

A Crack in the Firmament: Restoring Soundness of the Orion Proof System and More

Thomas den Hollander*

Daniel Slamanig[†]

Abstract

Orion (Xie et al. CRYPTO’22) is a recent plausibly post-quantum zero-knowledge argument system with a linear time prover. It improves over Brakedown (Golovnev et al. ePrint’21 and CRYPTO’23) by reducing proof size and verifier complexity to be polylogarithmic and additionally adds the zero-knowledge property. The argument system is demonstrated to be concretely efficient with a prover time being the fastest among all existing succinct proof systems and a proof size that is an order of magnitude smaller than Brakedown. Since its publication in CRYPTO 2022, two revisions have been made to the zk-SNARK. First, there was an issue with how zero-knowledge was handled. Second, Orion was discovered to be unsound, which was then repaired through the use of a commit-and-prove SNARK as an “outer” SNARK.

As we will show in this paper, unfortunately, Orion in its current revision is still unsound (with and without the zero-knowledge property) and we will demonstrate practical attacks on it. We then show how to repair Orion without additional assumptions, which requires non-trivial fixes when aiming to preserve the linear time prover complexity. The proposed fixes lead to an even improved efficiency, i.e., smaller proof size and verifier time, over the claimed efficiency of the initial version of Orion. Moreover, we provide the first rigorous security proofs and explicitly consider multi-point openings and non-interactivity. While revisiting Orion we make some additional contributions which might be of independent interest, most notable an improved code randomization technique that retains the minimum relative distance.

1 Introduction

Zero-knowledge proof systems are a fascinating concept and were introduced by Goldwasser, Micali, and Rackoff [GMR85] in the 1980s. The basic idea is that a prover who holds a witness to an NP statement can convince any verifier, who only knows the statement, of the validity of the statement without revealing any information about the witness. Besides being an interesting theoretical object of study on its own, it has shown to have numerous practical cryptographic applications, e.g., as a tool to prevent malicious behaviour in multi-party protocols, enabling privacy-friendly authentication or to construct digital signatures. While for many years most of these concepts were only of theoretical interest, nowadays zero-knowledge proofs see widespread adoption. One famous (implicit) long-term use of zero-knowledge proofs are Schnorr signatures [Sch90], which can be seen as a particular example of Sigma protocols made non-interactive via the Fiat-Shamir heuristic [FS87]. However, also their explicit construction has gained in popularity with the increased need for post-quantum replacements for existing signatures. One particularly popular approach is via the multi-party computation in the head (MPCitH) paradigm [Ish+09], enabling schemes that do not require structured hardness assumptions, e.g., Picnic [Cha+17] or FAEST [Bau+23], but meanwhile turned into a popular paradigm also for structured hardness assumptions.

The arguably most quickly developing domain related to zero-knowledge is that of zk-SNARKs (succinct non-interactive arguments of knowledge) [Bit+12]. Those are computationally sound zero-knowledge proofs that are knowledge sound (i.e., if a proof is accepted by a verifier it is guaranteed that the prover knows the witness), non-interactive (i.e., there is only a single message sent from the prover to the verifier) and they are succinct, i.e., the size of their proofs is sublinear in the size of the

*Research Institute CODE, Universität der Bundeswehr München; thomasdh@unibw.de

[†]Research Institute CODE, Universität der Bundeswehr München; daniel.slamanig@unibw.de

witness to be proven. A lot of the developments are spurred by practical applications in blockchains and cryptocurrencies.

There are various techniques to construct zk-SNARKS that come with different trade-offs in proof sizes as well as prover and verifier complexities (cf. [Nit19; Tha22] for a comprehensive overview). In the following we will focus on the construction of SNARKS that achieve linear time prover complexity and square root (or even polylogarithmic) proof complexity. While the first schemes that fall into this category [Boo+17; BCG20; BCL22] seemed to be only of theoretical interest, recent works [Gol+23; XZS22c] have shown how following these ideas can yield to practically efficient schemes. We recall that a popular paradigm to construct SNARKs is to rely on an information theoretic primitive often abstracted in the form of a polynomial IOP (PIOP) [BFS20] which is then combined with a polynomial commitment scheme [KZG10] and made non-interactive using the Fiat-Shamir heuristic. Brakedown [Gol+23] observes that [BCG20] implicitly describe a polynomial commitment scheme with a linear time commitment phase (relying on linear-time encodable codes), which they make explicit and combine them with the PIOP for R1CS from Spartan [Set20] to obtain a SNARK. The so obtained SNARK is not yet zero-knowledge, but can be made so using one layer of recursive proof composition, i.e., using an “outer SNARK” that is zero-knowledge to prove knowledge of a proof of the “inner” SNARK (that does not need to be zero-knowledge).

In this paper we are focusing on Orion [XZS22c], which improves over Brakedown [Gol+23]. Due to the used building blocks Orion like Breakdown is plausibly post-quantum secure and is transparent, i.e., does not require a trusted setup, both being highly desirable properties in practice. Moreover, to the best of our knowledge, under those SNARKs that are plausibly post-quantum, Orion has the fastest prover time among all existing schemes focusing on linear-time provers in the literature.

High-level overview of Orion. The Orion proof system [XZS22c] improves on Brakedown in two ways. First, by improving the linear-time encodable code used, to have a good minimum distance except with negligible probability. Second, by employing a “code-switching” proof composition technique, the proof size and verifier time go from $O(\sqrt{N})$ to $O(\log^2(N))$. Since the verifier work for Brakedown is $O(\sqrt{N})$, it is possible to use an outer zk-SNARK that is quasilinear while maintaining the overall linear-time prover. Specifically, Virgo [Zha+20] is used as the outer zk-SNARK.

To commit, the prover first parses the polynomial coefficients as a $k \times k$ matrix W . The matrix is then encoded by first encoding every row using the linear code E_C . To achieve zero-knowledge, a random row \vec{r}_i is sampled by the prover, which is then added as well as appended to every code-word. This yields matrix C . After that, every column of the resulting matrix is again encoded using E_C , giving matrix D . The prover finally creates a Merkle tree commitment to D to obtain the commitment \mathcal{R} to the polynomial.

In the evaluation procedure, the prover attempts to convince the verifier of a certain evaluation of the committed polynomial.¹ The verifier sends a random vector $\vec{\gamma}$ to the prover, which the prover uses to compute $\vec{y}_{\vec{\gamma}}$, a linear combination of the rows of W . In Brakedown, this row is sent to the verifier, who checks correctness by opening a number of columns at random and verifying that for every selected column index j this linear combination has been computed correctly. Due to the use of linear codes, any inconsistencies will be caught with overwhelming probability. Orion uses the same idea, but performs these checks in the outer zk-SNARK. In particular, the prover computes

$$\vec{c}_{\vec{\gamma}} = \sum_i \tilde{\gamma}_i C_{ij}, \quad \vec{y}_{\vec{\gamma}} = \sum_i \tilde{\gamma}_i W_{ij}, \quad \vec{r}_{\vec{\gamma}} = \sum_i \tilde{\gamma}_i \vec{r}_i,$$

and sends a commitment to $\vec{c}_{\vec{\gamma}}$ to the verifier. At this point the verifier samples a random column set J , after which the prover additionally commits to $\vec{y}_{\vec{\gamma}}$ and $\vec{r}_{\vec{\gamma}}$, as well as $C_{\cdot,j}$ for all columns $j \in J$. For this it uses the outer SNARK, which is in the form of a commit-and-prove SNARK (CP-SNARK), i.e., a part of the witness is hidden in a commitment and the zk-SNARK proves the relation and the consistency of the witness in the commitments. The verifier then provides a row index set I , after which the prover generates the zk-SNARK proof. Here, \mathcal{P} checks that $\langle \vec{\gamma}, \vec{C}_{\cdot,j} \rangle = E_C(\vec{y}_{\vec{\gamma}} + \vec{r}_{\vec{\gamma}}) \parallel \vec{r}_{\vec{\gamma}}$ at columns $j \in J$. This is identical to the check done in Brakedown, modulo the added zero-knowledge. The zk-SNARK also outputs $\vec{c}_{\vec{\gamma},j}$ so that the verifier can check consistency with

¹In [XZS22c], Eval is conceptually split into a prove and a verify phase. We merge these, since \mathcal{V} should only be convinced if the challenges were sampled honestly: in neither the interactive nor the non-interactive setting does Verify suffice as a standalone verification procedure.

the earlier commitment. Lastly, the proof outputs the column encoding using $E_C(D.j)_i$ at indices $i \in I$, with the goal that these can be checked with the commitment to D without having to open the columns entirely. Since we have a constant number of openings, the proof size is asymptotically dominated by the proof size of the outer SNARK.

Issues with the soundness of Orion. Since its publication in CRYPTO 2022 [XZS22c], two revisions have been made to the zk-SNARK in the paper.² First, there was an issue with how zero-knowledge was handled. In short, the polynomial $\phi(\vec{x})$ was masked by sampling a random polynomial $r(\vec{x})$ and proving evaluations of $(\phi + r)(\vec{x})$ as well as $r(\vec{x})$. However, since encoding the polynomials added redundancy, this could not be simulated. This was fixed in a revised version [XZS22a].

Second, due to an issue in the challenge order, Orion was discovered to be unsound. In particular, the row index set I would be provided too early on in the proof generation, causing the prover to be able to cheat. The authors of Hyperplonk [Che+23] are credited with discovering the problem. Indeed the same issue is not present in their polynomial commitment scheme Orion+, which has a much smaller proof size but is not post-quantum secure because of its reliance on pairing friendly elliptic curve groups. The mentioned soundness issue was then repaired by the authors of Orion in another revised version [XZS22b] through the use of a CP-SNARK, which forces the prover to commit to the witness before the index set I is provided by the verifier. This indeed prevents some ways in which the prover could cheat. Unless noted otherwise, when we henceforth talk about Orion we refer to this latest revision [XZS22b].

As we will demonstrate, unfortunately the Orion polynomial commitment in its current form, and by extension the Orion proof system, is still unsound. This is due to a second issue in the challenge-response order. In this paper, we will demonstrate that this issue indeed exists, and that it is practical to forge a proof such that *any* committed polynomial can be opened to *any* desired evaluation. The issue is inherent to the protocol and persists when considering its non-zero-knowledge version. We also provide an implementation of this attack, where we use additional techniques to make it efficient enough to be practical.

The issue lies in the fact that unlike the row set I , the column set J is known to the prover before committing to the CP-SNARK witness, which enables the prover to adaptively choose a witness that will pass the subsequent verification procedure. In fact, this is required for efficiency: if this index set was not known at this point, the prover would need to commit to C in its entirety.

Consequently, when using Virgo as the outer zk-SNARK as Orion does, this would cost $O(N \log(N))$ prover time, meaning that the prover no longer runs in overall linear time and thus invalidating the claims of having a linear time prover.

Resurrecting Orion. To remedy the unsoundness and simultaneously retain succinctness, we find a way to split the column index set J between rounds. This way, we can use one challenge set J that is unknown to the prover when committing to the witness of the CP-SNARK, while using another challenge set \hat{J} such that the prover can commit to $C_{\hat{J}}$ instead of C in its entirety.

However, this is not possible without breaking the zero-knowledge property of the scheme, which is why we also provide a new method of adding zero-knowledge while retaining the linear-time prover. For this we use zero-knowledge codes [BCL22], which we instantiate concretely using masking polynomials of constant degree. This way, we can retain the linear-time prover. In the process we obtain a new polynomial commitment scheme with smaller proof sizes and a more efficient verifier, that also works for polynomials of lower degree (beyond just padding up to a supported degree).

1.1 Summarizing our Contributions

Subsequently, we will briefly summarize the contributions of this work.

- We take another look at the Orion proof system [XZS22b] and show its unsoundness, which we also support by giving explicit and efficient attacks.

²Actually, there were three revisions. The third revision, however, is concerned with the testing whether a random bipartite graph is a lossless expander (required for the generalized Spielman codes) and is not relevant to our treatment of Orion in this paper.

- We show how to overcome these existing soundness issues and thus resurrect Orion. In order to not introduce additional assumptions and to maintain a linear time prover though, this requires fixes that are non-trivial. Notably, with our fixes we achieve a smaller proof size and verifier time while keeping prover time comparable. We also provide rigorous formal proofs for the modified Orion construction to show that our fixes indeed solve the soundness issues. We provide a new simulator and the first rigorous analysis of the zero-knowledge property. This is necessary since the current simulator used in [XZS22b] does not satisfy their own definition, since it requires knowledge of the evaluation point at the time of committing. Lastly, we provide an explicit treatment of how to make the Orion polynomial commitment scheme non-interactive using the Fiat-Shamir transformation [FS87].
- In the course of fixing Orion, we obtain a new polynomial commitment scheme with smaller proof sizes and a more efficient verifier, that also works for polynomials of lower degree and does not require padding up to a supported degree. We extend the polynomial commitment to open at multiple points by taking a random linear combination of all evaluation vectors while retaining zero-knowledge. Moreover, zero-knowledge can be shown even if the evaluation point is not known at the time of committing.
- We provide a general transformation to perfectly zero-knowledge codes [BCL22], which may be of independent interest. In brief, instead of a random vector to mask the witness we use a randomized polynomial. This retains the same relative minimum distance instead of reducing it by half as in the vector masking approach (which is used by Orion), while retaining linear encodability.

2 Preliminaries

For security parameter λ , we write $\text{negl}(\lambda)$ for the negligible function. We also write $a \approx_\lambda b$ as shorthand for $|a - b| \leq \text{negl}(\lambda)$ and $a||b$ for the concatenation of strings a and b .

We write $w(\vec{a}, \vec{b})$ to denote the Hamming weight of vectors \vec{a} and \vec{b} of equal length n , and $\Delta(\vec{a}, \vec{b}) = \{i \in [n] \mid \vec{a}_i \neq \vec{b}_i\}$ as well as $d(\vec{a}, \vec{b}) = |\Delta(\vec{a}, \vec{b})|$. For vector space L , we define $w(\vec{a}, L)$ as $w(\vec{a}, \vec{c})$ for the $c \in L$ that minimizes the weight— $\Delta(\vec{a}, L)$ and $d(\vec{a}, L)$ are defined analogously.

2.1 Randomized Linear Codes

Following [BCL22], we make use of *randomized linear codes*, which are encoding functions $\tilde{E}_C(\vec{m}; \vec{r}) = E_C(\vec{m} || \vec{r})$, where E_C is some linear code $\mathbb{F}^{k_m+k_r} \rightarrow \mathbb{F}^{n_m+n_r}$ and the message $m \in \mathbb{F}^{k_m}$ and randomness \mathbb{F}^{k_r} are logically separated. We simply write $\tilde{E}_C(\vec{m})$ if $\vec{r} \in_R \mathbb{F}^{k_r}$.

Definition 2.1. A randomized linear code \tilde{E}_C is *q-query uniform* if for any message $m \in \mathbb{F}^{k_m}$, any q adaptively chosen locations of $E_C(\tilde{m})$ are distributed uniformly at random.

An example of a q -query uniform linear code is the Reed-Solomon code $\text{RS}_{[n, k_m+q, d]}$ with $n = k_m + q + d - 1$. This follows from the fact that there is exactly one codeword that conforms to any $k_m + q$ evaluations.

2.2 Collision-Resistant Hashing and Merkle Trees

Let us consider a family $\{H_s\}_{s \in I}$ of functions for which there exists a PPT key generation algorithm $s \leftarrow \text{Gen}(1^\lambda)$. Given the security parameter λ , it outputs an index $s \in I$, or equivalently a function $H_s : \{0, 1\}^{\ell(\lambda)} \rightarrow \{0, 1\}^{\ell'(\lambda)}$ with $\ell(\lambda) > \ell'(\lambda)$. We require that there exists a PPT evaluation algorithm that on input $s \in I$ and $x \in \{0, 1\}^{\ell(\lambda)}$ computes $H_s(x)$. The collision-resistance property for such a family states that every PPT algorithm \mathcal{A} that is given $s \leftarrow \text{Gen}(1^\lambda)$, and outputs a pair $x \neq x'$ such that $H_s(x) = H_s(x')$ only succeeds with negligible probability.

The following presentation is based on the one in [Ame+23]. For a collision-resistant hash function³ H , a Merkle tree hash [Mer90] is a succinct commitment to a list $k = 2^d$ of values represented

³We will usually omit the key s for the sake of simplicity.

by a single hash value h (the root hash) such that opening the commitment on any of the k values requires to reveal $O(d)$ hash values.

Let m_1, \dots, m_k be k messages, then a Merkle tree is represented by a binary tree of depth d where the messages m_i are assigned to the leaves of the tree from left to right. The values assigned to the internal nodes of the tree are computed as $H(u||v)$ where u and v are the left and right children of the respective node. Consequently, the root of the tree represents a commitment to messages m_1, \dots, m_k . To open the commitment to a message m_i , the opening proof is represented by all the values assigned to nodes on the path from the root to m_i , and the values assigned to the siblings of all these nodes. Verifying the opening proof for m_i recomputes the entire path from m_i to the root bottom-up and compares the so obtained root hash to the commitment value.

The binding of a Merkle hash tree follows from the collision-resistance of the hash function. Moreover, when replacing the hash function H used to compute the Merkle tree by a random oracle, statistical binding follows from the hardness of finding a collision in this model. Additionally, in the random oracle model setting the leaf for every m_i to $m_i||r_i$ where r_i is chosen uniformly at random from $\{0, 1\}^\lambda$, the Merkle tree hash is also a statistically hiding commitment.

We will use the notation from [Ame+23] and denote committing to messages m_1, \dots, m_k by $h \leftarrow \text{Commit}_M(m_1, \dots, m_k)$ and opening of a message m_i by $(m_i, \text{path}(i)) \leftarrow \text{Open}_M(h, i)$. Moreover, we denote verification of an opening $(m_i, \text{path}(i))$ by $b \leftarrow \text{Verify}_M(h, m_i, \text{path}(i))$.

2.3 (Commit-and-Prove) SNARKs

In the following we will formally introduce SNARKs as well as commit-and-prove SNARKs (CP-SNARKs).

Succinct non-interactive argument of knowledge (SNARK).

Definition 2.2 (SNARK). A succinct non-interactive argument of knowledge (SNARK) for relation generator RGen is a tuple of algorithms $\Pi = (\text{Setup}, \text{Prove}, \text{Verify})$:

- $\text{crs} := (\text{pk}, \text{vk}) \leftarrow \text{Setup}(R, \text{aux}_R)$: given a relation and auxiliary information $(R, \text{aux}_R) \leftarrow \text{RGen}(1^\lambda)$ as input; outputs a common reference string crs consisting of a proving key pk and a verification key vk .
- $\pi \leftarrow \text{Prove}(\text{pk}, R, \text{aux}_R, u, w)$ given a proving key pk , a relation R , auxiliary information aux_R , a statement u and witness w with $(u, w) \in R$; outputs a proof π .
- $b \leftarrow \text{Verify}(\text{vk}, R, \text{aux}_R, u, \pi)$ given a verification key vk , a relation R , auxiliary information aux_R , a statement u and a proof π ; outputs a bit b indicating reject ($b = 0$) or accept ($b = 1$).

A SNARK satisfies completeness, knowledge soundness, and succinctness. If it is a zk-SNARK it in addition satisfies zero-knowledge. We formally introduce the properties below:

Completeness. Π is complete for RGen , if for all λ , all $(R, \text{aux}_R) \leftarrow \text{RGen}(1^\lambda)$ and $(u, w) \in R$

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{vk}) \leftarrow \text{Setup}(R, \text{aux}_R), \pi \leftarrow \text{Prove}(\text{pk}, R, \text{aux}_R, u, w) : \\ \text{Verify}(\text{vk}, R, \text{aux}_R, u, \pi) = 1 \end{array} \right] = 1.$$

Knowledge soundness. Π is knowledge sound if for all λ , all $(R, \text{aux}_R) \leftarrow \text{RGen}(1^\lambda)$ and any PPT adversary \mathcal{A} there exists a PPT extractor \mathcal{E} with full access to \mathcal{A} such that

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{vk}) \leftarrow \text{Setup}(R, \text{aux}_R), ((u, \pi); w) \leftarrow \mathcal{A} || \mathcal{E}(\text{pk}, \text{vk}, R, \text{aux}_R) : \\ (u, w) \notin R \wedge \text{Verify}(\text{vk}, R, \text{aux}_R, u, \pi) = 1 \end{array} \right] \approx_\lambda 0$$

Succinctness. Π is succinct if the size of the proof π is $\text{poly}(\lambda + \log |w|)$ and Verify runs in time $\text{poly}(\lambda + |u| + \log |w|)$.

Zero-knowledge. Π is zero-knowledge (a zk-SNARK) if for all λ , all $(R, \text{aux}_R) \leftarrow \text{RGen}(1^\lambda)$ and $(u, w) \in R$ there exists a PPT simulator $\mathcal{S} = (\mathcal{S}_0, \mathcal{S}_1)$ with $(R, \text{aux}_R, \tau) \leftarrow \mathcal{S}_0(1^\lambda)$ such that for all PPT adversaries \mathcal{A} it holds that

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{vk}) \leftarrow \text{Setup}(R, \text{aux}_R), \pi \leftarrow \text{Prove}(\text{pk}, R, \text{aux}_R, u, w) : \\ \mathcal{A}(R, \text{aux}_R, \text{pk}, \text{vk}, w, \pi) = 1 \end{array} \right] \approx_\lambda$$

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{vk},) \leftarrow \text{Setup}(R, \text{aux}_R), \pi \leftarrow \mathcal{S}_1(\text{pk}, R, \text{aux}_R, \text{u}, \tau) : \\ \mathcal{A}(R, \text{aux}_R, \text{pk}, \text{vk}, \text{w}, \pi) = 1 \end{array} \right].$$

Commit-and-Prove SNARKs. In a commit-and-prove SNARK (CP-SNARK) one wants to prove knowledge of (u, w) such that $(\text{u}, \text{w}) \in R$ holds with respect to a witness $\text{w} = (w, \omega)$ and w is the value committed to in a commitment c_w . We follow the formalization of [CFQ19] Campanelli et al. so that only a part of the witness w , i.e., w , is committed to and w splits over some subdomains. Before we start defining CP-SNARKs, we need to introduce commitment schemes.

Definition 2.3 (Commitment scheme). A commitment scheme is a triple of efficient algorithms $(\text{Setup}, \text{Commit}, \text{Open})$, which are defined as follows:

- $\text{ck} \leftarrow \text{Setup}(1^\lambda)$: takes as input a security parameter λ ; outputs the commitment key ck which implicitly contains a description of the input space \mathcal{D} , commitment space \mathcal{C} and opening space \mathcal{O} .
- $(c, o) \leftarrow \text{Commit}(\text{ck}, m)$: takes as input the commitment key ck and a value m ; outputs a commitment c and opening o .
- $b \leftarrow \text{Open}(\text{ck}, c, m, o)$: takes as input the commitment key ck , a commitment c , value m and opening o ; outputs a bit b indicating reject ($b = 0$) or accept ($b = 1$).

A commitment scheme is required to provide correctness, binding and hiding. We omit a formal definition of correctness as it is straightforward. The remaining properties are defined as follows.

Binding. A commitment scheme is binding, if for all PPT adversaries \mathcal{A}

$$\Pr \left[\begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda), (c^*, m^*, o^*, m'^*, o'^*) \leftarrow \mathcal{A}(\text{ck}), 1 \leftarrow \text{Open}(\text{ck}, c^*, m^*, o^*), \\ 1' \leftarrow \text{Open}(\text{ck}, c^*, m'^*, o'^*) : m^* \neq m'^* \end{array} \right] \approx_\lambda 0.$$

Hiding. A commitment scheme is hiding, if for all PPT adversaries \mathcal{A}

$$\Pr \left[\begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda), (m_0, m_1, \text{state}) \leftarrow \mathcal{A}(\text{ck}), b \xleftarrow{R} \{0, 1\}, \\ (c, o) \leftarrow \text{Commit}(\text{ck}, m_b), b^* \leftarrow \mathcal{A}(\text{ck}, c, \text{state}) : b = b^* \end{array} \right] \approx_\lambda \frac{1}{2}.$$

We now present the formal definition of a CP-SNARK, which is adapted from [CFQ19].

Definition 2.4 (CP-SNARK [CFQ19]). Let RGen be a relation generator for relations R over $\mathcal{D}_u \times \mathcal{D}_w \times \mathcal{D}_\omega$ such that \mathcal{D}_w splits over ℓ arbitrary domains $(\mathcal{D}_1 \times \dots \times \mathcal{D}_\ell)$ for some arity parameter $\ell \geq 1$. Let $\text{Com} = (\text{Setup}, \text{Commit}, \text{Open})$ be a commitment scheme (as per Definition 2.3) whose input space \mathcal{D} is such that $\mathcal{D}_i \subset \mathcal{D}$ for all $i \in [\ell]$. A commit and prove SNARK for Com and RGen is a zkSNARK for a family of relations RGen^{Com} such that:

- every $R^{\text{Com}} \in \text{RGen}^{\text{Com}}$ is represented by a pair (ck, R) where $\text{ck} \in \text{Com.Setup}(1^\lambda)$ and $R \in \text{RGen}$;
- R^{Com} is over pairs (u, w) where the statement is $\text{u} := (\text{u}, (c_j)_{j \in [\ell]}) \in \mathcal{D}_x \times \mathcal{C}^\ell$, the witness is $\text{w} := ((w_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, \omega) \in \mathcal{D}_1 \times \dots \times \mathcal{D}_\ell \times \mathcal{O}^\ell \times \mathcal{D}_\omega$, and the relation R^{Com} holds iff

$$\bigwedge_{j \in [\ell]} \text{Open}(\text{ck}, c_j, w_j, o_j) = 1 \wedge R(\text{u}, (w_j)_{j \in [\ell]}, \omega) = 1.$$

Furthermore, when we say that the CP-SNARK is knowledge-sound for a relation generator RGen , we mean it is a knowledge-sound SNARK for the relation generator RGen^{Com} that runs $\text{ck} \leftarrow \text{Com.Setup}(1^\lambda)$ and $(R, \text{aux}_R) \leftarrow \text{RGen}(1^\lambda)$, and returns $((\text{ck}, R), \text{aux}_R)$.

We denote a CP-SNARK as a triple of algorithms $CP = (\text{Setup}, \text{Prove}, \text{Verify})$:

- $\text{crs} := (\text{pk}, \text{vk}) \leftarrow \text{Setup}(\text{ck}, R, \text{aux}_R)$ generates a common reference string crs consisting of a proving key pk and a verification key vk .
- $\pi \leftarrow \text{Prove}(\text{pk}, R, \text{aux}_R, \text{u}, (c_j)_{j \in [\ell]}, (w_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, \omega)$ outputs the proof.
- $b \leftarrow \text{Verify}(\text{vk}, R, \text{aux}_R, \text{u}, (c_j)_{j \in [\ell]}, \pi)$ outputs a bit b indicating reject ($b = 0$) or accept $b = 1$.

2.4 Polynomial Commitments

We adapt the definition from [Gol+23], and add zero-knowledge. A polynomial commitment scheme for multivariate polynomials is a tuple of four (interactive) algorithms $\text{PC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Eval})$:

- $\text{pp} \leftarrow \text{Gen}(1^\lambda, \mu)$: takes as input μ (the number of variables in a multilinear polynomial); produces public parameters pp .
- $\mathcal{R} \leftarrow \text{Commit}(\text{pp}, \phi; r)$: takes as input a μ -variate polynomial ϕ over a finite field \mathbb{F} and randomization r ; produces a commitment \mathcal{R} . We omit r if it is sampled uniformly at random.
- $b \leftarrow \text{Open}(\text{pp}, \mathcal{R}, \phi, r)$: given randomization r , verifies the opening of commitment \mathcal{R} to the μ -variate polynomial ϕ ; outputs $b \in \{0, 1\}$.
- $b \leftarrow \text{Eval}(\text{pp}, \mathcal{R}, \vec{x}, y, \mu, \phi)$ is an interactive algorithm between a PPT prover \mathcal{P} and verifier \mathcal{V} . Both \mathcal{V} and \mathcal{P} know a commitment \mathcal{R} , the number of variables μ , the evaluation point \vec{x} and claimed evaluation y . \mathcal{P} additionally knows a μ -variate polynomial ϕ . \mathcal{P} attempts to convince \mathcal{V} that $\phi(\vec{x}) = y$. At the end of the protocol, \mathcal{V} outputs $b \in \{0, 1\}$.

Definition 2.5. A tuple of four (interactive) algorithms $(\text{Gen}, \text{Commit}, \text{Open}, \text{Eval})$ is a *zero-knowledge extractable polynomial commitment scheme* for μ -variate polynomials over a finite field \mathbb{F} if the following conditions hold.

Completeness. For any μ -variate polynomial ϕ and all randomness r ,

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{Gen}(1^\lambda, \mu), \mathcal{R} \leftarrow \text{Commit}(\text{pp}, \phi) : \\ \text{Eval}(\text{pp}, \mathcal{R}, \vec{x}, y, \mu, \phi) = 1 \wedge y = \phi(\vec{x}) \end{array} \right] = 1.$$

Binding. For any PPT adversary \mathcal{A} , size parameter $\mu \geq 1$,

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{Gen}(1^\lambda, \mu), (\mathcal{R}, \phi, \phi', r, r') = \mathcal{A}(\text{pp}), \\ b_0 \leftarrow \text{Open}(\text{pp}, \mathcal{R}, \phi, r), b_1 \leftarrow \text{Open}(\text{pp}, \mathcal{R}, \phi', r') : \\ b_0 = b_1 = 1 \wedge \phi \neq \phi' \end{array} \right] \approx_\lambda 0$$

Knowledge Soundness. For any PPT adversary \mathcal{A} , size parameter $\mu \geq 1$, there exists a PPT extractor \mathcal{E} with full access to \mathcal{A} such that:

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{Gen}(1^\lambda, \mu), (\mathcal{R}, \vec{x}, y) \leftarrow \mathcal{A}(\text{pp}), 1 = \langle \mathcal{A}, \mathcal{V}_{\text{Eval}} \rangle(\text{pp}, \mathcal{R}, \vec{x}, y, \mu, \phi), \\ (\phi, r) \leftarrow \mathcal{E}^{\mathcal{A}}(\text{pp}, (\mathcal{R}, \vec{x}, y)) : \phi(\vec{x}) \neq y \vee \text{Open}(\text{pp}, \mathcal{R}, \phi, r) \neq 1 \end{array} \right] \approx_\lambda 0$$

Zero-knowledge. For any μ -variate polynomial ϕ and any PPT adversary \mathcal{A} , there exists a simulator $\mathcal{S} = (\mathcal{S}_0, \mathcal{S}_1)$ for which it holds that

$$\Pr [\text{Real}_{\mathcal{A}, \phi}(\text{pp})] \approx_\lambda \Pr [\text{Ideal}_{\mathcal{A}, \phi, \mathcal{S}^{\mathcal{A}}}(\text{pp})],$$

where the experiments Real and Ideal are defined as follows:

$\text{Real}_{\mathcal{A}, \phi}(\text{pp})$: <ol style="list-style-type: none"> 1. $\mathcal{R} \leftarrow \text{Commit}(\text{pp}, \phi)$ 2. $\vec{x} \leftarrow \mathcal{A}(\mathcal{R}, \text{pp})$ 3. $y \leftarrow \phi(\vec{x})$ 4. $b \leftarrow \langle \mathcal{P}_{\text{Eval}}, \mathcal{A} \rangle(\text{pp}, \mathcal{R}, \vec{x}, y, \mu, \phi)$ 5. Output b 	$\text{Ideal}_{\mathcal{A}, \phi, \mathcal{S}^{\mathcal{A}}}(\text{pp})$: <ol style="list-style-type: none"> 1. $\mathcal{R} \leftarrow \mathcal{S}_0(1^\lambda, \text{pp})$ 2. $\vec{x} \leftarrow \mathcal{A}(\mathcal{R}, \text{pp})$ 3. $y \leftarrow \phi(\vec{x})$ 4. $b \leftarrow \langle \mathcal{S}, \mathcal{A} \rangle(\text{pp}, \mathcal{R}, \vec{x}, y, \mu)$ 5. Output b
--	---

Here $\langle \mathcal{A}, \mathcal{B} \rangle$ indicates an interaction between parties \mathcal{A} and \mathcal{B} , and $\mathcal{P}_{\text{Eval}}$ a prover running the Eval protocol.

3 Orion and its Unsoundness

The full protocol for the Orion polynomial commitment is given in Figure 1, and the CP-SNARK statement is given in Figure 2. A CP-SNARK is used in the construction with the intention of making the prover commit to all values before the verifier provides the challenges. Unfortunately, there exists issue with the challenge order. The row index set I is provided at step 6 in Figure 1, which is after the CP-SNARK commitment at step 4. However, the column challenge set J is provided *before* this commitment. This gives the prover the opportunity to choose the CP-SNARK witness based on J . For the columns of C this is not an issue, as these are already bound by the Merkle commitment \mathcal{R} to D . However, the vectors \vec{y}_γ , \vec{y}_1 , \vec{r}_1 and \vec{r}_1 are all committed to only after the verifier has sent their challenge set, when the prover executes the commitment phase of the CP-SNARK. Although their encodings \vec{c}_γ and \vec{c}_1 are committed to beforehand, we will see that the prover can nonetheless cheat. In fact, it turns out that this enables the prover to open any committed polynomial in \mathcal{R} to any point y . The commitment does not need to be forged specifically with this goal—any pre-existing commitment can be opened at any point \vec{x} to any desired evaluation y .

Suppose the prover has some commitment \mathcal{R} to a polynomial. In steps 1-4 of the Eval phase, \mathcal{P} follows the protocol honestly, except for the fact that \mathcal{P} sends the desired evaluation y in place of the honest evaluation corresponding to $\phi(x)$. At step 5, \mathcal{P} needs to provide the CP-SNARK witness. Here the prover is able to forge a \vec{y}_1 such that for any I sent by the verifier in step 6, the prover will be able to complete the outer SNARK proof. This is sufficient for the verifier to accept this invalid evaluation proof.

To see that this is always possible, first note that the outer SNARK is the only place where \vec{y}_1 is validated. The vectors are also not committed to until the commit phase of the CP-SNARK in Step 5. The CP-SNARK procedure constructs \vec{c}_1 directly from \vec{y}_1 and \vec{r}_1 , so the goal of the prover is to find \vec{y}_1' and \vec{r}_1' such that together with the derived $\vec{c}_1' = \text{E}_C(\vec{y}_1') + \vec{r}_1' \parallel \vec{r}_1'$, the following conditions hold:

1. $\forall j \in J : \vec{c}_{1j}' = \sum_{i \in [k]} \vec{x}_{0,i} C_{ij}$.
2. $y = \sum_{j \in [k]} \vec{x}_{1j} \vec{y}_{1j}'$
3. $\forall j \in J : \mathcal{R}_{\vec{c}_{1j}}$ opens to \vec{c}_{1j}'

The first two conditions are checked in the CP-SNARK, while the last condition is checked directly by the verifier. Since \vec{c}_1 contained in the commitment $\mathcal{R}_{\vec{c}_1}$ has been honestly computed from C by the prover, conditions (1) and (3) can in fact be reduced to the single condition that $\forall j \in J : \vec{c}_{1j}' = \vec{c}_{1j}$, where \vec{c}_{1j}' is the vector derived from \vec{y}_1 and \vec{r}_1 in the outer SNARK. The protocol can be broken as follows: we first select any vector \vec{y}_1 that satisfies condition (2).⁴ Then we need that

$$\begin{aligned} \forall j \in J \cap 1 \dots n : \quad & \vec{r}_{1j}' = \text{E}_C(\vec{y}_1)_j - \text{E}_C(\vec{y}_1')_j + \vec{r}_{1j}, \\ \forall j \in J \cap n + 1 \dots 2n : \quad & \vec{r}_{1j}' = \vec{r}_{1j}. \end{aligned}$$

To retain zero-knowledge, the challenge space is restricted so that these two sets are precisely disjoint. As such, we can just choose any \vec{r}_1 that satisfies these conditions.

One might then assume that the unsoundness is due to zero-knowledge, but in fact it is also possible to break (knowledge) soundness for the non-zero-knowledge version of the protocol. Even if we omit all \vec{r}_i so that $C_i = \text{E}_C(\vec{y}_1)$ and J is just sampled uniformly random from $[n]$, the protocol can still be broken. Take the generator matrix G of E_C . The condition $\vec{c}_{\gamma_j}' = \vec{c}_{\gamma_j}$ for $j \in J$ can be expressed by choosing the matrix $B_{ij} \equiv \mathbb{1}\{J_i = j\} \in \mathbb{F}^{t \times n}$, where J_i denotes the i th element of the set J in some order. B simply selects the t entries in J . Condition (2) can be appended as a column to the resulting matrix to obtain the linear system

$$\begin{bmatrix} \vec{c}_{1(j_1)}' \\ \vec{c}_{1(j_2)}' \\ \dots \\ \vec{c}_{1(j_t)}' \\ y \end{bmatrix} = \begin{bmatrix} BG^T \\ \vec{x}_1 \end{bmatrix} \begin{bmatrix} \vec{y}_{11}' \\ \vec{y}_{12}' \\ \dots \\ \vec{y}_{1k}' \end{bmatrix}. \quad (1)$$

⁴The only situation in which this is not possible, is if we structurally restrict ϕ to have a zero constant term, and we evaluate at $\vec{x} = \vec{0}$. Such an evaluation can only be 0 and as such a verifier cannot be convinced otherwise.

$\text{Commit}(\text{pp}, \phi; r)$: The prover \mathcal{P} performs the following operations, where any randomness is deterministically derived from r :

1. Parse ϕ as a $k \times k$ matrix W of its coefficients.
2. For each $i \in [k]$, sample row $\vec{r}_i \in \mathbb{F}^n$ uniformly at random.
3. Compute $C \in \mathbb{F}^{k \times 2n}$ by encoding every row of W using the encoding function E_C and hiding with \vec{r}_i as follows: $\forall i \in [k] : C_i = E_C(W_i) + \vec{r}_i \parallel \vec{r}_i$.
4. Compute $D \in \mathbb{F}^{n \times 2n}$ by encoding every column of C using E_C .
5. Commit to columns: $\mathcal{R}_j \leftarrow \text{Commit}_M(D_{\cdot j})$.
6. Output final commitment $\mathcal{R} \leftarrow \text{Commit}_M(\mathcal{R}_1, \dots, \mathcal{R}_{2n})$.

$\text{Eval}(\text{pp}, \mathcal{R}, X = \vec{x}_0 \otimes \vec{x}_1, y, \mu, \phi)$: The prover and verifier execute the following protocol:

1. \mathcal{V} sends a uniformly random vector $\vec{\gamma} \in \mathbb{F}^k$ to \mathcal{P} .
2. The prover computes

$$\begin{aligned} \vec{c}_1 &\leftarrow \sum_{i=1}^k \vec{x}_{0i} C_i, & \vec{y}_1 &\leftarrow \sum_{i=1}^k \vec{x}_{0i} W_i, & \vec{r}_1 &\leftarrow \sum_{i=1}^k \vec{x}_{0i} \vec{r}_i, & \mathcal{R}_{\vec{c}_1} &\leftarrow \text{Commit}_M(\vec{c}_1), \\ \vec{c}_\gamma &\leftarrow \sum_{i=1}^k \vec{\gamma}_i C_i, & \vec{y}_\gamma &\leftarrow \sum_{i=1}^k \vec{\gamma}_i W_i, & \vec{r}_\gamma &\leftarrow \sum_{i=1}^k \vec{\gamma}_i \vec{r}_i, & \mathcal{R}_{\vec{c}_\gamma} &\leftarrow \text{Commit}_M(\vec{c}_\gamma). \end{aligned}$$

3. \mathcal{P} computes $y = \sum_{i=1}^k \vec{y}_{1j} \vec{x}_{1j}$ and sends $\mathcal{R}_{\vec{c}_1}, \mathcal{R}_{\vec{c}_\gamma}$ and y to the verifier.
4. \mathcal{V} samples and sends column indices $J \subset [2n]$ uniformly at random, under the condition that

$$|J| = 2t, \quad |J \cap [n]| = t, \quad |J \cap [n+1 \dots 2n]| = t, \quad \neg \exists j : j, j+n \in \hat{J}.$$

5. The prover commits to the witness of the CP-SNARK (Figure 2), consisting of

$$\vec{y}_1, \vec{y}_\gamma, \forall (i,j) \in [k] \times J : C_{ij}, \vec{r}_1, \vec{r}_\gamma.$$

6. The verifier samples row index set $I \subset [n]$, $|I| = t$ and sends it to the prover.
7. The prover computes the CP-SNARK argument π and sends it with

$$\begin{aligned} \forall j \in \hat{J} : & \text{Open}_M(\mathcal{R}_{\vec{c}_1}, j), \text{Open}_M(\mathcal{R}_{\vec{c}_\gamma}, j), \\ \forall j \in \hat{J} : & \text{Open}_M(\mathcal{R}, j), \quad \forall (i,j) \in I \times \hat{J} : \text{Open}_M(\mathcal{R}_j, i). \end{aligned}$$

8. \mathcal{V} checks the proof π using the CP-SNARK verification procedure.
9. \mathcal{V} checks the Merkle tree proofs of D_{ij} in \mathcal{R} for all $(i,j) \in I \times J$.
10. \mathcal{V} checks the Merkle tree proofs of \vec{c}_{1j} and $\vec{c}_{\gamma j}$ in $\mathcal{R}_{\vec{c}_1}$ and $\mathcal{R}_{\vec{c}_\gamma}$.

The verifier accepts if all checks are successful.

Figure 1: The full description of the Orion polynomial commitment protocol.

There are two reasons why there might not exist a solution. First, it is possible that there is no other codeword equal to \vec{c}_1 at J . Second, it may be possible that all candidate codewords have messages that are orthogonal to \vec{x}_1 . To address both these possibilities: consider the vector space $A = \{\vec{m} : (BG^T \vec{m}) = \vec{0}\}$.

Lemma 1. Let $0 < l = k - t$. Then $\dim(A) \geq k - t$.

Proof. Any t entries of a codeword cannot uniquely determine the message, since $t < k$. That means

Committed witness: $\vec{y}_1, \vec{y}_\gamma, \vec{r}_1, \vec{r}_\gamma, \forall j \in J : C_{.j}$.

Public input:

- $\vec{\gamma}, x_0, x_1, y, I, J$
- $\forall j \in J : \vec{c}_{1j}, \vec{c}_{\gamma j}$
- $\forall (i, j) \in I \times J : D_{ij}$

Proof circuit: The CP-SNARK circuit performs the following checks and computations:

- Build encodings of row vector combinations:
 - $\vec{c}_\gamma \leftarrow E_C(\vec{y}_\gamma) + \vec{r}_\gamma \parallel \vec{r}_\gamma$.
 - $\vec{c}_1 \leftarrow E_C(\vec{y}_1) + \vec{r}_1 \parallel \vec{r}_1$.
 - Check for $j \in J$ that $\vec{c}_{\gamma j}$ and \vec{c}_{1j} match the input.
- Check that these are valid linear combinations:
 - Check $\forall j \in J : \vec{c}_{\gamma j} \stackrel{?}{=} \sum_{i \in [k]} \vec{\gamma}_i C_{ij}$.
 - Check $\forall j \in J : \vec{c}_{1j} \stackrel{?}{=} \sum_{i \in [k]} \vec{x}_{0i} C_{ij}$.
- Encode columns and compare:
 - $\forall j \in J : D_{.j} = E_C(C_{.j})$.
 - Check for $(i, j) \in I \times J$ that D_{ij} matches the input.
- Check claimed evaluation:
 - Check $y \stackrel{?}{=} \sum_{j \in [k]} \vec{x}_{1j} \vec{y}_{1j}$.

Figure 2: The procedure for the CP-SNARK. The committed witness contains the values that are part of the CP-SNARK commitment. Any other witness values arise only from the intermediate values due to the arithmetization.

that A , the subspace where entries are 0 at t points $j \in J$, is at least one-dimensional. Let \vec{m}_1 be a non-zero codeword in A . We can define $L = \{i\}$ where i is one index at which $(BG^T \vec{m}_1)_i \neq 0$. If $k = t + 1$, we are done. If $k > t + 1$, the vector space $\{\vec{m} : \forall j \in J : (G^T \vec{m})_j = 0 \wedge \forall i \in L : (G^T \vec{m})_i = 0\}$ does not contain \vec{m}_1 and is also one-dimensional, since $t + 1$ entries do not uniquely determine \vec{m} by assumption. We can keep adding indices for non-zero codewords to L until $|L| = l$. We now have l orthogonal messages m_i for $i \in L$, which form a basis for A . \square

It follows that the first possibility is never applicable if $k > t$. This is a very weak assumption that is always satisfied in practice: if $k \leq t$, the CP-SNARK does more work than it would do by evaluating the polynomial directly. As such, in this situation one should always opt to commit to the polynomial using the CP-SNARK commitment, and evaluate the polynomial in its circuit.

For the second condition, the system is solvable if and only if there exists $\vec{m} \in A$, $\vec{m} \notin \ker(\vec{x}_1)$. Here, $\ker(\vec{x}_1)$ has dimension $k - 1$ by the rank-nullity theorem (or k if $\vec{x}_1 = \vec{0}$, but then $y = 0$ would be the only possible evaluation regardless of ϕ). As k increases, so does $\dim(A) \geq k - t$ and therefore it becomes more likely that there exists such a message. We have the following lemma:

Lemma 2. Let $\vec{x}_1 \in_R \mathbb{F}_p^{1 \times k}$ and $k > t$. Then Eq. (1) has a solution except with negligible probability.

Proof. We know that the solutions to Eq. (1) are a subset of $\vec{y}_1 + \vec{m}$ for $\vec{m} \in A$ and $\vec{y}_1 = \sum_{i \in [k]} \vec{x}_{0i} W_i$ for the original witness. If there exists one message $\vec{m} \in A$ such that $\sum_{j \in [n]} \vec{x}_{1j} \vec{m}_j = \alpha \neq 0$, then $\vec{y}_1 + \left(y - \sum_{j \in [n]} \vec{x}_{1j} \vec{y}_{1j}\right) \alpha^{-1} \vec{m}$ will be a solution to the system. Consider one message $\vec{m}' \in A$ and

index j' where $\vec{m}'_{j'} \neq 0$. Then the probability that there exists no solution is given by

$$\Pr \left[\forall \vec{m} \in A : \langle \vec{x}_{1j}, \vec{m}_j \rangle = 0 \right] \leq \Pr \left[\sum_{j \in [n]} \vec{x}_{1j} \vec{m}'_j = 0 \right] = \Pr \left[\vec{x}_{1j'} = -\vec{m}'_{j'}^{-1} \sum_{j \in [n] \setminus \{j'\}} \vec{x}_{1j} \vec{m}'_j \right] = \frac{1}{|\mathbb{F}|}$$

□

Similarly, the case where $\vec{x}_{1j} = x^j$ for some uniformly random $x \in \mathbb{F}$ has negligible failure probability by the Schwartz-Zippel lemma. Unfortunately it is not possible to give a general success probability without knowing a distribution of the evaluation point \vec{x}_1 . It is likely that the adversary succeeds with high probability for evaluation points that are not adversarially chosen. This is corroborated by our experimental results, which have always yielded a solution for $k > t$.⁵ Note that even a small probability is already sufficient in the non-interactive setting, since the prover can simply try for $\text{poly}(\lambda)$ different challenge sets J .

3.1 Attack Implementation

Since Orion uses a Generalized Spielman code, the encoding algorithm can be expressed much more efficiently than through its generator matrix G , both computationally as well as in terms of memory usage. Fortunately, we do not need to construct G explicitly for the attack either.

The encoding algorithm $E_C^{(d)}(\vec{x})$ consists of three parts. We use d to indicate the depth, since the algorithm is recursive. In the base case, $d = 0$ and we just copy the message directly. Otherwise, $\vec{m}_1 = \vec{x}A^{(d)}$ is computed for some expander graph $A^{(d)}$. We then compute $\vec{c}_1 = E_C^{(d-1)}(\vec{m}_1)$ through a recursive call. Lastly $\vec{c}_2 = \vec{c}_1A^{(d)'} is computed for another expander graph. The final codeword is $\vec{x} \parallel \vec{c}_1 \parallel \vec{c}_2$.$

Instead of computing G from this process through encoding basis vectors for the message space, we apply G^T to B from the right. Note that B has size $t \times n$, much smaller than G . We start by inverting $A^{(d)'}$ to the right-most columns of B . We can view the result as a matrix that maps $\vec{x} \parallel \vec{c}_1$ to the output at the chosen indices. The next step is a recursive call, creating a matrix that takes $\vec{x} \parallel \vec{m}_1$ as a right input. Lastly, we invert $A^{(d)}$ and arrive at the final matrix that maps \vec{x} to the correct evaluation of E_C at the selected indices.

While every lossless expander graph may be sparse, the resulting generator matrix is not. Therefore this method is much more memory efficient: B has only t rows and starts with n columns, and BG^T will have dimension $t \times k$. During this process, the size of B has only shrunk from its original size. We therefore use only $O(|B|)$ memory. Computing is also efficient: inverting expander graphs is as efficient as applying them.

After having computed the resulting matrix BA , a valid codeword can be found by computing the pseudoinverse Q^\dagger of $Q = \begin{bmatrix} BA \\ \vec{x}_1 \end{bmatrix}$. We then find a solution to the system by setting

$$\vec{y}'_\gamma = Q^\dagger \begin{bmatrix} \vec{c}_{\gamma(J_1)} & \vec{c}_{\gamma(J_2)} & \dots & \vec{c}_{\gamma(J_t)} & y \end{bmatrix}^T.$$

4 Restoring Soundness of Orion

The order in which challenges and commitments must be sent is naturally constrained by the protocol. First of all, we do not want to commit to C entirely and so the column set J needs to be known at this time. Second, to check that C is consistent with the commitment to D , we should sample I only after committing to these columns. Third, the vector \vec{y}'_γ needs to be compared to C , but here the column vector J must be sampled afterwards to ensure \vec{y}'_γ cannot be chosen to open correctly at only these points.

⁵An implementation of this attack can be found at <https://github.com/ThomasdenH/orion-attack>.

4.1 One More Round of Commitments

One way to reconcile this, is by committing to \vec{y}_γ before \hat{I} is sampled, and opening this commitment in the SNARK. This can be generically achieved by creating a Merkle commitment and opening it inside the CP-SNARK. This is undesirable however, since hash-computations are expensive to perform inside the SNARK circuit. Alternatively, we could switch to a CP-SNARK that supports multiple rounds of commitments: the vectors $\vec{y}_1, \vec{y}_\gamma, \vec{r}_1$ and \vec{r}_γ would be committed to first, and when J is sampled, then the prover commits to $\vec{c}_{1,j}$ as the second part of the witness.

Done naively, this would incur significant overhead on the CP-SNARK, since the committed witness now uses two polynomial commitments of half the size instead of one. If the verifier time and proof size incurred by this step are polylogarithmic, these would almost double, with the opening of polynomial commitments representing the majority of the total work. This naive approach can be improved upon through the use of batching. One option is to use FRI [Ben+18], which can be efficiently batched with minimal overhead.

4.2 Sampling a Different Column Set

In this paper we will instead focus on regaining soundness of the original construction, where we do not want to put additional requirements on the outer SNARK. As we will see, we can fix the protocol without requiring that the CP-SNARK supports efficient batching or even multiple rounds of commitments at all. The only cost associated with our approach is that the verifier samples a new column set \hat{J} . Additionally, we have to switch the zero-knowledge scheme, but this will actually improve verifier time and proof size, while keeping the prover time comparable.

If it were possible to sample such a second, independent column set \hat{J} , this set could be used for the columns $C_{j'}$, while the other parts of the protocol could depend on a different J that is only provided by \mathcal{V} after \mathcal{P} has already committed to the CP-SNARK witness. This way, we can have the column set \hat{J} available when we commit to selected columns of C , while not allowing the prover to select \vec{y}_γ such that it opens correctly precisely at J , which would be sampled later.

Note that this approach has several advantages: first, we do not increase the witness size, just the statement size of the CP-SNARK. We have one additional challenge, which does not need to be sent with the proof. We don't create any new Merkle commitments or open existing commitments at additional points at all.

Unfortunately, this approach still has an issue. To retain zero-knowledge, the second column set J cannot be sampled independently from \hat{J} . It is necessary that $\neg \exists j : j, j+n \in J \cup \hat{J}$ or part of the witness could be learned. But this requirement implies that before committing to the CP-SNARK witness, \mathcal{P} knows t entries that will never be queried by the verifier in J .

4.3 Using Zero-knowledge Linear Codes

To remedy this issue, the zero-knowledge scheme is changed to use a randomized polynomial instead of a vector to mask the witness. This has the desirable property that entries are distributed uniformly at random as long as the degree is at least $|J \cup \hat{J}| - 1$. This construction is an instance of a q -query uniform linear code, which in turn is a specific case of a zero-knowledge linear code. In fact, this general transformation of any code to a perfectly zero-knowledge code may be of independent interest. It retains the same relative minimum distance instead of reducing it by half as in the vector masking approach, while retaining linear encodability. Crucial is the observation that we require $E_{C,ZK}$ to be linearly encodable in k , not in q .

Lemma 3. Let E_C be a linear code $[n, k, d]$ and let $q < 2n - 4d$. Then $E_{C,ZK}^q$, defined for $(\vec{y}, \vec{r}) \in \mathbb{F}^k \times \mathbb{F}^q$ as

$$\forall i \in [2n] : E_{C,ZK}^q(\vec{y}; \vec{r})_i = (E_C(\vec{y}) \parallel E_C(\vec{y}))_i + \sum_{j=0}^{q-1} \vec{r}_j i^j,$$

is the randomized linear code $[2n, k + q, 2d]$. Furthermore, if E_C is encodable in $O(n)$ operations, then $E_{C,ZK}^q$ is encodable in $O(q(n + q))$ operations.

Proof. The fact that $E_{C,ZK}^q$ is a linear code follows directly from the definition.

To see that the minimum distance is $2d$, note that the right hand side is just $\text{RS}_{[2n, q, 2n-q+1]}$. For $(\vec{y} \parallel \vec{r}) \neq \vec{0}$ it holds that that

$$\begin{aligned} w(E_{C, \text{ZK}}(\vec{y}; \vec{r})) &\geq \left| 2w(E_C(\vec{y})) - \text{RS}_{[2n, q, 2n-q+1]}(\vec{r}) \right| \\ &\geq \min(2d, 2n - q + 1, 2n - 2d - q - 1) \geq 2d \end{aligned}$$

by the restriction that $q < 2n - 4d$.

To see that it is q -query uniform, note that for any message \vec{y} and any q queried locations (i_j) and corresponding values (e_j) with $j \in [q]$, there is exactly one \vec{r} such that $E_{C, \text{ZK}}(\vec{y}; \vec{r})_{i_j} = e_j$ for all $j \in [q]$, since q points uniquely determine all coefficients of a degree $q - 1$ polynomial. This means that for \vec{r} uniformly random, any q queries are also uniformly random.

Lastly, encoding a codeword $E_{C, \text{ZK}}(\vec{y}; \vec{r})$ takes $O(k)$ operations for encoding E_C , and then $q(n + q + d - 1) = O(q(n + q))$ operations for computing the polynomial. \square

Using $E_{C, \text{ZK}}$ as the row encoding means that the challenge space can be simplified: \mathcal{V} can simply select \hat{J} and J independently uniformly at random. Additionally, for the same message space, our randomized linear code has twice the minimum distance as the code $E_C(\vec{y}) + \vec{r} \parallel \vec{r}$. This means that we get the same security level for smaller proof sizes.

Interestingly, only the prover every needs to run $E_{C, \text{ZK}}$. This means that it might be possible to encode randomness in the (secret) generator matrix instead of having a random contribution, potentially shrinking codeword size and thereby proof size. One potential issue is that it must be efficiently verifiable that the code has a good minimum relative distance. It may be possible that an existing randomly sampled linear-time encodable code already has the desired properties. This leaves an interesting research question with potential to optimize this zk-SNARK further.

4.4 A New Polynomial Commitment Scheme

Our new polynomial commitment scheme can be found in Figure 3, and the corresponding CP-SNARK relation in Figure 4.

Theorem 4. *Let E_C be a linear code $[n, k, d]$ with $\left(1 - \frac{d}{3n}\right)^t \leq \text{negl}(\lambda)$ and let $E_{C, \text{ZK}}^{2t}$ be a $2t$ -query uniform linear code $[\tilde{n}, \tilde{k}, \tilde{d}]$ with $\left(1 - \frac{\tilde{d}}{3\tilde{n}}\right)^t \leq \text{negl}(\lambda)$. Then the protocol in Figure 3, using a complete, knowledge-sound, zero-knowledge CP-SNARK to instantiate the outer SNARK relation in Figure 4, is a zero-knowledge polynomial commitment scheme.*

We provide a proof in Section 5.

In addition to our new polynomial commitment scheme, we also make an improvement to the simulator. The one currently used in [XZS22b] does not follow its own definition, since it requires knowledge of the evaluation point at the time of committing. This is likely an issue with the current simulator and proof as opposed to the scheme itself—it seems possible to describe a correct simulator. This difference has practical relevance too, however. The evaluation point may realistically not be known beforehand, even by the verifier, and in this case we still want zero-knowledge. The definition therefore also describes a more realistic scenario than what is currently achieved. We therefore describe a new simulator that can open \mathcal{R} to any evaluation point, anywhere.

4.5 Efficiency

It is clear that we achieve the same asymptotics as [XZS22b]. However, it turns out that our solution actually improves verifier efficiency as well as proof size, and is only slightly worse in terms of prover time.

First consider the impact on the CP-SNARK circuit. There, the only penalty in performance of our approach is the $2t \cdot 2n = O(\sqrt{N})$ multiplications to evaluate the randomized polynomial, instead of n additions. Since we already perform t encodings in $O(\sqrt{N})$ operations (with a higher constant), the circuit does not grow significantly. Since the prover time is $O(\sqrt{N})$ and the proof size and verifier time are $O(\log^2(N))$, we only expect the prover to incur a noticeable (but still minor)

$\text{Commit}(\text{pp}, \phi; r)$: The prover \mathcal{P} performs the following operations, where any randomness is deterministically derived from r :

1. Parse ϕ as a $k \times k$ matrix W of its coefficients.
2. Compute $C \in \mathbb{F}^{k \times \tilde{n}}$ by encoding every row of W using the encoding function $E_{C, \text{ZK}}^{2t}(W_i; \vec{r}_i)$ for $\vec{r}_i \in_R \mathbb{F}^q$.
3. Compute $D \in \mathbb{F}^{n \times \tilde{n}}$ by encoding every column of C using E_C .
4. Commit to columns: $\mathcal{R}_j \leftarrow \text{Commit}_M(D, j)$.
5. Output final commitment $\mathcal{R} \leftarrow \text{Commit}_M(\mathcal{R}_1, \dots, \mathcal{R}_{\tilde{n}})$.

$\text{Eval}(\text{pp}, \mathcal{R}, X = \vec{x}_0 \otimes \vec{x}_1, y, \mu, \phi)$: The prover and verifier execute the following protocol:

1. \mathcal{V} sends a uniformly random vector $\vec{\gamma} \in \mathbb{F}^k$ to \mathcal{P} .
2. The prover computes

$$\begin{aligned} \vec{c}_1 &\leftarrow \sum_{i=1}^k \vec{x}_{0i} C_i, & \vec{y}_1 &\leftarrow \sum_{i=1}^k \vec{x}_{0i} W_i, & \vec{r}_1 &\leftarrow \sum_{i=1}^k \vec{x}_{0i} \vec{r}_i, & \mathcal{R}_{\vec{c}_1} &\leftarrow \text{Commit}_M(\vec{c}_1), \\ \vec{c}_{\vec{\gamma}} &\leftarrow \sum_{i=1}^k \vec{\gamma}_i C_i, & \vec{y}_{\vec{\gamma}} &\leftarrow \sum_{i=1}^k \vec{\gamma}_i W_i, & \vec{r}_{\vec{\gamma}} &\leftarrow \sum_{i=1}^k \vec{\gamma}_i \vec{r}_i, & \mathcal{R}_{\vec{c}_{\vec{\gamma}}} &\leftarrow \text{Commit}_M(\vec{c}_{\vec{\gamma}}). \end{aligned}$$

3. \mathcal{P} computes $y = \sum_{i=1}^k \vec{y}_{1j} \vec{x}_{1j}$ and sends $\mathcal{R}_{\vec{c}_1}$, $\mathcal{R}_{\vec{c}_{\vec{\gamma}}}$ and y to the verifier.
4. \mathcal{V} samples and sends column indices $\hat{J} \subset [\tilde{n}]$ uniformly at random.
5. The prover commits to the witness of the CP-SNARK (Figure 4), consisting of

$$\vec{y}_1, \vec{y}_{\vec{\gamma}}, \vec{r}_1, \vec{r}_{\vec{\gamma}}, \forall_{j \in \hat{J}}: C_{\cdot j}.$$

6. The verifier samples and sends row index set $I \subset [n]$, $|I| = t$, as well as a new column index set $J \subset \tilde{n}$, $|J| = t$, uniformly at random.
7. The prover computes the CP-SNARK argument π and sends

$$\begin{aligned} \forall j \in J: & \text{Open}_M(\mathcal{R}_{\vec{c}_1}, j), \text{Open}_M(\mathcal{R}_{\vec{c}_{\vec{\gamma}}}, j), \\ \forall j \in J: & \text{Open}_M(\mathcal{R}, j), \quad \forall (i, j) \in I \times J: \text{Open}_M(\mathcal{R}, i). \end{aligned}$$

8. \mathcal{V} checks the proof π using the CP-SNARK verification procedure.
9. \mathcal{V} checks the Merkle tree proofs of D_{ij} in \mathcal{R} for all $(i, j) \in I \times J$.
10. \mathcal{V} checks the Merkle tree proofs of \vec{c}_{1j} and $\vec{c}_{\vec{\gamma}j}$ in $\mathcal{R}_{\vec{c}_1}$ and $\mathcal{R}_{\vec{c}_{\vec{\gamma}}}$.

The verifier accepts if all checks are successful.

Figure 3: The protocol for our fixed version of Orion.

overhead from the new CP-SNARK circuit, and the verifier time and proof size to not be noticeably impacted.

Outside of the CP-SNARK, the prover additionally needs to do N constant size polynomial evaluations. We do not expect this to make the prover significantly less efficient either, since it already needs to perform $5k$ encodings of E_C in $O(N)$ time, as well as compute $2N$ hash functions for the Merkle commitment.

As we have seen before, for the verifier time and proof size, the impact of the new CP-SNARK circuit is tiny. At the same time, the minimum distance of the used code becomes larger due to the new randomization technique. This means that the number of opened columns and rows, t , more than halves, meaning that there are much fewer elements in the proof. The verifier time depends on

Committed witness: $\vec{y}_1, \vec{y}_\gamma, \vec{r}_1, \vec{r}_\gamma, \forall j \in \hat{J} : C_{.j}$.

Public input:

- $\gamma, x_0, x_1, y, \hat{J}, I, J$.
- $\forall j \in J : \vec{c}_{1j}, \vec{c}_{\gamma j}$.
- $\forall (i, j) \in I \times J : D_{ij}$.

Proof circuit: The CP-SNARK circuit performs the following checks and computations:

- **Build encodings of row vector combinations:**
 - $\vec{c}_\gamma \leftarrow E_{C,ZK}(\vec{y}_\gamma; \vec{r}_\gamma)$
 - $\vec{c}_1 \leftarrow E_{C,ZK}(\vec{y}_1; \vec{r}_1)$
 - Check for $j \in J$ that $\vec{c}_{\gamma j}$ and \vec{c}_{1j} match the input.
- Check that these are valid linear combinations:
 - Check $\forall j \in \hat{J} : \vec{c}_{\gamma j} \stackrel{?}{=} \sum_{i \in [k]} \tilde{\gamma}_i C_{ij}$.
 - Check $\forall j \in \hat{J} : \vec{c}_{1j} \stackrel{?}{=} \sum_{i \in [k]} \tilde{x}_{0i} C_{ij}$.
- Encode columns:
 - $\forall j \in \hat{J} : D_{.j} = E_C(C_{.j})$.
 - Check for $(i, j) \in I \times J$ that D_{ij} matches the input.
- Check claimed evaluation:
 - Check $y \stackrel{?}{=} \sum_{j \in [k]} \tilde{x}_{1j} \tilde{y}_{1j}$.

Figure 4: The procedure for the fixed CP-SNARK. The committed witness contains the values that are part of the CP-SNARK commitment. Any other witness values arise only from the intermediate values due to the arithmetization.

verifying π and is otherwise linear in t , so it is expected to decrease significantly as well.

4.6 Supporting Smaller Polynomials

Orion requires that $4k = n \geq 2t$ to be able to satisfy the constraints on the challenge space. For the relative minimum distance $\delta = 0.055$, we set $t = 4795$ in order that $\left(1 - \frac{\delta}{3}\right)^t \leq 2^{-\lambda}$.⁶ Note that our code can be used for $n > \frac{t}{1-2\delta}$, or $k > \frac{t}{4(1-2\delta)} \approx 1347$ for the generalized Spielman code used in Orion. This means that we directly support polynomials of less than half the size.

As we have noted in Section 3 already, for polynomials where $k < t$ it would make sense to use the CP-SNARK directly. There are however some additional ways of making the proof system more efficient even for small k . For one, as was already noted in [XZS22b], at some point it is faster to avoid using codes for the columns and instead open these directly. This removes the cost of encoding and shrinks the amount of data to be committed to. For each column, k elements are now opened instead of t , and so this approach is always better when $k < t$. Second, for the rows it is possible to switch to Reed-Solomon codes in place of $E_{C,ZK}$ in situations where k is small enough that the $O(n \log(n))$ cost is acceptable. This will allow for much lower values of t due to the larger relative minimum distance of the code. For example, for $k = 2000$ we can pick the code $RS_{[7486, 2000+2t, 4847]}$ with $t = 365$ and for $k = 1000$ the code $RS_{[9896, 1000+2t, 8585]}$ with $t = 156$ suffices. Since proof size and verifier time

⁶Although Orion selects $t = 1568$, checking the linear combination gives rise to a success probability $\left(1 - \frac{\delta}{3 \cdot 2^n}\right)^t$, implying a security level of $\lambda \approx 14$. For $\lambda = 128$, a value of $t = 9635$ would have been required instead.

depend mostly on t , this will give a small and efficient proof system, although it comes at the cost of prover speed.

4.7 Open at Multiple Points

Since one of our goals is to support zero-knowledge, we cannot run the Eval procedure multiple times, even though this would be desirable in many protocols. Fortunately, the scheme is easily extended to open at multiple points by taking a random linear combination of all evaluation vectors. This protocol can be found in Figure 5.

Lemma 5. The protocol in Figure 5 is complete, binding and zero-knowledge. In addition it is knowledge sound, i.e. for any adversary \mathcal{A} , size parameter $\mu \geq 1$, there exists a PPT extractor \mathcal{E} with full access to \mathcal{A} such that:

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{Gen}(1^\lambda, \mu), (\text{pp}, \mathcal{R}, x_{0,0}^{\vec{r}} \otimes x_{0,1}^{\vec{r}}, \dots, \vec{x}y_0, \dots, y_{m-1}^{\vec{r}}) \leftarrow \mathcal{A}(\text{pp},) \\ 1 = \langle \mathcal{A}, \mathcal{V}_{\text{Eval}} \rangle(\text{pp}, \mathcal{R}, \vec{x}, y, \mu, \phi), (\phi, r) \leftarrow \mathcal{E}^{\mathcal{A}}(\text{pp}, (\mathcal{R}, \vec{x}, y)) : \\ \phi(\vec{x}_1) \neq y_1 \vee \dots \vee \phi(x_{m-1}^{\vec{r}}) \neq y_{m-1} \vee \text{Open}(\text{pp}, \mathcal{R}, \phi, r) \neq 1 \end{array} \right] \approx_\lambda 0.$$

We provide a proof in Section 5.

$\text{Eval}_{\text{multi}}(\text{pp}, \mathcal{R}, (X_0 = x_{1,0}^{\vec{r}} \otimes x_{1,1}^{\vec{r}}, \dots, X_{m-1} = x_{m-1,0}^{\vec{r}} \otimes x_{m-1,1}^{\vec{r}}), (y_0, \dots, y_{m-1}), \mu, \phi)$: The prover and verifier execute the following protocol:

- \mathcal{V} selects a random vector $\vec{\alpha} \in \mathbb{F}^m$ and sends it to the prover.
- \mathcal{P} and \mathcal{V} compute $X' = \sum_{i=0}^{m-1} \vec{\alpha}_i X_i$.
- \mathcal{P} and \mathcal{V} compute $y' = \sum_{i=0}^{m-1} \vec{\alpha}_i y_i$.
- \mathcal{P} and \mathcal{V} run the protocol $\text{Eval}(\text{pp}, \mathcal{R}, X', y', \mu, \phi)$.

Figure 5: A procedure to open the polynomial commitment at multiple points.

4.8 On (Weak) Simulation Extratability

Simulation Extractability (SE) is a strengthened knowledge-soundness property, which guarantees that the witness is extractable even if the adversary has access to the simulated proofs. This is a property often desirable in applications and intensively studied recently generically for PIOPs [Gan+22; KPT23; Fao+23] as well as for various concrete proof systems (e.g, [DG23; Lib24]). At least two variants can be distinguished: strong SE, where the adversary wins by providing a proof that was not simulated before, or weak SE, where the adversary wins by providing a proof for statement that was not simulated before.

Unfortunately, Polynomial commitment schemes based on [BCG20], such as those used in Orion, Brakedown [Gol+23] and Ligerio [Ame+17; Ame+23] cannot be (strong) SE. The consistency check relies on the fact that a linear combination of vectors far from all codewords, will also be far from any codeword. This means that the closest codeword is correct. However, it cannot be enforced that this vector (i.e. \vec{c}_γ in Orion) is equal to a codeword. In fact, if the prover changes just a single index they have a decent probability of passing all verifier checks. This means the prover can do something that resembles rerandomization and thus breaking strong SE. Note that the message corresponding to the closest codeword cannot be changed and therefore it is still possible that Orion satisfies weak SE. In fact, this seems likely if the outer CP-SNARK is at least weak SE: if we include the transcript in the outer CP-SNARK statement, we will never be able to re-use a simulated outer proof. Therefore \mathcal{P} is always forced to go through the verification procedure. Proving that the Orion PC is weak SE, or more importantly that the Orion zk-SNARK is weak SE, is an interesting question for future work.

5 Security

Like Orion we use the result that, if a single row of matrix C is far away from a codeword, then a random linear combination will also be far away with high probability. Specifically, we use the following claim (which is a reformulation of [Ame+23, Lemma A.1]), where d denotes the minimum distance of the code E_C .

Lemma 6. [Gol+23, Claim 1] Let $\hat{U} = (\hat{u}_1, \dots, \hat{u}_m)^T \in \mathbb{F}^{m \times n}$ and for each $i \in [m]$ let \vec{c}_i be the closest codeword in E_C to \hat{u}_i . Let E with $|E| \leq d/3$ be a subset of the columns $j \in [n]$ of \hat{U} on which there is even one row $i \in [m]$ such that $\hat{u}_{ij} \neq \vec{c}_{ij}$. With probability at least $1 - (|E| + 1)/|\mathbb{F}|$ over the choice of $\vec{\gamma} \in \mathbb{F}^m$, $\sum_{i=1}^m \vec{\gamma}_i \hat{u}_i$ has distance at least $|E|$ from any codeword in E_C .

Proof of Theorem 4. Below we give our prove that the construction in Figure 3 is a polynomial commitment scheme.

Proof. (Theorem 4)

Completeness. The protocol in Figure 3 is perfectly complete by construction. The CP-SNARK is complete by assumption so we conclude that the polynomial commitment scheme as a whole is complete.

Binding. Define the algorithm $\text{Open}(\text{pp}, \mathcal{R}, \phi, r)$ as the canonical one, i.e., by simply running $\mathcal{R}' \leftarrow \text{Commit}(\text{pp}, \phi; r)$ with the provided randomness r . Now the algorithm runs deterministically and we can return $b = 1$ if $\mathcal{R}' = \mathcal{R}$ and $b = 0$ otherwise.

Suppose there exists a PPT adversary \mathcal{A} that can find ϕ, r, ϕ', r' such that $\phi \neq \phi'$ and both $\text{Open}(\text{pp}, \mathcal{R}, \phi, r) = 1$ and $\text{Open}(\text{pp}, \mathcal{R}, \phi', r') = 1$. Since the coefficients of ϕ and ϕ' differ in at least one location, this means that their encoding is also distinct and therefore there must exist a node in the Merkle tree for \mathcal{R} for which we have found a hash collision. By assumption this can only occur with negligible probability.

Knowledge Soundness. For knowledge soundness, we construct an extractor \mathcal{E} that is able to extract a polynomial ϕ and randomness r such that $\phi(\vec{x}) = y \wedge \text{Open}(\text{pp}, \mathcal{R}, \phi, r) = 1$. We show knowledge-soundness in four steps:

- (1) We show that $\vec{y}_\gamma, \vec{y}_1, \vec{r}_\gamma$ and \vec{r}_1 are uniquely determined by $\mathcal{R}_{\vec{c}_\gamma}$ and $\mathcal{R}_{\vec{c}_1}$.
- (2) We show that C is uniquely determined by \mathcal{R} .
- (3) We show that W and thereby ϕ is uniquely determined by \mathcal{R} .
- (4) We show that $y = \phi(\vec{x})$.

(1): The adversary sends the commitment to the CP-SNARK at Step 5 of the protocol. In the next step, we sample I and J uniformly at random, after which \mathcal{P} provides proof π as well as openings \vec{c}_{γ_j} for t indices $j \in J$. We can run the outer SNARK extractor to obtain the CP-SNARK witness, which includes $\vec{y}_\gamma, \vec{r}_\gamma$.

First, note that \vec{y}_γ and \vec{r}_γ are uniquely defined after committing to the CP-SNARK witness by its binding property. Note also that each index in $\mathcal{R}_{\vec{c}_\gamma}$ can be opened to at most one point by the adversary with non-negligible probability, by the binding property of the Merkle commitment. We can therefore define \vec{c}_γ as this unique opening of $\mathcal{R}_{\vec{c}_\gamma}$, where we set $\vec{c}_{\gamma_j} = 0$ at any index j that the adversary is unable to open.

Suppose then that these \vec{y}_γ and \vec{r}_γ are not uniquely determined by $\mathcal{R}_{\vec{c}_\gamma}$, i.e. that the adversary is able to successfully complete the protocol after supplying $\mathcal{R}_{\vec{c}_\gamma}$ with non-negligible probability for two distinct $(\vec{y}_\gamma, \vec{r}_\gamma), (\vec{y}'_\gamma, \vec{r}'_\gamma)$ obtained by the CP-SNARK extractor. By the minimum distance of $E_{C, \text{ZK}}$, for at least one of the two, $\Delta(\vec{c}_\gamma, E_{C, \text{ZK}}(\vec{y}_\gamma; \vec{r}_\gamma)) \geq \frac{d}{2}$. The probability of opening all entries correctly is $\left(1 - \frac{d}{2\tilde{n}}\right)^t < \left(1 - \frac{d}{3\tilde{n}}\right)^t \leq \text{negl}(\lambda)$, where d and \tilde{n} are the minimum distance and codeword length of $E_{C, \text{ZK}}$, respectively. This is a contradiction and so \vec{y}_γ and \vec{r}_γ are uniquely determined by $\mathcal{R}_{\vec{c}_\gamma}$. The exact same logic implies that \vec{y}_1 and \vec{r}_1 are uniquely determined by $\mathcal{R}_{\vec{c}_1}$.

(2): For the columns of C we use a similar reasoning, but the protocol does not treat C in exactly the same way and so we consider it separately. Note that every index of a column committed to in \mathcal{R}_j

can be opened to at most one point with non-negligible probability, again by the binding property of the Merkle commitment. Define $D_{\cdot j}$ as the values to which the verifier can open, setting $D_{ij} = 0$ at any $i \in [n]$ where the verifier cannot open with non-negligible probability.

Suppose the CP-SNARK extractor obtains $C_{\cdot j}$ such that $\Delta(E_C(C_{\cdot j}), D_{\cdot j}) \geq \frac{d}{2}$. Then the probability of success for the adversary (C is fixed after the CP-SNARK commitment) of answering for challenge I is at most $\left(1 - \frac{d}{2n}\right)^t \leq \text{negl}(\lambda)$. Since \mathcal{R} can only be opened to at most one \mathcal{R}_j at any j with non-negligible probability, this means that \mathcal{R} uniquely determines $C_{\cdot j}$ for any column j where the adversary is able to succeed with non-negligible probability. For the remaining column indices, we can just set $C_{\cdot j} = 0$.

(3): We have seen that C is uniquely determined by the polynomial commitment \mathcal{R} . Suppose that there exists a row index i such that $C_{i\cdot}$ is at least $\frac{d}{3}$ away from any codeword in $E_{C,ZK}$. By Lemma 6, it holds that $\sum_{i \in [n]} \tilde{\gamma}_i C_{i\cdot}$ is also at least $\frac{d}{3}$ away from any codeword. However, we sample a uniformly random column set \hat{J} (after $\sum_{i \in [n]} \tilde{\gamma}_i C_{i\cdot}$ is uniquely determined) and test equality to codeword $E_{C,ZK}(\vec{y}_{\tilde{\gamma}}; \vec{r}_{\tilde{\gamma}})$ at $j \in \hat{J}$. The adversary's success probability is at most $\left(1 - \frac{d}{3n}\right)^t \leq \text{negl}(\lambda)$. This is a contradiction and so every row of the C has distance less than $\frac{d}{3}$ to some codeword. Call this codeword \tilde{C}_i . We define $W_i \parallel \vec{r}_i$ as the value such that $E_{C,ZK}(W_i; \vec{r}_i) = \tilde{C}_i$.

Finally, suppose that $E_{C,ZK}(\vec{y}_{\tilde{\gamma}}; \vec{r}_{\tilde{\gamma}}) \neq \sum_{i \in [n]} \tilde{\gamma}_i \tilde{C}_i$. Then their Hamming distance is at least d and so the adversary has at most probability $\left(1 - \frac{d}{n}\right)^t \leq \text{negl}(\lambda)$ of success. We therefore have

$$E_{C,ZK}(\vec{y}_{\tilde{\gamma}}; \vec{r}_{\tilde{\gamma}}) = \sum_{i \in [n]} \tilde{\gamma}_i \tilde{C}_i = \tilde{\gamma}_i \sum_{i \in [n]} E_{C,ZK}(W_i; \vec{r}_i) = E_{C,ZK}\left(\sum_{i \in [n]} \tilde{\gamma}_i W_i; \sum_{i \in [n]} \tilde{\gamma}_i \vec{r}_i\right)$$

and so $\vec{y}_{\tilde{\gamma}} = \sum_{i \in [n]} \tilde{\gamma}_i W_i$. Since this argument only follows from the fact that \hat{J} is selected uniformly at random and the linearity of $E_{C,ZK}$, by the same argument $\vec{y}_1 = \sum_{i \in [n]} \tilde{\gamma}_i W_i$. It follows from soundness of the CP-SNARK that $y = \vec{x}_0 W \vec{x}_1$.

Since we do not assume that E_C or $E_{C,ZK}$ are efficiently decodable, we still have to show how to recover W and the randomness. Note that by running the CP-SNARK extractor, we obtain $\vec{y}_{\tilde{\gamma}}$ for a random vector $\tilde{\gamma}$. We can therefore rewind the adversary $O(k)$ times to the start of Step 1 of Eval, until we have asked for k linearly independent vectors $\tilde{\gamma}$ as well as $\vec{y}_{\tilde{\gamma}} = \sum_{i \in [k]} \tilde{\gamma}_i W_i$ and $\vec{r}_{\tilde{\gamma}} = \sum_{i \in [k]} \tilde{\gamma}_i \vec{r}_i$. We can solve this linear system to obtain W and all \vec{r}_i . Note that there is no exponential blowup due to the protocol being multi-round: it suffices to complete the protocol once for k linearly independent challenges $\tilde{\gamma}$, which is sampled at the first step of Eval. Furthermore, we expect k challenges to be linearly independent except with negligible probability. To see this, suppose that in round r , we have sampled challenges $\vec{\gamma}_1, \dots, \vec{\gamma}_{r-1}$. A new uniformly random challenge can be decomposed as $\sum_{i=1}^{r-1} \alpha_i \vec{\gamma}_i + \sum_{i=r}^k \alpha_i \vec{\beta}_i$, where all $\vec{\beta}_i$ are orthogonal to each other and to all $\vec{\gamma}_i$, and α_i are sampled uniformly at random. The resulting vector will be orthogonal unless $\alpha_i = 0$ for all $r \leq i \leq k$. Since $\alpha_i \in \mathbb{F}$, this occurs with negligible probability.

Zero-Knowledge: The simulator can be found in Figure 6. First note that \mathcal{S} is complete by construction: $\hat{J} \cup J$ determines at most $2t$ coefficients of $\vec{r}_{\tilde{\gamma}}$ and \vec{r}_1 , and so it is always possible to find such polynomials. To see that no PPT adversary \mathcal{A} can distinguish between two executions, we show that the communication is distributed identically in both $\text{Real}_{\mathcal{A},\phi}$ and $\text{Ideal}_{\mathcal{A},\phi,\mathcal{S},A}$.

The elements sent by the prover are:

- Merkle commitments $\mathcal{R}, \mathcal{R}_{\vec{c}_{\tilde{\gamma}}}, \mathcal{R}_{\vec{c}_1}$ and its opening proofs. The unopened values remain statistically hidden.
- Evaluation y , which is identical in both games.
- The CP-SNARK proof π and corresponding commitment, which we compute honestly and which are in both cases indistinguishable from the output of the CP-SNARK simulator by zero-knowledge of the CP-SNARK.
- The openings D_{ij} for $I \times \hat{J}$, and $\vec{c}_{1j}, \vec{c}_{\tilde{\gamma}j}$ for $j \in \hat{J}$. Note that these are all linear combinations of columns $j \in \hat{J}$ of C . Since each row of C is a codeword of $E_{C,ZK}$ and we open $|\hat{J}| \leq 2t$ entries, these openings are distributed uniformly at random in both cases.

□

Function $S_0^A(1^\lambda, \text{pp})$: To simulate a commitment, S_0 simply runs $\text{Commit}(\text{pp}, \phi)$ for $\phi \equiv 0$.

Function $S_1^A(\text{pp}, \mathcal{R}, \vec{x}, y, \mu)$: The simulator engages in a protocol with the verifier \mathcal{V} .

- \mathcal{V} sends a uniformly random vector $\vec{\gamma} \in \mathbb{F}^k$.
- S learns \hat{J} and J from the verifier's random tape.
- S sets $\vec{y}_\gamma = \vec{0}$ and chooses \vec{r}_γ uniformly at random under the restriction that

$$\forall j \in \hat{J} \cup J: \quad \mathbb{E}_{\text{C,ZK}}(\vec{y}_\gamma; \vec{r}_\gamma)_j = \sum_{i \in [k]} \tilde{\gamma}_i C_{ij}.$$

- S chooses any \vec{y}_1 such that $y = \sum_{j \in [k]} \vec{x}_{1j} \vec{y}_{1j}$. It then selects \vec{r}_1 uniformly at random under the restriction that

$$\forall j \in \hat{J} \cup J: \quad \mathbb{E}_{\text{C,ZK}}(\vec{y}_1; \vec{r}_1)_j = \sum_{i \in [k]} \vec{x}_{0i} C_{ij} - (\mathbb{E}_{\text{C}}(\vec{y}_1) \parallel \vec{0})_j$$

- S continues the rest of the protocol honestly, starting with Step 3.

Figure 6: The zero-knowledge Simulator for the protocol in Figure 3.

Proof of Lemma 5. We now show that the construction in Figure 3 is a polynomial commitment scheme even for openings at multiple adversarially chosen points.

Proof. (Lemma 5):

Completeness, binding, zero-knowledge. These follow directly from Thm. 4.

Knowledge-Soundness. By Thm. 4, any PPT adversary has negligible probability of succeeding without the extractor obtaining $\phi(\cdot)$ that opens the polynomial commitment such that $\sum_{i=0}^m \alpha_i \phi(\vec{x}_i) = \sum_{i=0}^m \alpha_i y_i = \tilde{y}$.

Suppose that this extractor would have non-negligible probability of obtaining more than one ϕ in different executions, then this would mean that it would break the binding definition, so ϕ must be unique.

Suppose that there exists one i' for which $\phi(\vec{x}_{i'}) \neq y_{i'}$. Then

$$\begin{aligned} \Pr \left[\sum_{i=0}^{m-1} \alpha_i \phi(\vec{x}_i) = \sum_{i=0}^{m-1} \alpha_i y_i \right] &= \Pr \left[\alpha_{i'} (\phi(\vec{x}_{i'} - y_{i'})) = - \sum_{i=0, i \neq i'}^{m-1} \alpha_i (\phi(\vec{x}_i) - y_i) \right] \\ &= \Pr \left[\alpha_{i'} = -(\phi(\vec{x}_{i'} - y_{i'}))^{-1} \sum_{i=0, i \neq i'}^{m-1} \alpha_i (\phi(\vec{x}_i) - y_i) \right] = \frac{1}{|\mathbb{F}|}, \end{aligned}$$

and so the prover has negligible success probability. □

5.1 Non-interactivity through Fiat-Shamir

Since the protocol in Figure 3 is public coin, we can apply the Fiat-Shamir transform [FS87]. In the proof of Theorem 4 we show that rewinding k times to the start of `Eval` will fully extract ϕ from its commitment. This means that the non-interactive version of the polynomial commitment scheme is also knowledge sound.

To see this, let \mathcal{A} be an adversary that is able to finish the protocol with \mathcal{V} except with non-negligible probability, performing at most Q random oracle queries. We let \mathcal{A} complete the protocol and in addition run the CP-SNARK extractor. By definition the CP-SNARK extractor succeeds with overwhelming probability. We obtain a transcript that verifies except with negligible probability as well as the CP-SNARK witness. We then reset the random oracle and start another, independent run of \mathcal{A} . After k runs, we have k verifying transcripts and \vec{y}_γ for linearly independent $\vec{\gamma}$, and we can reconstruct W as described in the proof for Theorem 4. The adversary has performed at most kQ ROM queries in the process, in addition to the overhead of running the CP-SNARK extractor k times.

As a result we obtain a non-interactive polynomial commitment scheme. Orion uses the polynomial IOP from [Set20; Gol+23]. Both works apply the Fiat-Shamir transformation but do not provide a full analysis. We leave such a detailed analysis of using Fiat-Shamir to obtain a zk-SNARK and the associated security loss to future work.

Acknowledgements

We want to thank the authors of Orion for their comments on this work. In particular for mentioning the switch from Virgo to a FRI based outer SNARK that supports batching with minimum overhead as a way to reduce the overhead for the our fix that requires multiple rounds of commitments (as discussed in Section 4.1).

References

- [Ame+17] Scott Ames et al. “Ligero: Lightweight Sublinear Arguments Without a Trusted Setup”. In: *ACM CCS 2017: 24th Conference on Computer and Communications Security*. Ed. by Bhavani M. Thuraisingham et al. Dallas, TX, USA: ACM Press, Oct. 2017, pp. 2087–2104. DOI: [10.1145/3133956.3134104](https://doi.org/10.1145/3133956.3134104).
- [Ame+23] Scott Ames et al. “Ligero: lightweight sublinear arguments without a trusted setup”. In: *Designs, Codes and Cryptography* 91.11 (2023), pp. 3379–3424. DOI: [10.1007/s10623-023-01222-8](https://doi.org/10.1007/s10623-023-01222-8).
- [Bau+23] Carsten Baum et al. “Publicly Verifiable Zero-Knowledge and Post-Quantum Signatures from VOLE-in-the-Head”. In: *Advances in Cryptology – CRYPTO 2023, Part V*. Ed. by Helena Handschuh and Anna Lysyanskaya. Vol. 14085. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Cham, Switzerland, Aug. 2023, pp. 581–615. DOI: [10.1007/978-3-031-38554-4_19](https://doi.org/10.1007/978-3-031-38554-4_19).
- [BCG20] Jonathan Bootle, Alessandro Chiesa, and Jens Groth. “Linear-Time Arguments with Sublinear Verification from Tensor Codes”. In: *TCC 2020: 18th Theory of Cryptography Conference, Part II*. Ed. by Rafael Pass and Krzysztof Pietrzak. Vol. 12551. Lecture Notes in Computer Science. Durham, NC, USA: Springer, Cham, Switzerland, Nov. 2020, pp. 19–46. DOI: [10.1007/978-3-030-64378-2_2](https://doi.org/10.1007/978-3-030-64378-2_2).
- [BCL22] Jonathan Bootle, Alessandro Chiesa, and Siqi Liu. “Zero-Knowledge IOPs with Linear-Time Prover and Polylogarithmic-Time Verifier”. In: *Advances in Cryptology – EUROCRYPT 2022, Part II*. Ed. by Orr Dunkelman and Stefan Dziembowski. Vol. 13276. Lecture Notes in Computer Science. Trondheim, Norway: Springer, Cham, Switzerland, May 2022, pp. 275–304. DOI: [10.1007/978-3-031-07085-3_10](https://doi.org/10.1007/978-3-031-07085-3_10).
- [Ben+18] Eli Ben-Sasson et al. “Fast Reed-Solomon Interactive Oracle Proofs of Proximity”. In: *ICALP 2018: 45th International Colloquium on Automata, Languages and Programming*. Ed. by Ioannis Chatzigiannakis et al. Vol. 107. LIPIcs. Prague, Czech Republic: Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, July 2018, 14:1–14:17. DOI: [10.4230/LIPIcs.ICALP.2018.14](https://doi.org/10.4230/LIPIcs.ICALP.2018.14).
- [BFS20] Benedikt Bünz, Ben Fisch, and Alan Szepieniec. “Transparent SNARKs from DARK Compilers”. In: *Advances in Cryptology – EUROCRYPT 2020, Part I*. Ed. by Anne Canoute and Yuval Ishai. Vol. 12105. Lecture Notes in Computer Science. Zagreb, Croatia: Springer, Cham, Switzerland, May 2020, pp. 677–706. DOI: [10.1007/978-3-030-45721-1_24](https://doi.org/10.1007/978-3-030-45721-1_24).

- [Bit+12] Nir Bitansky et al. “From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again”. In: *ITCS 2012: 3rd Innovations in Theoretical Computer Science*. Ed. by Shafi Goldwasser. Cambridge, MA, USA: Association for Computing Machinery, Jan. 2012, pp. 326–349. DOI: [10.1145/2090236.2090263](https://doi.org/10.1145/2090236.2090263).
- [Boo+17] Jonathan Bootle et al. “Linear-Time Zero-Knowledge Proofs for Arithmetic Circuit Satisfiability”. In: *Advances in Cryptology – ASIACRYPT 2017, Part III*. Ed. by Tsuyoshi Takagi and Thomas Peyrin. Vol. 10626. Lecture Notes in Computer Science. Hong Kong, China: Springer, Cham, Switzerland, Dec. 2017, pp. 336–365. DOI: [10.1007/978-3-319-70700-6_12](https://doi.org/10.1007/978-3-319-70700-6_12).
- [CFQ19] Matteo Campanelli, Dario Fiore, and Anaïs Querol. “LegoSNARK: Modular Design and Composition of Succinct Zero-Knowledge Proofs”. In: *ACM CCS 2019: 26th Conference on Computer and Communications Security*. Ed. by Lorenzo Cavallaro et al. London, UK: ACM Press, Nov. 2019, pp. 2075–2092. DOI: [10.1145/3319535.3339820](https://doi.org/10.1145/3319535.3339820).
- [Cha+17] Melissa Chase et al. “Post-Quantum Zero-Knowledge and Signatures from Symmetric-Key Primitives”. In: *ACM CCS 2017: 24th Conference on Computer and Communications Security*. Ed. by Bhavani M. Thuraisingham et al. Dallas, TX, USA: ACM Press, Oct. 2017, pp. 1825–1842. DOI: [10.1145/3133956.3133997](https://doi.org/10.1145/3133956.3133997).
- [Che+23] Binyi Chen et al. “HyperPlonk: Plonk with Linear-Time Prover and High-Degree Custom Gates”. In: *Advances in Cryptology – EUROCRYPT 2023, Part II*. Ed. by Carmit Hazay and Martijn Stam. Vol. 14005. Lecture Notes in Computer Science. Lyon, France: Springer, Cham, Switzerland, Apr. 2023, pp. 499–530. DOI: [10.1007/978-3-031-30617-4_17](https://doi.org/10.1007/978-3-031-30617-4_17).
- [DG23] Quang Dao and Paul Grubbs. “Spartan and Bulletproofs are Simulation-Extractable (for Free!)” In: *Advances in Cryptology – EUROCRYPT 2023, Part II*. Ed. by Carmit Hazay and Martijn Stam. Vol. 14005. Lecture Notes in Computer Science. Lyon, France: Springer, Cham, Switzerland, Apr. 2023, pp. 531–562. DOI: [10.1007/978-3-031-30617-4_18](https://doi.org/10.1007/978-3-031-30617-4_18).
- [Fao+23] Antonio Faonio et al. “From Polynomial IOP and Commitments to Non-malleable zk-SNARKs”. In: *TCC 2023: 21st Theory of Cryptography Conference, Part III*. Ed. by Guy N. Rothblum and Hoeteck Wee. Vol. 14371. Lecture Notes in Computer Science. Taipei, Taiwan: Springer, Cham, Switzerland, Nov. 2023, pp. 455–485. DOI: [10.1007/978-3-031-48621-0_16](https://doi.org/10.1007/978-3-031-48621-0_16).
- [FS87] Amos Fiat and Adi Shamir. “How to Prove Yourself: Practical Solutions to Identification and Signature Problems”. In: *Advances in Cryptology – CRYPTO’86*. Ed. by Andrew M. Odlyzko. Vol. 263. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Berlin, Heidelberg, Germany, Aug. 1987, pp. 186–194. DOI: [10.1007/3-540-47721-7_12](https://doi.org/10.1007/3-540-47721-7_12).
- [Gan+22] Chaya Ganesh et al. “What Makes Fiat-Shamir zkSNARKs (Updatable SRS) Simulation Extractable?” In: *SCN 22: 13th International Conference on Security in Communication Networks*. Ed. by Clemente Galdi and Stanislaw Jarecki. Vol. 13409. Lecture Notes in Computer Science. Amalfi, Italy: Springer, Cham, Switzerland, Sept. 2022, pp. 735–760. DOI: [10.1007/978-3-031-14791-3_32](https://doi.org/10.1007/978-3-031-14791-3_32).
- [GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. “The Knowledge Complexity of Interactive Proof-Systems (Extended Abstract)”. In: *17th Annual ACM Symposium on Theory of Computing*. Providence, RI, USA: ACM Press, May 1985, pp. 291–304. DOI: [10.1145/22145.22178](https://doi.org/10.1145/22145.22178).
- [Gol+23] Alexander Golovnev et al. “Brakedown: Linear-Time and Field-Agnostic SNARKs for R1CS”. In: *Advances in Cryptology – CRYPTO 2023, Part II*. Ed. by Helena Handschuh and Anna Lysyanskaya. Vol. 14082. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Cham, Switzerland, Aug. 2023, pp. 193–226. DOI: [10.1007/978-3-031-38545-2_7](https://doi.org/10.1007/978-3-031-38545-2_7).
- [Ish+09] Yuval Ishai et al. “Zero-Knowledge Proofs from Secure Multiparty Computation”. In: *SIAM J. Comput.* 39.3 (2009), pp. 1121–1152. DOI: [10.1137/080725398](https://doi.org/10.1137/080725398). URL: <https://doi.org/10.1137/080725398>.

- [KPT23] Markulf Kohlweiss, Mahak Panholi, and Akira Takahashi. “How to Compile Polynomial IOP into Simulation-Extractable SNARKs: A Modular Approach”. In: *TCC 2023: 21st Theory of Cryptography Conference, Part III*. Ed. by Guy N. Rothblum and Hoeteck Wee. Vol. 14371. Lecture Notes in Computer Science. Taipei, Taiwan: Springer, Cham, Switzerland, Nov. 2023, pp. 486–512. DOI: [10.1007/978-3-031-48621-0_17](https://doi.org/10.1007/978-3-031-48621-0_17).
- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. “Constant-Size Commitments to Polynomials and Their Applications”. In: *Advances in Cryptology – ASIACRYPT 2010*. Ed. by Masayuki Abe. Vol. 6477. Lecture Notes in Computer Science. Singapore: Springer, Berlin, Heidelberg, Germany, Dec. 2010, pp. 177–194. DOI: [10.1007/978-3-642-17373-8_11](https://doi.org/10.1007/978-3-642-17373-8_11).
- [Lib24] Benoît Libert. “Simulation-Extractable KZG Polynomial Commitments and Applications to HyperPlonk”. In: *Public-Key Cryptography - PKC 2024 - 27th IACR International Conference on Practice and Theory of Public-Key Cryptography, Sydney, NSW, Australia, April 15-17, 2024, Proceedings, Part II*. Ed. by Qiang Tang and Vanessa Teague. Vol. 14602. Lecture Notes in Computer Science. Springer, 2024, pp. 68–98. DOI: [10.1007/978-3-031-57722-2_3](https://doi.org/10.1007/978-3-031-57722-2_3). URL: https://doi.org/10.1007/978-3-031-57722-2_3.
- [Mer90] Ralph C. Merkle. “A Certified Digital Signature”. In: *Advances in Cryptology – CRYPTO’89*. Ed. by Gilles Brassard. Vol. 435. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, New York, USA, Aug. 1990, pp. 218–238. DOI: [10.1007/0-387-34805-0_21](https://doi.org/10.1007/0-387-34805-0_21).
- [Nit19] Anca Nitulescu. “A Gentle Introduction to SNARKs”. In: 2019. URL: <https://api.semanticscholar.org/CorpusID:209520686>.
- [Sch90] Claus-Peter Schnorr. “Efficient Identification and Signatures for Smart Cards”. In: *Advances in Cryptology – CRYPTO’89*. Ed. by Gilles Brassard. Vol. 435. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, New York, USA, Aug. 1990, pp. 239–252. DOI: [10.1007/0-387-34805-0_22](https://doi.org/10.1007/0-387-34805-0_22).
- [Set20] Srinath Setty. “Spartan: Efficient and General-Purpose zkSNARKs Without Trusted Setup”. In: *Advances in Cryptology – CRYPTO 2020, Part III*. Ed. by Daniele Micciancio and Thomas Ristenpart. Vol. 12172. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Cham, Switzerland, Aug. 2020, pp. 704–737. DOI: [10.1007/978-3-030-56877-1_25](https://doi.org/10.1007/978-3-030-56877-1_25).
- [Tha22] Justin Thaler. “Proofs, Arguments, and Zero-Knowledge”. In: *Foundations and Trends® in Privacy and Security 4.2–4 (2022)*, pp. 117–660. ISSN: 2474-1558. DOI: [10.1561/33000000030](https://dx.doi.org/10.1561/33000000030). URL: <http://dx.doi.org/10.1561/33000000030>.
- [XZS22a] Tiancheng Xie, Yupeng Zhang, and Dawn Song. “Orion: Zero Knowledge Proof with Linear Prover Time”. Cryptology ePrint Archive, Paper 2022/1010. Oct. 2022. URL: <https://eprint.iacr.org/archive/2022/1010/20231027:195131>.
- [XZS22b] Tiancheng Xie, Yupeng Zhang, and Dawn Song. “Orion: Zero Knowledge Proof with Linear Prover Time”. Cryptology ePrint Archive, Paper 2022/1010. Nov. 2022. URL: <https://eprint.iacr.org/archive/2022/1010/20231102:185740>.
- [XZS22c] Tiancheng Xie, Yupeng Zhang, and Dawn Song. “Orion: Zero Knowledge Proof with Linear Prover Time”. In: *Advances in Cryptology – CRYPTO 2022, Part IV*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Vol. 13510. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Cham, Switzerland, Aug. 2022, pp. 299–328. DOI: [10.1007/978-3-031-15985-5_11](https://doi.org/10.1007/978-3-031-15985-5_11).
- [Zha+20] Jiaheng Zhang et al. “Transparent Polynomial Delegation and Its Applications to Zero Knowledge Proof”. In: *2020 IEEE Symposium on Security and Privacy*. San Francisco, CA, USA: IEEE Computer Society Press, May 2020, pp. 859–876. DOI: [10.1109/SP40000.2020.00052](https://doi.org/10.1109/SP40000.2020.00052).