



Orion: Zero Knowledge Proof with Linear Prover Time

Tiancheng Xie¹, Yupeng Zhang², and Dawn Song^{1(✉)}

¹ University of California, Berkeley, Berkeley, USA
`{tianc.x, dawnsong}@berkeley.edu`

² Texas A&M University, College Station, USA
`zhangyp@tamu.edu`

Abstract. Zero-knowledge proof is a powerful cryptographic primitive that has found various applications in the real world. However, existing schemes with succinct proof size suffer from a high overhead on the proof generation time that is super-linear in the size of the statement represented as an arithmetic circuit, limiting their efficiency and scalability in practice. In this paper, we present **Orion**, a new zero-knowledge argument system that achieves $O(N)$ prover time of field operations and hash functions and $O(\log^2 N)$ proof size. **Orion** is concretely efficient and our implementation shows that the prover time is 3.09s and the proof size is 1.5 MB for a circuit with 2^{20} multiplication gates. The prover time is the fastest among all existing succinct proof systems, and the proof size is an order of magnitude smaller than a recent scheme proposed in Golovnev et al. 2021.

In particular, we develop two new techniques leading to the efficiency improvement. (1) We propose a new algorithm to test whether a random bipartite graph is a lossless expander graph or not based on the densest subgraph algorithm. It allows us to sample lossless expanders with an overwhelming probability. The technique improves the efficiency and/or security of all existing zero-knowledge argument schemes with a linear prover time. The testing algorithm based on densest subgraph may be of independent interest for other applications of expander graphs. (2) We develop an efficient proof composition scheme, code switching, to reduce the proof size from square root to polylogarithmic in the size of the computation. The scheme is built on the encoding circuit of a linear code and shows that the witness of a second zero-knowledge argument is the same as the message in the linear code. The proof composition only introduces a small overhead on the prover time.

1 Introduction

Zero-knowledge proof (ZKP) allows a *prover* to convince a *verifier* that a statement is valid, without revealing any additional information about the prover's secret witness of the statement. Since it was first introduced in the seminal paper by Goldwasser, Micali and Rackoff [GMR89], ZKP has evolved from a purely theoretical interest to a concretely efficient cryptographic primitive,

leading to many real-world applications in practice. It has been widely used in blockchains and cryptocurrencies to achieve privacy (Zcash [BCG+14, zca]) and to improve scalability (zkRollup [zkr]). More recently, it also found applications in zero-knowledge machine learning [ZFZS20, LKKO20, LXZ21, FQZ+21, WYX+21], zero-knowledge program analysis [FDNZ21], and zero-knowledge middlebox [GAZ+22].

There are three major efficiency measures in ZKP: the overhead of the prover to generate the proof, which is referred to as the *prover time*; the total communication between the prover and the verifier, which is called the *proof size*; and the time to verify the proof, which is called the *verifier time*. Despite its recent progress, the efficiency of ZKP is still not good enough for many applications. In particular, the prover time is one of the major bottlenecks preventing existing ZKP schemes from scaling to large statements. As pointed out by Golovnev et al. in [GLS+], to prove a statement that can be modeled as an arithmetic circuit with N gates, existing schemes with succinct proof size either perform a fast Fourier transform (FFT) due to the Reed-Solomon code encodings or polynomial interpolations, or a multi-scalar exponentiation due to the use of discrete-logarithm assumptions or bilinear maps, over a vector of size $O(N)$. The former takes $O(N \log N)$ field additions and multiplications and the latter takes $O(N \log |\mathbb{F}|)$ field multiplications, where $|\mathbb{F}|$ is the size of the finite field. With the Pippenger's algorithm [Pip], the complexity of the multi-scalar exponentiation can be improved to $O(N \log |\mathbb{F}| / \log N)$, which is still super-linear as $\log |\mathbb{F}| = \omega(\log N)$ to ensure security. These operations are indeed the dominating cost of the prover time both asymptotically and concretely. See Sect. 1.3 for more discussions about existing ZKP schemes categorized by the underlying cryptographic techniques.

The only exceptions in the literature are schemes in [BCG+17, BCG20, BCL22, GLS+]. Bootle et al. [BCG+17] proposed the first ZKP scheme with a prover time of $O(N)$ field operations and a proof size of $O(\sqrt{N})$ using a linear-time encodable error-correcting code. The proof size is later improved to $O(N^{1/c})$ for any constant c via a tensor code in [BCG20], and then to $\text{polylog}(N)$ via a generic proof composition with a probabilistic checkable proof (PCP) in [BCL22]. These schemes are mainly for theoretical interests and do not have implementations with good concrete efficiency. Recently, Golovnev et al. [GLS+] proposed a ZKP scheme based on the techniques in [BCG20] by instantiating the linear-time encodable code with a randomized construction. However, the security guarantee (soundness error) is only inverse polynomial in the size of the circuit, instead of negligible. Moreover, the proof size of the implemented scheme is $O(\sqrt{N})$ (more details are presented in Sect. 1.3). Therefore, the following question still remains open:

Can we construct a concretely efficient ZKP scheme with $O(N)$ prover time and $\text{polylog}(N)$ proof size?

Table 1. Comparison to existing ZKP schemes with linear prover time. N is the size of the circuit/R1CS and $c \geq 2$ is a constant.

	Prover time	Proof size	Verifier time	Soundness error	Concrete efficiency
[BCG+17]	$O(N)$	$O(\sqrt{N})$	$O(N)$	$\text{negl}(N)$	✗
[BCG20]	$O(N)$	$O(N^{1/c})$	$O(N)$	$\text{negl}(N)$	✗
[BCL22]	$O(N)$	$\text{polylog}(N)$	$O(N)$	$\text{negl}(N)$	✗
[GLS+]	$O(N)$	$O(\sqrt{N})$	$O(N)$	$O(\frac{1}{\text{poly}(N)})$	✓
our scheme	$O(N)$	$O(\log^2 N)$	$O(N)$	$\text{negl}(N)$	✓

1.1 Our Contributions

We answer the question above positively in this paper by proposing a new ZKP scheme. In particular, our contributions include:

- First, we propose a random construction of the linear-time encodable code that has a constant relative distance with overwhelming probability. Such a code was used in all existing linear-time ZKP schemes [BCG+17, BCG20, BCL22, GLS+] and thus our new construction also improves their efficiency. The key technique is a new algorithm to test whether a random graph is a good expander graph based on the densest sub-graph algorithm, which may be of independent interest for other applications of expander graphs [SZT02].
- Second, we propose a new reduction that achieves a proof size of $O(\log^2 N)$ efficiently. Our new technique is a proof composition named “code switching” that reduces the proof size of the schemes in [BCG20, GLS+] from $O(\sqrt{N})$ to $O(\log^2 N)$ with a small overhead on the prover time.
- Finally, we implement our new ZKP scheme, **Orion**, and evaluate it experimentally. On a circuit with 2^{20} gates (rank-1-constraint-system (R1CS) with 2^{20} constraints), the prover time is 3.09 s, the proof size is 1.5 MBs and the verifier time is 70 ms. **Orion** has the fastest prover time among all existing ZKP schemes in the literature. The proof size is $6.5 \times$ smaller than the system in [GLS+]. The scheme is plausibly post-quantum secure and can be made non-interactive via the Fiat-Shamir heuristic [FS86].

Table 1 shows the comparison between our scheme and existing schemes with linear prover time and succinct proof size.

1.2 Technical Overview

Testing Expander Graphs via Densest Sub-graph. All existing ZKP schemes with linear prover time and succinct proof size [BCG+17, BCG20, BCL22, GLS+] use linear-time encodable codes with a constant relative distance proposed in [Spi96, DI14, GLS+], which in turn all rely on the existence of good expander graphs. In a good expander graph, any subset of vertices expands to a large number of neighbors. Figure 1 shows an example of a bipartite graph where any subset of vertices on the left of size 2 expands to at least 5 vertices on the right. See Sect. 2.1 for formal definitions and constructions. However, how to

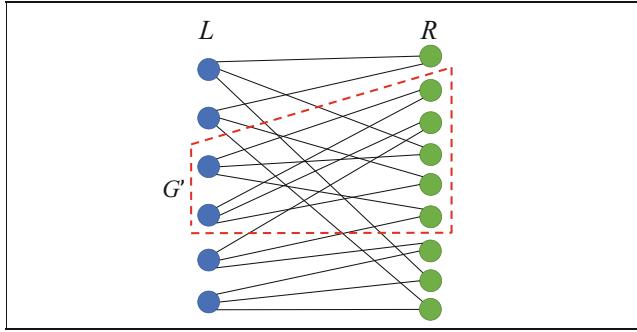


Fig. 1. An example of lossless expander. $k = 6, k' = 9, g = 3, \delta = 1, \epsilon = \frac{1}{6}$.

construct such good expanders remain unclear in practice. Explicit constructions [CRVW02] have large hidden constants in the complexity and thus are not practical. A random graph tends to have good expansion, but the probability that a random graph is not a good expander is inverse polynomial in the size of the graph. The code constructed from such a non-expanding graph does not have a good minimum distance, making the ZKP scheme insecure. Therefore, a randomly sampled graph is not good for cryptographic applications.

In this paper, we propose a new algorithm to efficiently test whether a random graph is a good expander or not. With the new testing algorithm, we are able to re-sample the random graph until it passes the test, obtaining a good expander with an overwhelming probability and boosting the soundness error of the ZKP scheme to be negligible. The testing algorithm is based on the densest sub-graph algorithm [Gol84]. The density of a graph $G = (V, E)$ is defined as the number of edges divided by the number of vertices $\frac{|E|}{|V|}$, and the densest sub-graph is simply the sub-graph in a graph with the maximum density. We observe that a good expander graph tends to have a small maximum density. This is because assuming the degree g of each vertex is a constant, e.g. $g = 3$ for all vertices on the left in Fig. 1, given any subset of vertices of size s in the graph, the total number of edges is fixed as $|E| = gs$ in the sub-graph defined by this subset and its neighbors. For example, any two vertices on the left in Fig. 1 as highlighted always have 6 outgoing edges. Then we differentiate two cases:

- In a good expander graph, any subset expands to a large number of neighbors, thus the total number of vertices in this sub-graph is large. Therefore, the density of any sub-graph is small;
- In contrast, if the graph is not a good expander, there is at least one subset that does not expand. Taking the sub-graph defined by this subset and its neighbors, again the number of edges is fixed, while the number of vertices is small. Therefore, the density of this sub-graph is large, which will be detected by the densest subgraph algorithm.

This observation gives us a way to differentiate good expanders. To the best of our knowledge, we are the first to make the connection between expander and the densest subgraph problem.

The real testing algorithm involves random sampling and repeating the densest sub-graph algorithm because of additional conditions of the expander. The formal algorithm, theorem and proofs are presented in Sect. 3.

Proof Composition via Code-Switching. With the expander graph sampled above and the corresponding linear code, we are able to build efficient ZKP schemes following the approaches in [BCG+17, BCG20, GLS+]. However, the proof size is $O(N^{1/c})$ instead of $\text{polylog}(N)$. To reduce the proof size, a common technique in the literature is proof composition. Instead of sending the proof directly to the verifier, the prover uses a second ZKP scheme to show that the proof of the first ZKP is indeed valid. In particular, in [BCG+17, BCG20, GLS+], the proof is a codeword of the linear-time encodable code, and the checks can be represented as several inner products between the message in the codeword of the proof and some public vectors.

Unfortunately, we do not have a second ZKP scheme with a $\text{polylog}(N)$ proof size to prove inner products. If we had it, we would already be able to build a ZKP scheme with $\text{polylog}(N)$ proof size in the first place. Instead, we rely on the fact that the proof is a codeword of the linear code and construct the second ZKP scheme as follows. One component of the second ZKP scheme is the *encoding circuit* of the linear-time encodable code. It takes the witness of the second ZKP scheme, encodes it and outputs several random locations of the codeword. The verifier checks that these random locations are the same as the proof of the first ZKP scheme, without receiving the entire proof. By the distance of the linear-time encodable code, we show that the witness of the second ZKP must be the same as the message in the proof of the first ZKP with overwhelming probability. After that, the other component of the second ZKP checks the inner product relationship modeled as an arithmetic circuit.

With this idea, we can use any general-purpose ZKP scheme on arithmetic circuits with a $\text{polylog}(N)$ proof size as the second ZKP scheme in the proof composition. The size of this circuit is only $O(\sqrt{N})$, thus the second ZKP does not introduce any overhead on the prover time as long as its prover time is no more than quadratic. In our construction, we use the ZKP scheme in [ZXZS20] as the second ZKP. The scheme is based on the interactive oracle proofs (IOP) and the witness is encoded using the Reed-Solomon code. Therefore, the technique is called code switching. The formal protocols are presented in Sect. 4.

1.3 Related Work

Zero-knowledge proof was introduced in [GMR89] and generic constructions based on PCPs were proposed by Kilian [Kil92] and Micali [Mic00] in the early days. Driven by various applications mentioned in the introduction, there has been significant progress in efficient ZKP protocols and systems. Categorized by their underlying techniques, there are ZKP systems based on bilinear maps [PHGR13, BSCG+13, BFR+13, BSCTV14, CFH+15, WSR+15, FFG+16, GKM+18, MBKM19, GWC19, CHM+20, KPPS20], MPC-in-the-head [GMO16, CDG+17, AHIV17, KKW18], interactive proofs [ZGK+17a, ZGK+17b, WTS+18,

[ZGK+18](#), [XZZ+19](#), [ZLW+21](#)], discrete logarithm [\[BBB+18, BFS20, Set20, SL20\]](#), interactive oracle proofs (IOP) [\[BSCR+19, BSBHR19, ZXZS20, BFH+20, COS20, BDFG20\]](#), and lattices [\[BBC+18, ESLL19, BLNS20, ISW21\]](#). As mentioned in the introduction, these schemes perform either an FFT (such as schemes based on MPC-in-the-head and IOP) or a multi-scalar exponentiation (such as schemes based on discrete-log and bilinear pairing), making the complexity of the prover time super-linear in the size of the circuit.

With the techniques proposed in [\[XZZ+19, ZLW+21\]](#), the prover time of the schemes based on the interactive proofs (the GKR protocol [\[GKR08\]](#)) is linear if the size of the input is significantly smaller than the size of the circuit. However, the goal of this paper is to make the prover time strictly linear without such a requirement, and our polynomial commitment scheme can also be plugged into these schemes to improve their efficiency.

Schemes with Linear Prover Time. As mentioned before, schemes in [\[BCG+17, BCG20, BCL22, GLS+\]](#) are the only candidates in the literature with linear prover time and succinct proof size. They all use linear-time encodable codes based on expander graphs and our first contribution applies to all of them. Moreover, our ZKP scheme is based on the polynomial commitment in [\[GLS+\]](#) and the tensor IOP in [\[BCG20\]](#), and we improve the proof size to $O(\log^2 N)$ through a proof composition. In fact, the scheme in [\[BCL22\]](#) also proposes a proof composition with the PCP in [\[Mie09\]](#). However, the complexity of the PCP is polynomial time. That is why the scheme in [\[BCL22\]](#) has to be built on the scheme in [\[BCG20\]](#) with a proof size of $O(N^{1/c})$ and is not concretely efficient, while our scheme can be built on top of the efficient scheme in [\[GLS+\]](#) with a proof size of $O(\sqrt{N})$. A similar proof composition with PCP was also used in [\[RZR20\]](#) for a different purpose. We view our approach using the encoding circuit as a variant of the proof composition that is efficient in practice, and we inherit the name “code switching” from [\[RZR20\]](#).

Finally, the scheme in [\[GLS+\]](#) samples a random graph to build the linear-time encodable code. The scheme achieves a soundness error of $O(\frac{1}{\text{poly}(N)})$ and the authors spent great efforts to calculate parameters to achieve a concrete failure probability of 2^{-100} for large circuits in practice [\[GLS+, Claim 2 and Fig. 2\]](#). Our sampling algorithm provides the provable security guarantee for a negligible soundness error for their scheme. Moreover, we improve the proof size from $O(\sqrt{N})$ to $O(\log^2 N)$ efficiently, solving an open problem left in [\[GLS+\]](#).

Schemes with Linear Proof Size. Recently, there is a line of work constructing ZKP based on secure multiparty computation (MPC) techniques [\[WYKW20, DIO21, BMRS21, YSWW21\]](#) and these schemes have demonstrated fast prover time in practice. If one treats a block cipher (e.g., AES) as a constant-time operation because of the CPU instruction, these schemes indeed have a linear time prover (we are using a similar CPU instruction for the hash function SHA-256 in our scheme to achieve linear prover time). However, they have linear proof size in the size of the circuit, are inherently interactive, and are not publicly verifiable, which are not desirable in many applications. We mainly focus on non-interactive ZKP with succinct proof size and public verifiability in this paper.

Expander Testing. Testing the properties of expander graphs is a deeply explored area in computer science. Many works [NS07, CS07, GR11] have proposed efficient testing algorithms without accessing the whole graph. However, these algorithms do not directly apply to our testing of lossless expander. For example, the algorithm in [NS07] based on random walks can differentiate good expanders from graphs that are far from expanders, while our scheme can differentiate whether a graph is a lossless expander or not with overwhelming probability. Of course our algorithm accesses the entire graph, which is fine in our application of linear-time encodable code. To the best of our knowledge, we are not aware of any testing algorithm with such properties.

There are also impossibility results on expander testing [KS16]. Due to different definitions of expansion, our testing algorithm cannot distinguish the cases in [KS16, Theorem 1.1] and thus it does not violate the impossibility results.

2 Preliminary

We use $[N]$ to denote the set $\{0, 1, 2, \dots, N - 1\}$. $\text{poly}(N)$ means a function upper bounded by a polynomial in N with a constant degree. We use $\lambda = \omega(\log N)$ to denote the security parameter, and $\text{negl}(N)$ to denote the negligible function in N , i.e. $\text{negl}(N) \leq \frac{1}{\text{poly}(N)}$ for all sufficiently large N and any polynomial. Some papers define $\text{negl}(\lambda)$ as the negligible function. As λ is a function of N , they are essentially the same and $\text{negl}(N) \leq \frac{1}{2^\lambda}$. “PPT” stands for probabilistic polynomial time. $\langle A(x), B(y) \rangle(z)$ denotes an interactive protocol between algorithms A, B with x as the input of A , y as the input of B and z as the common input.

2.1 Linear Time Encodable Linear Code

Definition 1 (Linear Code). A linear error-correcting code with message length k and codeword length n is a linear subspace $C \in \mathbb{F}^n$, such that there exists an injective mapping from message to codeword $E_C : \mathbb{F}^k \rightarrow C$, which is called the encoder of the code. Any linear combination of codewords is also a codeword. The rate of the code is defined as $\frac{k}{n}$. The distance between two codewords u, v is the hamming distance denoted as $\Delta(u, v)$. The minimum distance is $d = \min_{u, v} \Delta(u, v)$. Such a code is denoted as $[n, k, d]$ linear code, and we also refer to $\frac{d}{n}$ as the relative distance of the code.

Generalized Spielman Code. In our construction, we use a family of linear codes that can be encoded in linear time and has a constant relative distance [Spi96, DI14, GLS+]. The code was first proposed by Daniel Spielman in [Spi96] over the Boolean alphabet. Druk and Ishai [DI14] generalized it to a finite field \mathbb{F} , and introduced a distance boosting technique to achieve the Gilbert-Varshamov bound [Gil52, Var57]. We only use the basic construction over \mathbb{F} without the distance boosting, and thus refer to it as the generalized Spielman code in this paper. The code relies on the existence of lossless expander graphs, which is defined below:

Definition 2 (Lossless Expander [Spi96]). Let $G = (L, R, E)$ be a bipartite graph. $0 < \epsilon < 1$ and $0 < \delta$ be some constants. The vertex set consists of L and R , two disjoint subsets, henceforth the left and right vertex set. Let $\Gamma(S)$ be the neighbor set of some vertex set S . We say G is an $(k, k'; g)$ -lossless expander if $|L| = k$, $|R| = k' = \alpha k$ for some constant α , and the following property hold:

1. *Degree:* The degree of every vertex in L is g .
2. *Expansion:* $|\Gamma(S)| \geq (1 - \epsilon)g|S|$ for every $S \subseteq L$ with $|S| \leq \frac{\delta|L|}{g}$.

Intuitively speaking, a lossless expander has very strong expansion. As the degree of each left vertex is g , a set of $|S|$ left vertices have at most $g|S|$ neighbors, while the second condition requires that every set expands to at least $(1 - \epsilon)g|S|$ vertices for a small constant ϵ . Meanwhile, as the right vertex set has $|R| = \alpha k$ vertices, such an expansion is not possible if $|S| > \frac{\alpha k}{(1 - \epsilon)g}$, thus there is a condition $|S| \leq \frac{\delta k}{g}$ bounding the size of S . An example is shown in Fig. 1.

Construction of Generalized Spielman Code. With the lossless expander, we give a brief description of the generalized Spielman code. Let $G = (L, R, E)$ be a lossless expander with $|L| = 2^t$, $|R| = 2^{t-1}$. Let A_t be a $2^t \times 2^{t-1}$ matrix where $A_t[i][j] = 1$ if there is an edge i, j in G for $i \in [2^t], j \in [2^{t-1}]$; otherwise $A_t[i][j] = 0$. The generalized Spielman code is constructed as follows:

1. Let $E_C^t(x)$ be the encoder function of input length $|x| = 2^t$, and its output will be a codeword of size 2^{t+2} . We use E_C to denote the encoder function when length is clear.
2. If $|x| \leq n_0$ then directly output x , for some constant n_0 .
3. Compute $m_1 = xA_t$. Each entry of m_1 can be viewed as a vertex in R , and value of each vertex is the summation of its neighbors in L . The length of m_1 is 2^{t-1} .
4. Recursively apply the encoder E_C^{t-1} on m_1 , let $c_1 = E_C^{t-1}(m_1)$.
5. Compute $c_2 = c_1 A_{t+1}$.
6. Output $x \odot c_1 \odot c_2$ as the codeword of size 2^{t+2} . \odot denotes concatenation.

Lemma 1 (Generalized Spielman code, [DI14]). Given a family of lossless expander, that achieves $(1 - \epsilon)g|S|$ expansion with $|S| \leq \frac{\delta|L|}{g}$, for input size k , the generalized Spielman code is a $[4k, k, \frac{\delta}{8}k]$ linear code over \mathbb{F} .

The code in [GLS+] is a variant of generalized Spielman code. In their construction, random weights are assigned to each edge of lossless expander at line 3, 5. And randomize the output at line 6: $(x \otimes r) \odot c_1 \odot c_2$, here \otimes is element-wise multiply, r is a random vector.

Definition 3 (Tensor code). Let C be a $[n, k, d]$ linear code, the tensor code $C^{\otimes 2}$ of dimension 2 is the linear code in \mathbb{F}^{n^2} with message length k^2 , codeword length n^2 , and distance nd . We can view the codeword as a $n \times n$ matrix. We define the encoding function below:

1. A message of length $k \times k$ is parsed as a $k \times k$ matrix. Each row of the matrix is encoded using E_C , resulting in a codeword C_1 of size $k \times n$.
2. Each column of C_1 is then encoded again using E_C . The result C_2 of size $n \times n$ is the codeword of the tensor code.

2.2 Collision-Resistant Hash Functions and Merkle Tree

Let $H : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$ be a hash function. A Merkle Tree is a data structure that allows one to commit to $l = 2^{\text{dep}}$ messages by a single hash value h , such that revealing any bit of the message require $\text{dep} + 1$ hash values.

A Merkle hash tree is represented by a binary tree of depth dep where l messages elements m_1, m_2, \dots, m_l are assigned to the leaves of the tree. The values assigned to internal nodes are computed by hashing the value of its two child nodes. To reveal m_i , we need to reveal m_i together with the values on the path from m_i to the root. We denote the algorithm as follows:

1. $h \leftarrow \text{Merkle.Commit}(m_1, \dots, m_l)$.
2. $(m_i, \pi_i) \leftarrow \text{Merkle.Open}(m, i)$.
3. $\{\text{accept}, \text{reject}\} \leftarrow \text{Merkle.Verify}(\pi_i, m_i, h)$.

2.3 Zero-Knowledge Arguments

An argument system for an NP relation R is a protocol between a computationally bounded prover \mathcal{P} and a verifier \mathcal{V} . At the end of the protocol \mathcal{V} will be convinced that there exists a witness w such that $(x, w) \in R$ for some public input x . We focus on arguments of knowledge which require the prover know the witness w . We formally define zero-knowledge as follows:

Definition 4 (View). We denote by $\text{View}(\langle \mathcal{P}, \mathcal{V} \rangle(x))$ the view of \mathcal{V} in an interactive protocol with \mathcal{P} . Namely, it is the random variable $(r, b_1, b_2, \dots, b_n, v_1, v_2, \dots, v_m)$ where r is \mathcal{V} 's randomness, b_1, \dots, b_n are messages from \mathcal{V} to \mathcal{P} , and v_1, \dots, v_m are messages from \mathcal{P} to \mathcal{V} .

Definition 5. Let \mathcal{R} be an NP relation. A tuple of algorithm $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ is a zero-knowledge argument of knowledge for \mathcal{R} if the following holds.

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

$$\Pr[\langle \mathcal{P}(w), \mathcal{V}() \rangle(\text{pp}, x) = \text{accept}] = 1.$$

- **Knowledge Soundness.** For any PPT adversary \mathcal{P}^* , there exists a PPT extractor ε such that for every pp output by $\mathcal{G}(1^\lambda)$ and any x , the following probability is $\text{negl}(N)$:

$$\Pr[\langle \mathcal{P}^*(\cdot), \mathcal{V}() \rangle(\text{pp}, x) = \text{accept}, (x, w) \notin \mathcal{R} | w \leftarrow \varepsilon(\text{pp}, x, \text{View}(\langle \mathcal{P}^*(\cdot), \mathcal{V}() \rangle(\text{pp}, x)))]$$

- **Zero knowledge.** There exists a PPT simulator \mathcal{S} such that for any PPT algorithm \mathcal{V}^* , $(x, w) \in R$, pp output by $\mathcal{G}(1^\lambda)$, it holds that

$$\text{View}(\langle \mathcal{P}(w), \mathcal{V}^*() \rangle(x)) \approx \mathcal{S}^{\mathcal{V}^*}(\text{pp}, x)$$

Where $\mathcal{S}^{\mathcal{V}^*}(x)$ denotes that \mathcal{S} is given oracle accesses to \mathcal{V}^* ’s random tape.

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

Definition 6 (Arithmetic circuit). An arithmetic circuit C over \mathbb{F} and a set of variables x_1, \dots, x_N is a directed acyclic graph as follows:

1. Each vertex is called a “gate”. A gate with in-degree zero is an input gate and is labeled as a variable x_i or a constant field element in \mathbb{F} .
2. Other gates have 2 incoming edges. It calculates the addition or multiplication over the two inputs and output the result.
3. The size of the circuit is defined as the number of gates N .

2.4 Polynomial Commitment

A polynomial commitment consists of three algorithms:

- $\text{PC.Commit}(\phi(\cdot))$: the algorithm outputs a commitment \mathcal{R} of the polynomial $\phi(\cdot)$.
- $\text{PC.Prove}(\phi, \vec{x}, \mathcal{R})$: given an evaluation point $\phi(\vec{x})$, the algorithm outputs a tuple $\langle \vec{x}, \phi(\vec{x}), \pi_{\vec{x}} \rangle$, where $\pi_{\vec{x}}$ is the proof.
- $\text{PC.VerifyEval}(\pi_{\vec{x}}, \vec{x}, \phi(\vec{x}), \mathcal{R})$: given $\pi_{\vec{x}}, \vec{x}, \phi(\vec{x}), \mathcal{R}$, the algorithm checks if $\phi(\vec{x})$ is the correct evaluation. The algorithm outputs accept or reject.

Definition 7 ((Multivariate) Polynomial commitment). A polynomial commitment scheme has the following properties:

- **Correctness.** For every polynomial ϕ and evaluation point \vec{x} , the following probability holds:

$$\Pr \left(\begin{array}{c} \text{PC.Commit}(\phi) \rightarrow \mathcal{R} \\ \text{PC.Prove}(\phi, \vec{x}, \mathcal{R}) \rightarrow \vec{x}, y, \pi \\ y = \phi(\vec{x}) \\ \text{PC.VerifyEval}(\pi, \vec{x}, y, \mathcal{R}) \rightarrow \text{accept} \end{array} \right) = 1$$

- **Knowledge Soundness.** For any PPT adversary \mathcal{P}^* with PC.Commit^* , PC.Prove^* , there exists a PPT extractor \mathcal{E} such that the probability below is negligible:

$$\Pr \left(\begin{array}{c} \text{PC.Commit}^*(\phi^*) \rightarrow \mathcal{R}^* \\ \text{PC.Prove}^*(\phi^*, \vec{x}, \mathcal{R}^*) \rightarrow \vec{x}, y^*, \pi^* \\ \text{PC.VerifyEval}(\pi^*, \vec{x}, y^*, \mathcal{R}^*) \rightarrow \text{accept} \end{array} \middle| \phi^* \leftarrow \mathcal{E}(\mathcal{R}^*, \vec{x}, \pi^*, y^*) \wedge y^* \neq \phi^*(\vec{x}) \right)$$

- **Zero-knowledge.** For security parameter λ , polynomial ϕ , any PPT adversary \mathcal{A} , there exists a simulator $\mathcal{S} = [\mathcal{S}_0, \mathcal{S}_1]$, we consider following two experiments:

$\text{Real}_{\mathcal{A}, \phi}(\text{pp})$:

1. $\mathcal{R} \leftarrow \text{Commit}(\text{pp}, \phi)$
2. $\vec{x} \leftarrow \mathcal{A}(\mathcal{R}, \text{pp})$
3. $(\vec{x}, y, \pi) \leftarrow \text{Prove}(\phi, \vec{x}, \mathcal{R})$
4. $b \leftarrow \mathcal{A}(\pi, \vec{x}, y, \mathcal{R})$
5. Output b

$\text{Ideal}_{\mathcal{A}, \mathcal{S}^A}(\text{pp})$:

1. $\mathcal{R} \leftarrow \mathcal{S}_0(1^\lambda, \text{pp})$
2. $\vec{x} \leftarrow \mathcal{A}(\mathcal{R}, \text{pp})$
3. $(\vec{x}, y, \pi) \leftarrow \mathcal{S}_1^A(\vec{x}, \text{pp})$, given oracle access to $y = \phi(\vec{x})$
4. $b \leftarrow \mathcal{A}(\pi, \vec{x}, y, \mathcal{R})$
5. Output b

For any PPT adversary \mathcal{A} , two experiments are identically distributed:

$$\Pr[|\text{Real}_{\mathcal{A}, f}(\text{pp}) - \text{Ideal}_{\mathcal{A}, \mathcal{S}^A}(\text{pp})| = 1] \leq \text{negl}(N)$$

3 Testing Algorithm for Lossless Expander

As explained above, the generalized Spielman code relies on the existence of lossless expanders. On one hand, there are explicit constructions of lossless expanders in the literature [CRVW02]. However, there are large hidden constants in the complexity and the constructions are not practical. On the other hand, a random bipartite graph is a lossless expander with a high probability of $1 - O(\frac{1}{\text{poly}(k)})$, where k is the size of the left vertex set in the bipartite graph. However, this is not good enough for cryptographic applications.

In this section, we propose a new approach to sample a lossless expander with a negligible failure probability. The key ingredient of our approach is a new algorithm to test whether a randomly sampled bipartite graph is a lossless expander or not. We begin the section by introducing the classical randomized construction of a lossless expander and its analysis.

3.1 Random Construction of Lossless Expander

As defined in Definition 2, a lossless expander graph is a g -left-regular bipartite graph $G = (L, R, E)$. Wigderson et al. [HLW06, Lemma 1.9] showed that a random bipartite graph is a lossless expander with a high probability. In particular, we have the following lemma:

Lemma 2 ([HLW06]). *For fixed constant parameters $g, \delta, \alpha, \epsilon$, a random g -left-regular bipartite graph is a $(k, k'; g)$ -lossless-expander with probability $1 - O(\frac{1}{\text{poly}(k)})$.*

Proof. Let $G = (L, R, E)$ be a random bipartite graph with k vertices on the left and $k' = O(k)$ vertices on the right, where each left vertex connects to a randomly chosen set of g vertices on the right.

Let $s = |S|$ be the cardinality of a left subset of vertices $S \subseteq L$ such that $s \leq \frac{\delta k}{g}$, and let $t = |T|$ be the cardinality of a right subset of vertices $T \subseteq R$ such that $t \leq (1 - \epsilon)gs$. Let $X_{S, T}$ be an indicator random variable for the event that all

the edges from S connect to T . Then for a particular S , if $\sum_{T \in R} X_{S,T} = 0$, then the number of neighboring vertices of S must be larger than $(1-\epsilon)gs$. Otherwise, if there exists a $T \in R$ such that $X_{S,T} = 1$, i.e., all edges from S connect to T , the graph is not a lossless expander. As the edges are sampled randomly, the probability of this *non-expanding* event is $(\frac{t}{k'})^{sg}$. Therefore, summing over all S and by the union bound, the probability of a non-expanding graph is:

$$\begin{aligned} \Pr[(\sum_{S,T} X_{S,T}) > 0] &\leq \sum_{S,T} \Pr[X_{S,T} = 1] = \sum_{S,T} \left(\frac{t}{k'}\right)^{sg} \\ &\leq \sum_{s=2}^{\frac{\delta k}{g}} \binom{k}{s} \binom{k'}{t} \left(\frac{t}{k'}\right)^{sg} \leq \sum_{s=2}^{\frac{\delta k}{g}} \binom{k}{s} \binom{k'}{(1-\epsilon)gs} \left(\frac{(1-\epsilon)gs}{k'}\right)^{sg} \end{aligned}$$

Using the inequality $\binom{k}{s} \leq \left(\frac{ke}{s}\right)^s$, the probability above is

$$\begin{aligned} &\leq \sum_{s=2}^{\frac{\delta k}{g}} \left(\frac{ke}{s}\right)^s \left(\frac{k'e}{(1-\epsilon)gs}\right)^{(1-\epsilon)gs} \left(\frac{(1-\epsilon)gs}{k'}\right)^{sg} \\ &= \sum_{s=2}^{\frac{\delta k}{g}} \left(\frac{ke}{s}\right)^s e^{(1-\epsilon)gs} \left(\frac{(1-\epsilon)gs}{k'}\right)^{egs} \\ &= \sum_{s=2}^{\frac{\delta k}{g}} e^{(1-\epsilon)gs+s} \cdot \left(\frac{k}{s}\right)^s \cdot \left(\frac{(1-\epsilon)gs}{k'}\right)^{egs} \end{aligned} \tag{1}$$

When s, ϵ, g are constants and $k' = O(k)$, $e^{(1-\epsilon)gs+s}$ is a constant, $(\frac{k}{s})^s$ is $O(\text{poly}(k))$, and $(\frac{(1-\epsilon)gs}{k'})^{egs}$ is $O(\frac{1}{\text{poly}(k)})$. Therefore, the overall upper bound is at least $O(\frac{1}{\text{poly}(k)})$.

The derivation above shows that the probability that a random graph is not a lossless expander is upper-bounded by $O(\frac{1}{\text{poly}(k)})$, which is not negligible. Furthermore, we show that the lower-bound of the non-expanding probability is also not negligible through a simple argument here.

We focus on the case where s is a constant. The number of all possible subgraphs induced by a left subset of vertices S is at most $k'^{sg} = O(\text{poly}(k))$. That is, the size of the entire probability space is bounded by a polynomial. The number of non-expanding graphs is at least 1 (e.g., all edges from S connect to a single vertex in R). Therefore, the non-expanding probability is at least $O(\frac{1}{\text{poly}(k)})$.

Lossless Expander in [GLS+]. As explained in Sect. 2.1, in [GLS+], the authors extended the generalized Spielman code by adding random weights to the edges in the bipartite graph. However, the graph still needs to be a lossless expander in order to achieve a constant relative distance, and the same issue above applies to

their construction. In particular, as shown by [GLS+, Claim 2], the probability of *not* sampling a lossless expander is

$$2^{kH(15/k) + \alpha k H(19.2/(\alpha k)) - 15g \log \frac{\alpha k}{19.2}},$$

where $H(x) = -x \log x - (1-x) \log(1-x)$. We show that the probability above is not negligible. First, for any constant const ,

$$\begin{aligned} xH(\text{const}/x) &= x\left(-\frac{\text{const}}{x} \log \frac{\text{const}}{x} - (1 - \frac{\text{const}}{x}) \log(\frac{x - \text{const}}{x})\right) \\ &= (\text{const} \log(x) - \text{const} \log \text{const}) + (1 - \frac{\text{const}}{x}) \log(\frac{x - \text{const}}{x}). \end{aligned}$$

By taking the limit, we have $\lim_{x \rightarrow \infty} xH(\text{const}/x) = (\text{const} \log(x) - \text{const} \log \text{const}) + 1 \times 0$. Therefore, $xH(\text{const}/x) = O(\log x)$. Applying this fact to the equation above, $kH(15/k) + \alpha k H(19.2/(\alpha k)) = O(\log k)$, and $-15g \log \frac{\alpha k}{19.2} = -O(\log k)$. Therefore, $2^{kH(15/k) + \alpha k H(19.2/(\alpha k)) - 15g \log \frac{\alpha k}{19.2}}$ is at least $2^{-O(\log k)} = \frac{1}{\text{poly}(k)}$. The failure probability is similar to the upper bound in Eq. 1.

3.2 Algorithm Based on Densest Sub-graph

To reduce the non-expanding probability of the random construction, we take a closer look at the equations above. Equation 1 shows that the probability that a random bipartite graph is a not lossless expander is upper bounded by $\frac{1}{\text{poly}(k)}$. However, we observe that within the summation, the probability is actually negligible when s is large. In particular, if we decompose the summation in Eq. 1 into two sums, one for $2 \leq s \leq \log \log k$, and the other for $s \geq \log \log k$, the second part is

$$\sum_{s=\log \log k}^{\frac{\delta k}{g}} e^{(1-\epsilon)gs+s} \cdot \left(\frac{k}{s}\right)^s \cdot \left(\frac{(1-\epsilon)gs}{k'}\right)^{\epsilon gs}. \quad (2)$$

Lemma 3. *Equation 2 is negligible if the following conditions are met:*

1. $(1-\epsilon)\delta + \frac{\delta}{g} + \frac{\delta}{g} \log(\frac{g}{\delta}) + \log(\frac{\delta}{\alpha})\epsilon\delta < -0.001$,
2. $\epsilon d > 2$.

Here -0.001 is just any small constant that is less than 0. We give a proof in the full version of the paper. To provide an intuition on how these parameters are set, we give an example here: $\delta = \frac{1}{11}$, $\epsilon = \frac{7}{16}$, $g = 16$, $k' = \frac{1}{2}k$. We can verify the condition:

1. $\epsilon g = 7 > 2$.
2. $(1-\epsilon)\delta + \frac{\delta}{g} + \frac{\delta}{g} \log(\frac{g}{\delta}) + \log(\frac{\delta}{\alpha})\epsilon\delta = -0.009 < -0.001$.

Sampling Lossless Expander with Negligible Failure Probability. The observation above shows that the non-expanding probability is dominated by small sub-graphs with size $2 \leq s \leq \log \log k$. This actually matches our lower bound in Sect. 3.1, as there are only polynomially many such sub-graphs and there exist ones that do not expand. Therefore, in order to reduce the non-expanding probability, we propose a new algorithm that detects small sub-graphs of size $s \leq \log \log k$ that do not expand. The algorithm is based on the densest sub-graph problem, and we are the first to make the connection between the densest sub-graph and the lossless expander.

Definition 8 (Densest Sub-graph Problem). Let $G = (V, E)$ be an undirected graph, and let $S = (E_S, V_S)$ be a subgraph of G . The density of S is defined to be $\text{den}(S) = \frac{|E_S|}{|V_S|}$. The densest sub-graph problem is to find S such that it maximizes $\text{den}(S)$. We denote the maximum density by $\text{Den}(G)$.

Theorem 1 [Gol84]. For any graph $G = (V, E)$, there is a polynomial time algorithm that find the densest sub-graph $G' = (V', E')$ such that $V' \subseteq V$ and G' is the sub-graph. And $\frac{|E'|}{|V'|}$ is maximized. The running time of the algorithm is $O(|V||E| \log |E| \log |V|)$.

We will use this algorithm as a building block of our testing algorithm. First, we define a notion of perfect expander, and then derive the density of a perfect expander.

Definition 9 (Perfect expander). Let $G = (L, R, E)$ be a bipartite graph. We say G is an $(k^*, k'; g)$ -**perfect expander** if $|L| = k^*$, $|R| = k'$, the following property holds (where $\Gamma(S)$ denotes the set of neighbors of a set S in G):

1. *Degree:* every vertex $a \in L$, it has constant degree g .
2. *Expansion:* $|\Gamma(S)| \geq (1 - \varepsilon)g|S|$ for every $S \subseteq L$.

Compared to lossless expander, the perfect expander does not have the upper bound on $|S|$ in the expansion property. Therefore, k' has to be much larger than k^* , unlike the case of lossless expander where $k' = O(k)$. Now we show that the density of a perfect expander is low:

Theorem 2. If a bipartite graph is a perfect expander, its density is at most $\frac{g}{1+(1-\varepsilon)g}$; otherwise, the density of the graph is larger than $\frac{g}{1+(1-\varepsilon)g}$.

Proof. We first show that the density of a perfect expander is at most $\frac{g}{1+(1-\varepsilon)g}$. For any subset $L' \subseteq L$, we prove that among all sub-graphs that L' is the left vertex set, the graph induced by $(L', \Gamma(L'))$ has the maximum density.

To see this, suppose $V' = (L', R')$, $R' \neq \Gamma(L')$ has density $\frac{|E'|}{|V'|}$ that is the densest sub-graph with L' as its left vertex set.

Case 1: If there exists a vertex $y \in R'$, $y \notin \Gamma(L')$, then there is no edge between y and L' . We can increase the density by removing y from R' , as $\frac{|E'|}{|V'|-1} > \frac{|E'|}{|V'|}$. This is a contradiction. Therefore, $R' \subseteq \Gamma(L')$.

Case 2: If there exists an element $y \in \Gamma(L')$, $y \notin R'$, let $c \geq 1$ be the number of edges between y and L' , by adding y to R' , the density becomes $\frac{|E'|+c}{|V'|+1} > \frac{|E'|}{|V'|}$. This is a contradiction again and thus $\Gamma(L') \subseteq R'$.

Therefore, we have $\Gamma(L') = R'$ and $V' = (L', \Gamma(L'))$ maximizes the density among all sub-graphs with L' as the left vertex set. Let that sub-graph be G' . By the expansion property of the perfect expander, $\text{den}(G') = \frac{|E'|}{|V'|} \leq \frac{|L'|g}{|L'|+(1-\epsilon)g|L'|} = \frac{g}{1+(1-\epsilon)g}$. Therefore, the maximum density $\text{Den}(G) = \max_{L' \subseteq L} \text{den}(G') \leq \frac{g}{1+(1-\epsilon)g}$.

Next, we show that if a bipartite graph is not a perfect expander, its density is larger than $\frac{g}{1+(1-\epsilon)g}$. Let S^* be the set such that $|\Gamma(S^*)| < (1-\epsilon)g|S^*|$, then the density of the sub-graph $G' = (V' = (S^*, \Gamma(S^*)), E')$ is $\frac{|E'|}{|V'|} > \frac{g|S^*|}{|S^*|+(1-\epsilon)g|S^*|} = \frac{g}{1+(1-\epsilon)g}$, so $\text{Den}(G) \geq \text{den}(G') > \frac{g}{1+(1-\epsilon)g}$.

3.3 Testing Random Lossless Expander

Theorem 2 suggests a way to test whether a random graph is a lossless expander. As discussed in Lemma 3, when $s \geq \log \log k$ the non-expanding probability is negligible. Thus, it suffices to test whether there is a sub-graph of size $s < \log \log k$ that does not expand. In particular, we are trying to distinguish the following two cases:

1. **Yes case:** For $G = (L, R, E)$, $\forall S \subseteq L, |S| \leq \log \log k$, we have $|\Gamma(S)| \geq (1-\epsilon)g|S|$.
2. **No case:** For $G = (L, R, E)$, there exists a subset $S^* \subseteq L, |S^*| \leq \log \log k$, such that $|\Gamma(S^*)| < (1-\epsilon)g|S_0|$.

To distinguish these two cases, we cannot directly apply the densest sub-graph algorithm on the entire bipartite graph, because the expansion property only holds for $|S| \leq \frac{\delta k}{g}$ by Definition 2 of the lossless expander. The densest sub-graph algorithm would return a large sub-graph with $|S| > \frac{\delta k}{g}$ even if it is a lossless expander, as the density of the large sub-graph could be larger than $\frac{g}{1+(1-\epsilon)g}$ by Theorem 2.

Instead, we randomly sample sub-graphs $G^* = ((L', \Gamma(L')), E')$ with $\frac{\delta k}{g}$ vertexes in the left vertex set. If there exists a small non-expanding sub-graph with at most $\log \log k$ vertices on the left, the density of this small sub-graph is larger than $\frac{g}{1+(1-\epsilon)g}$ and the probability of it is in the sub-graph G^* is at least $(\frac{\delta}{g})^{\log \log k}$. Once it is contained in G' , the densest-sub-graph algorithm will output a sub-graph with density larger than $\frac{g}{1+(1-\epsilon)}$. We will sample G^* $\frac{g}{\delta} \log \log k$ times to amplify the probability. The formal algorithm is presented in Algorithm 1.

Theorem 3 (Distinguisher). *Algorithm 1 achieves the following properties:*

1. *If G is a Yes case, then the algorithm will return SUCC with probability 1.*

Algorithm 1. Distinguisher

```

1: Let  $G = (L, R, E)$  be the random bipartite graph.
2:
3: for  $i \in [(\frac{g}{\delta})^{\log \log k}]$  do
4:   Sample a random set  $L'$ , where  $|L'| = \frac{\delta k}{g}$ .
5:   Run densest graph algorithm in [Gol84] on the subgraph induced by  $L'$ :  $G^* = ((L', \Gamma(L')), E')$  to find its densest subgraph.
6:   if  $\text{Den}(G^*) > \frac{g}{1+(1-\varepsilon)g}$  then
7:     return FAIL
8: return SUCC

```

2. If G is a **No case**, then the algorithm will return FAIL with probability at least $1 - \frac{1}{e}$.

Proof. By Theorem 2, if the random graph is in **Yes case**, then the distinguisher will always return SUCC, since for every induced sub-graph G^* , it is a perfect expander. Otherwise, if the random graph contains a subset $S_0 \subseteq L, |S_0| \leq \log \log k$ such that $|\Gamma(S_0)| < (1 - \varepsilon)g|S_0|$, then with probability at least $(\frac{\delta k}{g})^{\log \log k} = (\frac{\delta}{g})^{\log \log k}$, S_0 will be a subset of L' sampled by the algorithm. In this case, L' is not a perfect expander graph and by Theorem 2, $\text{Den}(G^*) > \frac{g}{1+(1-\varepsilon)g}$ and the algorithm will return FAIL. Since we repeat it $\frac{g \log \log n}{\delta}$ times, the probability that we did not successfully sample S_0 is $(1 - (\frac{\delta}{g})^{\log \log k})^{(\frac{g}{\delta})^{\log \log k}}$. By the inequality $(1 - \frac{1}{n})^n \leq \frac{1}{e}$, we have $(1 - (\frac{\delta}{g})^{\log \log k})^{(\frac{g}{\delta})^{\log \log k}} \leq \frac{1}{e}$.

By repeating the distinguisher λ times, we can amplify the detection probability of the No case to $1 - \frac{1}{e^\lambda}$. Finally, we re-sample the random graph until the distinguisher returns SUCC. The successful probability of one sampling is $1 - O(\frac{1}{\text{poly}(k)})$, so the expected number of sampling is a constant. The algorithm runs $\lambda(\frac{g}{\delta})^{\log \log k}$ instances of the densest sub-graph algorithm, and each instance involves a graph with at most $\delta \frac{k}{g}$ vertices and δk edges, so the total running time is $O(\lambda(\frac{g}{\delta})^{\log \log k} k^2 \log^2 k) = O(\lambda \text{polylog}(k) k^2)$. The same algorithm can also apply to the lossless expander graph in [GLS+]. Our sampling algorithm is very efficient in practice. First, it does not involve any cryptographic operations and is done once. Second, $k = \sqrt{N}$ in our protocol of the polynomial commitment in the next section, so the complexity is actually quasi-linear in the size of the zero-knowledge argument instance. Finally, the complexity of the densest sub-graph algorithm in Theorem 1 is for arbitrary graphs. As observed in our experiments, the algorithm is faster on random bipartite graphs and we conjecture that there is a better complexity analysis, which is left as an interesting future work.

4 Our New Zero-Knowledge Argument

In this section, we present the construction of our zero-knowledge argument scheme. Many existing papers show that one can build zero-knowledge

arguments from polynomial commitments [WTS+18, ZXZS20, CHM+20, Set20, GWC19, BFS20, GLS+]. We adopt the same technique and focus on constructing a polynomial commitment because of its simplicity and efficiency, but our approach can be applied directly to the zero-knowledge arguments for R1CS in [BCG20, BCL22] to improve the prover time and the proof size. We start the section by describing the polynomial commitment scheme in [GLS+] based on the tensor IOP protocol in [BCG20] with a proof size of $O(\sqrt{N})$.

4.1 Polynomial Commitment from Tensor Query

In [GLS+], Golovnev et al. observed that a polynomial evaluation can be expressed as a tensor product. Here we only consider multilinear polynomial commitments, which can be used to construct zero-knowledge arguments based on the approaches in [ZGK+17b, WTS+18, XZZ+19, ZXZS20, Set20], but our scheme can be extended to univariate polynomials. In particular, given a multilinear polynomial ϕ , its evaluation on input vector $x_0, x_1, \dots, x_{\log N-1}$ is:

$$\phi(x_0, x_1, \dots, x_{\log N-1}) = \sum_{i_0=0}^1 \sum_{i_1=0}^1 \dots \sum_{i_{\log N-1}=0}^1 w_{i_0 i_1 \dots i_{\log N-1}} x_0^{i_0} x_1^{i_1} \dots x_{\log N-1}^{i_{\log N-1}}.$$

The degree of each variable is either 0 or 1 by the definition of a multilinear polynomial, and thus there are N monomials and coefficients with $\log N$ variables. We let $i = \sum_{j=0}^{\log N-1} 2^j i_j$, that is, $i_0 i_1 \dots i_{\log N-1}$ is the binary representation of number i . We use w to denote the coefficients where $w[i] = w_{i_0 i_1 \dots i_{\log N-1}}$. Similarly we define $X_i = x_0^{i_0} x_1^{i_1} \dots x_{\log N-1}^{i_{\log N-1}}$. Let $k = \sqrt{N}$, $r_0 = \{X_0, X_1, \dots, X_{k-1}\}$, $r_1 = \{X_{0 \times k}, X_{1 \times k}, X_{2 \times k}, \dots, X_{(k-1) \times k}\}$. Then we have $X = r_0 \otimes r_1$. The polynomial evaluation is reduced to a tensor product $\phi(x_0, x_1, \dots, x_{\log N-1}) = \langle w, r_0 \otimes r_1 \rangle$. Using the tensor IOP protocol in [BCG20], one can build a polynomial commitment [GLS+] and we present the protocol in Protocol 2 for completeness. Here we reuse the notation k as it is exactly the message length of the linear code.

As shown in the protocol, to commit to a polynomial, PC.Commit parses the coefficients w as a $k \times k$ matrix and encodes it using the tensor code with dimension 2 as defined in Definition 3. Then the algorithm constructs a Merkle tree commitment for every column $C_2[:, i]$ of the $n \times n$ codeword C_2 , and finally builds another Merkle tree on top of their roots as the final commitment.

To answer the tensor query, there are two checks in the protocol: a proximity check and a consistency check. The proximity check ensures that the matrix in the commitment is indeed close to a codeword of the tensor code. The consistency check ensures that $y = \langle r_0 \otimes r_1, w \rangle$ assuming \mathcal{R} is a commitment of a codeword.

Proximity Check. The proximity check has two steps. First, the verifier sends a random vector γ_0 to the prover, and the prover computes the linear combination of all rows of C_1 and w with γ_0 , as in Step 8 in Protocol 2. Because of the property of a linear code, c_{γ_0} is a codeword with message y_{γ_0} , and this step is referred to as the “fold” operation in [BCG20]. Second, the prover shows that c_{γ_0} is indeed computed from the committed tensor codeword. To do so, the verifier randomly

Protocol 2. Polynomial commitment from [BCG20, GLS+]

Public input: The evaluation point \vec{x} , parsed as a tensor product $r = r_0 \otimes r_1$;

Private input: the polynomial ϕ , the coefficient of ϕ is denoted by w .

Let C be the $[n, k, d]$ -linear code, $E_C : \mathbb{F}^k \rightarrow \mathbb{F}^n$ be the encoding function, $N = k \times k$. If N is not a perfect square, we can pad it to the next perfect square.

We use a python style notation to select the i -th column of a matrix $\text{mat}[:, i]$.

- 1: **function** PC.COMMIT(ϕ)
- 2: Parse w as a $k \times k$ matrix. The prover computes the tensor code encoding C_1, C_2 locally as defined in Definition 3. Here C_1 is a $k \times n$ matrix and C_2 is a $n \times n$ matrix.
- 3: **for** $i \in [n]$ **do**
- 4: Compute the Merkle tree root $\text{Root}_i = \text{Merkle.Commit}(C_2[:, i])$.
- 5: Compute a Merkle tree root $\mathcal{R} = \text{Merkle.Commit}([\text{Root}_0, \dots, \text{Root}_{n-1}])$ and output \mathcal{R} as the commitment.
- 6: **function** PC.PROVE($\phi, \vec{x}, \mathcal{R}$)
- 7: The prover receives a random vector $\gamma_0 \in \mathbb{F}^k$ from the verifier.
- 8: $c_{\gamma_0} = \sum_{i=0}^{k-1} \gamma_0[i] C_1[i], y_{\gamma_0} = \sum_{i=0}^{k-1} \gamma_0[i] w[i]$. ▷ Proximity
- 9: $c_1 = \sum_{i=0}^{k-1} r_0[i] C_1[i], y_1 = \sum_{i=0}^{k-1} r_0[i] w[i]$. ▷ Consistency
- 10: Prover sends $c_1, y_1, c_{\gamma_0}, y_{\gamma_0}$ to the verifier.
- 11: Verifier randomly samples $t \in [n]$ indexes as an array \hat{I} and send it to prover.
- 12: **for** $\text{idx} \in \hat{I}$ **do**
- 13: Prover sends $C_1[:, \text{idx}]$ and the Merkle tree proof of Root_{idx} for $C_2[:, \text{idx}]$ under \mathcal{R} to verifier
- 14: **function** PC.VERIFY EVAL($\pi_{\vec{x}}, \vec{x}, y = \phi(\vec{x}), \mathcal{R}$)
- 15: $\forall \text{idx} \in \hat{I}, c_{\gamma_0}[\text{idx}] == \langle \gamma_0, C_1[:, \text{idx}] \rangle$ and $E_C(y_{\gamma_0}) == c_{\gamma_0}$. ▷ Proximity
- 16: $\forall \text{idx} \in \hat{I}, c_1[\text{idx}] == \langle r_0, C_1[:, \text{idx}] \rangle$ and $E_C(y_1) == c_1$. ▷ Consistency
- 17: $y == \langle r_1, y_1 \rangle$. ▷ Tensor product
- 18: $\forall \text{idx} \in \hat{I}, E_C(C_1[:, \text{idx}])$ is consistent with Root_{idx} , and Root_{idx} 's Merkle tree proof is valid.
- 19: Output **accept** if all conditions above holds. Otherwise output **reject**.

selects t columns and the prover opens them with their Merkle tree proofs. The verifier checks that the inner product between each column and the random vector γ_0 is equal to the corresponding element of c_{γ_0} (Step 15). As shown in [BCG+17, BCG20], if the linear code has a constant relative distance, the committed matrix is close to a tensor codeword with overwhelming probability.

Consistency Check. The consistency check follows exactly the same steps of the proximity check. Instead of using a random vector from the verifier, the linear combination is done with r_0 of the tensor query $r_0 \otimes r_1$. Similarly, c_1 is a codeword of the linear code with message y_1 , and $\phi(x) = \langle y_1, r_1 \rangle$ by the definition of tensor product and polynomial evaluation. As shown in [BCG20], by the check in Step 16, if the committed matrix in \mathcal{R} is close to a tensor codeword, then $y = \phi(x)$ with overwhelming probability. In particular, there exist an extractor to extract a polynomial ϕ from the commitment such that $y = \phi(x)$.

Theorem 4 (Polynomial commitment [BCG20, GLS+]). *Protocol 2 is a polynomial commitment that is correct and sound as defined in Definition 7.*

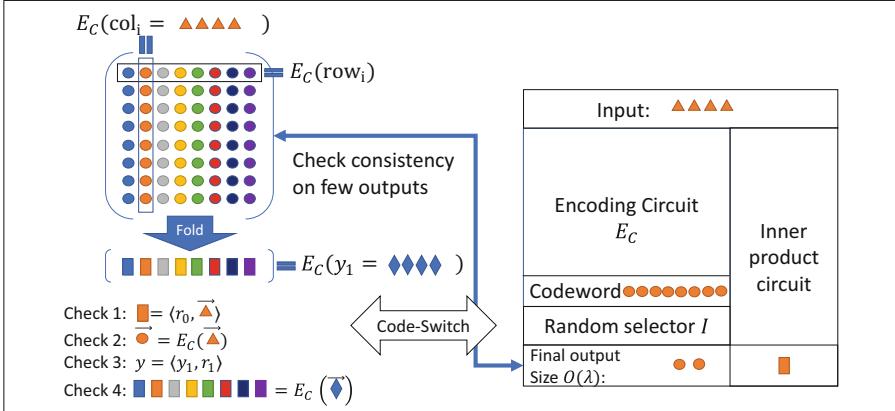


Fig. 2. An illustration of code switching. The circuit on the right for Check 1, 2 and Check 3, 4 are the same.

Efficiency. The prover’s computation is dominated by encoding the tensor code, which takes $O(N)$ time using a linear-time encodable code such as the generalized Spielman code. The proof size is $O(t\sqrt{N})$, as the prover opens t random columns of size \sqrt{N} to the verifier. The verifier time is also $O(t\sqrt{N})$ to check the inner products and to encode t columns.

4.2 Efficient Proof Composition via Code Switching

The proof size of the polynomial commitment in Protocol 2 is $O(\sqrt{N})$ (the complexity hides a security parameter t). There are three steps that incur $O(\sqrt{N})$ proof size in Protocol 2: Step 8, 9, and 13. In this section, we present a new protocol that reduces the proof size to $O(\log^2 N)$ via the technique of proof composition. The idea is to use a second proof system to prove that the checks of these three steps are satisfied, without sending the proofs of these steps to the verifier directly.

To design the second proof system efficiently, our key observation is that the values sent by the prover in these three steps are messages of the linear-time encodable code. That is, y_{γ_0} is the message of c_{γ_0} in Step 8, y_1 is the message of c_1 in Step 9 and $C_1[:, \text{idx}]$ is the message of $C_2[:, \text{idx}]$ for every idx in Step 13. Therefore, the second proof system takes y_{γ_0}, y_1 and $C_1[:, \text{idx}]$ for $\text{idx} \in I$ as the witness, and performs the following computations:

1. It encodes the witness using the encoding circuit of the linear-time encodable code.
2. It outputs a subset of random indices of the codewords chosen by the verifier. By checking whether the values of these indices are consistent with the commitments by the prover via the Merkle tree, it guarantees that the witness is indeed the same as the messages specified above with overwhelming probability because of the minimum distance property of the code.

Protocol 3. Code Switching Statement C_{CS}

Witness: $y_{\gamma_0}, y_1, C_1[:, \text{idx}] \forall \text{idx} \in \hat{I}$ in Protocol 2.

Public input: γ_0, r_0, r_1, y .

Public information: \hat{I} and I chosen by the verifier.

- 1: Encode $c_{\gamma_0} := E_C(y_{\gamma_0})$, $c_1 := E_C(y_1)$.
- 2: **for** $\text{idx} \in \hat{I}$ **do**
- 3: Encode $C_2[:, \text{idx}] := E_C(C_1[:, \text{idx}])$
- 4: **for** $\text{idx} \in \hat{I}$ **do**
- 5: Check if $c_{\gamma_0}[\text{idx}] == \langle \gamma_0, C_1[:, \text{idx}] \rangle$. ▷ Proximity
- 6: Check if $c_1[\text{idx}] == \langle r_0, C_1[:, \text{idx}] \rangle$. ▷ Consistency
- 7: Check if $\langle r_1, y_1 \rangle == y$. ▷ Tensor product
- 8: **for** $0 \leq j < |I|$ **do** ▷ Encoder check
- 9: Output $c_1[I[j]], c_{\gamma_0}[I[j]]$.
- 10: **for** $\text{idx} \in \hat{I}$ **do**
- 11: Output $C_2[I[j], \text{idx}]$

3. Finally, it checks that these messages and their codewords satisfy the conditions in line 15, 16 and 17 of Protocol 2.

The idea is illustrated in Fig. 2, and we formally present the statement of the second proof system in Protocol 3. Note that \hat{I} is the random set chosen by the verifier in Protocol 2, and is only used as a notation for the subscripts in Protocol 3. I is the random set chosen by the verifier for the code switching. In this way, we switch the message encoded using the linear-time encodable code to the witness of the second proof system. In our implementation, we are using an IOP-based zero-knowledge argument with the Reed-Solomon code, we use the name “code switching” as in [RZR20].

We apply any zero-knowledge argument scheme \mathcal{ZK} on the statement and then check the consistency between the output and the Merkle tree commitment \mathcal{R} of the codeword of the linear-time encodable code. We present the new protocol in Protocol 4 and highlight the differences from Protocol 2 in blue. As shown in the protocol, instead of sending $c_1, y_1, c_{\gamma_0}, y_{\gamma_0}$, the prover commits to c_1 and c_{γ_0} in Step 8 and 9. The codeword C_2 was already committed column-wise in \mathcal{R} . The prover then proves the constraints of $c_1, y_1, c_{\gamma_0}, y_{\gamma_0}$ and $C_1[:, \text{idx}]$ using the code switching technique in Step 13. In this way, we are able to reduce the proof size and the verifier time of Protocol 2 to $O(\log^2 N)$.

Theorem 5. *Protocol 4 is a polynomial commitment as defined in Definition 7.*

The proof is presented in the full version of the paper.

Complexity of Protocol 4. The prover time remains $O(N)$. This is because in Step 8 and 9, the prover additionally commits to c_1, c_{γ_0} , which only takes $O(n) = O(\sqrt{N})$ time. In Step 13, the prover invokes another zero-knowledge argument on C_{CS} . C_{CS} consists of $t + 2$ encoding circuits E_C of the linear-time encodable code and $t + 2$ inner products. As the encoding circuit is of size $O(k)$, we will

Protocol 4. Polynomial commitment with code-switching

Public input: The evaluation point \vec{x} , parsed as a tensor product $r = r_0 \otimes r_1$;
Private input: the polynomial ϕ with coefficients w .

- 1: **function** COMMIT(ϕ)
- 2: Parse w as a $k \times k$ matrix. The prover computes the tensor code encoding C_1, C_2 locally as defined in Definition 3.
- 3: **for** $i \in [n]$ **do**
- 4: Compute the Merkle tree root $\text{Root}_i = \text{Merkle.Commit}(C_2[:, i])$.
- 5: Compute a Merkle tree root $\mathcal{R} = \text{Merkle.Commit}([\text{Root}_0, \dots, \text{Root}_{n-1}])$ and output \mathcal{R} as the commitment.
- 6: **function** PROVE($\phi, \vec{x}, \mathcal{R}$)
- 7: The prover receives a random vector $\gamma_0 \in \mathbb{F}^k$ from the verifier.
- 8: $c_1 = \sum_{i=0}^{k-1} r_0[i]C_1[i], y_1 = \sum_{i=0}^{k-1} r_0[i]w[i], \mathcal{R}_{c_1} = \text{Merkle.Commit}(c_1)$
- 9: $c_{\gamma_0} = \sum_{i=0}^{k-1} \gamma_0[i]C_1[i], y_{\gamma_0} = \sum_{i=0}^{k-1} \gamma_0[i]w[i], \mathcal{R}_{\gamma_0} = \text{Merkle.Commit}(c_{\gamma_0})$
- 10: The prover computes the answer $y := \langle y_0, r_1 \rangle$. Prover sends $\mathcal{R}_{c_1}, \mathcal{R}_{\gamma_0}, y$ to the verifier.
- 11: The verifier randomly samples $t \in [n]$ indexes as an array \hat{I} and send it to prover.
- 12: The verifier randomly samples another index set $I \subseteq [k], |I| = t$ and sends it to the prover.
- 13: The prover calls the zero-knowledge argument protocol $\mathcal{ZK.P}$ on C_{CS} . Let π_{z_k} be the proof of the zero-knowledge argument. The prover sends the output of C_{CS} : $C_2[I[j], \text{idx}] \forall \text{idx} \in \hat{I}, c_1[I[j]], c_{\gamma_0}[I[j]]$ and π_{z_k} to the verifier.
- 14: The prover sends the Merkle tree proofs of $C_2[I[j], \text{idx}] \forall \text{idx} \in \hat{I}$ under Root_{idx} .
- 15: The prover sends the Merkle tree proofs of $\text{Root}_{\text{idx}} \forall \text{idx} \in \hat{I}$ under \mathcal{R} .
- 16: The prover sends the Merkle tree proofs of $c_1[I[j]], c_{\gamma_0}[I[j]]$ under $\mathcal{R}_{c_1}, \mathcal{R}_{c_{\gamma_0}}$.
- 17: **function** VERIFY EVAL($\pi_{\vec{x}}, \vec{x}, y = \phi(\vec{x}), \mathcal{R}$)
- 18: The verifier calls the zero-knowledge argument protocol $\mathcal{ZK.V}$ on C_{CS} .
- 19: The verifier checks the Merkle tree proofs of $C_2[I[j], \text{idx}] \forall \text{idx} \in \hat{I}$.
- 20: The verifier checks the Merkle tree proofs of $\text{Root}_{\text{idx}} \forall \text{idx} \in \hat{I}$ using \mathcal{R} .
- 21: The verifier checks the Merkle tree proofs of $c_1[I[j]], c_{\gamma_0}[I[j]]$ using $\mathcal{R}_{c_1}, \mathcal{R}_{c_{\gamma_0}}$.
- 22: Output **accept** if all checks pass. Otherwise output **reject**.

present the analysis in the full version of the paper, and the circuit to compute an inner product is of size $O(k)$, the overall circuit size is $O(t \cdot k)$. By using any zero-knowledge argument scheme with a quasi-linear prover time, such as [ZXZS20], the prover time of this step is $O(t \cdot k \log k)$. Since $k = \sqrt{N}$, the prover time is still $O(N)$ dominated by the encoding and the commitment of the $k \times k$ matrix in COMMIT(). With the code switching technique, the proof size and the verifier time becomes $O(t \log^2 k) = O(t \log^2 N)$.

4.3 Putting Everything Together

In this section, we show how to achieve zero-knowledge on top of our new polynomial commitment in Protocol 4, and sketch how to build a zero-knowledge argument using the polynomial commitment.

Protocol 5. zk-Polynomial commitment

Public input: The evaluation point \vec{x} , parsed as a tensor product $r = r_0 \otimes r_1$;

Private input: the polynomial ϕ with coefficients w .

- 1: **function** ZKCOMMIT(ϕ_w)
- 2: The prover randomly samples $m \in \mathbb{F}^{|w|}$.
- 3: Output $\mathcal{R}_{w+m} = \text{COMMIT}(w+m)$, $\mathcal{R}_m = \text{COMMIT}(m)$.
- 4: **function** ZKPROVE($\phi, \vec{x}, \mathcal{R}$)
- 5: Let ϕ_m be the masking polynomial, ϕ_{m+w} be the masked polynomial.
- 6: Run $\text{Prove}(\phi_{m+w}, \vec{x}, \mathcal{R}_{m+w})$. Let the random index set used during the protocol be \hat{I}_0, I_0 .
- 7: Run $\text{Prove}(\phi_m, \vec{x}, \mathcal{R}_m)$. In this step, the verifier samples the random index set \hat{I}_1, I_1 used during the protocol such that $\hat{I}_0 \cap \hat{I}_1 = \emptyset \wedge I_0 \cap I_1 = \emptyset$.
- 8: **function** ZKVERIFY($\pi_{\vec{x}}^{w+m}, \pi_{\vec{x}}^m, \vec{x}, y_{w+m}, y_m, \mathcal{R}_{w+m}, \mathcal{R}_m$)
- 9: The final polynomial evaluation $\phi(\vec{x})$ should be $y_{w+m} - y_m$.
- 10: Execute $\text{VerifyEval}(\pi_{w+m}, \vec{x}, y_{w+m}, \mathcal{R}_{w+m})$.
- 11: Execute $\text{VerifyEval}(\pi_m, \vec{x}, y_m, \mathcal{R}_m)$.
- 12: Output accept if all checks above passes, otherwise output reject.

Achieving Zero-Knowledge. We apply a masking technique similar to the one in [BCG+17]. The codeword C_2 is masked by a codeword MSK of a masking polynomial with random coefficients m . We use our proof system to prove $y_{w+m} = \langle (w+m), r_0 \otimes r_1 \rangle$ and $y_m = \langle m, r_0 \otimes r_1 \rangle$ simultaneously, and the final answer of the polynomial evaluation is $y = y_{w+m} - y_m$. We present the protocol in Protocol 5.

Theorem 6. *Protocol 5 is a zero-knowledge polynomial commitment scheme by definition 7.*

We present the proof in the full version of the paper.

Zero-Knowledge Argument. Finally, we build our zero-knowledge argument system by combining the multivariate polynomial commitment with the sumcheck protocol as in [Set20, GLS+]. We state the theorem here and refer the readers to [Set20, GLS+] for the construction and the proof.

Theorem 7. *There exists a zero-knowledge argument scheme by definition 5 with $O(N)$ prover time, $O(\log^2 N)$ proof size and $O(N)$ verifier time.*

As we are using the IOP-based scheme in [ZXZS20] as the second zero-knowledge argument in the proof composition, our scheme is an IOP with a linear proof size and logarithmic query complexity. The scheme can be made non-interactive via the Fiat-Shamir [FS86] heuristic, and has plausible post-quantum security. Following the frameworks in [CHM+20, COS20, Set20, GLS+], our scheme can be turned into a holographic proof with a $\text{polylog}(N)$ verifier time in a straight-forward way.

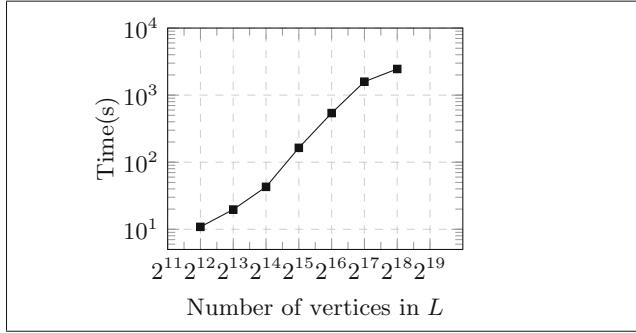


Fig. 3. Running time of our expander testing algorithm.

5 Experiments

We have implemented our scheme, Orion, and we present the evaluations of the system and the comparisons to existing ZKP schemes in this section.

Settings and Parameters. Our polynomial commitment scheme is implemented in C++ with 6000 lines of code. The proof composition uses Virgo in [ZXZS20] and its open-source implementation. We combine the polynomial commitment with a sumcheck protocol to get our zero-knowledge argument following the approach in [Set20] and we implement our own code for this part.

Expander Graph used in Our Implementation. We use a modified version of generalized Spielman code in [GLS+]. The code assigns a random weight to each edge of the expander graph, achieving a better minimum distance. We take a step further and fine-tune the dimensions more aggressively. With our testing algorithm, the failure probability of the expander sampling remains negligible. There are two types of expander graph used in our construction and the parameters are G_1 : $\alpha = 0.33, \delta = 0.6, \epsilon = 0.78, g = 6$; G_2 : $\alpha = 0.337, g = 6, \delta = g, \epsilon = 0.88$.

Parameters of the our Linear Code. With expanders above, the final relative distance is 0.055. We set the security parameter $\lambda = 128$. This leads to opening $t = \frac{128}{\log(1-0.055)} = 1568$ columns and locations in Protocol 4.

Hash Function and Finite Field. We use the SHA-256 hash function implemented by [arm]. We use the extension field of $GF((2^{61} - 1)^2)$ as our underlying field to be compatible with the zero-knowledge argument in [ZXZS20].

Environment and Method. We use an AWS m6i-32xlarge instance with Intel(R) Xeon(R) Platinum 8375C CPU @ 2.90 GHz CPU and 512 GB memory to execute all of our experiments. However, the largest instance in our experiment only utilize 16 GB of memory. All experiments are using a single thread except the expander testing algorithm. For each data point, we run the experiments 10 times and report the average.

5.1 Expander Testing

We first show the performance of our expander testing algorithm in Sect. 3. We implemented the densest sub-graph algorithm in [Gol84], which uses network-

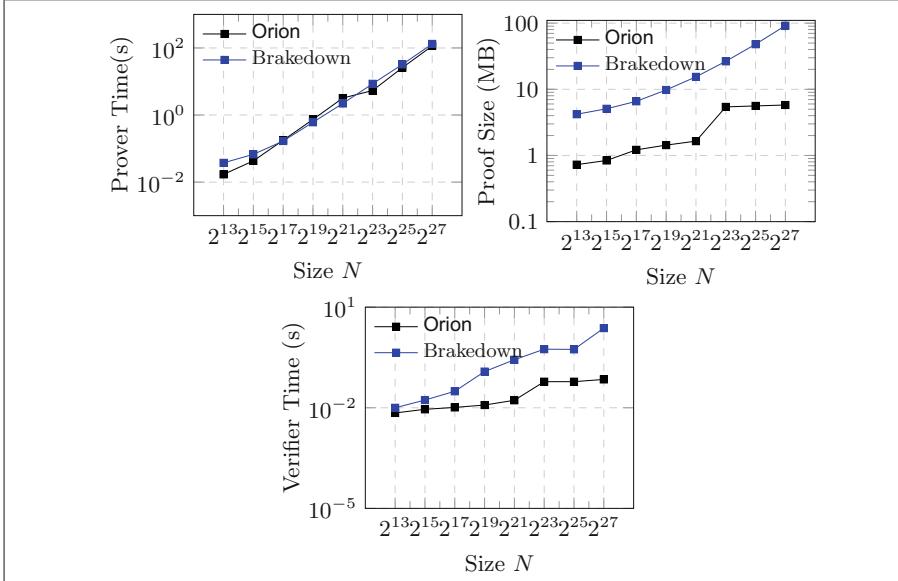


Fig. 4. Performance of polynomial commitments.

flow algorithm as a black-box. In our implementation, we use Dinic's algorithm [Din70], the complexity of which is $O(|V|^2|E|)$ on general graphs. However, on random bipartite graphs, the Dinic's algorithm runs significantly faster and as observed in our experiments, it scales almost linearly in the size of the graph.

Figure 3 shows the running time of the algorithm. We vary the size of left vertex set L in the random bipartite graph from 2^{12} to 2^{18} , and the size of R is set to be $|L| \times \alpha$. The implementation uses multi-threading utilizing all 128 CPU cores. As shown in the figure, it only takes 163 s to test whether a random bipartite graph with $|L| = 2^{15}$ vertices is a lossless expander with a failure probability $\text{negl}(N) = 2^{-128}$. The running time almost grows linearly in $|L|$. As $k = \sqrt{N}$ in our zero-knowledge argument, this is enough for our experiments. As the sampling of the lossless expander is done once, our testing algorithm is very practical.

5.2 Polynomial Commitment

In this section, we report the performance of our polynomial commitment scheme and compare it with the scheme Brakedown in [GLS+], which is the only implemented polynomial commitment scheme with a linear prover time. We use the open-source implementation of Brakedown at [Wla] in the comparison. Our current implementation is for the plain version of the polynomial commitment without zero-knowledge, which is the same as Brakedown.

Figure 4 shows the performance of our polynomial commitment and the polynomial commitment in Brakedown. We vary the size of the polynomials from 2^{15}

to 2^{29} and measure the prover time, the proof size and the verifier time. As shown in the figure, our prover time is even slightly faster than **Brakedown**. It only takes 115 s for a polynomial with 2^{27} coefficients, while it is 132 s in **Brakedown**. This is because we use more aggressive parameters of the expander code, while still achieving 128-bit of security thanks to our expander testing algorithm. Moreover, the additional proof composition in our scheme involves a second zero-knowledge argument on a circuit of size $O(\sqrt{N})$. In our experiments, this extra zero-knowledge argument takes less than 20% of the total prover time, justifying that our code switching technique only introduces a small overhead on the prover time.

Our proof size and verifier time is significantly smaller than **Brakedown**. The proof size is only 6 MBs for a polynomial of size 2^{27} , $16\times$ smaller than **Brakedown**. The verifier time is 70 ms for $N = 2^{27}$, $33\times$ faster than **Brakedown**. The result shows the improvement of the $O(\log^2 N)$ proof size in our scheme.

Note that there is a jump from $N = 2^{21}$ to $N = 2^{23}$ in the proof size and verifier time. This is because in our implementation, instead of directly parsing the coefficients into $\sqrt{N} \times \sqrt{N}$ matrix, we optimize the dimensions for better performance. When $N < 2^{23}$, it is not meaningful to do code-switching on the columns. The prover only does the code-switching on the row (Protocol 4 Step 8 and 9), but opens the columns directly. We observe that this gives the best prover time and the proof size. When $N \geq 2^{23}$, the prover does the code-switching for both the row and the columns (Protocol 4, Step 8–13). Therefore, the proof size and the verifier time have a big increase because of the larger column size and the additional code-switching protocol.

5.3 Zero-Knowledge Arguments

Finally, we present the performance of our zero-knowledge argument scheme for R1CS as a whole in this section. We focus the comparison to existing schemes that work on R1CS and have transparent setup and plausible post-quantum security. They include **Brakedown** [GLS+], **Aurora** [BSCR+19] and **Ligero** [AHIV17]. We use the implementation of **Brakedown** at [Wla], and the open-source code of **Ligero** and **Aurora** at [aur] in the experiments.

We randomly generate the R1CS instances and vary the number of constraints from 2^{15} to 2^{20} . As shown in Fig. 5, **Orion** has the fastest prover among all schemes. It only takes 3.09 s to generate the proof for $N = 2^{20}$. This is slightly faster than **Brakedown** for the same reason as explained in Sect. 5.2. It is $20\times$ faster than **Ligero** and $142\times$ faster than **Aurora** because of the linear prover time and the simplified reduction via polynomial commitments.

The proof size of **Orion** is significantly smaller than **Brakedown** and **Ligero**. It is only 1.5 MB for $N = 2^{20}$, $6.5\times$ smaller than **Brakedown** and $12.5\times$ smaller than **Ligero**. The proof size is even comparable to **Aurora**, which has $O(\log^2 N)$ proof size and uses the Reed-Solomon code with a much better minimum distance than our linear code. The result justifies the improvement of our code switching.

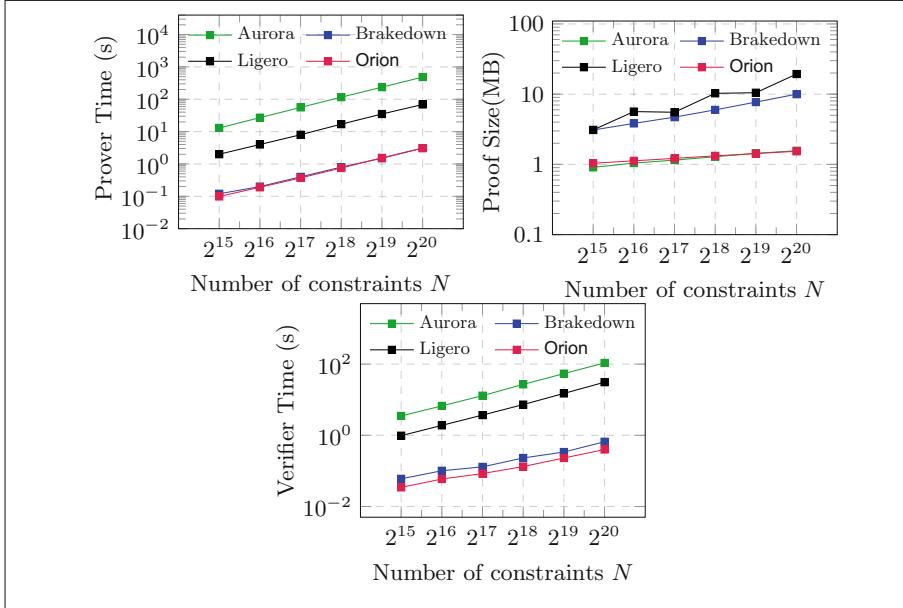


Fig. 5. Performance of zero-knowledge arguments on R1CS.

The verifier time of all schemes grow linearly with N and the comparisons are similar to the prover time. One can reduce the verifier time to sublinear in the holographic setting using the techniques in [CHM+20, COS20, Set20].

Other Related Schemes. There are several other existing transparent zero-knowledge argument schemes. Hyrax [WTS+18], Virgo [ZXZS20] and Virgo++ [ZLW+21] work on layered arithmetic circuits and STARK [BSBHR19] works on an algebraic intermediate representation that is close to a RAM program. It is hard to compare directly to R1CS, but we expect our prover time to be faster than these systems for similar computations based on the results shown in prior papers [ZXZS20, ZLW+21]. Spartan and schemes in [SL20] are using the same framework of polynomial commitment and sumcheck as in our scheme. However, they are based on discrete-log and bilinear pairing and thus are not post-quantum secure. As shown in [GLS+], their prover time is slower than Brakedown while the proof size is better (tens of KBs). Finally, Bulletproofs [BBB+18] and Supersonic [BFS20] are also based on discrete-log and group of unknown order. Their prover time is orders of magnitude slower than schemes mentioned above, while providing the smallest proof size (1–2 KBs) because of the underlying cryptographic techniques.

Acknowledgements. We thank Yuval Ishai for helpful discussions and valuable feedback on the paper. The material is supported by DARPA under Contract No. HR001120C0087, the NSF award #2144625 and the 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, NSF or CLTC.

References

- [AHIV17] Ames, S., Hazay, C., Ishai, Y., Venkitasubramaniam, M.: Lightweight sub-linear arguments without a trusted setup. In: CCS, Ligero (2017)
- [arm] armfazh. flo-shani-aesni. <https://github.com/armfazh/flo-shani-aesni>
- [aur] libIOP. <https://github.com/scipr-lab/libiop>
- [BBB+18] Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: short proofs for confidential transactions and more. In: IEEE S&P (2018)
- [BBC+18] Baum, C., Bootle, J., Cerulli, A., Del Pino, R., Groth, J., Lyubashevsky, V.: Sub-linear lattice-based zero-knowledge arguments for arithmetic circuits. In: CRYPTO (2018)
- [BCG+14] Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., Virza, M.: Decentralized anonymous payments from bitcoin. In: IEEE S&P, Zerocash (2014)
- [BCG+17] 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
- [BCG20] Bootle, J., Chiesa, A., Groth, J.: Linear-time arguments with sublinear verification from tensor codes. In: TCC (2020)
- [BCL22] Bootle, J., Chiesa, A., Liu, S.: Zero-knowledge IOPs with linear-time prover and polylogarithmic-time verifier. In: Dunkelman, O., Dziembowski, S. (eds.) Advances in Cryptology - EUROCRYPT 2022. EUROCRYPT 2022. Lecture Notes in Computer Science, vol. 13276, pp. 275–304. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07085-3_10
- [BDFG20] Boneh, D., Drake, J., Fisch, B., Gabizon, A.: Halo infinite: Recursive zk-SNARKs from any additive polynomial commitment scheme. Cryptology ePrint Archive, Report 2020/1536 (2020)
- [BFH+20] Bhaduria, R., Fang, Z., Hazay, C., Venkitasubramaniam, M., Xie, T., Zhang, Y.: Ligero++: a new optimized sublinear IOP. In: CCS (2020)
- [BFR+13] Braun, B., Feldman, A.J., Ren, Z., Setty, S.T.V., Blumberg, A.J., Walfish, M.: Verifying computations with state. In: SOSP (2013)
- [BFS20] Bünz, B., Fisch, B., Szepieniec, A.: Transparent SNARKs from DARK compilers. In: Eurocrypt (2020)
- [BLNS20] Bootle, J., Lyubashevsky, V., Nguyen, N.K., Seiler, G.: A non-PCP approach to succinct quantum-safe zero-knowledge. In: CRYPTO (2020)
- [BMRS21] Baum, C., Malozemoff, A.J., Rosen, M., Scholl, P.: Mac’n’cheese: zero-knowledge proofs for arithmetic circuits with nested disjunctions. In: CRYPTO (2021)
- [BSBHR19] Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable zero knowledge with no trusted setup. In: CRYPTO (2019)

- [BSCG+13] Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., Virza, M.: Verifying program executions succinctly and in zero knowledge. In: CRYPTO, SNARKs for C (2013)
- [BSCR+19] Ben-Sasson, E., Chiesa, A., Riabzev, M., Spooner, N., Virza, M., Ward, N.P.: Aurora: transparent succinct arguments for R1CS. In: Eurocrypt (2019)
- [BSCTV14] Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Scalable zero knowledge via cycles of elliptic curves. In: CRYPTO (2014)
- [CDG+17] Chase, M., et al.: Post-quantum zero-knowledge and signatures from symmetric-key primitives. In: CCS (2017)
- [CFH+15] Costello, C., Zahur, S.: Versatile verifiable computation. In: IEEE S&P, Geppetto (2015)
- [CHM+20] Chiesa, A., Yuncong, H., Maller, M., Mishra, P., Vesely, N., Ward, N.: Preprocessing zkSNARKs with universal and updatable SRS. In: Eurocrypt, Marlin (2020)
- [COS20] Chiesa, A., Ojha, D., Spooner, N.: Post-quantum and transparent recursive proofs from holography. In: Eurocrypt, Fractal (2020)
- [CRVW02] Capalbo, M., Reingold, O., Vadhan, S., Wigderson, A.: Randomness conductors and constant-degree lossless expanders. In: STOC (2002)
- [CS07] Czumaj, A., Sohler, C.: Testing expansion in bounded-degree graphs. In: IEEE FOCS (2007)
- [DI14] Druk, E., Ishai, Y.: Linear-time encodable codes meeting the gilbert-varshamov bound and their cryptographic applications. In: ITCS (2014)
- [Din70] Dinic, E.A.: Algorithm for solution of a problem of maximum flow in networks with power estimation. In: Soviet Math. Doklady (1970)
- [DIO21] Dittmer, S., Ishai, Y., Ostrovsky, R.: Line-point zero knowledge and its applications. In: ITC (2021)
- [ESLL19] Esgin, M.F., Steinfeld, R., Liu, J.K., Liu, D.: Lattice-based zero-knowledge proofs: new techniques for shorter and faster constructions and applications. In: CRYPTO (2019)
- [FDNZ21] Fang, Z., Darais, D., Near, J., Zhang, Y.: Zero knowledge static program analysis. In: CCS (2021)
- [FFG+16] Fiore, D., Fournet, C., Ghosh, E., Kohlweiss, M., Ohrimenko, O., Parno, B.: Hash first, argue later: adaptive verifiable computations on outsourced data. In: CCS (2016)
- [FQZ+21] Feng, B., Qin, L., Zhang, Z., Ding, Y., Chu, S.: ZEN: efficient zero-knowledge proofs for neural networks. Cryptology ePrint Archive, Report 2021/087 (2021)
- [FS86] Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: CRYPTO (1986)
- [GAZ+22] Grubbs, P., Arun, A., Bonneau, J., Walfish, M.: Zero-knowledge middleboxes. In: USENIX Security, Ye Zhang (2022)
- [Gil52] Gilbert, E.N.: A comparison of signalling alphabets. Bell Syst. Tech. J. **31**(3), 504–522 (1952)
- [GKM+18] Groth, J., Kohlweiss, M., Maller, M., Meiklejohn, S., Miers, I.: Updatable and universal common reference strings with applications to zk-SNARKs. In: CRYPTO (2018)
- [GKR08] Goldwasser, S., Kalai, Y.T., Rothblum, G.: Delegating computation: interactive proofs for muggles. In: STOC (2008)

- [GLS+] Golovnev, A., Lee, J., Setty, S., Thaler, J., Wahby, R.S.: Brakedown: linear-time and post-quantum snarks for r1cs. Cryptology ePrint Archive (2021). <https://ia.cr/2021/1043>
- [GMO16] Giacomelli, I., Madsen, J., Orlandi, C.: Faster zero-knowledge for boolean circuits. In: USENIX Security, ZKBoo (2016)
- [GMR89] Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM J. Comput. **18**(1), 186–208 (1989)
- [Gol84] Goldberg, A.V.: Finding a maximum density subgraph. University of California Berkeley (1984)
- [GR11] Goldreich, O., Ron, D.: On testing expansion in bounded-degree graphs. In: Goldreich, O. (ed.) Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation. LNCS, vol. 6650, pp. 68–75. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22670-0_9
- [GWC19] Gabizon, A., Williamson, Z.J., Ciobotaru, O.: Plonk: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953 (2019)
- [HLW06] Hoory, S., Linial, N., Wigderson, A.: Expander graphs and their applications. Bull. Amer. Math. Soc. **43**(4), 439–561 (2006)
- [ISW21] Ishai, Y., Su, H., Wu, D.J.: Shorter and faster post-quantum designated-verifier zkSNARKs from lattices. In: CCS (2021)
- [Kil92] Kilian, J.: A note on efficient zero-knowledge proofs and arguments (extended abstract). In: STOC (1992)
- [KKW18] Katz, J., Kolesnikov, V., Wang, X.: Improved non-interactive zero knowledge with applications to post-quantum signatures. In: CCS (2018)
- [KPPS20] Kosba, A.E., Papadopoulos, D., Papamanthou, C., Song, D.: MIRAGE: succinct arguments for randomized algorithms with applications to universal zk-SNARKs. In: USENIX Security (2020)
- [KS16] Khot, S., Saket, R.: Hardness of bipartite expansion. In: ESA (2016)
- [LKKO20] Lee, S., Ko, H., Kim, J., Oh, H.: vCNN: Verifiable convolutional neural network based on zk-SNARKs. Cryptology ePrint Archive, Report 2020/584 (2020)
- [LXZ21] Liu, T., Xie, X., Zhang, Y.: zkCNN: zero knowledge proofs for convolutional neural network predictions and accuracy. In: CCS (2021)
- [MBKM19] Maller, M., Bowe, S., Kohlweiss, M., Meiklejohn, S.: Sonic: zero-knowledge snarks from linear-size universal and updatable structured reference strings. In: CCS (2019)
- [Mic00] Micali, S.: Computationally sound proofs. SIAM J. Comput. **30**(4), 1253–1298 (2000)
- [Mie09] Mie, T.: Short PCPPs verifiable in polylogarithmic time with $O(1)$ queries. Ann. Math. Artif. Intell. **56**(3), 313–338 (2009)
- [NS07] Nachmias, A., Shapira, A.: Testing the expansion of a graph. Electr. Colloquium Comput. Complex. (ECCC) **14**, 01 (2007)
- [PHGR13] Parno, B., Howell, J., Gentry, C., Raykova, M.: Nearly practical verifiable computation. In: IEEE S&P, Pinocchio (2013)
- [Pip] Pippenger, N.: On the evaluation of powers and related problems. In: SFCS, IEEE Computer Society (1976)
- [RZR20] Ron-Zewi, N., Rothblum, R.D.: Local proofs approaching the witness length. In: FOCS (2020)
- [Set20] Setty, S.: Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In: CRYPTO (2020)

- [SL20] Setty, S., Lee, J.: Quarks: quadruple-efficient transparent zkSNARKs. Cryptology ePrint Archive, Report 2020/1275 (2020)
- [Spi96] Spielman, D.A.: Linear-time encodable and decodable error-correcting codes. IEEE Trans. Inf. Theor. **42**(6), 1723–1731 (1996)
- [SZT02] Song, D., Zuckerman, D., Tygar, J.D.: Expander graphs for digital stream authentication and robust overlay networks. In: S&P, IEEE (2002)
- [Var57] Varshamov, R.R.: Estimate of the number of signals in error correcting codes. Doklady Akad. Nauk, SSSR **117**, 739–741 (1957)
- [Wla] Wahby, R.S.: lcpc authors. lcpc. <https://github.com/conroi/lcpc>
- [WSR+15] Wahby, R.S., Setty, S.T.V., Ren, Z., Blumberg, A.J., Walfish, M.: Efficient RAM and control flow in verifiable outsourced computation. In: NDSS (2015)
- [WTS+18] Wahby, R.S., Tzialla, I., Shelat, A., Thaler, J., Walfish, M.: Doubly-efficient zkSNARKs without trusted setup. In: S&P (2018)
- [WYKW20] Weng, C./. Yang, K., Katz, J., Wang, X.: Wolverine: fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits. In: S&P (2020)
- [WYX+21] Weng, C., Yang, K., Xie, X., Katz, J., Wang, X.: Mystique: efficient conversions for zero-knowledge proofs with applications to machine learning. In: USENIX Security (2021)
- [XZZ+19] Xie, T., Zhang, J., Zhang, Y., Papamanthou, C., Song, D.: Succinct zero-knowledge proofs with optimal prover computation. In: CRYPTO, Libra (2019)
- [YSWW21] Yang, K., Sarkar, P., Weng, C., Wang, X.: Quicksilver: efficient and affordable zero-knowledge proofs for circuits and polynomials over any field. In: CCS (2021)
- [zca] Zcash. <https://z.cash/>
- [ZFZS20] Zhang, J., Fang, Z., Zhang, Y., Song, D.: Zero knowledge proofs for decision tree predictions and accuracy. In: CCS (2020)
- [ZGK+17a] Zhang, Y., Genkin, D., Katz, J., Papadopoulos, D., Papamanthou, C.: vSQL: verifying arbitrary SQL queries over dynamic outsourced databases. In: S&P (2017)
- [ZGK+17b] Zhang, Y., Genkin, D., Katz, J., Papadopoulos, D., Papamanthou, C.: A zero-knowledge version of vSQL. Cryptology ePrint Archive: Report 2017/1146 (2017)
- [ZGK+18] Zhang, Y., Genkin, D., Katz, J., Papadopoulos, D., Papamanthou, C.: vRAM: faster verifiable RAM with program-independent preprocessing. In: S&P (2018)
- [zkr] An incomplete guide to rollups. <https://vitalik.ca/general/2021/01/05/rollup.html>
- [ZLW+21] Zhang, J., et al.: Doubly efficient interactive proofs for general arithmetic circuits with linear prover time. In: CCS (2021)
- [ZXZS20] Zhang, J., Xie, T., Zhang, Y., Song, D.: Transparent polynomial delegation and its applications to zero knowledge proof. In: S&P, IEEE (2020)