circuit reduction, which justifies the expressiveness of such circuits.

## 1.2 Our Techniques

Our main technical contributions are a GKR protocol with linear prover time and an efficient approach to turn the GKR protocol into zero-knowledge. We summarize the key ideas behind these two contributions. The detailed protocols are presented in Sects. 3 and 4 respectively.

**GKR with Linear Prover.** Goldwasser et al. [30] showed an approach to model the evaluation of a layered circuit as a sequence of summations on polynomials defined by values in consecutive layers of the circuit. Using the famous sumcheck protocol (see Sect. 2.3), they developed a protocol (the GKR protocol) allowing the verifier to validate the circuit evaluation in logarithmic time with a logarithmic size proof. However, the polynomials in the protocol are multivariate with  $2s$  variables, where  $S$  is the number of gates in one layer of the circuit and  $s = \log S$ . Naively running the sumcheck protocol on these polynomials incurs  $S^2$  prover time, as there are at least  $2^{2s} = S^2$  monomials in a  $2s$ -variate polynomial. Later, Cormode et al. [23] observed that these polynomials are sparse, containing only  $S$  nonzero monomials and improved the prover time to  $S \log S$ .

In our new approach, we divide the protocol into two separate sumchecks. In each sumcheck, the polynomial only contains  $s$  variables, and can be expressed as the product of two multilinear polynomials. Utilizing the sparsity of the circuit, we develop new algorithms to scan through each gate of the circuit and compute the closed-form of all these multilinear polynomials explicitly, which takes  $O(S)$  time. With this new way of representation, the prover can deploy a dynamic programming technique to generate the proofs in each sumcheck in  $O(S)$  time, resulting in a total prover time of  $O(S)$ .

**Efficient Zero-Knowledge GKR.** The original GKR protocol is not zero-knowledge, since the messages in the proof can be viewed as weighed sums of the values in the circuit and leak information. In [48, 49], the authors proposed to turn the GKR protocol into zero-knowledge by hiding the messages in homomorphic commitments, which incurs a big overhead in the verification time. In [22], Chiesa et al. proposed an alternative approach by masking the protocol with random polynomials. However, the masking polynomials are as big as the original ones and the prover time becomes exponential, making the approach mainly of theoretical interest.

In our scheme, we first show that in order to make the sumcheck protocol zero-knowledge, the prover can mask it with a “small” polynomial. In particular, the masking polynomial only contains logarithmically many random coefficients. The intuition is that though the original polynomial has  $O(2^\ell)$  or more terms ( $\ell$  is the number of variables in the polynomial), the prover only sends  $O(\ell)$  messages in the sumcheck protocol. Therefore, it suffices to mask the original polynomial with a random one with  $O(\ell)$  coefficients to achieve zero-knowledge.

In particular, we set the masking polynomial as the sum of  $\ell$  univariate random polynomials with the same variable-degree. In Sect. 4.1, we show that the entropy of this mask exactly counters the leakage of the sumcheck, proving that it is sufficient and optimal.

Besides the sumcheck, the GKR protocol additionally leaks two evaluations of the polynomial defined by values in each layer of the circuit. To make these evaluations zero-knowledge, we mask the polynomial by a special low-degree random polynomial. In particular, we show that after the mask, the verifier in total learns 4 messages related to the evaluations of the masking polynomial and we can prove zero-knowledge by making these messages linearly independent. Therefore, the masking polynomial is of constant size: it consists of 2 variables with variable degree 2.

### 1.3 Related Work

In recent years there has been significant progress in efficient ZKP protocols and systems. In this section, we discuss related work in this area, with the focus on those with sublinear proofs.

**QAP-Based.** Following earlier work of Ishai [34], Groth [33] and Lipmaa [38], Gennaro et al. [28] introduced quadratic arithmetic programs (QAPs), which forms the basis of most recent implementations [10, 14, 19, 24, 27, 42, 47] including libSNARK [13]. The proof size in these systems is constant, and the verification time depends only on the input size. Both these properties are particularly appealing and have led to real-world deployments, e.g., ZCash [9]. One of the main bottlenecks, however, of QAP-based systems is the high overhead in the prover running time and memory consumption, making it hard to scale to large statements. In addition, a separate trusted setup for every statement is required.

**IOPs.** Based on “(MPC)-in-the-head” introduced in [21, 29, 35], Ames et al. [6] proposed a ZKP scheme called Liger. It only uses symmetric key operations and the prover time is fast in practice. However, it generates proofs of size  $O(\sqrt{C})$ , which is several megabytes in practice for moderate-size circuits. In addition, the verification time is quasi-linear to the size of the circuit. It is categorized as interactive PCP, which is a special case of interactive oracle proofs (IOPs). IOP generalizes the probabilistically checkable proofs (PCPs) where earlier works of Kilian [36] and Micali [41] are built on. In the IOP model, Ben-Sasson et al. built libstark [8], a zero-knowledge transparent argument of knowledge (zkSTARK). libstark does not rely on trusted setup and executes in the RAM model of computation. Their verification time is only linear to the description of the RAM program, and succinct (logarithmic) in the time required for program execution. Recently, Ben-Sasson et al. [11] proposed Aurora, a new ZKP system in the IOP model with the proof size of  $O(\log^2 C)$ .

**Discrete Log.** Before Bulletproof [16], earlier discrete-log based ZKP schemes include the work of Groth [32], Bayer and Groth [7] and Bootle et al. [17].

**Hash-Based.** Bootle et al. [18] proposed a ZKP scheme with linear prover time and verification time. The verification only requires  $O(C)$  field additions. However, the proof size is  $O(\sqrt{C})$  and the constants are large.

**Interactive Proofs.** The line of work that relates to our paper the most is based on interactive proofs [31]. In the seminal work of [30], Goldwasser et al. proposed an efficient interactive proof for layered arithmetic circuits. Later, Cormode et al. [23] improved the prover complexity of the interactive proof in [30] to  $O(C \log C)$  using multilinear extensions instead of low degree extensions. Several follow-up works further reduce the prover time assuming special structures of the circuit. For regular circuits where the wiring pattern can be described in constant space and time, Thaler [43] introduced a protocol with  $O(C)$  prover time; for data parallel circuits with many copies of small circuits with size  $C'$ , a  $O(C \log C')$  protocol is presented in the same work, later improved to  $O(C + C' \log C)$  by Wahby et al. in [46]; for circuits with many non-connected but different copies, Zhang et al. showed a protocol with  $O(C \log C')$  prover time.

In [50], Zhang et al. extended the GKR protocol to an argument system using a protocol for verifiable polynomial delegation. Zhang et al. [51] and Wahby et al. [48] make the argument system zero-knowledge by putting all the messages in the proof into homomorphic commitments, as proposed by Cramer and Damgard in [25]. This approach introduces a high overhead on the verification time compared to the plain argument system without zero-knowledge, as each addition becomes a multiplication and each multiplication becomes an exponentiation in the homomorphic commitments. The multiplicative overhead is around two orders of magnitude in practice. Additionally, the scheme of [48], Hyrax, removes the trusted setup of the argument system by introducing a new polynomial delegation, increasing the proof size and verification time to  $O(\sqrt{n})$ .

## 2 Preliminaries

### 2.1 Notation

In this paper, we use  $\lambda$  to denote the security parameter, and  $\text{negl}(\lambda)$  to denote the negligible function in  $\lambda$ . “PPT” stands for probabilistic polynomial time. We use  $f(), h()$  for polynomials,  $x, y, z$  for vectors of variables and  $g, u, v$  for vectors of values.  $x_i$  denotes the  $i$ -th variable in  $x$ . We use bold letters such as  $\mathbf{A}$  to represent arrays. For a multivariate polynomial  $f$ , its “variable-degree” is the maximum degree of  $f$  in any of its variables.

**Assumptions.** Our scheme uses bilinear pairing and relies on the  $q$ -Strong Bilinear Diffie-Hellman ( $q$ -SBDH) assumption and an extended version of the Power Knowledge of Exponent (PKE) assumption [49, 50]. We present bilinear pairing and the assumptions formally in the full version of the paper.

### 2.2 Interactive Proofs and Zero-Knowledge Arguments

**Interactive Proofs.** An interactive proof allows a prover  $\mathcal{P}$  to convince a verifier  $\mathcal{V}$  the validity of some statement. The interactive proof runs in several rounds,

allowing  $\mathcal{V}$  to ask questions in each round based on  $\mathcal{P}$ 's answers of previous rounds. We phrase this in terms of  $\mathcal{P}$  trying to convince  $\mathcal{V}$  that  $f(x) = 1$ . The proof system is interesting only when the running time of  $\mathcal{V}$  is less than the time of directly computing the function  $f$ . We give the formal definition of interactive proofs in the full version.

**Zero-Knowledge Arguments.** An argument system for an NP relationship  $R$  is a protocol between a computationally-bounded prover  $\mathcal{P}$  and a verifier  $\mathcal{V}$ . At the end of the protocol,  $\mathcal{V}$  is convinced by  $\mathcal{P}$  that there exists a witness  $w$  such that  $(x; w) \in R$  for some input  $x$ . We focus on arguments of knowledge which have the stronger property that if the prover convinces the verifier of the statement validity, then the prover must know  $w$ . We use  $\mathcal{G}$  to represent the generation phase of the public key  $\text{pk}$  and the verification key  $\text{vk}$ . Formally, consider the definition below, where we assume  $R$  is known to  $\mathcal{P}$  and  $\mathcal{V}$ .

**Definition 1.** *Let  $R$  be an NP relation. A tuple of algorithm  $(\mathcal{G}, \mathcal{P}, \mathcal{V})$  is a zero-knowledge argument of knowledge for  $R$  if the following holds.*

- **Correctness.** *For every  $(\text{pk}, \text{vk})$  output by  $\mathcal{G}(1^\lambda)$  and  $(x, w) \in R$ ,*

$$\langle \mathcal{P}(\text{pk}, w), \mathcal{V}(\text{vk}) \rangle(x) = \text{accept}$$

- **Soundness.** *For any PPT prover  $\mathcal{P}$ , there exists a PPT extractor  $\varepsilon$  such that for every  $(\text{pk}, \text{vk})$  output by  $\mathcal{G}(1^\lambda)$  and any  $x$ , it holds that*

$$\Pr[\langle \mathcal{P}(\text{pk}), \mathcal{V}(\text{vk}) \rangle(x) = \text{accept} \wedge (x, w) \notin R | w \leftarrow \varepsilon(\text{pk}, x)] \leq \text{negl}(\lambda)$$

- **Zero knowledge.** *There exists a PPT simulator  $\mathcal{S}$  such that for any PPT adversary  $\mathcal{A}$ , auxiliary input  $z \in \{0, 1\}^{\text{poly}(\lambda)}$ ,  $(x; w) \in R$ , it holds that*  

$$\Pr[\langle \mathcal{P}(\text{pk}, w), \mathcal{A} \rangle = \text{accept} : (\text{pk}, \text{vk}) \leftarrow \mathcal{G}(1^\lambda); (x, w) \leftarrow \mathcal{A}(z, \text{pk}, \text{vk})] =$$
  

$$\Pr[\langle \mathcal{S}(\text{trap}, z, \text{pk}), \mathcal{A} \rangle = \text{accept} : (\text{pk}, \text{vk}, \text{trap}) \leftarrow \mathcal{S}(1^\lambda); (x, w) \leftarrow \mathcal{A}(z, \text{pk}, \text{vk})]$$

We say that  $(\mathcal{G}, \mathcal{P}, \mathcal{V})$  is a **succinct** argument system if the running time of  $\mathcal{V}$  and the total communication between  $\mathcal{P}$  and  $\mathcal{V}$  (proof size) are  $\text{poly}(\lambda, |x|, \log |w|)$ .

### 2.3 GKR Protocol

In [30], Goldwasser et al. proposed an efficient interactive proof protocol for layered arithmetic circuits, which we use as a building block for our new zero-knowledge argument and is referred as the *GKR* protocol. We present the detailed protocol here.

**Sumcheck Protocol.** The sumcheck problem is a fundamental problem that has various applications. The problem is to sum a polynomial  $f : \mathbb{F}^\ell \rightarrow \mathbb{F}$  on the binary hypercube  $\sum_{b_1, b_2, \dots, b_\ell \in \{0, 1\}} f(b_1, b_2, \dots, b_\ell)$ . Directly computing the sum requires exponential time in  $\ell$ , as there are  $2^\ell$  combinations of  $b_1, \dots, b_\ell$ . Lund

et al. [39] proposed a *sumcheck* protocol that allows a verifier  $\mathcal{V}$  to delegate the computation to a computationally unbounded prover  $\mathcal{P}$ , who can convince  $\mathcal{V}$  that  $H$  is the correct sum. We provide a description of the sumcheck protocol in Protocol 1. The proof size of the sumcheck protocol is  $O(d\ell)$ , where  $d$  is the variable-degree of  $f$ , as in each round,  $\mathcal{P}$  sends a univariate polynomial of one variable in  $f$ , which can be uniquely defined by  $d + 1$  points. The verifier time of the protocol is  $O(d\ell)$ . The prover time depends on the degree and the sparsity of  $f$ , and we will give the complexity later in our scheme. The sumcheck protocol is complete and sound with  $\epsilon = \frac{d\ell}{|\mathbb{F}|}$ .

**Protocol 1 (Sumcheck)** *The protocol proceeds in  $\ell$  rounds.*

– In the first round,  $\mathcal{P}$  sends a univariate polynomial

$$f_1(x_1) \stackrel{\text{def}}{=} \sum_{b_2, \dots, b_\ell \in \{0,1\}} f(x_1, b_2, \dots, b_\ell),$$

$\mathcal{V}$  checks  $H = f_1(0) + f_1(1)$ . Then  $\mathcal{V}$  sends a random challenge  $r_1 \in \mathbb{F}$  to  $\mathcal{P}$ .

– In the  $i$ -th round, where  $2 \leq i \leq \ell - 1$ ,  $\mathcal{P}$  sends a univariate polynomial

$$f_i(x_i) \stackrel{\text{def}}{=} \sum_{b_{i+1}, \dots, b_\ell \in \{0,1\}} f(r_1, \dots, r_{i-1}, x_i, b_{i+1}, \dots, b_\ell),$$

$\mathcal{V}$  checks  $f_{i-1}(r_{i-1}) = f_i(0) + f_i(1)$ , and sends a random challenge  $r_i \in \mathbb{F}$  to  $\mathcal{P}$ .

– In the  $\ell$ -th round,  $\mathcal{P}$  sends a univariate polynomial

$$f_\ell(x_\ell) \stackrel{\text{def}}{=} f(r_1, r_2, \dots, r_{\ell-1}, x_\ell),$$

$\mathcal{V}$  checks  $f_{\ell-1}(r_{\ell-1}) = f_\ell(0) + f_\ell(1)$ . The verifier generates a random challenge  $r_\ell \in \mathbb{F}$ . Given oracle access to an evaluation  $f(r_1, r_2, \dots, r_\ell)$  of  $f$ ,  $\mathcal{V}$  will accept if and only if  $f_\ell(r_\ell) = f(r_1, r_2, \dots, r_\ell)$ . The instantiation of the oracle access depends on the application of the sumcheck protocol.

**Definition 2 (Multi-linear Extension).** Let  $V : \{0, 1\}^\ell \rightarrow \mathbb{F}$  be a function. The multilinear extension of  $V$  is the unique polynomial  $\tilde{V} : \mathbb{F}^\ell \rightarrow \mathbb{F}$  such that  $\tilde{V}(x_1, x_2, \dots, x_\ell) = V(x_1, x_2, \dots, x_\ell)$  for all  $x_1, x_2, \dots, x_\ell \in \{0, 1\}^\ell$ .

$\tilde{V}$  can be expressed as:

$$\tilde{V}(x_1, x_2, \dots, x_\ell) = \sum_{b \in \{0,1\}^\ell} \prod_{i=1}^\ell [((1 - x_i)(1 - b_i) + x_i b_i) \cdot V(b)]$$

where  $b_i$  is  $i$ -th bit of  $b$ .

**Multilinear Extensions of Arrays.** Inspired by the close form equation of the multilinear extension given above, we can view an array  $\mathbf{A} = (a_0, a_1, \dots, a_{n-1})$  as a function  $A : \{0, 1\}^{\log n} \rightarrow \mathbb{F}$  such that  $\forall i \in [0, n - 1], A(i) = a_i$ . Therefore, in this paper, we abuse the use of multilinear extension on an array as the multilinear extension  $\tilde{A}$  of  $A$ .

**High Level Ideas of GKR.** Let  $C$  be a layered arithmetic circuit with depth  $d$  over a finite field  $\mathbb{F}$ . Each gate in the  $i$ -th layer takes inputs from two gates in



the  $(i + 1)$ -th layer; layer 0 is the output layer and layer  $d$  is the input layer. The protocol proceeds layer by layer. Upon receiving the claimed output from  $\mathcal{P}$ , in the first round,  $\mathcal{V}$  and  $\mathcal{P}$  run the sumcheck protocol to reduce the claim about the output to a claim about the values in the layer above. In the  $i$ -th round, both parties reduce a claim about layer  $i - 1$  to a claim about layer  $i$  through the sumcheck protocol. Finally, the protocol terminates with a claim about the input layer  $d$ , which can be checked directly by  $\mathcal{V}$ , or is given as an oracle access. If the check passes,  $\mathcal{V}$  accepts the claimed output.

**Notation.** Before describing the GKR protocol, we introduce some additional notations. We denote the number of gates in the  $i$ -th layer as  $S_i$  and let  $s_i = \lceil \log S_i \rceil$ . (For simplicity, we assume  $S_i$  is a power of 2, and we can pad the layer with dummy gates otherwise.) We then define a function  $V_i : \{0, 1\}^{s_i} \rightarrow \mathbb{F}$  that takes a binary string  $b \in \{0, 1\}^{s_i}$  and returns the output of gate  $b$  in layer  $i$ , where  $b$  is called the gate label. With this definition,  $V_0$  corresponds to the output of the circuit, and  $V_d$  corresponds to the input layer. Finally, we define two additional functions  $add_i, mult_i : \{0, 1\}^{s_{i-1}+2s_i} \rightarrow \{0, 1\}$ , referred as *wiring predicates* in the literature.  $add_i$  ( $mult_i$ ) takes one gate label  $z \in \{0, 1\}^{s_{i-1}}$  in layer  $i - 1$  and two gate labels  $x, y \in \{0, 1\}^{s_i}$  in layer  $i$ , and outputs 1 if and only if gate  $z$  is an addition (multiplication) gate that takes the output of gate  $x, y$  as input. With these definitions,  $V_i$  can be written as follows:

$$V_i(z) = \sum_{x,y \in \{0,1\}^{s_{i+1}}} (add_{i+1}(z, x, y)(V_{i+1}(x) + V_{i+1}(y)) + mult_{i+1}(z, x, y)(V_{i+1}(x)V_{i+1}(y))) \tag{1}$$

for any  $z \in \{0, 1\}^{s_i}$ .

In the equation above,  $V_i$  is expressed as a summation, so  $\mathcal{V}$  can use the sumcheck protocol to check that it is computed correctly. As the sumcheck protocol operates on polynomials defined on  $\mathbb{F}$ , we rewrite the equation with their multilinear extensions:

$$\begin{aligned} \tilde{V}_i(g) &= \sum_{x,y \in \{0,1\}^{s_{i+1}}} f_i(x, y) \\ &= \sum_{x,y \in \{0,1\}^{s_{i+1}}} (a\tilde{d}_{i+1}(g, x, y)(\tilde{V}_{i+1}(x) + \tilde{V}_{i+1}(y)) \\ &\quad + m\tilde{u}l_{i+1}(g, x, y)(\tilde{V}_{i+1}(x)\tilde{V}_{i+1}(y))), \end{aligned} \tag{2}$$

where  $g \in \mathbb{F}^{s_i}$  is a random vector.

**Protocol.** With Eq. 2, the GKR protocol proceeds as follows. The prover  $\mathcal{P}$  first sends the claimed output of the circuit to  $\mathcal{V}$ . From the claimed output,  $\mathcal{V}$  defines polynomial  $\tilde{V}_0$  and computes  $\tilde{V}_0(g)$  for a random  $g \in \mathbb{F}^{s_0}$ .  $\mathcal{V}$  and  $\mathcal{P}$  then invoke a sumcheck protocol on Eq. 2 with  $i = 0$ . As described in Sect. 2.3, at the end of the sumcheck,  $\mathcal{V}$  needs an oracle access to  $f_i(u, v)$ , where  $u, v$  are randomly selected in  $\mathbb{F}^{s_{i+1}}$ . To compute  $f_i(u, v)$ ,  $\mathcal{V}$  computes  $add_{i+1}(u, v)$  and  $mult_{i+1}(u, v)$  locally (they only depend on the wiring pattern of the circuit, but not on the values),

asks  $\mathcal{P}$  to send  $\tilde{V}_1(u)$  and  $\tilde{V}_1(v)$  and computes  $f_i(u, v)$  to complete the sumcheck protocol. In this way,  $\mathcal{V}$  and  $\mathcal{P}$  reduces a claim about the output to two claims about values in layer 1.

**Combining Two Claims: Condensing to One Claim.** In [30], Goldwasser et al. presented a protocol to reduce two claims  $\tilde{V}_i(u)$  and  $\tilde{V}_i(v)$  to one as following.  $\mathcal{V}$  defines a line  $\gamma : \mathbb{F} \rightarrow \mathbb{F}^{s_i}$  such that  $\gamma(0) = u, \gamma(1) = v$ .  $\mathcal{V}$  sends  $\gamma(x)$  to  $\mathcal{P}$ . Then  $\mathcal{P}$  sends  $\mathcal{V}$  a degree  $s_i$  univariate polynomial  $h(x) = \tilde{V}_i(\gamma(x))$ .  $\mathcal{V}$  checks that  $h(0) = \tilde{V}_i(u), h(1) = \tilde{V}_i(v)$ . Then  $\mathcal{V}$  randomly chooses  $r \in \mathbb{F}$  and computes a new claim  $h(r) = \tilde{V}_i(\gamma(r)) = \tilde{V}_i(w)$  on  $w = \gamma(r) \in \mathbb{F}^{s_i}$ .  $\mathcal{V}$  sends  $r, w$  to  $\mathcal{P}$ . In this way, the two claims are reduced to one claim  $\tilde{V}_i(w)$ . Combining this protocol with the sumcheck protocol on Eq. 2,  $\mathcal{V}$  and  $\mathcal{P}$  can reduce a claim on layer  $i$  to one claim on layer  $i + 1$ , and eventually to a claim on the input.

**Combining Two Claims: Random Linear Combination.** In [22], Chiesa et al. proposed an alternative approach using random linear combinations. Upon receiving the two claims  $\tilde{V}_i(u)$  and  $\tilde{V}_i(v)$ ,  $\mathcal{V}$  selects  $\alpha_i, \beta_i \in \mathbb{F}$  randomly and computes  $\alpha_i \tilde{V}_i(u) + \beta_i \tilde{V}_i(v)$ . Based on Eq. 2, this can be written as

$$\begin{aligned}
 & \alpha_i \tilde{V}_i(u) + \beta_i \tilde{V}_i(v) \\
 = & \alpha_i \sum_{x, y \in \{0, 1\}^{s_{i+1}}} (\tilde{add}_{i+1}(u, x, y)(\tilde{V}_{i+1}(x) + \tilde{V}_{i+1}(y)) + \tilde{mult}_{i+1}(u, x, y)(\tilde{V}_{i+1}(x)\tilde{V}_{i+1}(y))) \\
 & + \beta_i \sum_{x, y \in \{0, 1\}^{s_{i+1}}} (\tilde{add}_{i+1}(v, x, y)(\tilde{V}_{i+1}(x) + \tilde{V}_{i+1}(y)) + \tilde{mult}_{i+1}(v, x, y)(\tilde{V}_{i+1}(x)\tilde{V}_{i+1}(y))) \\
 = & \sum_{x, y \in \{0, 1\}^{s_{i+1}}} ((\alpha_i \tilde{add}_{i+1}(u, x, y) + \beta_i \tilde{add}_{i+1}(v, x, y))(\tilde{V}_{i+1}(x) + \tilde{V}_{i+1}(y))) \\
 & + (\alpha_i \tilde{mult}_{i+1}(u, x, y) + \beta_i \tilde{mult}_{i+1}(v, x, y))(\tilde{V}_{i+1}(x)\tilde{V}_{i+1}(y)) \tag{3}
 \end{aligned}$$

$\mathcal{V}$  and  $\mathcal{P}$  then execute the sumcheck protocol on Eq. 3 instead of Eq. 2. At the end of the sumcheck protocol,  $\mathcal{V}$  still receives two claims about  $\tilde{V}_{i+1}$ , computes their random linear combination and proceeds to an layer above recursively. In our new ZKP scheme, we will mainly use the second approach.

**Theorem 1.** [23, 30, 43, 45]. *Let  $C : \mathbb{F}^n \rightarrow \mathbb{F}^k$  be a depth- $d$  layered arithmetic circuit. The GKR protocol is an interactive proof for the function computed by  $C$  with soundness  $O(d \log |C|/|\mathbb{F}|)$ . It uses  $O(d \log |C|)$  rounds of interaction and running time of the prover  $\mathcal{P}$  is  $O(|C| \log |C|)$ . Let the optimal computation time for all  $\tilde{add}_i$  and  $\tilde{mult}_i$  be  $T$ , the running time of  $\mathcal{V}$  is  $O(n + k + d \log |C| + T)$ . For log-space uniform circuits it is  $T = \text{polylog } |C|$ .*

### 2.4 Zero-Knowledge Verifiable Polynomial Delegation Scheme

Let  $\mathbb{F}$  be a finite field,  $\mathcal{F}$  be a family of  $\ell$ -variate polynomial over  $\mathbb{F}$ , and  $d$  be a variable-degree parameter. A zero-knowledge verifiable polynomial delegation scheme (zkVPD) for  $f \in \mathcal{F}$  and  $t \in \mathbb{F}^\ell$  consists of the following algorithms:

- (pp, vp)  $\leftarrow$  KeyGen( $1^\lambda, \ell, d$ ),

---

**Algorithm 1.**  $\mathcal{F} \leftarrow \text{FunctionEvaluations}(f, \mathbf{A}, r_1, \dots, r_\ell)$

---

**Input:** Multilinear  $f$  on  $\ell$  variables, initial bookkeeping table  $\mathbf{A}$ , random  $r_1, \dots, r_\ell$ ;  
**Output:** All function evaluations  $f(r_1, \dots, r_{i-1}, t, b_{i+1}, \dots, b_\ell)$ ;

- 1: **for**  $i = 1, \dots, \ell$  **do**
- 2:     **for**  $b \in \{0, 1\}^{\ell-i}$  **do**     //  $b$  is both a number and its binary representation.
- 3:         **for**  $t = 0, 1, 2$  **do**
- 4:             Let  $f(r_1, \dots, r_{i-1}, t, b) = \mathbf{A}[b] \cdot (1 - t) + \mathbf{A}[b + 2^{\ell-i}] \cdot t$
- 5:              $\mathbf{A}[b] = \mathbf{A}[b] \cdot (1 - r_i) + \mathbf{A}[b + 2^{\ell-i}] \cdot r_i$
- 6: Let  $\mathcal{F}$  contain all function evaluations  $f(\cdot)$  computed at Step 4
- 7: **return**  $\mathcal{F}$

---

- $\text{com} \leftarrow \text{Commit}(f, r_f, \text{pp})$ ,
- $\{\text{accept}, \text{reject}\} \leftarrow \text{CheckComm}(\text{com}, \text{vp})$ ,
- $(y, \pi) \leftarrow \text{Open}(f, t, r_f, \text{pp})$ ,
- $\{\text{accept}, \text{reject}\} \leftarrow \text{Verify}(\text{com}, t, y, \pi, \text{vp})$ .

A zkVPD scheme satisfies correctness, soundness and zero knowledge. We give the formal definitions in the full version.

### 3 GKR Protocol with Linear Prover Time

In this section we present a new algorithm for the prover of the GKR protocol [30] that runs in linear time for *arbitrary layered circuits*. Before that, we present some necessary building blocks.

#### 3.1 Linear-Time Sumcheck for a Multilinear Function [43]

In [43], Thaler proposed a linear-time algorithm for the prover of the sumcheck protocol on a multilinear function  $f$  on  $\ell$  variables (the algorithm runs in  $O(2^\ell)$  time). We review this algorithm here. Recall that in the  $i$ -th round of the sumcheck protocol the prover sends the verifier the univariate polynomial on  $x_i$

$$\sum_{b_{i+1}, \dots, b_\ell \in \{0,1\}} f(r_1, \dots, r_{i-1}, x_i, b_{i+1}, \dots, b_\ell),$$

where  $r_1, \dots, r_{i-1}$  are random values chosen by the verifier in previous rounds. Since  $f$  is multilinear, it suffices for the prover to send two evaluations of the polynomial at points  $t = 0$  and  $t = 1$ , namely the evaluations

$$\sum_{b_{i+1}, \dots, b_\ell \in \{0,1\}} f(r_1, \dots, r_{i-1}, 0, b_{i+1}, \dots, b_\ell) \tag{4}$$

and

$$\sum_{b_{i+1}, \dots, b_\ell \in \{0,1\}} f(r_1, \dots, r_{i-1}, 1, b_{i+1}, \dots, b_\ell). \tag{5}$$

To compute the above sums the prover maintains a *bookkeeping table*  $\mathbf{A}$  for  $f$ . This table, at round  $i$ , has  $2^{\ell-i+1}$  entries storing the values

---

**Algorithm 2.**  $\{a_1, \dots, a_\ell\} \leftarrow \text{SumCheck}(f, \mathbf{A}, r_1, \dots, r_\ell)$

---

**Input:** Multilinear  $f$  on  $\ell$  variables, initial bookkeeping table  $\mathbf{A}$ , random  $r_1, \dots, r_\ell$ ;

**Output:**  $\ell$  sumcheck messages for  $\sum_{x \in \{0,1\}^\ell} f(x)$ . Each message  $a_i$  consists of 3 elements  $(a_{i0}, a_{i1}, a_{i2})$ ;

```

1:  $\mathcal{F} \leftarrow \text{FunctionEvaluations}(f, \mathbf{A}, r_1, \dots, r_\ell)$ 
2: for  $i = 1, \dots, \ell$  do
3:   for  $t \in \{0, 1, 2\}$  do
4:      $a_{it} = \sum_{b \in \{0,1\}^{\ell-i}} f(r_1, \dots, r_{i-1}, t, b)$  // All evaluations needed are in  $\mathcal{F}$ .
5: return  $\{a_1, \dots, a_\ell\}$ ;

```

---

$f(r_1, \dots, r_{i-1}, b_i, b_{i+1}, \dots, b_\ell)$  for all  $b_i, \dots, b_\ell \in \{0, 1\}$  and is initialized with evaluations of  $f$  on the hypercube. For every entry of  $\mathbf{A}$ , the prover subsequently computes, as in Step 4 of Algorithm 1 `FunctionEvaluations`<sup>2</sup> two values  $f(r_1, \dots, r_{i-1}, 0, b_{i+1}, \dots, b_\ell)$  and  $f(r_1, \dots, r_{i-1}, 1, b_{i+1}, \dots, b_\ell)$ . Once these function evaluations are in place, the prover can easily sum over them and compute the required sumcheck messages as required by Relations 4 and 5. This is done in Algorithm 2.

**Complexity Analysis.** Both Algorithms 1 and 2 run in  $O(2^\ell)$  time: The first iteration takes  $O(2^\ell)$ , the second  $O(2^{\ell-1})$  and so on. Therefore the bound holds.

### 3.2 Linear-Time Sumcheck for Products of Multilinear Functions [43]

The linear-time sumcheck in the previous section can be generalized to a product of two multilinear functions. Let now  $f$  and  $g$  be two multilinear functions on  $\ell$  variables each, we describe a linear-time algorithm to compute the messages of the prover for the sumcheck on the product  $f \cdot g$ , as proposed in [43]. Note that we cannot use Algorithm 2 here since  $f \cdot g$  is not multilinear. However, similarly with the single-function case, the prover must now send, at round  $i$ , the following evaluations at points  $t = 0$ ,  $t = 1$  and  $t = 2$

$$\sum_{b_{i+1}, \dots, b_\ell \in \{0,1\}} f(r_1, \dots, r_{i-1}, t, b_{i+1}, \dots, b_\ell) \cdot g(r_1, \dots, r_{i-1}, t, b_{i+1}, \dots, b_\ell)$$

The above can be easily computed by computing evaluations for functions  $f$  and  $g$  *separately* using Algorithm 1 and the combining the results using our new Algorithm 3 `SumCheckProduct`. We now have the following lemma:

**Lemma 1.** *Algorithm SumCheckProduct runs in time  $O(2^\ell)$*

---

<sup>2</sup> To be compatible with other protocols later, we use three values  $t = 0, 1, 2$  in our evaluations instead of just two.

### 3.3 Linear-Time Sumcheck for GKR Functions

Let us now consider the sumcheck problem on a particular class of functions that are relevant for the GKR protocol (that is why we call them GKR functions). In particular we want to compute the sumcheck

$$\sum_{x,y \in \{0,1\}^\ell} f_1(g, x, y) f_2(x) f_3(y), \tag{6}$$

for a fixed point  $g \in \mathbb{F}^\ell$ , where  $f_2(x), f_3(x) : \mathbb{F}^\ell \rightarrow \mathbb{F}$  are multilinear extensions of arrays  $\mathbf{A}_{f_2}, \mathbf{A}_{f_3}$  of size  $2^\ell$ , and function  $f_1 : \mathbb{F}^{3\ell} \rightarrow \mathbb{F}$  is the multilinear extension of a sparse array with  $O(2^\ell)$  (out of  $2^{3\ell}$  possible) nonzero elements. It is not hard to see that the sumcheck polynomials in GKR given by Eqs. 2 and 3 satisfy these properties.

---

**Algorithm 3.**  $\{a_1, \dots, a_\ell\} \leftarrow \text{SumCheckProduct}(f, \mathbf{A}_f, g, \mathbf{A}_g, r_1, \dots, r_\ell)$

---

**Input:** Multilinear  $f$  and  $g$ , initial bookkeeping tables  $\mathbf{A}_f$  and  $\mathbf{A}_g$ , random  $r_1, \dots, r_\ell$ ;  
**Output:**  $\ell$  sumcheck messages for  $\sum_{x \in \{0,1\}^\ell} f(x)g(x)$ . Each message  $a_i$  consists of 3 elements  $(a_{i0}, a_{i1}, a_{i2})$ ;

- 1:  $\mathcal{F} \leftarrow \text{FunctionEvaluations}(f, \mathbf{A}_f, r_1, \dots, r_\ell)$
- 2:  $\mathcal{G} \leftarrow \text{FunctionEvaluations}(g, \mathbf{A}_g, r_1, \dots, r_\ell)$
- 3: **for**  $i = 1, \dots, \ell$  **do**
- 4:     **for**  $t \in \{0, 1, 2\}$  **do**
- 5:          $a_{it} = \sum_{b \in \{0,1\}^{\ell-i}} f(r_1, \dots, r_{i-1}, t, b) \cdot g(r_1, \dots, r_{i-1}, t, b)$      // All evaluations needed are in  $\mathcal{F}$  and  $\mathcal{G}$ .
- 6: **return**  $\{a_1, \dots, a_\ell\}$ ;

---

We note here that applying Algorithm 1 FunctionEvaluations for this particular class of polynomials would lead to quadratic prover time. This is because  $f_1$  has  $2^{2\ell}$  variables to sum on yielding  $O(2^{2\ell})$  complexity. However, one could take advantage of the sparsity of  $f_1$ : the prover can store only the  $O(2^\ell)$  non-zero values of the bookkeeping table  $\mathbf{A}$ . This is exactly the approach used in many prior work [23, 46, 51]. However, with this approach, the number of nonzero values that must be considered in Step 2 is always at most  $2^\ell$ , since it is not guaranteed that this number will reduce to half (i.e., to  $2^{\ell-i}$ ) after every update in Step 5 because it is sparse. Therefore, the overall complexity becomes  $O(\ell \cdot 2^\ell)$ .

In this section we effectively reduce this bound to  $O(2^\ell)$ . Our protocol divides the sumcheck into two phases: the first  $\ell$  rounds bounding the variables of  $x$  to a random point  $u$ , and the last  $\ell$  rounds bounding the variables of  $y$  to a random point  $v$ . The central idea lies in rewriting Eq. 6 as follows

$$\begin{aligned} \sum_{x,y \in \{0,1\}^\ell} f_1(g, x, y) f_2(x) f_3(y) &= \sum_{x \in \{0,1\}^\ell} f_2(x) \sum_{y \in \{0,1\}^\ell} f_1(g, x, y) f_3(y) \\ &= \sum_{x \in \{0,1\}^\ell} f_2(x) h_g(x), \end{aligned}$$

where  $h_g(x) = \sum_{y \in \{0,1\}^\ell} f_1(g, x, y) f_3(y)$ .

---

**Algorithm 4.**  $\mathbf{A}_{h_g} \leftarrow \text{Initialize\_PhaseOne}(f_1, f_3, \mathbf{A}_{f_3}, g)$

---

**Input:** Multilinear  $f_1$  and  $f_3$ , initial bookkeeping tables  $\mathbf{A}_{f_3}$ , random  $g = g_1, \dots, g_\ell$ ;

**Output:** Bookkeeping table  $\mathbf{A}_{h_g}$ ;

```

1: procedure  $\mathbf{G} \leftarrow \text{Precompute}(g)$  //  $\mathbf{G}$  is an array of size  $2^\ell$ .
2:   Set  $\mathbf{G}[0] = 1$ 
3:   for  $i = 0, \dots, \ell - 1$  do
4:     for  $b \in \{0, 1\}^i$  do
5:        $\mathbf{G}[b, 0] = \mathbf{G}[b] \cdot (1 - g_{i+1})$ 
6:        $\mathbf{G}[b, 1] = \mathbf{G}[b] \cdot g_{i+1}$ 
7:  $\forall x \in \{0, 1\}^\ell$ , set  $\mathbf{A}_{h_g}[x] = 0$ 
8: for every  $(z, x, y)$  such that  $f_1(z, x, y)$  is non-zero do
9:    $\mathbf{A}_{h_g}[x] = \mathbf{A}_{h_g}[x] + \mathbf{G}[z] \cdot f_1(z, x, y) \cdot \mathbf{A}_{f_3}[y]$ 
10: return  $\mathbf{A}_{h_g}$ ;

```

---

**Phase One.** With the formula above, in the first  $\ell$  rounds, the prover and the verifier are running exactly a sumcheck on a product of two multilinear functions  $f_2 \cdot h_g$ , since functions  $f_2$  and  $h_g$  can be viewed as functions only in  $x$ — $y$  can be considered constant (it is always summed on the hypercube). To compute the sumcheck messages for the first  $\ell$  rounds, given their bookkeeping tables, we can call

$$\text{SumCheckProduct}(h_g(x), \mathbf{A}_{h_g}, f_2(x), \mathbf{A}_{f_2}, u_1, \dots, u_\ell)$$

in Algorithm 3. By Lemma 1 this will take  $O(2^\ell)$  time. We now show how to initialize the bookkeeping tables in linear time.

Initializing the Bookkeeping Tables:

Initializing the bookkeeping table for  $f_2$  in  $O(2^\ell)$  time is trivial, since  $f_2$  is a multilinear extension of an array and therefore the evaluations on the hypercube are known. Initializing the bookkeeping table for  $h_g$  in  $O(2^\ell)$  time is more challenging but we can leverage the sparsity of  $f_1$ . Consider the following lemma.

**Lemma 2.** *Let  $\mathcal{N}_x$  be the set of  $(z, y) \in \{0, 1\}^{2\ell}$  such that  $f_1(z, x, y)$  is non-zero. Then for all  $x \in \{0, 1\}^\ell$ , it is  $h_g(x) = \sum_{(z,y) \in \mathcal{N}_x} I(g, z) \cdot f_1(z, x, y) \cdot f_3(y)$ , where  $I(g, z) = \prod_{i=1}^\ell ((1 - g_i)(1 - z_i) + g_i z_i)$ .*

*Proof.* As  $f_1$  is a multilinear extension, as shown in [43], we have  $f_1(g, x, y) = \sum_{z \in \{0,1\}^\ell} I(g, z) f_1(z, x, y)$ , where  $I$  is the multilinear extension of the identity polynomial, i.e.,  $I(w, z) = 1$  iff  $w = z$  for all  $w, z \in \{0, 1\}^\ell$ . Therefore, we have

$$\begin{aligned} h_g(x) &= \sum_{y \in \{0,1\}^\ell} f_1(g, x, y) f_3(y) = \sum_{z,y \in \{0,1\}^\ell} I(g, z) f_1(z, x, y) f_3(y) \\ &= \sum_{(z,y) \in \mathcal{N}_x} I(g, z) \cdot f_1(z, x, y) \cdot f_3(y) \end{aligned}$$

Moreover,  $I(w, z) = \prod_{i=1}^\ell ((1 - w_i)(1 - z_i) + w_i z_i)$  is the unique polynomial that evaluates to 1 iff  $w = z$  for all  $w, z \in \{0, 1\}^\ell$ . As the multilinear extension is unique, we have  $I(g, z) = \prod_{i=1}^\ell ((1 - g_i)(1 - z_i) + g_i z_i)$ . □

---

**Algorithm 5.**  $\mathbf{A}_{f_1} \leftarrow \text{Initialize\_PhaseTwo}(f_1, g, u)$ 


---

**Input:** Multilinear  $f_1$ , random  $g = g_1, \dots, g_\ell$  and  $u = u_1, \dots, u_\ell$ ;

**Output:** Bookkeeping table  $\mathbf{A}_{f_1}$ ;

- 1:  $\mathbf{G} \leftarrow \text{Precompute}(g)$
  - 2:  $\mathbf{U} \leftarrow \text{Precompute}(u)$
  - 3:  $\forall y \in \{0, 1\}^\ell$ , set  $\mathbf{A}_{f_1}[y] = 0$
  - 4: **for** every  $(z, x, y)$  such that  $f_1(z, x, y)$  is non-zero **do**
  - 5:      $\mathbf{A}_{f_1}[y] = \mathbf{A}_{f_1}[y] + \mathbf{G}[z] \cdot \mathbf{U}[x] \cdot f_1(z, x, y)$
  - 6: **return**  $\mathbf{A}_{f_1}$ ;
- 

**Lemma 3.** *The bookkeeping table  $\mathbf{A}_{h_g}$  can be initialized in time  $O(2^\ell)$ .*

*Proof.* As  $f_1$  is sparse,  $\sum_{x \in \{0, 1\}^\ell} |\mathcal{N}_x| = O(2^\ell)$ . From Lemma 2, given the evaluations of  $I(g, z)$  for all  $z \in \{0, 1\}^\ell$ , the prover can iterate all  $(z, y) \in \mathcal{N}_x$  for all  $x$  to compute  $\mathbf{A}_{h_g}$ . The full algorithm is presented in Algorithm 4.

Procedure  $\text{Precompute}(g)$  is to evaluate  $\mathbf{G}[z] = I(g, z) = \prod_{i=1}^\ell ((1 - g_i)(1 - z_i) + g_i z_i)$  for  $z \in \{0, 1\}^\ell$ . By the closed-form of  $I(g, z)$ , the procedure iterates each bit of  $z$ , and multiplies  $1 - g_i$  for  $z_i = 0$  and multiplies  $g_i$  for  $z_i = 1$ . In this way, the size of  $\mathbf{G}$  doubles in each iteration, and the total complexity is  $O(2^\ell)$ .

Step 8–9 computes  $h_g(x)$  using Lemma 2. When  $f_1$  is represented as a map of  $(z, x, y)$ ,  $f_1(z, x, y)$  for non-zero values, the complexity of these steps is  $O(2^\ell)$ . In the GKR protocol, this is exactly the representation of a gate, where  $z, x, y$  are labels of the gate, its left input and its right input, and  $f_1(z, x, y) = 1$ .  $\square$

With the bookkeeping tables, the prover runs  $\text{SumCheckProduct}(h_g(x), \mathbf{A}_{h_g}, f_2(x), \mathbf{A}_{f_2}, u_1, \dots, u_\ell)$  in Algorithm 3 and the complexity for phase one is  $O(2^\ell)$ .

**Phase Two.** At this point, all variables in  $x$  have been bounded to random numbers  $u$ . In the second phase, the equation to sum on becomes

$$\sum_{y \in \{0, 1\}^\ell} f_1(g, u, y) f_2(u) f_3(y)$$

Note here that  $f_2(u)$  is merely a single value which we already computed in phase one. Both  $f_1(g, u, y)$  and  $f_3(y)$  are polynomials on  $y$  with  $\ell$  variables. Similar to phase one, to compute the messages for the last  $\ell$  rounds we can call

$$\text{SumCheckProduct}(f_1(g, u, y), \mathbf{A}_{f_1}, f_3(y) \cdot f_2(u), \mathbf{A}_{f_3} \cdot f_2(u), v_1, \dots, v_\ell).$$

Note here that  $\mathbf{A}_{f_1}$  is the bookkeeping table for  $f_1(g, u, y)$ , not the original sparse function  $f_1(g, x, y)$ .

*Initializing the Bookkeeping Table for  $f_1$ :*

It now remains to initialize the bookkeeping table for  $f_1(g, u, y)$  efficiently. Similar to phase one, we have the following lemma:

**Lemma 4.** *Let  $\mathcal{N}_y$  be the set of  $(z, x) \in \{0, 1\}^{2\ell}$  such that  $f_1(z, x, y)$  is non-zero. Then for all  $y \in \{0, 1\}^\ell$ , it is  $f_1(g, u, y) = \sum_{(z, x) \in \mathcal{N}_y} I(g, z) \cdot I(u, x) \cdot f_1(z, x, y)$ .*

*Proof.* This immediately follows from the fact that  $f_1$  is a multilinear extension. We have  $f_1(g, u, y) = \sum_{z,y \in \{0,1\}^\ell} I(g, z) \cdot I(u, x) \cdot f_1(z, x, y)$ , where the closed form of  $I$  is given in Lemma 2.  $\square$

**Lemma 5.** *The bookkeeping table  $A_{f_1}$  can be initialized in time  $O(2^\ell)$ .*

*Proof.* Similar to Algorithm 4, the prover again iterates all non-zero indices of  $f_1$  to compute it using Lemma 4. The full algorithm is presented in Algorithm 5.  $\square$

### 3.4 Putting Everything Together

The sumcheck protocol in GKR given by Eq. 2 can be decomposed into several instances that have the form of Eq. 6 presented in the previous section. The term

$$\sum_{x,y \in \{0,1\}^{s_{i+1}}} \tilde{mult}_{i+1}(g, x, y)(\tilde{V}_{i+1}(x)\tilde{V}_{i+1}(y))$$

is exactly the same as Eq. 6. The term  $\sum_{x,y \in \{0,1\}^{s_{i+1}}} \tilde{add}_{i+1}(g, x, y)(\tilde{V}_{i+1}(x) + \tilde{V}_{i+1}(y))$  can be viewed as:

$$\sum_{x,y \in \{0,1\}^{s_{i+1}}} \tilde{add}_{i+1}(g, x, y)\tilde{V}_{i+1}(x) + \sum_{x,y \in \{0,1\}^{s_{i+1}}} \tilde{add}_{i+1}(g, x, y)\tilde{V}_{i+1}(y)$$

The first sum can be computed using the same protocol in Sect. 3.3 without  $f_3(y)$ , and the second sum can be computed without  $f_2(x)$ . The complexity for both cases remains linear. Due to linearity of the sumcheck protocol, the prover can execute these 3 instances simultaneously in every round, and sum up the individual messages and send them to the verifier.

**Combining Two Claims.** After the sumcheck in the GKR protocol is completed, as described in Sect. 2.3, the prover and the verifier need to combine the two claims about  $\tilde{V}_{i+1}$  received at the end of the sumcheck protocol to one to avoid the exponential blow-up. There are two ways to combine the two claims and we show how to do each of them in linear time.

The second approach using random linear combinations is rather straight forward. After the output layers,  $\mathcal{P}$  and  $\mathcal{V}$  execute sumcheck protocol on Eq. 3 instead of Eq. 2, which still satisfies the properties of Eq. 6. One could view it as 6 instances of Eq. 6 and the prover time is still linear. Moreover, there is a better way to further improve the efficiency. Taking  $\sum_{x,y \in \{0,1\}^{s_{i+1}}} (\alpha_i \tilde{mult}_{i+1}(u, x, y) + \beta_i \tilde{mult}_{i+1}(v, x, y))\tilde{V}_{i+1}(x)\tilde{V}_{i+1}(y)$  as an example, in Algorithm 4, the prover runs Precompute twice on  $u$  and  $v$  to generate two arrays ( $\mathbf{G}_1$  and  $\mathbf{G}_2$ ), and sets  $\mathbf{G}[b] = \alpha_i \mathbf{G}_1[b] + \beta_i \mathbf{G}_2[b]$  for all  $b$ . The rest of the algorithms remains the same. This only incurs a small overhead in practice in our implementation, compared to the original algorithm on Eq. 6.

Though with the approach above we already have a linear prover GKR protocol, the technique to condense two points to one proposed in the original GKR protocol [30] may still be interesting in some scenarios (e.g., in our implementation, we use this approach in the last layer and only make one query to the multi-linear extension of the input, which is more efficient practice). We present an algorithm to reduce the prover time of this approach to linear in the full version of the paper.



## 4 Zero Knowledge Argument Protocols

In this section, we present the construction of our new zero-knowledge argument system. In [50], Zhang et al. proposed to combine the GKR protocol with a verifiable polynomial delegation protocol, resulting in an argument system. Later, in [48, 49], the construction was extended to zero-knowledge, by sending all the messages in the GKR protocol in homomorphic commitments and performing all the checks by zero-knowledge equality and product testing. This incurs a high overhead for the verifier compared to the plain version without zero-knowledge, as each multiplication becomes an exponentiation and each equality check becomes a  $\Sigma$ -protocol, which is around  $100\times$  slower in practice.

In this paper, we follow the same blueprint of combining GKR and VPD to obtain an argument system, but instead show how to extend it to be zero-knowledge efficiently. In particular, the prover masks the GKR protocol with special random polynomials so that the verifier runs a “randomized” GKR that leaks no extra information and her overhead is small. A similar approach was used by Chiesa et al. in [22]. In the following, we present the zero-knowledge version of each building block, followed by the whole zero-knowledge argument.

### 4.1 Zero Knowledge Sumcheck

As a core step of the GKR protocol,  $\mathcal{P}$  and  $\mathcal{V}$  execute a sumcheck protocol on Eq. 2, during which  $\mathcal{P}$  sends  $\mathcal{V}$  evaluations of the polynomial at several random points chosen by  $\mathcal{V}$ . These evaluations leak information about the values in the circuit, as they can be viewed as weighted sums of these values.

To make the sumcheck protocol zero-knowledge, we take the approach proposed by Chiesa et al. in [22], which is masking the polynomial in the sumcheck protocol by a random polynomial. In this approach, to prove

$$H = \sum_{x_1, x_2, \dots, x_\ell \in \{0,1\}} f(x_1, x_2, \dots, x_\ell),$$

the prover generates a random polynomial  $g$  with the same variables and individual degrees of  $f$ . She commits to the polynomial  $g$ , and sends the verifier a claim  $G = \sum_{x_1, x_2, \dots, x_\ell \in \{0,1\}} g(x_1, x_2, \dots, x_\ell)$ . The verifier picks a random number  $\rho$ , and execute a sumcheck protocol with the prover on

$$H + \rho G = \sum_{x_1, x_2, \dots, x_\ell \in \{0,1\}} (f(x_1, x_2, \dots, x_\ell) + \rho g(x_1, x_2, \dots, x_\ell)).$$

At the last round of this sumcheck, the prover opens the commitment of  $g$  at  $g(r_1, \dots, r_\ell)$ , and the verifier computes  $f(r_1, \dots, r_\ell)$  by subtracting  $\rho g(r_1, \dots, r_\ell)$  from the last message, and compares it with the oracle access of  $f$ . It is shown that as long as the commitment and opening of  $g$  are zero-knowledge, the protocol is zero-knowledge. Intuitively, this is because all the coefficients of  $f$  are masked by those of  $g$ . The soundness still holds because of the random linear combination of  $f$  and  $g$ .

Unfortunately, the masking polynomial  $g$  is as big as  $f$ , and opening it to a random point later is expensive. In [22], the prover sends a PCP oracle of  $g$ , and executes a zero-knowledge sumcheck to open it to a random point, which incurs an exponential complexity for the prover. Even replacing it with the zkVPD protocol in [49], the prover time is slow in practice.

In this paper, we show that it suffices to mask  $f$  with a small polynomial to achieve zero-knowledge. In particular, we set  $g(x_1, \dots, x_\ell) = a_0 + g_1(x_1) + g_2(x_2) + \dots + g_\ell(x_\ell)$ , where  $g_i(x_i) = a_{i,1}x_i + a_{i,2}x_i^2 + \dots + a_{i,d}x_i^d$  is a random univariate polynomial of degree  $d$  ( $d$  is the variable degree of  $f$ ). Note here that the size of  $g$  is only  $O(d\ell)$ , while the size of  $f$  is exponential in  $\ell$ .

The intuition of our improvement is that the prover sends  $O(d\ell)$  messages in total to the verifier during the sumcheck protocol, thus a polynomial  $g$  with  $O(d\ell)$  random coefficients is sufficient to mask all the messages and achieve zero-knowledge. We present the full protocol in Construction 1.

The completeness of the protocol holds obviously. The soundness follows the soundness of the sumcheck protocol and the random linear combination in step 2 and 3, as proven in [22]. We give a proof of zero knowledge in the full version.

**Theorem 2 (Zero knowledge).** *For every verifier  $\mathcal{V}^*$  and every  $\ell$ -variate polynomial  $f : \mathbb{F}^\ell \rightarrow \mathbb{F}$  with variable degree  $d$ , there exists a simulator  $\mathcal{S}$  such that given access to  $H = \sum_{x_1, x_2, \dots, x_\ell \in \{0,1\}} f(x_1, x_2, \dots, x_\ell)$ ,  $\mathcal{S}$  is able to simulate the partial view of  $\mathcal{V}^*$  in step 1–4 of Construction 1.*

**Construction 1.** *We assume the existence of a zkVPD protocol defined in Sect. 2.4. For simplicity, we omit the randomness  $r_f$  and public parameters  $\text{pp}, \text{vp}$  without any ambiguity. To prove the claim  $H = \sum_{x_1, x_2, \dots, x_\ell \in \{0,1\}} f(x_1, x_2, \dots, x_\ell)$ :*

1.  $\mathcal{P}$  selects a polynomial  $g(x_1, \dots, x_\ell) = a_0 + g_1(x_1) + g_2(x_2) + \dots + g_\ell(x_\ell)$ , where  $g_i(x_i) = a_{i,1}x_i + a_{i,2}x_i^2 + \dots + a_{i,d}x_i^d$  and all  $a_{i,j}$ s are uniformly random.  $\mathcal{P}$  sends  $H = \sum_{x_1, x_2, \dots, x_\ell \in \{0,1\}} f(x_1, x_2, \dots, x_\ell)$ ,  $G = \sum_{x_1, x_2, \dots, x_\ell \in \{0,1\}} g(x_1, x_2, \dots, x_\ell)$  and  $\text{com}_g = \text{Commit}(g)$  to  $\mathcal{V}$ .
2.  $\mathcal{V}$  uniformly selects  $\rho \in \mathbb{F}^*$ , computes  $H + \rho G$  and sends  $\rho$  to  $\mathcal{P}$ .
3.  $\mathcal{P}$  and  $\mathcal{V}$  run the sumcheck protocol on

$$H + \rho G = \sum_{x_1, x_2, \dots, x_\ell \in \{0,1\}} (f(x_1, x_2, \dots, x_\ell) + \rho g(x_1, x_2, \dots, x_\ell))$$

4. At the last round of the sumcheck protocol,  $\mathcal{V}$  obtains a claim  $h_\ell(r_\ell) = f(r_1, r_2, \dots, r_\ell) + \rho g(r_1, r_2, \dots, r_\ell)$ .  $\mathcal{P}$  and  $\mathcal{V}$  opens the commitment of  $g$  at  $r = (r_1, \dots, r_\ell)$  by  $(g(r), \pi) \leftarrow \text{Open}(g, r), \text{Verify}(\text{com}_g, g(r), r, \pi)$ . If  $\text{Verify}$  outputs reject,  $\mathcal{V}$  aborts.
5.  $\mathcal{V}$  computes  $h_\ell(r_\ell) - \rho g(r_1, \dots, r_\ell)$  and compares it with the oracle access of  $f(r_1, \dots, r_\ell)$ .

## 4.2 Zero Knowledge GKR

To achieve zero-knowledge, we replace the sumcheck protocol in GKR with the zero-knowledge version described in the previous section. However, the protocol still leaks additional information. In particular, at the end of the zero-knowledge sumcheck,  $\mathcal{V}$  queries the oracle to evaluate the polynomial on a random point. When executed on Eq. 2, this reveals two evaluations of the polynomial  $\check{V}_i$  defined by the values in the  $i$ -th layer of the circuit:  $\check{V}_i(u)$  and  $\check{V}_i(v)$ .

To prevent this leakage, Chiesa et al. [22] proposed to replace the multi-linear extension  $\check{V}_i$  with a low degree extension, such that learning  $\check{V}_i(u)$  and  $\check{V}_i(v)$  does not leak any information about  $V_i$ . Define a low degree extension of  $V_i$  as

$$\dot{V}_i(z) \stackrel{\text{def}}{=} \check{V}_i(z) + Z_i(z) \sum_{w \in \{0,1\}^\lambda} R_i(z, w), \quad (7)$$

where  $Z(z) = \prod_{i=1}^{s_i} z_i(1 - z_i)$ , i.e.,  $Z(z) = 0$  for all  $z \in \{0,1\}^{s_i}$ .  $R_i(z, w)$  is a random low-degree polynomial and  $\lambda$  is the security parameter. With this low degree extension, Eq. 2 becomes

$$\begin{aligned} \dot{V}_i(g) &= \sum_{x,y \in \{0,1\}^{s_{i+1}}} \tilde{mult}_{i+1}(g, x, y) (\dot{V}_{i+1}(x) \dot{V}_{i+1}(y)) \\ &\quad + \tilde{add}_{i+1}(g, x, y) (\dot{V}_{i+1}(x) + \dot{V}_{i+1}(y)) + Z_i(g) \sum_{w \in \{0,1\}^\lambda} R_i(g, w) \end{aligned} \quad (8)$$

$$\begin{aligned} &= \sum_{x,y \in \{0,1\}^{s_{i+1}}, w \in \{0,1\}^\lambda} (I(\mathbf{0}, w) \cdot \tilde{mult}_{i+1}(g, x, y) (\dot{V}_{i+1}(x) \dot{V}_{i+1}(y)) \\ &\quad + \tilde{add}_{i+1}(g, x, y) (\dot{V}_{i+1}(x) + \dot{V}_{i+1}(y)) + I((x, y), \mathbf{0}) Z_i(g) R_i(g, w)) \end{aligned} \quad (9)$$

where  $I(\mathbf{a}, \mathbf{b})$  is an identity polynomial  $I(\mathbf{a}, \mathbf{b}) = 0$  iff  $\mathbf{a} = \mathbf{b}$ . The first equation holds because  $\check{V}_i$  agrees with  $\dot{V}_i$  on the Boolean hyper-cube  $\{0,1\}^{s_i}$ , as  $Z_i(z) = 0$  for binary inputs. The second equation holds because the mask in  $\dot{V}_i$  is in the form of a “sum” and can be moved into the sumcheck equation.

When executing the zero-knowledge sumcheck protocol on Eq. 8, at the end of the protocol,  $\mathcal{V}$  receives  $\dot{V}_{i+1}(u)$  and  $\dot{V}_{i+1}(v)$  for random points  $u, v \in \mathbb{F}^{s_{i+1}}$  chosen by  $\mathcal{V}$ . They no longer leak information about  $V_{i+1}$ , as they are masked by  $Z_{i+1}(z) \sum_{w \in \{0,1\}^\lambda} R_{i+1}(z, w)$  for  $z = u$  and  $z = v$ .  $\mathcal{V}$  computes  $\tilde{mult}_{i+1}(g, u, v)$  and  $\tilde{add}_{i+1}(g, u, v)$  as before, computes  $Z_i(g)$ ,  $I(\mathbf{0}, c)$ ,  $I((u, v), \mathbf{0})$  where  $c \in \mathbb{F}^\lambda$  is a random point chosen by  $\mathcal{V}$  for variable  $w$ , opens  $R_i(g, w)$  at  $c$  with  $\mathcal{P}$  through a polynomial commitment, and checks that together with  $\dot{V}_{i+1}(u)$ ,  $\dot{V}_{i+1}(v)$  received from  $\mathcal{P}$  they are consistent with the last message of the sumcheck.  $\mathcal{V}$  then uses  $\dot{V}_{i+1}(u)$ ,  $\dot{V}_{i+1}(v)$  to proceed to the next round.

Unfortunately, similar to the zero-knowledge sumcheck, the masking polynomial  $R_i$  is very large in [22]. Opening  $R_i$  at a random point takes exponential time for  $\mathcal{P}$  either using a PCP oracle as in [22] or potentially using a zkVPD, as  $R$  has  $s_i + 2s_{i+1} + \lambda$  variables.

In this section, we show that we can set  $R_i$  to be a small polynomial to achieve zero-knowledge. In particular,  $R_i$  has only two variables with variable degree 2. This is because in the  $(i - 1)$ -th round,  $\mathcal{V}$  receives two evaluations of  $V_i$ ,  $\check{V}_i(u)$  and  $\check{V}_i(v)$ , which are masked by  $\sum_w R_i(u, w)$  and  $\sum_w R_i(v, w)$ ; in the

$i$ -th sumcheck,  $\mathcal{V}$  opens  $R_i$  at  $R_i(u, c)$  and  $R_i(v, c)$ . It suffices to make these four evaluations linearly independent, assuming the commitment and opening of  $R_i$  are using a zkVPD. Therefore, we set the low-degree term in Eq. 7 as  $Z_i(z) \sum_{w \in \{0,1\}} R_i(z_1, w)$ , i.e.  $R_i$  only takes two variables, the first variable  $z_1$  of  $z$  and an extra variable  $w \in \{0, 1\}$  instead of  $\{0, 1\}^\lambda$ , with variable degree 2.

The full protocol is presented in Construction 2. Here we use superscriptions (e.g.,  $u^{(i)}$ ) to denote random numbers or vectors for the  $i$ -th layer of the circuit.

**Construction 2.** 1. On a layered arithmetic circuit  $C$  with  $d$  layers and input in, the prover  $\mathcal{P}$  sends the output of the circuit out to the verifier  $\mathcal{V}$ .

2.  $\mathcal{P}$  randomly selects polynomials  $R_1(z_1, w), \dots, R_d(z_1, w) : \mathbb{F}^2 \rightarrow \mathbb{F}$  with variable degree 2.  $\mathcal{P}$  commits to these polynomials by sending  $\text{com}_i \leftarrow \text{Commit}(R_i)$  to  $\mathcal{V}$  for  $i \in [1, d]$ .
3.  $\mathcal{V}$  defines  $\dot{V}_0(z) = \tilde{V}_0(z)$ , where  $\tilde{V}_0(z)$  is the multilinear extension of out.  $\dot{V}_0(z)$  can be viewed as a special case with  $R_0(z_1, w)$  being the 0 polynomial.  $\mathcal{V}$  evaluates it at a random point  $\dot{V}_0(g^{(0)})$  and sends  $g^{(0)}$  to  $\mathcal{P}$ .
4.  $\mathcal{P}$  and  $\mathcal{V}$  execute the zero knowledge sumcheck protocol presented in Construction 1 on

$$\begin{aligned} \dot{V}_0(g^{(0)}) = & \sum_{x,y \in \{0,1\}^{s_1}} \tilde{mult}_1(g^{(0)}, x, y)(\dot{V}_1(x)\dot{V}_1(y)) \\ & + \tilde{add}_1(g^{(0)}, x, y)(\dot{V}_1(x) + \dot{V}_1(y)) \end{aligned}$$

If  $u_1^{(1)} = v_1^{(1)}$ ,  $\mathcal{P}$  aborts. At the end of the protocol,  $\mathcal{V}$  receives  $\dot{V}_1(u^{(1)})$  and  $\dot{V}_1(v^{(1)})$ .  $\mathcal{V}$  computes  $\tilde{mult}_1(g^{(0)}, u^{(1)}, v^{(1)})$ ,  $\tilde{add}_1(g^{(0)}, u^{(1)}, v^{(1)})$  and checks that

$$\tilde{mult}_1(g^{(0)}, u^{(1)}, v^{(1)})\dot{V}_1(u^{(1)})\dot{V}_1(v^{(1)}) + \tilde{add}_1(g^{(0)}, u^{(1)}, v^{(1)})(\dot{V}_1(u^{(1)}) + \dot{V}_1(v^{(1)}))$$

equals to the last message of the sumcheck (evaluation oracle).

1. For layer  $i = 1, \dots, d - 1$ :
  - (a)  $\mathcal{V}$  randomly selects  $\alpha^{(i)}, \beta^{(i)} \in \mathbb{F}$  and sends them to  $\mathcal{P}$ .
  - (b) Let  $Mult_{i+1}(x, y) = \alpha^{(i)}\tilde{mult}_{i+1}(u^{(i)}, x, y) + \beta^{(i)}\tilde{mult}_{i+1}(v^{(i)}, x, y)$  and  
 $Add_{i+1}(x, y) = \alpha^{(i)}\tilde{add}_{i+1}(u^{(i)}, x, y) + \beta^{(i)}\tilde{add}_{i+1}(v^{(i)}, x, y)$ .  $\mathcal{P}$  and  $\mathcal{V}$  run the zero knowledge sumcheck on the equation  
 $\alpha^{(i)}\dot{V}_i(u^{(i)}) + \beta^{(i)}\dot{V}_i(v^{(i)}) =$

$$\begin{aligned} & \sum_{\substack{x,y \in \{0,1\}^{s_{i+1}} \\ w \in \{0,1\}}} (I(\mathbf{0}, w) \cdot Mult_{i+1}(x, y)(\dot{V}_{i+1}(x)\dot{V}_{i+1}(y)) \\ & + Add_{i+1}(x, y)(\dot{V}_{i+1}(x) + \dot{V}_{i+1}(y)) \\ & + I((x, y), \mathbf{0})(\alpha^{(i)}Z_i(u^{(i)})R_i(u_1^{(i)}, w) + \beta^{(i)}Z_i(v^{(i)})R_i(v_1^{(i)}, w))) \end{aligned}$$

If  $u_1^{(i+1)} = v_1^{(i+1)}$ ,  $\mathcal{P}$  aborts.

- (c) At the end of the zero-knowledge sumcheck protocol,  $\mathcal{P}$  sends  $\mathcal{V}$   $\dot{V}_{i+1}(u^{(i+1)})$  and  $\dot{V}_{i+1}(v^{(i+1)})$ .
- (d)  $\mathcal{V}$  computes

$$a_{i+1} = \alpha^{(i)} \tilde{\text{mult}}_{i+1}(u^{(i)}, u^{(i+1)}, v^{(i+1)}) + \beta^{(i)} \tilde{\text{mult}}_{i+1}(v^{(i)}, u^{(i+1)}, v^{(i+1)})$$

and

$$b_{i+1} = \alpha^{(i)} \tilde{\text{add}}_{i+1}(u^{(i)}, u^{(i+1)}, v^{(i+1)}) + \beta^{(i)} \tilde{\text{add}}_{i+1}(v^{(i)}, u^{(i+1)}, v^{(i+1)})$$

locally.  $\mathcal{V}$  computes  $Z_i(u^{(i)})$ ,  $Z_i(v^{(i)})$ ,  $I(\mathbf{0}, c^{(i)})$ ,  $I((u^{(i+1)}, v^{(i+1)}), \mathbf{0})$  locally.

- (e)  $\mathcal{P}$  and  $\mathcal{V}$  open  $R_i$  at two points  $R_i(u_1^{(i)}, c^{(i)})$  and  $R_i(v_1^{(i)}, c^{(i)})$  using **Open and Verify**.
- (f)  $\mathcal{V}$  computes the following as the evaluation oracle and uses it to complete the last step of the zero-knowledge sumcheck.

$$\begin{aligned} & I(\mathbf{0}, c^{(i)})(a_{i+1}(\dot{V}_{i+1}(u^{(i+1)})\dot{V}_{i+1}(v^{(i+1)})) + \\ & b_{i+1}(\dot{V}_{i+1}(u^{(i+1)}) + \dot{V}_{i+1}(v^{(i+1)}))) + \\ & I((u^{(i+1)}, v^{(i+1)}), \mathbf{0})(\alpha^{(i)} Z_i(u^{(i)})R_i(u_1^{(i)}, c^{(i)}) + \beta^{(i)} Z_i(v^{(i)})R_i(v_1^{(i)}, c^{(i)})) \end{aligned}$$

If all checks in the zero knowledge sumcheck and **Verify** passes,  $\mathcal{V}$  uses  $\dot{V}_{i+1}(u^{(i+1)})$  and  $\dot{V}_{i+1}(v^{(i+1)})$  to proceed to the  $(i + 1)$ -th layer. Otherwise,  $\mathcal{V}$  outputs reject and aborts.

6. At the input layer  $d$ ,  $\mathcal{V}$  has two claims  $\dot{V}_d(u^{(d)})$  and  $\dot{V}_d(v^{(d)})$ .  $\mathcal{V}$  opens  $R_d$  at 4 points  $R_d(u_1^{(d)}, 0)$ ,  $R_d(u_1^{(d)}, 1)$ ,  $R_d(v_1^{(d)}, 0)$ ,  $R_d(v_1^{(d)}, 1)$  and checks that  $\dot{V}_d(u^{(d)}) = \tilde{V}_d(u^{(d)}) + Z_d(u^{(d)}) \sum_{w \in \{0,1\}} R_d(u_1^{(d)}, w)$  and  $\dot{V}_d(v^{(d)}) = \tilde{V}_d(v^{(d)}) + Z_d(v^{(d)}) \sum_{w \in \{0,1\}} R_d(v_1^{(d)}, w)$ , given oracle access to two evaluations of  $\tilde{V}_d$  at  $u^{(d)}$  and  $v^{(d)}$ . If the check passes, output **accept**; otherwise, output **reject**.

**Theorem 3.** Construction 2 is an interactive proof protocol, for a function  $f$  defined by a layered arithmetic circuit  $C$  such that  $f(\text{in}, \text{out}) = 1$  iff  $C(\text{in}) = \text{out}$ . In addition, for every verifier  $\mathcal{V}^*$  and every layered circuit  $C$ , there exists a simulator  $\mathcal{S}$  such that given oracle access to **out**,  $\mathcal{S}$  is able to simulate the partial view of  $\mathcal{V}^*$  in step 1–5 of Construction 2.

The completeness follows from the construction explained above and the completeness of the zero knowledge sumcheck. The soundness follows the soundness of the GKR protocol with low degree extensions, as proven in [30] and [22]. We defer the proof of zero knowledge to the full version.

### 4.3 Zero Knowledge VPD

In this section, we present the instantiations of the zkVPD protocol, as described in Sect. 2.4. For every intermediate layer  $i$ , we use the same zkVPD protocol as proposed by Zhang et al. in [49] to commit and open the masking polynomials  $g_i(x), R_i(z_1, w)$ . In fact, as we show in the previous sections, these polynomials are very small ( $g_i$  is the sum of univariate polynomials and  $R_i$  has 2 variables with variable degree 2), the zkVPD protocols become very simple. The complexity of KeyGen, Commit, Open, Verify and proof size are all  $O(s_i)$  for  $g_i$  and are all  $O(1)$  for  $R_i$ . We omit the full protocols due to space limit.

For the zkVPD used for the input layer, we design a customized protocol based on the zkVPD protocol in [49]. Recall that at the end of the GKR protocol,  $\mathcal{P}$  sends two evaluations of  $\dot{V}_d(z) = \tilde{V}_d(z) + Z_d(z) \sum_{w \in \{0,1\}} R_d(z_1, w)$  at  $z = u^{(d)}$  and  $z = v^{(d)}$ . In our zero knowledge proof protocol, which will be presented in Sect. 4.4,  $\mathcal{P}$  commits to  $\dot{V}_d(z)$  using the zkVPD at the beginning, and opens it to the two points selected by  $\mathcal{V}$ .

The protocol in [49] works for any polynomial with  $\ell$  variables and any variable degree, and is particularly efficient for multilinear polynomials. We modify the protocol for our zero-knowledge proof scheme and preserve the efficiency. Note that though  $\dot{V}_d(z)$  is a low degree extension of the input, it can be decomposed to the sum of  $\tilde{V}_d(z)$ , a multilinear polynomial, and  $Z_d(z) \sum_{w \in \{0,1\}} R_d(z_1, w)$ . Moreover,  $Z_d(u^{(d)})$  and  $Z_d(v^{(d)})$  can be computed directly by  $\mathcal{V}$ . Therefore, in our construction,  $\mathcal{P}$  commits to  $\tilde{V}_d(z)$  and  $\sum_{w \in \{0,1\}} R_d(z_1, w)$  separately, and later opens the sum together given  $Z_d(u^{(d)})$  and  $Z_d(v^{(d)})$ , which is naturally supported because of the homomorphic property of the commitment. Another optimization is that unlike other layers of the circuit,  $R_d(z_1, w)$  itself is not opened at two points ( $\mathcal{V}$  does not receive  $R_d(u^{(d)}, c^{(d)})$  and  $R_d(v^{(d)}, c^{(d)})$  in Construction 2). Therefore, it suffices to set  $\dot{V}_d(z) = \tilde{V}_d(z) + Z_d(z)R_d(z_1)$ , where  $R_d$  is a univariate linear polynomial. We will give the full protocol in the full version.

### 4.4 Putting Everything Together

In this section, we present our zero knowledge argument scheme. At a high level, similar to [48–50],  $\mathcal{V}$  can use the GKR protocol to verify the correct evaluation of a circuit  $C$  on input  $x$  and a witness  $w$ , given an oracle access to the evaluation of a polynomial defined by  $x, w$  on a random point. We instantiate the oracle using the zkVPD protocol. Formally, we present the construction in Construction 3, which combines our zero knowledge GKR and zkVPD protocols. Similar to the protocols in [48, 49], Step 6 and 7 are to check that  $\mathcal{P}$  indeed uses  $x$  as the input to the circuit.

**Construction 3.** Let  $\lambda$  be the security parameter,  $\mathbb{F}$  be a prime field,  $n$  be an upper bound on input size, and  $S$  be an upper bound on circuit size. We use  $\text{VPD}_1, \text{VPD}_2, \text{VPD}_3$  to denote the zkVPD protocols for input layer, masking polynomials  $g_i$  and  $R_i$  described in Construction 2.

- $\mathcal{G}(1^\lambda, n, S)$ : run  $(\text{pp}_1, \text{vp}_1) \leftarrow \text{VPD}_1.\text{KeyGen}(1^\lambda, \log n)$ ,  $(\text{pp}_2, \text{vp}_2) \leftarrow \text{VPD}_2.\text{KeyGen}(1^\lambda, \log S)$ ,  $(\text{pp}_3, \text{vp}_3) \leftarrow \text{VPD}_3.\text{KeyGen}(1^\lambda)$ . Output  $\text{pk} = (\text{pp}_1, \text{pp}_2, \text{pp}_3)$  and  $\text{vk} = (\text{vp}_1, \text{vp}_2, \text{vp}_3)$ .
- $\langle \mathcal{P}(\text{pk}, w), \mathcal{V}(\text{vk})(x) \rangle$ : Let  $C$  be a layered arithmetic circuit over  $\mathbb{F}$  with  $d$  layers, input  $x$  and witness  $w$  such that  $|x| + |w| \leq n$ ,  $|C| \leq S$  and  $C(x; w) = 1$ . Without loss of generality, assume  $|w|/|x| = 2^m - 1$  for some  $m \in \mathbb{N}$ .
  1.  $\mathcal{P}$  selects a random bivariate polynomial  $R_d$  with variable degree 2 and commits to the input of  $C$  by sending  $\text{com}_d \leftarrow \text{VPD}_1.\text{Commit}(\tilde{V}_d, r_V, r_R, \text{pp}_1)$  to  $\mathcal{V}$ , where  $\tilde{V}_d$  is the multilinear extension of array  $(x; w)$  and  $\tilde{V}_d = \tilde{V}_d + R_d$
  2.  $\mathcal{V}$  runs  $\text{VPD}_1.\text{CheckComm}(\text{com}_d, \text{vp}_1)$ . If it outputs reject,  $\mathcal{V}$  aborts and outputs reject.
  3.  $\mathcal{P}$  and  $\mathcal{V}$  execute Step 1–5 of the zero knowledge GKR protocol in Construction 2, with the zkVPDs instantiated with  $\text{VPD}_2$  and  $\text{VPD}_3$ . If Construction 2 rejects,  $\mathcal{V}$  outputs reject and aborts. Otherwise, by the end of this step,  $\mathcal{V}$  receives two claims of  $\tilde{V}_d$  at  $u^{(d)}$  and  $v^{(d)}$ .
  4.  $\mathcal{P}$  runs  $(y_1, \pi_1) \leftarrow \text{VPD}_1.\text{Open}(\tilde{V}, r_V, r_R, u^{(d)}, \text{pp}_1)$ ,  $(y_2, \pi_2) \leftarrow \text{VPD}_1.\text{Open}(\tilde{V}, r_V, r_R, v^{(d)}, \text{pp}_1)$  and sends  $y_1, \pi_1, y_2, \pi_2$  to  $\mathcal{V}$ .
  5.  $\mathcal{V}$  runs  $\text{Verify}(\text{com}_d, u^{(d)}, y_1, \pi_1, \text{vp}_1)$  and  $\text{Verify}(\text{com}_d, v^{(d)}, y_2, \pi_2, \text{vp}_1)$  and output reject if either check fails. Otherwise,  $\mathcal{V}$  checks  $\tilde{V}_d(u^{(d)}) = y_1$  and  $\tilde{V}_d(v^{(d)}) = y_2$ , and rejects if either fails.
  6.  $\mathcal{V}$  computes the multilinear extension of input  $x$  at a random point  $r_x \in \mathbb{F}^{\log |x|}$  and sends  $r_x$  to  $\mathcal{P}$ .
  7.  $\mathcal{P}$  pads  $r_x$  to  $r'_x \in \mathbb{F}^{\log |x|} \times \mathbb{0}^{\log |w|}$  with  $\log |w|$  0s and sends  $\mathcal{V} (y_x, \pi_x) \leftarrow \text{VPD}_1.\text{Open}(\tilde{V}_d, r_V, r_R, r'_x, \text{pp}_1)$ .  $\mathcal{V}$  checks  $\text{Verify}(\text{com}_d, r'_x, y_x, \pi_x, \text{vp}_1)$  and  $y_x$  equals the evaluation of the multilinear extension on  $x$ .  $\mathcal{V}$  outputs reject if the checks fail. Otherwise,  $\mathcal{V}$  outputs accept.

**Theorem 4.** For an input size  $n$  and a finite field  $\mathbb{F}$ , Construction 3 is a zero knowledge argument for the relation

$$\mathcal{R} = \{(C, x; w) : C \in \mathcal{C}_{\mathbb{F}} \wedge |x| + |w| \leq n \wedge C(x; w) = 1\},$$

as defined in Definition 1, under the  $q$ -SBDH and the extended PKE assumptions. Moreover, for every  $(C, x; w) \in \mathcal{R}$ , the running time of  $\mathcal{P}$  is  $O(|C|)$  field operations and  $O(n)$  multiplications in the base group of the bilinear map. The running time of  $\mathcal{V}$  is  $O(|x| + d \cdot \log |C|)$  if  $C$  is log-space uniform with  $d$  layers.  $\mathcal{P}$  and  $\mathcal{V}$  interact  $O(d \log |C|)$  rounds and the total communication (proof size) is  $O(d \log |C|)$ . In case  $d$  is polylog( $|C|$ ), the protocol is a succinct argument.

**Proof Sketch.** The correctness and the soundness follow from those of the two building blocks, zero knowledge GKR and zkVPD.

To prove zero knowledge, consider a simulator  $\mathcal{S}$  that calls the simulator  $\mathcal{S}_{GKR}$  of zero knowledge GKR given in Sect. 4.2 as a subroutine, which simulates

the partial view up to the input layer. At the input layer, the major challenge is that  $\mathcal{S}$  committed to (a randomly chosen)  $\hat{V}_d^*$  at the beginning of the protocol, before knowing the points  $u^{(d)}, v^{(d)}$  to evaluate on. If  $\mathcal{S}$  opens the commitment honestly, with high probability the evaluations are not consistent with the last message of the GKR (sumcheck in layer  $d-1$ ) and a malicious  $\mathcal{V}^*$  can distinguish the ideal world from the real world. In our proof, we resolve this issue by using the simulator  $\mathcal{S}_{VPD}$  of our zkVPD protocol. Given the trapdoor `trap` used in `KeyGen`,  $\mathcal{S}_{VPD}$  is able to open the commitment to any value in zero knowledge, and in particular it opens to those messages that are consistent with the GKR protocol in our scheme, which completes the construction of  $\mathcal{S}$ .

The complexity of our zero knowledge argument scheme follows from our new GKR protocol with linear prover time, and the complexity of the zkVPD protocol for the input layer analyzed in Sect. 4.3. The masking polynomials  $g_i, R_i$  and their commitments and openings introduce no asymptotic overhead and are efficient in practice.

**Removing Interactions.** Our construction can be made non-interactive in the random oracle model using Fiat—Shamir heuristic [26]. Though GKR protocol is not constant round, recent results [12, 20] show that applying Fiat-Shamir only incurs a polynomial soundness loss in the number of rounds in GKR. In our implementation, the GKR protocol is on a 254-bit prime field matching the bilinear group used in the zkVPD. The non-interactive version of our system provides a security level of 100+ bits.

## 5 Implementation and Evaluation

**Software.** We fully implement `Libra`, our new zero knowledge proof system in C++. There are around 3000 lines of code for the zkGKR protocol, 1000 lines for the zkVPD protocol and 700 lines for circuit generators. Our system provides an interface to take a generic layered arithmetic circuit and turn it into a zero knowledge proof. We implement a new class for large integers named `u512`, and use it together with the GMP [2] library for large numbers and field arithmetic. We use the ate-pairing [1] library on a 254-bit elliptic curve for the bilinear map used in zkVPD. We plan to open-source our system.

**Hardware.** We run all of the experiments on Amazon EC2 `c5.9xlarge` instances with 70 GB of RAM and Intel Xeon platinum 8124m CPU with 3 GHz virtual core. Our current implementation is not parallelized and we only use a single CPU core in the experiments. We report the average running time of 10 executions.

In the implementation, we developed a concrete optimization to support various types of gates with no extra overhead, instead of only addition and multiplication. It may be of independent interest and is presented in the full version.



**Table 2.** Prover time of our linear GKR and previous GKR variants.

Matrix multiplication	Matrix size	$4 \times 4$	$16 \times 16$	$64 \times 64$	$256 \times 256$
	[43]	0.0003 s	0.006 s	0.390 s	29.0 s
	Ours	0.0004 s	0.014 s	0.788 s	50.0 s
Image scaling	#pixels	$112 \times 112$	$176 \times 176$	$560 \times 560$	$1072 \times 1072$
	[46]	0.445 s	0.779 s	7.54 s	29.2 s
	Ours	0.337 s	1.25 s	19.8 s	79.2 s
Image scaling with different parameters	#pixels	$112 \times 112$	$176 \times 176$	$560 \times 560$	$1072 \times 1072$
	[50]	5.45 s	21.8 s	348 s	1441 s
	Ours	0.329 s	1.22 s	19.3 s	77.2 s
Random circuit	#gates per layer	$2^8$	$2^{12}$	$2^{16}$	$2^{20}$
	[23]	0.008 s	0.179 s	3.79 s	83.1 s
	Ours	0.002 s	0.039 s	0.635 s	10.8 s

## 5.1 Improvements on GKR Protocols

In this section, we compare the performance of our new GKR protocol with linear prover time with all variants of GKR in the literature on different circuits.

**Methodology and Benchmarks.** For fair comparisons, we re-implement all of these variants in C++ with the same libraries. The variants include: (1)  $O(C)$  for regular circuits, proposed in [43], where the two inputs of a gate can be described by two mapping functions with constant size in constant time. See [43] for the formal definition of regular circuits. (2)  $O(C + C' \log C')$  for data-parallel circuits with a small copy of size  $C'$ , proposed in [46]. (3)  $O(C \log C')$  for circuits with non-connected different copies of size  $C'$ , proposed in [51]. (4)  $O(C \log C)$  for arbitrary circuits, proposed in [23].

We compare our GKR protocol to these variants on the benchmarks below:

- **Matrix multiplication:**  $\mathcal{P}$  proves to  $\mathcal{V}$  that it knows two matrices whose product equals a public matrix. The representation of this function with an arithmetic circuit is highly regular<sup>3</sup>. We evaluate on different dimensions from  $4 \times 4$  to  $256 \times 256$  and the elements in the matrices are 32-bit integers.
- **Image scaling:** It computes a low-resolution image by scaling from a high-resolution image. We use the classic Lanczos re-sampling [44] method. It computes each pixel of the output as the convolution of the input with a sliding window and a kernel function defined as:  $k(x) = \text{sinc}(x)/\text{sinc}(ax)$ , if  $-a < x < a$ ;  $k(x) = 0$ , otherwise, where  $a$  is the scaling parameter and  $\text{sinc}(x) = \sin(x)/x$ . This function is data parallel, where each sub-circuit computes the same function to generate one pixel of the output image. We evaluate by

<sup>3</sup> We use the circuit representation of matrix multiplication with  $O(n^3)$  gates for fair comparisons, not the special protocol proposed in [43].

fixing the window size as  $16 \times 16$  and increase the image size from  $112 \times 112$  to  $1072 \times 1072$ . The pixels are 8-bit integers for greyscale images.

- **Image scaling of different parameters:** It is the same computation as above with different scaling parameters in the kernel function for different pixels. The circuit of this function consists of different sub-copies. We evaluate it with the same image sizes as above.
- **Random circuit:** It is randomly generated layered circuit. We randomly sample the type of each gate, input value and the wiring patterns. We fix the depth as 3 and increase the number of gates per layer from  $2^8$  to  $2^{20}$ .

To be consistent with the next section, all the protocols are executed on a 254-bit prime field. This does not affect the comparison at all, as all the protocols are in the same field. In Table 2, we report the prover time of the protocols. The proof size and the verification time of all the variants are similar.

**Results.** As shown in Table 2, the performance of our GKR protocol is comparable to those special protocols for structured circuits, and much better than the state-of-the-art on generic circuits. For example, for matrix multiplication, our protocol is slower by 1.3–2.4 $\times$ , because the protocol in [43] writes the wiring of matrix multiplication explicitly and does not need to compute *add* and *mult*. For image scaling, our protocol is slower by 2.5–4 $\times$ . This gap would become even smaller when the size of each sub-copy is larger. Here we use a small  $16 \times 16$  block, while the number of copies is 49–4489.


For image scaling with different parameters and generic random circuits, our protocol has a speedup of 4–8 $\times$ , and the speedup will increase with the scale of the circuits, as indicated by the complexity.

Besides the speedup on complicated circuits, a significant advantage of our new GKR protocol is on the prover interface of the system. In prior work such as [46, 51], as the protocols are particularly efficient for structured circuits, the circuits must be represented as small copies and the numbers of each copy. Even worse, the structure is explored per layer of the circuit, making the numbers of each copy potentially different in different layers. (E.g., 6 gates may be considered 3 copies with 2 gates and 2 copies with 3 gates in two different layers for efficiency purposes.) This constraint makes the interface of these systems hard to use and generalize. Our result gives a unified solution for arbitrary circuits, and it is the main reason that our prover can take the description of any layered arithmetic circuit potentially generated by other tools like Verilog.

## 5.2 Comparing to Other ZKP Schemes

In this section, we show the performance of Libra as a whole and compare it with several state-of-the-art zero knowledge proof systems.

**Methodology.** We compare with the following systems: libSNARK [13], Ligerio [6], libSTARK [8], Hyrax [48], Bulletproofs [16] and Aurora [11]. See Sect. 1 for more explanations of these systems and their asymptotic.



**Fig. 1.** Comparisons of prover time, proof size and verification time between Libra and existing zero knowledge proof systems.

- libSNARK: We use jsnark [4] to write the circuits (rank one constraint system (R1CS)), which compiles them to ZKP using the libSNARK backend [5].
- Ligero: As the system is not open-source, we use the same number reported in [6] on computing hashes.
- libSTARK: After communications with the authors of [8], we obtain numbers for proving the same number of hashes in the 3rd benchmark below from the authors. The experiments are executed on a server with 512GB of DDR3 RAM (1.6GHz) and 16 cores (2 threads per core) at speed of 3.2GHz.
- Hyrax: We use the open-source implementation of the system at [3].
- Bulletproofs: We use the system re-implemented by [48] at [3].
- Aurora: As a recently accepted paper, the system is not available and we extrapolate its performance using the numbers reported in the paper [11] for circuits with  $2^{10} - 2^{20}$  R1CS constrains.

**Benchmarks.** We evaluate the systems on three benchmarks: matrix multiplication, image scaling and Merkle Tree [40], which are used in [48]. Matrix

multiplication and image scaling are the same as explained in Sect. 5.1. In the third benchmark,  $\mathcal{P}$  proves to  $\mathcal{V}$  that it knows the value of the leaves of a Merkle tree [40] that computes to a public root value [15]. We use SHA-256 for the hash function. We implement it with a flat circuit where each sub-computation is one instance of the hash function. The consistency of the input and output of corresponding hashes are then checked by the circuit. There are  $2M - 1$  SHA256 invocations for a Merkle tree with  $M$  leaves. We increase the number of leaves from 16 to 256. We use the SHA-256 implemented by jsnark [4] in R1CS format to run libSNARK and estimate Aurora, and we use the SHA-256 arithmetic circuit implemented by Hyrax to run Hyrax, Bulletproofs and Libra. We only show the performance of Ligerio and libSTARK on the third benchmark.

We report the prover time, proof size and verification time in Fig. 1.

**Prover Time.** As shown in Fig. 1(a), (b) and (c), the prover in Libra is the fastest among all systems in all three benchmarks we tested. Ligerio is one of the best existing ZKP systems on prover time as it is purely based on symmetric key operations. Comparing to Ligerio, the prover time of Libra is  $1.15\times$  faster on a Merkle tree with 2 leaves and  $2\times$  faster with 256 leaves. Comparing to other systems, Libra improves the prover time by  $3.4\text{--}8.9\times$  vs. Hyrax,  $7.1\text{--}16.1\times$  vs. Aurora,  $10.1\text{--}12.4\times$  vs. libSTARK and  $65\text{--}166\times$  vs. Bulletproof.

Libra is also faster than libSNARK on general circuits by  $5\text{--}10\times$ , as shown in Fig. 1(a) and (b). The performance of Libra is comparable to libSNARK on Merkle trees in Fig. 1(c). This is because (1) most values in the circuit of SHA256 are binary, which is friendly to the prover of libSNARK as the time of exponentiation is proportional to the bit-length of the values; (2) The R1CS of SHA256 is highly optimized by jsnark [4] and real world products like Zcash [9]. There are only 26,000 constrains in one hash. In the arithmetic circuit used by Libra, there are 60,000 gates with 38,000 of them being multiplication gates. Even so, Libra is still as fast as libSNARK on a Merkle tree with 2 leaves and  $2\times$  faster with 256 leaves. We plan to further optimize the implementation of SHA256 as an arithmetic circuit in the future.

The gap between Libra and other systems will become bigger as the size of the circuit grows, as the prover time in these systems (other than Bulletproof) scales quasi-linearly with the circuit size. The evaluations justify that the prover time in Libra is both optimal asymptotically, and efficient in practice.

**Verification Time.** Figure 1(d), (e) and (f) show the verification time. Our verifier is much slower than libSNARK and libSTARK, which runs in 1.8 ms and 28–44 ms respectively in all the benchmarks.

Other than these two systems, the verification time of Libra is faster, as it grows sub-linearly with the circuit size. In particular, our verification time ranges from 0.08–1.15 s in the benchmarks we consider. In Fig. 1(f), the verification time of Libra is  $8\times$  slower than Aurora when  $M = 2$ , and  $15\times$  faster when  $M = 256$ . Libra is  $2.5\times$  slower than Ligerio with  $M = 2$  and  $4\times$  faster with  $M = 256$ . Comparing to Hyrax and Bulletproof, our verification is  $1.2\text{--}9\times$  and  $27\text{--}900\times$  faster respectively. Again, the gap increases with the scale of the circuits as our verification is succinct.

**Proof Size.** We report the proof size in Fig. 1(g), (h) and (i). Our proof size is much bigger than libSNARK, which is 128 bytes for all circuits, and Bulletproof, which ranges in 2–5.5 KBs. The proof size in Libra is in the range of 30–60 KBs, except for the matrix multiplications where it reduces to 5–9 KBs. This is better than Aurora, Hyrax and libSTARK, which also have poly-logarithmic proof size to the circuit. Finally, the proof size in Ligerio is  $O(\sqrt{C})$  and grows to several MBs.

**Setup Time.** Among all the systems, only Libra and libSNARK require trusted setup. Thanks to the optimization described in the beginning of this section, it only takes 202s to generate the public parameters in our largest instance with  $n = 2^{24}$ . Libra only needs to perform this setup once and it can be used for all benchmarks and all circuits with no more inputs. libSNARK requires a per-circuit setup. For example, it takes 1027s for the Merkle tree with 256 leaves, and takes 210s for  $64 \times 64$  matrix multiplications.

**Acknowledgments.** This material is in part based upon work supported by DARPA under Grant No. N66001-15-C-4066 and Center for Long-Term Cybersecurity (CLTC). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA or CLTC. Charalampos Papamanthou’s work was supported by NSF grants #1652259 and #1514261 and by a NIST grant.

## References

1. Ate-pairing. <https://github.com/herumi/ate-pairing>
2. The GNU multiple precision arithmetic library. <https://gmplib.org/>
3. Hyrax reference implementation. <https://github.com/hyraxZK/hyraxZK>
4. jsnark. <https://github.com/akosba/jsnark>
5. libsnark. <https://github.com/scipr-lab/libsnark>
6. Ames, S., Hazay, C., Ishai, Y., Venkitasubramaniam, M.: Ligerio: lightweight sub-linear arguments without a trusted setup. In: Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (2017)
7. Bayer, S., Groth, J.: Efficient zero-knowledge argument for correctness of a shuffle. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 263–280. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-29011-4\\_17](https://doi.org/10.1007/978-3-642-29011-4_17)
8. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable, transparent, and post-quantum secure computational integrity. Cryptology ePrint (2018)
9. Ben-Sasson, E., et al.: Zerocash: decentralized anonymous payments from bitcoin. In: Proceedings of the Symposium on Security and Privacy SP (2014)
10. Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., Virza, M.: SNARKs for C: verifying program executions succinctly and in zero knowledge. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 90–108. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-40084-1\\_6](https://doi.org/10.1007/978-3-642-40084-1_6)
11. Ben-Sasson, E., Chiesa, A., Riabzev, M., Spooner, N., Virza, M., Ward, N.P.: Aurora: transparent succinct arguments for R1CS. Cryptology ePrint (2018)

12. Ben-Sasson, E., Chiesa, A., Spooner, N.: Interactive oracle proofs. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9986, pp. 31–60. Springer, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-53644-5\\_2](https://doi.org/10.1007/978-3-662-53644-5_2)
13. Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Succinct non-interactive zero knowledge for a von Neumann architecture. In: Proceedings of the USENIX Security Symposium (2014)
14. Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Scalable zero knowledge via cycles of elliptic curves. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8617, pp. 276–294. Springer, Heidelberg (2014). [https://doi.org/10.1007/978-3-662-44381-1\\_16](https://doi.org/10.1007/978-3-662-44381-1_16)
15. Blum, M., Evans, W., Gemmell, P., Kannan, S., Naor, M.: Checking the correctness of memories. *Algorithmica* **12**(2–3), 225–244 (1994)
16. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: short proofs for confidential transactions and more. In: Proceedings of the Symposium on Security and Privacy (SP), pp. 319–338 (2018)
17. Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 327–357. Springer, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-49896-5\\_12](https://doi.org/10.1007/978-3-662-49896-5_12)
18. Bootle, J., Cerulli, A., Ghadafi, E., Groth, J., Hajiabadi, M., Jakobsen, S.K.: Linear-time zero-knowledge proofs for arithmetic circuit satisfiability. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10626, pp. 336–365. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-70700-6\\_12](https://doi.org/10.1007/978-3-319-70700-6_12)
19. Braun, B., Feldman, A.J., Ren, Z., Setty, S.T.V., Blumberg, A.J., Walfish, M.: Verifying computations with state. In: ACM SIGOPS 24th Symposium on Operating Systems Principles, SOSP (2013)
20. Canetti, R., Chen, Y., Holmgren, J., Lombardi, A., Rothblum, G.N., Rothblum, R.D.: Fiat-Shamir from simpler assumptions. IACR Cryptology ePrint Archive 2018:1004 (2018)
21. Chase, M., et al.: Post-quantum zero-knowledge and signatures from symmetric-key primitives. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 1825–1842. ACM (2017)
22. Chiesa, A., Forbes, M.A., Spooner, N.: A zero knowledge sumcheck and its applications. CoRR abs/1704.02086 (2017). <http://arxiv.org/abs/1704.02086>
23. Cormode, G., Mitzenmacher, M., Thaler, J.: Practical verified computation with streaming interactive proofs. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS 2012 (2012)
24. Costello, C., et al.: Geppetto: Versatile verifiable computation. In: S&P (2015)
25. Cramer, R., Damgård, I.: Zero-knowledge proofs for finite field arithmetic, or: can zero-knowledge be for free? In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 424–441. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0055745>
26. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987). [https://doi.org/10.1007/3-540-47721-7\\_12](https://doi.org/10.1007/3-540-47721-7_12)
27. Fiore, D., Fournet, C., Ghosh, E., Kohlweiss, M., Ohrimenko, O., Parno, B.: Hash first, argue later: adaptive verifiable computations on outsourced data. In: ACM SIGSAC Conference on Computer and Communications Security (2016)

28. Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct NIZKs without PCPs. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 626–645. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-38348-9\\_37](https://doi.org/10.1007/978-3-642-38348-9_37)
29. Giacomelli, I., Madsen, J., Orlandi, C.: ZKBoo: faster zero-knowledge for boolean circuits. In: USENIX Security Symposium, pp. 1069–1083 (2016)
30. Goldwasser, S., Kalai, Y.T., Rothblum, G.N.: Delegating computation: interactive proofs for Muggles. *J. ACM* **62**(4), 27:1–27:64 (2015)
31. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. *SIAM J. Comput.* **18**(1), 186–208 (1989)
32. Groth, J.: Linear algebra with sub-linear zero-knowledge arguments. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 192–208. Springer, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-03356-8\\_12](https://doi.org/10.1007/978-3-642-03356-8_12)
33. Groth, J.: Short pairing-based non-interactive zero-knowledge arguments. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 321–340. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-17373-8\\_19](https://doi.org/10.1007/978-3-642-17373-8_19)
34. Ishai, Y., Kushilevitz, E., Ostrovsky, R.: Efficient arguments without short PCPs. In: 22nd Annual IEEE Conference on Computational Complexity (CCC 2007) (2007)
35. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Zero-knowledge from secure multiparty computation. In: Proceedings of the Annual ACM Symposium on Theory of Computing, pp. 21–30. ACM (2007)
36. Kilian, J.: A note on efficient zero-knowledge proofs and arguments (extended abstract). In: Proceedings of the ACM Symposium on Theory of Computing (1992)
37. Kosba, A., Miller, A., Shi, E., Wen, Z., Papamanthou, C.: Hawk: the blockchain model of cryptography and privacy-preserving smart contracts. In: Proceedings of Symposium on Security and Privacy (SP) (2016)
38. Lipmaa, H.: Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 169–189. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-28914-9\\_10](https://doi.org/10.1007/978-3-642-28914-9_10)
39. Lund, C., Fortnow, L., Karloff, H., Nisan, N.: Algebraic methods for interactive proof systems. *J. ACM* **39**(4), 859–868 (1992)
40. Merkle, R.C.: A digital signature based on a conventional encryption function. In: Pomerance, C. (ed.) CRYPTO 1987. LNCS, vol. 293, pp. 369–378. Springer, Heidelberg (1988). [https://doi.org/10.1007/3-540-48184-2\\_32](https://doi.org/10.1007/3-540-48184-2_32)
41. Micali, S.: Computationally sound proofs. *SIAM J. Comput.* (2000)
42. Parno, B., Howell, J., Gentry, C., Raykova, M.: Pinocchio: nearly practical verifiable computation. In: S&P 2013, pp. 238–252 (2013)
43. Thaler, J.: Time-optimal interactive proofs for circuit evaluation. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 71–89. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-40084-1\\_5](https://doi.org/10.1007/978-3-642-40084-1_5)
44. Turkowski, K.: Filters for common resampling tasks. In: Graphics Gems, pp. 147–165. Academic Press Professional, Inc. (1990)
45. Vu, V., Setty, S., Blumberg, A.J., Walfish, M.: A hybrid architecture for interactive verifiable computation. In: Proceedings of the 2013 IEEE Symposium on Security and Privacy, SP 2013 (2013)
46. Wahby, R.S., et al.: Full accounting for verifiable outsourcing. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. ACM (2017)
47. Wahby, R.S., Setty, S.T., Ren, Z., Blumberg, A.J., Walfish, M.: Efficient RAM and control flow in verifiable outsourced computation. In: NDSS (2015)

48. Wahby, R.S., Tzialla, I., Shelat, A., Thaler, J., Walfish, M.: Doubly-efficient zkSNARKs without trusted setup. In: 2018 IEEE Symposium on Security and Privacy (SP), pp. 926–943. IEEE (2018)
49. Zhang, Y., Genkin, D., Katz, J., Papadopoulos, D., Papamanthou, C.: A zero-knowledge version of vSQL. Cryptology ePrint (2017)
50. Zhang, Y., Genkin, D., Katz, J., Papadopoulos, D., Papamanthou, C.: vSQL: verifying arbitrary SQL queries over dynamic outsourced databases. In: 2017 IEEE Symposium on Security and Privacy (SP), pp. 863–880. IEEE (2017)
51. Zhang, Y., Genkin, D., Katz, J., Papadopoulos, D., Papamanthou, C.: vRAM: faster verifiable RAM with program-independent preprocessing. In: Proceeding of IEEE Symposium on Security and Privacy (S&P) (2018)