

Proof-Carrying Data from Multi-folding Schemes

Zibo Zhou¹, Zongyang Zhang¹, and Jin Dong²

¹ School of Cyber Science and Technology, Beihang University, Beijing, China
{zbzhou, zongyangzhang}@buaa.edu.cn

² Beijing Academy of Blockchain and Edge Computing, BABEC
dongjin@baec.org.cn

Abstract. Proof-carrying data (PCD) is a powerful cryptographic primitive that allows mutually distrustful parties to perform distributed computation defined on directed acyclic graphs in an efficiently verifiable manner. Important efficiency parameters include prover’s cost at each step and the recursion overhead that measures the additional cost apart from proving the computation.

In this paper, we construct a PCD scheme having the smallest prover’s cost and recursion overhead in the literature. Specifically, the prover’s cost at each step is dominated by only one $O(|C|)$ -sized multi-scalar multiplication (MSM), and the recursion overhead is dominated by only one $2r$ -sized MSM, where $|C|$ is the computation size and r is the number of incoming edges at certain step. In contrast, the state-of-the-art PCD scheme requires $4r + 12$ $O(|C|)$ -sized MSMs w.r.t. the prover’s cost and six $2r$ -sized MSMs, one $6r$ -sized MSM w.r.t. the recursion overhead. In addition, our PCD scheme supports more expressive constraint system for computations—*customizable constraint system (CCS)* that supports high-degree constraints efficiently, in contrast with *rank-1 constraint system (R1CS)* that supports only quadratic constraints used in existing PCD schemes.

Underlying our PCD scheme is a multi-folding scheme that reduces the task of checking multiple instances into the task of checking one. We generalize existing construction to support arbitrary number of instances.

Keywords: proof-carrying data · folding schemes · recursive zero-knowledge proofs.

1 Introduction

Proof-carrying data (PCD), introduced by Chiesa and Tromer [13], is a powerful cryptographic primitive that enables mutually distrustful parties to perform distributed computations defined on directed acyclic graphs, while every intermediate state of the computation can be verified efficiently. It generalizes incrementally verifiable computation (IVC) [32] which enables a possibly infinite computation defined on path graphs such that the correctness can be verified efficiently at any point. PCD has found numerous applications to enforcing language semantics [15], verifiable MapReduce computations [14], verifiable photo editing [27], verifiable registries [31], and blockchains [4,11,20].

There has been tremendous interest and progress in designing efficient PCD schemes. A classic method for constructing PCD is via *recursive composition* of succinct non-interactive argument of knowledge (SNARK) [2,1,12]. Informally, at each step i , the prover uses a SNARK to prove that the i -th step of the computation is executed correctly and the SNARK verifier expressed as an arithmetic circuit has accepted the SNARK proof from step $i - 1$. Important efficiency parameters include the prover’s cost at each step and the *recursion overhead*, i.e., the verifier’s operations expressed as an arithmetic circuit that the prover must prove at each step besides proving the computation. This method yields a secure PCD construction but the concrete efficiency is limited by the use of cycles of expensive pairing-friendly elliptic curves for pairing-based SNARKs [1] or heavy use of cryptographic hash functions for hash-based SNARKs [12].

Bünz et al. [8], denoted as BCMS20, introduced an alternative method for constructing PCD by formalizing a novel notion—atomic accumulation scheme from Halo [5]. Instead of expressing the whole SNARK verifier as a circuit, BCMS20 only requires expressing the verifier of the atomic accumulation scheme as a circuit, whose size is much smaller. However, the prover at each step has to perform expensive fast Fourier transforms (FFTs) and the recursion overhead is still dominated by group operations that scale logarithmically with the size of the computation. Later, Bünz et al. [7], denoted as BCLMS21, extended the notion of atomic accumulation scheme to split accumulation scheme. By designing a non-interactive argument of knowledge with a split accumulation scheme for *rank-1 constraint system (R1CS)* [19], they constructed a PCD scheme whose recursion overhead is dominated by constant number of group operations. But the prover at each step still has to generate a proof, resulting in slightly high prover’s cost and recursion overhead. Kothapalli and Setty [23] avoided generating a proof by introducing and constructing a multi-folding scheme for *customizable constraint system (CCS)* [29]. Nevertheless, their multi-folding scheme is tailored to fold two instances and thus results in an IVC scheme that is only a special case of PCD.

1.1 Our Contributions

In this paper, we generalize the multi-folding scheme of Kothapalli and Setty [23] and then construct a PCD scheme having the smallest prover’s cost and recursion overhead in the literature. Table 1 shows the comparison between the state-of-the-art PCD scheme and ours. We elaborate our contributions below.

A Multi-folding Scheme for Arbitrary Number of Instances. We generalize the multi-folding scheme of Kothapalli and Setty [23] to support folding arbitrary number of instances. Our scheme could reduce the task of checking multiple instances into the task of checking a single instance. By using the Fiat-Shamir transformation [18] we could make it non-interactive, which is the basis of our PCD scheme.

An Efficient PCD Scheme Supporting *Customizable Constraint System*. We construct an efficient PCD scheme that supports a more expressive constraint system—*customizable constraint system (CCS)*. The recursion overhead of our PCD scheme is dominated by only one multi-scalar multiplication (MSM)³ of size $2r$, and the prover’s cost at each step is dominated by only one MSM of size $O(n)$, where r is the number of incoming edges of a node at certain step and n is the number of variables in the constraint system. The proof size is logarithmic and the verifier’s cost is dominated by only one MSM of size n . Therefore, our PCD scheme not only has smaller prover’s cost and recursion overhead, but also reduces the proof size exponentially and reduces the verifier’s cost significantly. Besides, the CCS our PCD scheme supports can generalize not only R1CS without overhead but also Plonkish which supports constraints with degree larger than two and thereby represents certain program executions more succinctly. In particular, when $t = 3, q = 2, d = 2$, the data showed in Table 1 measures the efficiency of our PCD scheme using R1CS.

However, achieving zero-knowledge in our PCD scheme requires the multi-folding scheme to be zero-knowledge and an additional invocation of a general zero-knowledge succinct non-interactive argument of knowledge (zkSNARK), which we leave in future work.

³ For field elements a_1, \dots, a_r and elliptic curve group elements G_1, \dots, G_r , the multi-scalar multiplication denotes the operation $a_1G_1 + \dots + a_rG_r$. The group scalar multiplication denotes the operation a_1G_1 .

1.2 Related Work

IVC/PCD from SNARKs. Chiesa and Tromer [13] introduced the notion of proof-carrying data (PCD) which generalizes incrementally verifiable computation (IVC) [32]. Bitansky et al. [2] proved that any SNARK for machine computation can be efficiently transformed into a corresponding PCD scheme for constant-depth graphs via recursive composition. But it was not realized in practice due to enormous computational cost. Ben-Sasson et al. [1] achieved the first implementation of PCD using pairing-based SNARKs instantiated via pairing-friendly cycles of elliptic curves. But these curves must be constructed over large (780-bit) fields due to their low embedding degrees, resulting in poor concrete efficiency. Chiesa et al. [12] obtained the first efficient realization of post-quantum transparent PCD, but the number of constraints needed to express the verifier’s computation in a circuit is dominated by numerous hash function invocations.

Table 1. Comparison of proof-carrying data schemes

| | BCLMS21 [7] | This work |
|---------------------------|--|--|
| Constraint system | R1CS | CCS |
| Recursion overhead | 6 $(2r)$ -MSM 1 $(6r)$ -MSM $5r$ \mathbb{G} $r + 3$ RO | $O(d \log m + r \cdot (\log m + t + dq))$ \mathbb{F} 1 $(2r)$ -MSM $r + 1$ H $2 \log_2 m + 2$ RO |
| Prover’s cost (each step) | $O(r \cdot (N + m + n))$ \mathbb{F} $4r + 12$ (m) -MSM 6 $(2r)$ -MSM 1 $(6r)$ -MSM $5r$ \mathbb{G} $r + 4$ RO | $O(r \cdot (N + tm + n + qmd \log^2 d))$ \mathbb{F} 1 $(2r)$ -MSM 1 $O(n)$ -MSM $2 \log_2 m + 2$ RO |
| Proof size | $O(m + n)$ \mathbb{F} 15 \mathbb{G} | $O(d \log m + t + \log n)$ \mathbb{F} $O(\log n)$ \mathbb{G} |
| Verifier’s cost | $O(N)$ \mathbb{F} 10 $O(m)$ -MSM 1 RO | $O(d \log m + t + dq + \log n)$ \mathbb{F} 1 $O(n)$ -MSM $O(\log m + \log n)$ RO |
| Zero-knowledge | yes | no |

Notations: R1CS denotes rank-1 constraint system with parameters (m, n, N) where $N = \Omega(\max(m, n))$. CCS denotes customizable constraint system with parameters (m, n, N, t, q, d) . r denotes the number of incoming edges of a node at certain step. \mathbb{G} denotes group scalar multiplications when measuring overhead/cost or group elements when measuring proof size. Similarly, \mathbb{F} denotes field operations or field elements. H denotes the invocation of hash functions. RO denotes the invocation of random oracles. 1 $(2r)$ -MSM denotes one multi-scalar multiplication of size $2r$.

IVC/PCD from accumulation/folding schemes. Bünz et al. [8] developed and formalized a novel approach from Bowe et al. [5] to construct PCD, i.e., PCD from atomic accumulation schemes. Following this line of works, Boneh et al. [3] formalized a method to construct PCD from additive polynomial commitment schemes with a (public/private) aggregation scheme. Meanwhile, Bünz et al. [7] improved [8] by introducing the notion of split accumulation schemes to construct PCD.

Later, many researchers focus on IVC instead of PCD for its conceptual simplicity and practical applications in verifiable delay functions and zkEVM. Kothapalli et al. [24] introduced the notion of folding schemes and constructed an IVC scheme—Nova whose recursion overhead is dominated by only two group scalar multiplications. Kothapalli and Setty [21] adapted Nova to SuperNova that realizes non-uniform IVC with the same recursion overhead. Nova and SuperNova both express each step of an incremental computation with R1CS, which is restricted to checking quadratic constraints in a specific form. In practice, practitioners often use custom constraint systems such as Plonkish that allows gates to compute custom high-degree polynomials. To overcome this limitation, Mohnblatt [26] introduced Sangria which adapted the folding scheme of Nova to handle Plonkish, but the number of group operations grows linearly in the gate degree. Recently, Kothapalli and Setty [23] introduced HyperNova, an IVC scheme for CCS that simultaneously generalizes Plonkish, R1CS, and algebraic intermediate representation (AIR) without overheads. Bünz and Chen [10] introduced ProtoStar, a non-uniform IVC scheme for Plonkish/CCS. Eagen and Gabizon [17] continued the work of ProtoStar and constructed an efficient folding scheme—ProtoGalaxy for multiple instances. Kothapalli and Setty [22] introduced CycleFold—a conceptually simple approach to instantiate folding-scheme-based IVCs over a cycle of elliptic curves.

2 Preliminaries

For $t \in \mathbb{N}$, let $[t]$ denote the set $\{1, 2, \dots, t\}$. We use \mathbb{F} to denote a finite field, \mathbb{F}^t to denote a vector space of dimension t over \mathbb{F} , and $\mathbb{F}[\ell]$ to denote the family of ℓ -variate multilinear polynomials over \mathbb{F} . For $x \in \mathbb{F}^t$, let x_i denote the i -th element. We use $y \stackrel{\$}{\leftarrow} S$ to denote the assignment of a uniformly random element in set S to y . We use $y := c$ to denote the assignment of the constant value c to y . When A is an algorithm, we use $y \leftarrow A(x)$ to denote the assignment of the output of A on input x to y . We use λ to denote the security parameter and will drop it from the notation when it is implicit. We use $\text{negl}(\lambda)$ to denote a negligible function in λ . Let PPT denote probabilistic polynomial time. A *multiset* is an extension of the concept of a set where every element has a positive multiplicity. For a tuple containing a semicolon, those variables listed before the semicolon are public (known to both the prover and the verifier), and those listed after it are secret (known only to the prover).

Lemma 1 (Multilinear Extension [16]). *For every function $f : \{0, 1\}^\ell \rightarrow \mathbb{F}$, there is a unique ℓ -variate multilinear polynomial $\tilde{f} : \mathbb{F}^\ell \rightarrow \mathbb{F}$ such that $\tilde{f}(x) = f(x)$ for all $x \in \{0, 1\}^\ell$. We call \tilde{f} the multilinear extension (MLE) of function f , and \tilde{f} can be expressed as*

$$\tilde{f}(X) = \sum_{x \in \{0, 1\}^\ell} f(x) \cdot \tilde{eq}(X, x),$$

where $\tilde{eq}(X, x) = \prod_{i=1}^{\ell} (x_i X_i + (1 - x_i)(1 - X_i))$ and itself is the MLE of function

$$eq(X, x) : \{0, 1\}^\ell \rightarrow \mathbb{F}, \text{ and } eq(X, x) = \begin{cases} 1 & \text{if } X = x \\ 0 & \text{otherwise} \end{cases}.$$

Lemma 2 (Lagrange Interpolation of Multilinear Polynomials [30, Lemma 3.6]). For $\ell \in \mathbb{N}$, any ℓ -variate multilinear polynomial P could be expressed as $P(X) = \sum_{x \in \{0,1\}^\ell} \tilde{e}q(X, x) \cdot P(x)$.

Lemma 3 (Schwartz-Zippel Lemma [28]). Let $f : \mathbb{F}^\ell \rightarrow \mathbb{F}$ be a non-zero ℓ -variate polynomial of total degree d . Let S be any finite subset of \mathbb{F} . Then for $x \stackrel{\$}{\leftarrow} S^\ell$, $\Pr[f(x) = 0] \leq \frac{d}{|S|}$.

2.1 The Sum-Check Protocol

We adapt the description from Kothapalli and Setty [23]. The sum-check protocol [25] is an interactive protocol allowing a prover \mathcal{P} to convince a verifier \mathcal{V} of the validity of the statement

$$T = \sum_{x_1 \in \{0,1\}} \cdots \sum_{x_\ell \in \{0,1\}} g(x_1, \dots, x_\ell),$$

where $g : \mathbb{F}^\ell \rightarrow \mathbb{F}$ is an ℓ -variate polynomial over some finite field \mathbb{F} , and the degree of each variable is at most d . While \mathcal{V} could directly compute T using $O(2^\ell)$ evaluation of g , the sum-check protocol reduces \mathcal{V} 's work to be polynomial in ℓ . In the protocol, \mathcal{V} takes as input randomness $r \in \mathbb{F}^\ell$ and interacts with \mathcal{P} over a sequence of ℓ rounds. At the end of this interaction, \mathcal{V} outputs a claim about the evaluation $g(r)$. We denote the sum-check protocol as $c \leftarrow \langle \mathcal{P}, \mathcal{V}(r) \rangle (g, \ell, d, T)$, and it satisfies the following properties.

- **Completeness.** If $T = \sum_{x \in \{0,1\}^\ell} g(x)$, then for an honest \mathcal{P} and for all $r \in \mathbb{F}^\ell$, $\Pr_r [c \leftarrow \langle \mathcal{P}, \mathcal{V}(r) \rangle (g, \ell, d, T) \wedge g(r) = c] = 1$.
- **Soundness.** If $T \neq \sum_{x \in \{0,1\}^\ell} g(x)$, then for any adversary \mathcal{P}^* and for all $r \in \mathbb{F}^\ell$, $\Pr_r [c \leftarrow \langle \mathcal{P}^*, \mathcal{V}(r) \rangle (g, \ell, d, T) \wedge g(r) = c] \leq \ell \cdot d / |\mathbb{F}|$.
- **Succinctness.** The communication cost is $O(\ell \cdot d)$ elements in \mathbb{F} .

2.2 Polynomial Commitment Schemes

Definition 1 (Polynomial Commitment Schemes for Multilinear Polynomials [9]). A polynomial commitment scheme for multilinear polynomials is a tuple $\text{PCS} = (\text{Setup}, \text{Com}, \text{Open}, \text{Eval})$ where

- $\text{pp}_{\text{PCS}} \leftarrow \text{Setup}(1^\lambda, \ell)$ takes the security parameter λ with the unary form and the number of variables ℓ in a multilinear polynomial as input, and outputs some public parameter pp_{PCS} ;
- $C \leftarrow \text{Com}(\text{pp}_{\text{PCS}}, g)$ takes pp_{PCS} and an ℓ -variate multilinear polynomial over a finite field $g \in \mathbb{F}[\ell]$ as input, and outputs a commitment C ;
- $1/0 \leftarrow \text{Open}(\text{pp}_{\text{PCS}}, C, g)$ takes $\text{pp}_{\text{PCS}}, C, g$ as input, and outputs $1/0$ to denote that C is indeed a commitment to g or not.
- $1/0 \leftarrow \text{Eval}(\text{pp}_{\text{PCS}}, C, \ell, r, v; g)$ is an interactive protocol between a prover \mathcal{P} and a verifier \mathcal{V} . Both of them know $\text{pp}_{\text{PCS}}, C, \ell, r \in \mathbb{F}^\ell, v \in \mathbb{F}$. \mathcal{P} additionally knows g and attempts to convince \mathcal{V} that $C = \text{Com}(\text{pp}_{\text{PCS}}, g)$ and $g(r) = v$. The outputs $1/0$ denote that \mathcal{V} accepts or not.

A polynomial commitment scheme PCS should satisfy completeness, binding and knowledge soundness defined below.

- **Completeness.** PCS has completeness if for any ℓ -variate multilinear polynomial $g \in \mathbb{F}[\ell]$,

$$\Pr \left[\text{Eval}(\text{pp}_{\text{PCS}}, C, \ell, r, v; g) = 1 \mid \begin{array}{l} \text{pp}_{\text{PCS}} \leftarrow \text{Setup}(1^\lambda, \ell); \\ C \leftarrow \text{Com}(\text{pp}_{\text{PCS}}, g) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

- **Binding.** PCS has binding if for any PPT adversary \mathcal{A} , size parameter $\ell \geq 1$,

$$\Pr \left[\begin{array}{l} b_0 = b_1 \neq 0 \\ \wedge g_0 \neq g_1 \end{array} \middle| \begin{array}{l} \text{pp}_{\text{PCS}} \leftarrow \text{Setup}(1^\lambda, \ell); (C, g_0, g_1) \leftarrow \mathcal{A}(\text{pp}_{\text{PCS}}); \\ b_0 \leftarrow \text{Open}(\text{pp}_{\text{PCS}}, C, g_0); b_1 \leftarrow \text{Open}(\text{pp}_{\text{PCS}}, C, g_1) \end{array} \right] \leq \text{negl}(\lambda).$$

- **Knowledge soundness.** PCS has knowledge soundness if given $\text{pp}_{\text{PCS}} \leftarrow \text{Setup}(1^\lambda, \ell)$, Eval is a succinct argument of knowledge for NP relation

$$\mathcal{R}_{\text{Eval}}(\text{pp}_{\text{PCS}}) = \{(C, r, v; g) : g \in \mathbb{F}[\ell] \wedge g(r) = v \wedge \text{Open}(\text{pp}_{\text{PCS}}, C, g) = 1\}.$$

PCS is additively homomorphic if for all ℓ and public parameters pp_{PCS} , and for any ℓ -variate multilinear polynomial $g_1, g_2 \in \mathbb{F}[\ell]$, $\text{Com}(\text{pp}_{\text{PCS}}, g_1) + \text{Com}(\text{pp}_{\text{PCS}}, g_2) = \text{Com}(\text{pp}_{\text{PCS}}, g_1 + g_2)$.

A polynomial commitment to an ℓ -variate multilinear polynomial can be viewed as a commitment to a 2^ℓ -sized vector. In this paper, we instantiate the polynomial commitment scheme for multilinear polynomials with Bulletproofs [6], whose proof size is dominated by $O(\ell)$ group elements and verifier's cost is dominated by one MSM of size $O(2^\ell)$ in the Eval protocol.

2.3 Multi-folding Schemes

Definition 2 (Multi-folding Schemes [23]). Consider relations \mathcal{R}_1 and \mathcal{R}_2 over public parameters, structure, instance, and witness tuples, a predicate compat that structures in \mathcal{R}_1 and \mathcal{R}_2 must satisfy, and size parameters $\mu, \nu \in \mathbb{N}$. A multi-folding scheme for $(\mathcal{R}_1, \mathcal{R}_2, \text{compat}, \mu, \nu)$ is a tuple of algorithms $\text{MFS} = (\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ where

- $\text{fpp} \leftarrow \mathcal{G}(1^\lambda)$ is a PPT generator algorithm. It takes the security parameter λ with the unary form as input, and outputs some public parameter fpp ;
- $(\text{fpk}, \text{fvk}) \leftarrow \mathcal{K}(\text{fpp}, (\mathbf{s}_1, \mathbf{s}_2))$ is a deterministic encoder algorithm. It takes fpp and structure $\mathbf{s}_1, \mathbf{s}_2$ among the instances to be folded as input, and outputs a prover key fpk and a verifier key fvk ;
- $(\mathbf{u}, \mathbf{w}) \leftarrow \langle \mathcal{P}(\text{fpk}, \vec{\mathbf{w}}_1, \vec{\mathbf{w}}_2), \mathcal{V}(\text{fvk}) \rangle(\vec{\mathbf{u}}_1, \vec{\mathbf{u}}_2)$ denotes the interaction between a PPT prover \mathcal{P} and a PPT verifier \mathcal{V} . \mathcal{P} inputs fpk , a vector of instances $\vec{\mathbf{u}}_1$ in \mathcal{R}_1 of size μ with structure \mathbf{s}_1 , a vector of instances $\vec{\mathbf{u}}_2$ in \mathcal{R}_2 of size ν with structure \mathbf{s}_2 , and corresponding witnesses $\vec{\mathbf{w}}_1, \vec{\mathbf{w}}_2$. \mathcal{V} inputs $\text{fvk}, \vec{\mathbf{u}}_1, \vec{\mathbf{u}}_2$. At the end of the interaction, \mathcal{P} outputs a folded instance-witness pair (\mathbf{u}, \mathbf{w}) in \mathcal{R}_1 with structure \mathbf{s}_1 , and \mathcal{V} outputs a folded instance \mathbf{u} in \mathcal{R}_1 with structure \mathbf{s}_1 .

Let $\mathcal{R}^{(n)}$ be the relation so that $(\text{fpp}, \mathbf{s}, \vec{\mathbf{u}}, \vec{\mathbf{w}}) \in \mathcal{R}^{(n)}$ if and only if $(\text{fpp}, \mathbf{s}, \vec{\mathbf{u}}_i, \vec{\mathbf{w}}_i) \in \mathcal{R}$ for all $i \in [n]$. A multi-folding scheme MFS should satisfy perfect completeness and knowledge soundness defined below.

- **Perfect Completeness.** MFS has perfect completeness if for all PPT adversaries \mathcal{A} ,

$$\Pr \left[\begin{array}{l} (\text{fpp}, \mathbf{s}_1, \mathbf{u}, \mathbf{w}) \in \mathcal{R}_1 \end{array} \middle| \begin{array}{l} \text{fpp} \leftarrow \mathcal{G}(1^\lambda); \\ ((\mathbf{s}_1, \mathbf{s}_2), (\vec{\mathbf{u}}_1, \vec{\mathbf{w}}_1), (\vec{\mathbf{u}}_2, \vec{\mathbf{w}}_2)) \leftarrow \mathcal{A}(\text{fpp}); \\ \text{compat}(\mathbf{s}_1, \mathbf{s}_2) = \text{true}; \\ (\text{fpp}, \mathbf{s}_1, \vec{\mathbf{u}}_1, \vec{\mathbf{w}}_1) \in \mathcal{R}_1^{(\mu)}; (\text{fpp}, \mathbf{s}_2, \vec{\mathbf{u}}_2, \vec{\mathbf{w}}_2) \in \mathcal{R}_2^{(\nu)}; \\ (\text{fpk}, \text{fvk}) \leftarrow \mathcal{K}(\text{fpp}, (\mathbf{s}_1, \mathbf{s}_2)); \\ (\mathbf{u}, \mathbf{w}) \leftarrow \langle \mathcal{P}(\text{fpk}, \vec{\mathbf{w}}_1, \vec{\mathbf{w}}_2), \mathcal{V}(\text{fvk}) \rangle(\vec{\mathbf{u}}_1, \vec{\mathbf{u}}_2) \end{array} \right] = 1.$$

- **Knowledge Soundness.** *MFS has knowledge soundness if for any expected polynomial time adversary \mathcal{A} and \mathcal{P}^* , there exists an expected polynomial time extractor \mathcal{E} such that over all randomness ρ ,*

$$\Pr \left[\begin{array}{l} (\text{fpp}, s_1, \vec{u}_1, \vec{w}_1) \in \mathcal{R}_1^{(\mu)} \\ \wedge (\text{fpp}, s_2, \vec{u}_2, \vec{w}_2) \in \mathcal{R}_2^{(\nu)} \end{array} \middle| \begin{array}{l} \text{fpp} \leftarrow \mathcal{G}(1^\lambda); \\ ((s_1, s_2), (\vec{u}_1, \vec{u}_2), \text{st}) \leftarrow \mathcal{A}(\text{fpp}, \rho); \\ \text{compat}(s_1, s_2) = \text{true}; \\ (\vec{w}_1, \vec{w}_2) \leftarrow \mathcal{E}(\text{fpp}, \rho) \end{array} \right] \geq$$

$$\Pr \left[\begin{array}{l} (\text{fpp}, s_1, u, w) \in \mathcal{R}_1 \end{array} \middle| \begin{array}{l} \text{fpp} \leftarrow \mathcal{G}(1^\lambda); \\ ((s_1, s_2), (\vec{u}_1, \vec{u}_2), \text{st}) \leftarrow \mathcal{A}(\text{fpp}, \rho); \\ \text{compat}(s_1, s_2) = \text{true}; \\ (\text{fpk}, \text{fvk}) \leftarrow \mathcal{K}(\text{fpp}, (s_1, s_2)); \\ (u, w) \leftarrow \langle \mathcal{P}^*(\text{fpk}, \text{st}), \mathcal{V}(\text{fvk}) \rangle(\vec{u}_1, \vec{u}_2) \end{array} \right] - \text{negl}(\lambda).$$

- **Efficiency.** *We consider MFS non-trivial if the communication costs and \mathcal{V} 's computation are lower in the case where \mathcal{V} participates in the multi-folding scheme and then checks a witness sent by \mathcal{P} for the folded instance than the case where \mathcal{V} checks witnesses sent by \mathcal{P} for each of the original instances.*

Construction 1 (Non-interactive Multi-folding Schemes [23]). *A public-coin multi-folding scheme where all the messages sent from \mathcal{V} to \mathcal{P} are chosen uniformly at random could be transformed into a non-interactive multi-folding scheme $\text{NIMFS} = (\mathcal{G}', \mathcal{K}', \mathcal{P}', \mathcal{V}')$ using the Fiat-Shamir transformation [18] where*

- $\text{fpp}' \leftarrow \mathcal{G}'(1^\lambda)$ takes 1^λ as input, and outputs some public parameter fpp' ;
- $(\text{fpk}', \text{fvk}') \leftarrow \mathcal{K}'(\text{fpp}', (s_1, s_2))$ takes fpp' and structures s_1, s_2 as input, and outputs a prover key fpk' and a verifier key fvk' ;
- $(u, w, \pi) \leftarrow \mathcal{P}'(\text{fpk}', (\vec{u}_1, \vec{w}_1), (\vec{u}_2, \vec{w}_2))$ takes fpk' , (\vec{u}_1, \vec{w}_1) , (\vec{u}_2, \vec{w}_2) as input, and outputs a folded instance-witness pair (u, w) and a proof π ;
- $u \leftarrow \mathcal{V}'(\text{fvk}', \vec{u}_1, \vec{u}_2, \pi)$ takes fvk' , $\vec{u}_1, \vec{u}_2, \pi$ as input, and outputs a folded instance u .

2.4 (Linearized) Committed Customizable Constraint System

Setty et al. [29] introduced customizable constraint systems (CCS), a new constraint system that generalizes arithmetic circuits. Later, Kothapalli and Setty [23] described two variants of CCS—committed customizable constraint systems (CCCS) and linearized committed customizable constraint systems (LCCCS)—that are amenable to construct their multi-folding schemes. In particular, they let \mathcal{R}_1 be the LCCCS relation $\mathcal{R}_{\text{LCCCS}}$ and \mathcal{R}_2 be the CCCS relation $\mathcal{R}_{\text{CCCS}}$, a NP-complete relation. We follow their work and adapt their definition of CCCS and LCCCS.

For a matrice $M \in \mathbb{F}^{m \times n}$ where $m, n \in \mathbb{N}$ and let $s = \log_2 m$, $s' = \log_2 n$, interpret it as a function $\{0, 1\}^s \times \{0, 1\}^{s'} \rightarrow \mathbb{F}$. Then we could define its unique multilinear extension \widetilde{M} as a multilinear polynomial in $s + s'$ variables such that for all $x \in \{0, 1\}^s$, $y \in \{0, 1\}^{s'}$, $\widetilde{M}(x, y) = M(x, y)$. Similarly, for a vector $w \in \mathbb{F}^m$, interpret it as a function $\{0, 1\}^s \rightarrow \mathbb{F}$. Then we let \widetilde{w} denote its unique multilinear extension.

Definition 3 (Committed Customizable Constraint System [23]). *A CCCS relation $\mathcal{R}_{\text{CCCS}}$ consists of tuples containing public parameters, structure, instance and witness where*

- public parameters consist of size bounds $m, n, N, \ell, t, q, d \in \mathbb{N}$ and pp_{PCS} , where assume that $n = 2 \cdot (\ell + 1)$ for simplicity, $\text{pp}_{\text{PCS}} \leftarrow \text{Setup}(1^\lambda, \log_2 n - 1)$;
- the structure consists of:
 - a sequence of sparse multilinear polynomials in $s + s'$ variables $\widetilde{M}_1, \dots, \widetilde{M}_t$ such that they evaluate to a non-zero value in at most $N = \Omega(\max(m, n))$ locations over the Boolean hypercube $\{0, 1\}^s \times \{0, 1\}^{s'}$, where assume that $s = \log_2 m, s' = \log_2 n$, and for $i \in [t]$, \widetilde{M}_i is the unique multilinear extension of matrix $M_i \in \mathbb{F}^{m \times n}$;
 - a sequence of q multisets (S_1, \dots, S_q) , where an element in each multiset is from the set $[t]$ and the cardinality of each multiset is at most d ;
 - a sequence of q constants (c_1, \dots, c_q) , where each constant is from \mathbb{F} ;
- the instance is (C, \mathbf{x}) , where C is a commitment to a multilinear polynomial in $s' - 1$ variables and $\mathbf{x} \in \mathbb{F}^\ell$;
- the witness is a multilinear polynomial \widetilde{w} in $s' - 1$ variables, where \widetilde{w} is the unique multilinear extension of vector $w \in \mathbb{F}^{\ell+1}$.

Given public parameters, a $\mathcal{R}_{\text{CCCS}}$ structure-instance tuple is satisfied by a $\mathcal{R}_{\text{CCCS}}$ witness if $C = \text{Com}(\text{pp}_{\text{PCS}}, \widetilde{w})$ and for all $x \in \{0, 1\}^s$,

$$\sum_{i=1}^q c_i \cdot \left(\prod_{j \in S_i} \left(\sum_{y \in \{0, 1\}^{s'}} \widetilde{M}_j(x, y) \cdot \widetilde{z}(y) \right) \right) = 0,$$

where \widetilde{z} is the unique multilinear extension of vector $(1, \mathbf{x}, w) \in \mathbb{F}^n$, i.e., an s' -variate multilinear polynomial such that $\widetilde{z}(y) = \widetilde{(1, \mathbf{x}, w)}(y)$ for all $y \in \{0, 1\}^{s'}$.

Definition 4 (Linearized Committed Customizable Constraint System [23]). A $\mathcal{L}_{\text{CCCS}}$ relation $\mathcal{R}_{\text{LCCCS}}$ consists of tuples containing public parameters, structure, instance and witness where the public parameters and structure are the same as those in a CCCS relation. The instance and witness are as follows:

- the instance is $(C, u, \mathbf{x}, r, v_1, \dots, v_t)$, where $u \in \mathbb{F}, \mathbf{x} \in \mathbb{F}^\ell, r \in \mathbb{F}^s, v_i \in \mathbb{F}$ for all $i \in [t]$, and C is a commitment to a multilinear polynomial in $s' - 1$ variables;
- the witness is a multilinear polynomial \widetilde{w} in $s' - 1$ variables, where \widetilde{w} is the unique multilinear extension of vector $w \in \mathbb{F}^{\ell+1}$.

Given public parameters, a $\mathcal{R}_{\text{LCCCS}}$ structure-instance tuple is satisfied by a $\mathcal{R}_{\text{LCCCS}}$ witness if $C = \text{Com}(\text{pp}_{\text{PCS}}, \widetilde{w})$ and for all $i \in [t]$,

$$v_i = \sum_{y \in \{0, 1\}^{s'}} \widetilde{M}_i(r, y) \cdot \widetilde{z}(y),$$

where \widetilde{z} is the unique multilinear extension of vector $(u, \mathbf{x}, w) \in \mathbb{F}^n$, i.e., an s' -variate multilinear polynomial such that $\widetilde{z}(y) = \widetilde{(u, \mathbf{x}, w)}(y)$ for all $y \in \{0, 1\}^{s'}$.

2.5 Proof-Carrying Data

We refer to the definition of Bünz et al [7]. Define a transcript T as a directed acyclic graph where each vertex $v \in V(\mathsf{T})$ is labeled by local data $z_{\text{loc}}^{(v)}$ and each edge $e \in E(\mathsf{T})$ is labeled by a message

$z^{(e)} \neq \perp$. The output $\mathbf{o}(\mathsf{T})$ of a transcript T is $z^{(e)}$ where $e = (v, v')$ is the lexicographically-first edge such that v' is a sink. For a class of compliance predicates F , define that a vertex $v \in V(\mathsf{T})$ is φ -compliant for $\varphi \in \mathsf{F}$ if for all outgoing edges $e = (v, v') \in E(\mathsf{T})$:

- (base case) if v has no incoming edges, $\varphi(z^{(e)}, z_{\text{loc}}^{(v)}, \perp, \dots, \perp)$ accepts;
- (recursive case) if v has incoming edges e_1, \dots, e_r , $\varphi(z^{(e)}, z_{\text{loc}}^{(v)}, z^{(e_1)}, \dots, z^{(e_r)})$ accepts.

We say that T is φ -compliant if all of its vertices are φ -compliant.

Definition 5 (Proof-Carrying Data [7]). A proof-carrying data scheme for a class of compliance predicates F is a tuple of algorithms $\text{PCD} = (\mathsf{G}, \mathsf{K}, \mathsf{P}, \mathsf{V})$ where

- $\text{pp} \leftarrow \mathsf{G}(1^\lambda)$ takes the security parameter λ with the unary form as input, and outputs some public parameter pp ;
- $(\text{pk}, \text{vk}) \leftarrow \mathsf{K}(\text{pp}, \varphi)$ takes pp and a compliance predicate $\varphi \in \mathsf{F}$ as input, and outputs a prover key pk and a verifier key vk ;
- $\Pi \leftarrow \mathsf{P}(\text{pk}, z, z_{\text{loc}}, \{z_i, \Pi_i\}_{i=1}^r)$ takes pk , message z of the outgoing edge, local data z_{loc} , messages $\{z_i\}_{i=1}^r$ of incoming edges and their corresponding proofs $\{\Pi_i\}_{i=1}^r$ as input, and outputs a new proof Π to attest the correctness of z ;
- $0/1 \leftarrow \mathsf{V}(\text{vk}, z, \Pi)$ takes vk, z, Π as input, and outputs $0/1$ to reject or accept.

A proof-carrying data scheme PCD should satisfy perfect completeness and knowledge soundness defined below. Note that we do not require it to satisfy zero-knowledge here as our PCD scheme needs an auxiliary $zk\text{SNARK}$ to achieve zero-knowledge which is not the main focus of this paper.

- **Perfect Completeness.** PCD has perfect completeness if for every adversary \mathcal{A} ,

$$\Pr \left[\mathsf{V}(\text{vk}, z, \Pi) = 1 \mid \begin{array}{l} \text{pp} \leftarrow \mathsf{G}(1^\lambda); \\ (\varphi, z, z_{\text{loc}}, \{z_i, \Pi_i\}_{i=1}^r) \leftarrow \mathcal{A}(\text{pp}); \\ (\text{pk}, \text{vk}) \leftarrow \mathsf{K}(\text{pp}, \varphi); \\ \varphi \in \mathsf{F}; \varphi(z, z_{\text{loc}}, z_1, \dots, z_r) = 1; \\ \forall i \in [r], z_i = \perp \text{ or } \mathsf{V}(\text{vk}, z_i, \Pi_i) = 1; \\ \Pi \leftarrow \mathsf{P}(\text{pk}, z, z_{\text{loc}}, \{z_i, \Pi_i\}_{i=1}^r) \end{array} \right] = 1.$$

- **Knowledge Soundness.** PCD has knowledge soundness (w.r.t. an auxiliary input distribution \mathcal{D}) if for every expected polynomial time adversary P^* , there exists an expected polynomial time extractor $\mathcal{E}_{\mathsf{P}^*}$ such that for every set Z ,

$$\Pr \left[\begin{array}{l} \varphi \in \mathsf{F} \\ \wedge (\text{pp}, \text{ai}, \varphi, \mathbf{o}(\mathsf{T}), \text{ao}) \in Z \\ \wedge \mathsf{T} \text{ is } \varphi\text{-compliant} \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \mathsf{G}(1^\lambda); \\ \text{ai} \leftarrow \mathcal{D}(\text{pp}); \\ (\varphi, \mathsf{T}, \text{ao}) \leftarrow \mathcal{E}_{\mathsf{P}^*}(\text{pp}, \text{ai}) \end{array} \right] \geq$$

$$\Pr \left[\begin{array}{l} \varphi \in \mathsf{F} \\ \wedge (\text{pp}, \text{ai}, \varphi, \mathbf{o}, \text{ao}) \in Z \\ \wedge \mathsf{V}(\text{vk}, \mathbf{o}, \Pi) = 1 \end{array} \mid \begin{array}{l} \text{pp} \leftarrow \mathsf{G}(1^\lambda); \\ \text{ai} \leftarrow \mathcal{D}(\text{pp}); \\ (\varphi, \mathbf{o}, \Pi, \text{ao}) \leftarrow \mathsf{P}^*(\text{pp}, \text{ai}); \\ (\text{pk}, \text{vk}) \leftarrow \mathsf{K}(\text{pp}, \varphi) \end{array} \right] - \text{negl}(\lambda).$$

3 A Multi-folding Scheme for Arbitrary Number of Instances

Recall that a multi-folding scheme for $(\mathcal{R}_1, \mathcal{R}_2, \text{compat}, \mu, \nu)$ allow a prover and a verifier to reduce the task of checking μ instances in \mathcal{R}_1 with structure \mathbf{s}_1 and ν instances in \mathcal{R}_2 with structure \mathbf{s}_2 into the task of checking a single instance in \mathcal{R}_1 with structure \mathbf{s}_1 . In our case, we let $\mathcal{R}_1 = \mathcal{R}_{\text{LCCCS}}$, $\mathcal{R}_2 = \mathcal{R}_{\text{CCCS}}$, and $\text{compat}(\mathbf{s}_1, \mathbf{s}_2)$ require $\mathbf{s}_1 = \mathbf{s}_2$. Kothapalli and Setty [23] constructed a multi-folding scheme for $(\mathcal{R}_{\text{LCCCS}}, \mathcal{R}_{\text{CCCS}}, \text{compat}, 1, 1)$. We continue their work and construct a multi-folding scheme for arbitrary values of μ, ν , which is the basis of our proof-carrying data scheme.

3.1 Overview

Suppose that there are μ $\mathcal{R}_{\text{LCCCS}}$ instance-witness pairs $\{\phi_k\}_{k \in [\mu]}$, ν $\mathcal{R}_{\text{CCCS}}$ instance-witness pairs $\{\psi_{k'}\}_{k' \in [\nu]}$, where $\phi_k := (C_1, u, \mathbf{x}_1, r_x, v_1, \dots, v_t; \widetilde{w}_1)$, $\psi_{k'} := (C_2, \mathbf{x}_2; \widetilde{w}_2)$. Let $\phi_k \cdot \widetilde{z}_1 := (\phi_k \cdot u, \widetilde{\phi_k \cdot \mathbf{x}_1}, \phi_k \cdot w_1)$, $\psi_{k'} \cdot \widetilde{z}_2 := (1, \widetilde{\psi_{k'} \cdot \mathbf{x}_2}, \psi_{k'} \cdot w_2)$. The prover \mathcal{P} and the verifier \mathcal{V} both input μ $\mathcal{R}_{\text{LCCCS}}$ instances $\{\phi_k \cdot (C_1, u, \mathbf{x}_1, r_x, v_1, \dots, v_t)\}_{k \in [\mu]}$ and ν $\mathcal{R}_{\text{CCCS}}$ instances $\{\psi_{k'} \cdot (C_2, \mathbf{x}_2)\}_{k' \in [\nu]}$. \mathcal{P} additionally inputs the corresponding witnesses $\{\phi_k \cdot \widetilde{w}_1\}_{k \in [\mu]}$, $\{\psi_{k'} \cdot \widetilde{w}_2\}_{k' \in [\nu]}$. To obtain the folded instance and witness, we rely on the random linear combination technique.

For $k \in [\mu]$, according to $\mathcal{R}_{\text{LCCCS}}$, for $i \in [t]$, we have

$$\begin{aligned} \phi_k \cdot v_i &= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(\phi_k \cdot r_x, y) \cdot \phi_k \cdot \widetilde{z}_1(y) \\ &= \sum_{x \in \{0,1\}^s} \widetilde{e}q(\phi_k \cdot r_x, x) \cdot \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(x, y) \cdot \phi_k \cdot \widetilde{z}_1(y) \right) \quad \text{By Lemma 2} \end{aligned}$$

We first perform a random linear combination of these values $\{\phi_k \cdot v_i\}_{k \in [\mu], i \in [t]}$, i.e., for $\gamma \xleftarrow{\$} \mathbb{F}$,

$$\sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \phi_k \cdot v_i = \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot \sum_{x \in \{0,1\}^s} L_{k,i}(x), \quad (1)$$

where $L_{k,i}(X) := \widetilde{e}q(\phi_k \cdot r_x, X) \cdot \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(X, y) \cdot \phi_k \cdot \widetilde{z}_1(y) \right)$.

For $k' \in [\nu]$, according to $\mathcal{R}_{\text{CCCS}}$, we have for all $x \in \{0,1\}^s$,

$$\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(x, y) \cdot \psi_{k'} \cdot \widetilde{z}_2(y) \right) = 0.$$

Denoting the left-hand side of the above equation as a polynomial $q_{k'}(X)$, we have for $x \in \{0,1\}^s$, $q_{k'}(x) = 0$. Then the polynomial

$$G_{k'}(X) := \sum_{x \in \{0,1\}^s} \widetilde{e}q(X, x) \cdot q_{k'}(x)$$

is multilinear and vanishes on all $x \in \{0,1\}^s$. Hence, we have $G_{k'}(X)$ is a zero polynomial. For $\beta \xleftarrow{\$} \mathbb{F}^s$, $G_{k'}(\beta) = 0$. Let $Q_{k'}(X) := \widetilde{e}q(\beta, X) \cdot q_{k'}(X)$. We have for $k' \in [\nu]$, $\sum_{x \in \{0,1\}^s} Q_{k'}(x) =$

$G_{k'}(\beta) = 0$. Then based on Equation (1), we further perform a random linear combination that

$$\sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \phi_k \cdot v_i = \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot \sum_{x \in \{0,1\}^s} L_{k,i}(x) + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \cdot \sum_{x \in \{0,1\}^s} Q_{k'}(x).$$

Let

$$g(X) := \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot L_{k,i}(X) + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \cdot Q_{k'}(X).$$

We have

$$\sum_{x \in \{0,1\}^s} g(x) = \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \phi_k \cdot v_i,$$

which is exactly a statement that a sum-check protocol could prove. Therefore, the task of checking μ $\mathcal{R}_{\text{LCCCS}}$ instances and ν $\mathcal{R}_{\text{CCCS}}$ instances is reduced into the task of performing a sum-check protocol. With $r'_x \xleftarrow{\$} \mathbb{F}^s$, \mathcal{P} and \mathcal{V} run $c \leftarrow \langle \mathcal{P}, \mathcal{V}(r'_x) \rangle (g, s, d+1, \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \phi_k \cdot v_i)$. Now, \mathcal{V} has to check the equation $g(r'_x) = c$. We let \mathcal{P} first compute some intermediate values and then \mathcal{V} compute $g(r'_x)$ using these values. Specifically, \mathcal{P} computes and sends $\{\sigma_{k,i}, \theta_{k',i}\}_{k \in [\mu], k' \in [\nu], i \in [t]}$ to \mathcal{V} , where

$$\sigma_{k,i} := \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \phi_k \cdot \widetilde{z}_1(y), \quad \theta_{k',i} := \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \psi_{k'} \cdot \widetilde{z}_2(y). \quad (2)$$

Then \mathcal{V} computes

$$g(r'_x) := \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot \widetilde{e}q(\phi_k \cdot r_x, r'_x) \cdot \sigma_{k,i} + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \cdot \widetilde{e}q(\beta, r'_x) \cdot \left(\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \theta_{k',j} \right)$$

and compares it with c . Nevertheless, \mathcal{V} still has to check the correctness of $\{\sigma_{k,i}, \theta_{k',i}\}_{k \in [\mu], k' \in [\nu], i \in [t]}$. Observe that equations in (2) are exactly LCCCS relations. Thus, we could first perform a random linear combination of $\{\sigma_{k,i}, \theta_{k',i}\}_{k \in [\mu], k' \in [\nu]}$ for $i \in [t]$ and then prove the correctness of the folded $\mathcal{R}_{\text{LCCCS}}$ instance. Now, the task of checking μ $\mathcal{R}_{\text{LCCCS}}$ instances and ν $\mathcal{R}_{\text{CCCS}}$ instances is reduced into the task of checking a single $\mathcal{R}_{\text{LCCCS}}$ instance.

3.2 Formal Description

Construction 2 (A Multi-folding Scheme for Arbitrary Number of Instances). We formally present our multi-folding scheme as follows.

$\text{fpp} \leftarrow \mathcal{G}(1^\lambda)$:

1. Sample size bounds $m, n, N, \ell, t, q, d \in \mathbb{N}$ with $n = 2 \cdot (\ell + 1)$.
2. Compute $\text{pp}_{\text{pcs}} \leftarrow \text{Setup}(1^\lambda, \log_2 n - 1)$, and output $\text{fpp} := (m, n, N, \ell, t, q, d, \text{pp}_{\text{pcs}})$.

$(\text{fpk}, \text{fvk}) \leftarrow \mathcal{K}(\text{fpp}, (s_1, s_2))$:

1. Parse $s_1 = s_2$ as $((\widetilde{M}_1, \dots, \widetilde{M}_t), (S_1, \dots, S_q), (c_1, \dots, c_q))$.

2. Output $\text{fpk} := (\text{fpp}, \mathbf{s}_1)$, $\text{fvk} := (\text{fpp}, \mathbf{s}_1)$.

$$(\mathbf{u}, \mathbf{w}) \leftarrow \langle \mathcal{P}(\text{fpk}, \vec{\mathbf{w}}_1, \vec{\mathbf{w}}_2), \mathcal{V}(\text{fvk}) \rangle(\vec{\mathbf{u}}_1, \vec{\mathbf{u}}_2)$$

1. Parse $(\vec{\mathbf{u}}_1, \vec{\mathbf{w}}_1)$ as $\{\phi_k\}_{k \in [\mu]}$, $(\vec{\mathbf{u}}_2, \vec{\mathbf{w}}_2)$ as $\{\psi_{k'}\}_{k' \in [\nu]}$.

2. \mathcal{V} samples $\gamma \xleftarrow{\$} \mathbb{F}, \beta \xleftarrow{\$} \mathbb{F}^s$ and sends them to \mathcal{P} . \mathcal{V} then samples $r'_x \xleftarrow{\$} \mathbb{F}^s$.

3. \mathcal{P} and \mathcal{V} run the sum-check protocol

$$c \leftarrow \langle \mathcal{P}, \mathcal{V}(r'_x) \rangle(g, s, d + 1, \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \phi_k \cdot v_i).$$

\mathcal{V} aborts once he outputs “reject” in the sum-check protocol.

4. \mathcal{P} computes and sends $\{\sigma_{k,i}, \theta_{k',i}\}_{k \in [\mu], k' \in [\nu], i \in [t]}$ to \mathcal{V} .

5. \mathcal{V} computes $\{e_{k,1} := \tilde{e}q(\phi_k \cdot r_x, r'_x)\}_{k \in [\mu]}, e_2 := \tilde{e}q(\beta, r'_x)$, and aborts if

$$c \neq \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot e_{k,1} \cdot \sigma_{k,i} + \sum_{k'=1}^{\nu} \gamma^{\mu t + k'} \cdot e_2 \cdot \left(\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \theta_{k',j} \right).$$

6. \mathcal{V} samples $\rho \xleftarrow{\$} \mathbb{F}$ and sends it to \mathcal{P} .

7. \mathcal{P} and \mathcal{V} output the folded $\mathcal{R}_{\text{LCCCS}}$ instance $\mathbf{u} := (C', u', x', r'_x, v'_1, \dots, v'_t)$, where for all $i \in [t]$,

$$\begin{aligned} C' &:= \sum_{k=1}^{\mu} \rho^{k-1} \cdot \phi_k \cdot C_1 + \sum_{k'=1}^{\nu} \rho^{\mu-1+k'} \cdot \psi_{k'} \cdot C_2, \\ u' &:= \sum_{k=1}^{\mu} \rho^{k-1} \cdot \phi_k \cdot u + \sum_{k'=1}^{\nu} \rho^{\mu-1+k'} \cdot 1, \\ x' &:= \sum_{k=1}^{\mu} \rho^{k-1} \cdot \phi_k \cdot x_1 + \sum_{k'=1}^{\nu} \rho^{\mu-1+k'} \cdot \psi_{k'} \cdot x_2, \\ v'_i &:= \sum_{k=1}^{\mu} \rho^{k-1} \cdot \sigma_{k,i} + \sum_{k'=1}^{\nu} \rho^{\mu-1+k'} \cdot \theta_{k',i}. \end{aligned}$$

8. \mathcal{P} outputs the folded $\mathcal{R}_{\text{LCCCS}}$ witness $\mathbf{w} := \tilde{w}'$, where

$$\tilde{w}' := \sum_{k=1}^{\mu} \rho^{k-1} \cdot \phi_k \cdot \tilde{w}_1 + \sum_{k'=1}^{\nu} \rho^{\mu-1+k'} \cdot \psi_{k'} \cdot \tilde{w}_2.$$

Theorem 1. *Construction 2 is a public-coin multi-folding scheme for $(\mathcal{R}_{\text{LCCCS}}, \mathcal{R}_{\text{CCCS}}, \text{compat}, \mu, \nu)$ with perfect completeness and knowledge soundness.*

Proof. We now describe the proof of Theorem 1. We mainly refer to the proof of the multi-folding scheme of Kothapalli and Setty [23].

Perfect Completeness. For public parameters $\text{fpp} = (m, n, N, \ell, t, q, d, \text{pp}_{\text{PCS}})$ and $s = \log_2 m$, $s' = \log_2 n$, consider arbitrary adversarially chosen structure $\mathbf{s}_1 = \mathbf{s}_2 = ((\widetilde{M}_1, \dots, \widetilde{M}_t), (S_1, \dots, S_q), (c_1, \dots, c_q))$, μ satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pairs $\{\phi_k\}_{k \in [\mu]}$ and ν satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pairs $\{\psi_{k'}\}_{k' \in [\nu]}$. We show that \mathcal{V} will not abort and the folded $\mathcal{R}_{\text{LCCCS}}$ instance $(C', u', x', r'_x, v'_1, \dots, v'_t)$ is satisfied by the folded $\mathcal{R}_{\text{LCCCS}}$ witness \widetilde{w}' .

Firstly, since $\{\phi_k\}_{k \in [\mu]}$ are satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pairs, we have for $k \in [\mu], i \in [t]$ and $\phi_k.\widetilde{z}_1 = (\phi_k.u, \phi_k.x_1, \phi_k.w_1)$,

$$\begin{aligned} \phi_k.v_i &= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(\phi_k.r_x, y) \cdot \phi_k.\widetilde{z}_1(y) && \text{By precondition} \\ &= \sum_{x \in \{0,1\}^s} \widetilde{e}q(\phi_k.r_x, x) \cdot \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(x, y) \cdot \phi_k.\widetilde{z}_1(y) \right) && \text{By Lemma 2} \\ &= \sum_{x \in \{0,1\}^s} L_{k,i}(x) && \text{By construction.} \end{aligned}$$

Furthermore, since $\{\psi_{k'}\}_{k' \in [\nu]}$ are satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pairs, we have for $k' \in [\nu]$, $\psi_{k'}.\widetilde{z}_2 = (1, \psi_{k'}.x_2, \psi_{k'}.w_2)$ and $x \in \{0,1\}^s$,

$$\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(x, y) \cdot \psi_{k'}.\widetilde{z}_2(y) \right) = 0.$$

Denoting the left-hand side of the above equation as a polynomial $q_{k'}(X)$, we have for $x \in \{0,1\}^s$, $q_{k'}(x) = 0$. Then the polynomial $G_{k'}(X) := \sum_{x \in \{0,1\}^s} \widetilde{e}q(X, x) \cdot q_{k'}(x)$ is multilinear and vanishes on all $x \in \{0,1\}^s$. Hence, we have $G_{k'}(X)$ is a zero polynomial. For $\beta \xleftarrow{\$} \mathbb{F}^s$, $G_{k'}(\beta) = 0$. By construction, we have for $k' \in [\nu]$, $\sum_{x \in \{0,1\}^s} Q_{k'}(x) = G_{k'}(\beta) = 0$.

Therefore, for $\gamma \xleftarrow{\$} \mathbb{F}$, we have

$$\begin{aligned} \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \phi_k.v_i &= \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \left(\sum_{x \in \{0,1\}^s} L_{k,i}(x) \right) + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \sum_{x \in \{0,1\}^s} Q_{k'}(x) \\ &= \sum_{x \in \{0,1\}^s} \left(\sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} L_{k,i}(x) + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} Q_{k'}(x) \right) \\ &= \sum_{x \in \{0,1\}^s} g(x). \end{aligned}$$

Thus, by the completeness of the sum-check protocol, \mathcal{V} will not output “reject” inside it. Moreover, we have

$$\begin{aligned} c &= g(r'_x) = \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} L_{k,i}(r'_x) + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} Q_{k'}(r'_x) \\ &= \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot e_{k,1} \cdot \sigma_{k,i} + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \cdot e_2 \cdot \left(\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \theta_{k',j} \right). \end{aligned}$$

We finally have that \mathcal{V} will not abort.

Secondly, by construction, we have that for $k \in [\mu], k' \in [\nu]$,

$$\begin{aligned} & (\phi_k.C_1, \phi_k.u, \phi_k.x_1, r'_x, \sigma_{k,1}, \dots, \sigma_{k,t}; \phi_k.\widetilde{w}_1), \\ & (\psi_{k'}.C_2, 1, \psi_{k'}.x_2, r'_x, \theta_{k',1}, \dots, \theta_{k',t}; \psi_{k'}.\widetilde{w}_2) \end{aligned}$$

are all satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pairs. Therefore, for $\widetilde{z}' = (u', x', w')$ and $i \in [t]$, we have that

$$\begin{aligned} v'_i &= \sum_{k=1}^{\mu} \rho^{k-1} \cdot \sigma_{k,i} + \sum_{k'=1}^{\nu} \rho^{\mu-1+k'} \cdot \theta_{k',i} && \text{By construction} \\ &= \sum_{k=1}^{\mu} \rho^{k-1} \cdot \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \phi_k.\widetilde{z}_1(y) \right) + \sum_{k'=1}^{\nu} \rho^{\mu-1+k'} \cdot \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \psi_{k'}.\widetilde{z}_2(y) \right) \\ &= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \left(\sum_{k=1}^{\mu} \rho^{k-1} \cdot \phi_k.\widetilde{z}_1(y) + \sum_{k'=1}^{\nu} \rho^{\mu-1+k'} \cdot \psi_{k'}.\widetilde{z}_2(y) \right) \\ &= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \widetilde{z}'(y). \end{aligned}$$

By the additive homomorphism property of the polynomial commitment scheme, we have that $C' = \text{Com}(\text{pp}_{\text{PCS}}, \widetilde{w}')$. Therefore, $(C', u', x', r'_x, v'_1, \dots, v'_t)$ is a satisfied $\mathcal{R}_{\text{LCCCS}}$ instance and \widetilde{w}' is the corresponding witness.

We conclude that Construction 2 has perfect completeness.

Knowledge Soundness. Consider a malicious prover \mathcal{P}^* that succeeds with probability ϵ . For public parameters $\text{fpp} = (m, n, N, \ell, t, q, d, \text{pp}_{\text{PCS}})$ and $s = \log_2 m, s' = \log_2 n$. Consider an adversary \mathcal{A} that adaptively picks structures $\mathbf{s}_1 = \mathbf{s}_2 = ((\widetilde{M}_1, \dots, \widetilde{M}_t), (S_1, \dots, S_q), (c_1, \dots, c_q))$ that satisfy compat , μ $\mathcal{R}_{\text{LCCCS}}$ instances $\{\phi_k.u_1 := \phi_k.(C_1, u, x_1, r_x, v_1, \dots, v_t)\}_{k \in [\mu]}$, ν $\mathcal{R}_{\text{CCCS}}$ instances $\{\psi_{k'}.u_2 := \psi_{k'}.(C_2, x_2)\}_{k' \in [\nu]}$ and some auxiliary state st . We construct an expected polynomial time extractor \mathcal{E} that succeeds with probability $\epsilon - \text{negl}(\lambda)$ in obtaining satisfied witnesses for the original instances as follows.

$$\left(\{\phi_k.\widetilde{w}_1\}_{k \in [\mu]}, \{\psi_{k'}.\widetilde{w}_2\}_{k' \in [\nu]} \right) \leftarrow \mathcal{E}(\text{fpp}, \rho):$$

1. Invoke \mathcal{A} to obtain the output tuple: $((\mathbf{s}_1, \mathbf{s}_2), \{\phi_k.u_1\}_{k \in [\mu]}, \{\psi_{k'}.u_2\}_{k' \in [\nu]}, \text{st}) \leftarrow \mathcal{A}(\text{fpp}, \rho)$.
2. Compute $(\text{fpk}, \text{fvk}) \leftarrow \mathcal{K}(\text{fpp}, (\mathbf{s}_1, \mathbf{s}_2))$.
3. Run the interaction $(\mathbf{u}^{(1)}, \widetilde{w}^{(1)}) \leftarrow \langle \mathcal{P}^*(\text{fpk}, \text{st}), \mathcal{V}(\text{fvk}) \rangle (\{\phi_k.u_1\}_{k \in [\mu]}, \{\psi_{k'}.u_2\}_{k' \in [\nu]})$ *once* with the final verifier challenge $\rho^{(1)} \stackrel{\$}{\leftarrow} \mathbb{F}$.
4. Abort if $(\text{fpp}, \mathbf{s}_1, \mathbf{u}^{(1)}, \widetilde{w}^{(1)}) \notin \mathcal{R}_{\text{LCCCS}}$.
5. Rerun the interaction $(\mathbf{u}^{(1)}, \widetilde{w}^{(1)}) \leftarrow \langle \mathcal{P}^*(\text{fpk}, \text{st}), \mathcal{V}(\text{fvk}) \rangle (\{\phi_k.u_1\}_{k \in [\mu]}, \{\psi_{k'}.u_2\}_{k' \in [\nu]})$ with different verifier's final challenges while maintaining the same prior randomness. Keep doing so until it obtains $\mu + \nu - 1$ folded $\mathcal{R}_{\text{LCCCS}}$ instance-witness pairs $\{(\mathbf{u}^{(k'')}, \widetilde{w}^{(k'')})\}_{k'' \in [2, \dots, \mu + \nu]}$ such that $(\text{fpp}, \mathbf{s}_1, \mathbf{u}^{(k'')}, \widetilde{w}^{(k'')}) \in \mathcal{R}_{\text{LCCCS}}$ for all $k'' \in [2, \dots, \mu + \nu]$. Let $\{\rho^{(k'')}\}$ denote the corresponding verifier's final challenges.

6. Abort if there exists a collision in the verifier's final challenges.

7. Interpolate points $\{(\rho^{(k'')}, \widetilde{w}^{(k'')})\}_{k'' \in [\mu + \nu]}$ to retrieve witnesses $(\{\phi_k \cdot \widetilde{w}_1\}_{k \in [\mu]}, \{\psi_{k'} \cdot \widetilde{w}_2\}_{k' \in [\nu]})$ such that for $k'' \in [\mu + \nu]$,

$$\sum_{k=1}^{\mu} \rho^{(k'')^{k-1}} \cdot \phi_k \cdot \widetilde{w}_1 + \sum_{k'=1}^{\nu} \rho^{(k'')^{\mu-1+k'}} \cdot \psi_{k'} \cdot \widetilde{w}_2 = \widetilde{w}^{(k'')}. \quad (3)$$

8. Output $(\{\phi_k \cdot \widetilde{w}_1\}_{k \in [\mu]}, \{\psi_{k'} \cdot \widetilde{w}_2\}_{k' \in [\nu]})$.

We now demonstrate that the extractor \mathcal{E} runs in expected polynomial time and succeeds with probability $\epsilon - \text{negl}(\lambda)$.

Firstly, the extractor \mathcal{E} runs the interaction $\langle \mathcal{P}^*, \mathcal{V} \rangle$ once, and if it does not abort, keeps rerunning the interaction until it obtains $\mu + \nu - 1$ satisfied folded $\mathcal{R}_{\text{LCCCS}}$ instance-witness pairs. Thus, the expected number of times \mathcal{E} runs the interaction is

$$\begin{aligned} & \Pr[\text{First call to } \langle \mathcal{P}^*, \mathcal{V} \rangle \text{ fails}] \cdot 1 + \\ & \Pr[\text{First call to } \langle \mathcal{P}^*, \mathcal{V} \rangle \text{ succeeds}] \cdot \left(1 + \frac{\mu + \nu - 1}{\Pr[\langle \mathcal{P}^*, \mathcal{V} \rangle \text{ succeeds}]}\right) \\ & = (1 - \epsilon) \cdot 1 + \epsilon \cdot \left(1 + \frac{\mu + \nu - 1}{\epsilon}\right) = \mu + \nu. \end{aligned}$$

Assuming that μ, ν are polynomial in the security parameter, we have that \mathcal{E} runs in expected polynomial time.

Secondly, let E_1 denote the event that \mathcal{E} successfully produces some outputs in less than T times of running the interaction $\langle \mathcal{P}^*, \mathcal{V} \rangle$. Given E_1 , let E_2 denote the event that the outputs of \mathcal{E} are satisfied witnesses. We have that

$$\Pr[\mathcal{E} \text{ succeeds}] = \Pr[E_1] \cdot \Pr[E_2].$$

We now analyze $\Pr[E_1]$ and $\Pr[E_2]$. By the success probability of \mathcal{P}^* , we have that \mathcal{E} does not abort in step (4) with probability ϵ . Given that \mathcal{E} does not abort in step (4), by Markov's inequality, we have that \mathcal{E} runs the interaction $\langle \mathcal{P}^*, \mathcal{V} \rangle$ more than T times with probability $\frac{\mu + \nu}{T}$. Given that \mathcal{E} runs the interaction less than T times, which has probability $1 - \frac{\mu + \nu}{T}$, we have \mathcal{E} tests at most T values for ρ . Thus, the probability that \mathcal{E} does not abort in step (6) is $1 - \frac{T^2}{|\mathbb{F}|}$. Thus, we have

$$\Pr[E_1] = \left(1 - \frac{\mu + \nu}{T}\right) \cdot \epsilon \cdot \left(1 - \frac{T^2}{|\mathbb{F}|}\right).$$

Setting $T = \sqrt[3]{|\mathbb{F}|}$ and assuming $T \geq \mu + \nu$, we have that $\Pr[E_1] = \epsilon - \text{negl}(\lambda)$.

To analyze $\Pr[E_2]$, we first show that the retrieved witnesses are valid openings to the corresponding polynomial commitments in the instance, and then show that they satisfy the remaining algebraic relations with some probability. For $k'' \in [\mu + \nu]$, let $\mathbf{u}^{(k'')} = (C^{(k'')}, u^{(k'')}, \mathbf{x}^{(k'')}, r'_x, v_1^{(k'')}, \dots, v_t^{(k'')})$.

Since $\widetilde{w}^{(k'')}$ is a satisfied $\mathcal{R}_{\text{LCCCS}}$ witness, we have that

$$\begin{aligned}
& \sum_{k=1}^{\mu} \rho^{(k'')^{k-1}} \cdot \text{Com}(\text{pp}_{\text{PCS}}, \phi_k \cdot \widetilde{w}_1) + \sum_{k'=1}^{\nu} \rho^{(k'')^{\mu-1+k'}} \cdot \text{Com}(\text{pp}_{\text{PCS}}, \psi_{k'} \cdot \widetilde{w}_2) \\
&= \text{Com}\left(\text{pp}_{\text{PCS}}, \sum_{k=1}^{\mu} \rho^{(k'')^{k-1}} \cdot \phi_k \cdot \widetilde{w}_1 + \sum_{k'=1}^{\nu} \rho^{(k'')^{\mu-1+k'}} \cdot \psi_{k'} \cdot \widetilde{w}_2\right) \\
&\hspace{15em} \text{By additive homomorphism} \\
&= \text{Com}(\text{pp}_{\text{PCS}}, \widetilde{w}^{(k'')}) \hspace{15em} \text{By Equation (3)} \\
&= C^{(k'')} \hspace{15em} \text{Witness } \widetilde{w}^{(k'')} \text{ is a satisfied opening} \\
&= \sum_{k=1}^{\mu} \rho^{(k'')^{k-1}} \cdot \phi_k \cdot C_1 + \sum_{k'=1}^{\nu} \rho^{(k'')^{\mu-1+k'}} \cdot \psi_{k'} \cdot C_2 \text{ By the verifier's computation}
\end{aligned}$$

Treat the above equation as a univariate polynomial equation in $\rho^{(k'')}$. Since it holds for all $k'' \in [\mu + \nu]$, we must have that for $k \in [\mu], k' \in [\nu]$,

$$\phi_k \cdot C_1 = \text{Com}(\text{pp}_{\text{PCS}}, \phi_k \cdot \widetilde{w}_1), \quad \psi_{k'} \cdot C_2 = \text{Com}(\text{pp}_{\text{PCS}}, \psi_{k'} \cdot \widetilde{w}_2),$$

which means that $(\{\phi_k \cdot \widetilde{w}_1\}_{k \in [\mu]}, \{\psi_{k'} \cdot \widetilde{w}_2\}_{k' \in [\nu]})$ are valid openings.

Next, by the extractor's construction we have that $\{\sigma_{k,i}, \theta_{k',i}\}_{k \in [\mu], k' \in [\nu], i \in [t]}$ sent by the prover are identical across all executions of the interaction $\langle \mathcal{P}^*, \mathcal{V} \rangle$. By the verifier's computation, we have that for $k'' \in [\mu + \nu], i \in [t]$,

$$v_i^{(k'')} = \sum_{k=1}^{\mu} \rho^{(k'')^{k-1}} \cdot \sigma_{k,i} + \sum_{k'=1}^{\nu} \rho^{(k'')^{\mu-1+k'}} \cdot \theta_{k',i}. \quad (4)$$

Since $\{\widetilde{w}^{(k'')}\}_{k'' \in [\mu+\nu]}$ are satisfied $\mathcal{R}_{\text{LCCCS}}$ witnesses, we have that for $k'' \in [\mu + \nu], i \in [t]$,

$$v_i^{(k'')} = \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \widetilde{z}^{(k'')}(y), \quad (5)$$

where $\widetilde{z}^{(k'')} = (u^{(k'')}, x^{(k'')}, \widetilde{w}^{(k'')})$. By Equations (3), (4), (5) and the verifier's computation, we have that for $k'' \in [\mu + \nu], i \in [t]$,

$$\begin{aligned}
& \sum_{k=1}^{\mu} \rho^{(k'')^{k-1}} \cdot \sigma_{k,i} + \sum_{k'=1}^{\nu} \rho^{(k'')^{\mu-1+k'}} \cdot \theta_{k',i} \\
&= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \widetilde{z}^{(k'')}(y) \\
&= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \left(\sum_{k=1}^{\mu} \rho^{(k'')^{k-1}} \cdot \phi_k \cdot \widetilde{z}_1(y) + \sum_{k'=1}^{\nu} \rho^{(k'')^{\mu-1+k'}} \cdot \psi_{k'} \cdot \widetilde{z}_2(y) \right) \\
&= \sum_{k=1}^{\mu} \rho^{(k'')^{k-1}} \cdot \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \phi_k \cdot \widetilde{z}_1(y) + \sum_{k'=1}^{\nu} \rho^{(k'')^{\mu-1+k'}} \cdot \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \psi_{k'} \cdot \widetilde{z}_2(y),
\end{aligned}$$

where $\phi_k.\tilde{z}_1 = (\phi_k.u, \widetilde{\phi_k.x_1}, \phi_k.w_1)$ and $\psi_{k'}.\tilde{z}_2 = (1, \widetilde{\psi_{k'.x_2}}, \psi_{k'.w_2})$. Treat the above equation as a univariate polynomial equation in $\rho^{(k')}$. Since it holds for all $k'' \in [\mu + \nu]$, we must have that for $k \in [\mu], k' \in [\nu], i \in [t]$,

$$\sigma_{k,i} = \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \phi_k.\tilde{z}_1(y), \quad \theta_{k',i} = \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \psi_{k'}.\tilde{z}_2(y).$$

Then since the verifier does not abort, we have that

$$\begin{aligned} c &= \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot e_{k,1} \cdot \sigma_{k,i} + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \cdot e_2 \cdot \left(\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \theta_{k',j} \right) \\ &= \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot \tilde{e}q(\phi_k.r_x, r'_x) \cdot \sigma_{k,i} + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \cdot \tilde{e}q(\beta, r'_x) \cdot \left(\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \theta_{k',j} \right) \\ &= \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot L_{k,i}(r'_x) + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \cdot Q_{k'}(r'_x) \\ &= g(r'_x), \end{aligned}$$

which by the soundness of the sum-check protocol, implies that with probability $1 - O(s \cdot d)/|\mathbb{F}| = 1 - \text{negl}(\lambda)$ over the choice of r'_x ,

$$\begin{aligned} \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \phi_k.v_i &= \sum_{x \in \{0,1\}^s} g(x) \\ &= \sum_{x \in \{0,1\}^s} \left(\sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot L_{k,i}(x) + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \cdot Q_{k'}(x) \right) \\ &= \sum_{k=1}^{\mu} \sum_{i=1}^t \gamma^{(k-1)t+i} \cdot \sum_{x \in \{0,1\}^s} L_{k,i}(x) + \sum_{k'=1}^{\nu} \gamma^{\mu t+k'} \cdot \sum_{x \in \{0,1\}^s} Q_{k'}(x). \end{aligned}$$

By the Schwartz-Zippel lemma [28], this implies that with probability $1 - O(\mu \cdot t + \nu)/|\mathbb{F}| = 1 - \text{negl}(\lambda)$ over the choice of γ , we have for $k \in [\mu], i \in [t]$,

$$\phi_k.v_i = \sum_{x \in \{0,1\}^s} L_{k,i}(x), \tag{6}$$

and for $k' \in [\nu]$,

$$0 = \sum_{x \in \{0,1\}^s} Q_{k'}(x). \tag{7}$$

Expanding Equation (6), we have that

$$\begin{aligned} \phi_k.v_i &= \sum_{x \in \{0,1\}^s} L_{k,i}(x) \\ &= \sum_{x \in \{0,1\}^s} \tilde{e}q(\phi_k.r_x, x) \cdot \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(x, y) \cdot \phi_k.\tilde{z}_1(y) \right) \\ &= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(\phi_k.r_x, y) \cdot \phi_k.\tilde{z}_1(y). \end{aligned}$$

By Lemma 2

Since we have argued that $\{\phi_k \cdot \widetilde{w}_1\}_{k \in [\mu]}$ are valid openings, we have that they are satisfied $\mathcal{R}_{\text{LCCCS}}$ witnesses with probability $1 - \text{negl}(\lambda)$.

Expanding Equation (7), we have that

$$\begin{aligned} 0 &= \sum_{x \in \{0,1\}^s} Q_{k'}(x) \\ &= \sum_{x \in \{0,1\}^s} \widetilde{e}q(\beta, x) \cdot \left(\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(x, y) \cdot \psi_{k'} \cdot \widetilde{z}_2(y) \right) \right) \\ &= \sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(\beta, y) \cdot \psi_{k'} \cdot \widetilde{z}_2(y) \right). \end{aligned}$$

By the Schwartz-Zippel lemma, this implies that with probability $1 - O(s \cdot d)/|\mathbb{F}| = 1 - \text{negl}(\lambda)$ over the choice of β , we have that for all $x \in \{0,1\}^s$,

$$\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(x, y) \cdot \psi_{k'} \cdot \widetilde{z}_2(y) \right) = 0.$$

Since we have argued that $\{\psi_{k'} \cdot \widetilde{w}_2\}_{k' \in [\nu]}$ are valid openings, we have that they are satisfied $\mathcal{R}_{\text{CCCS}}$ witnesses with probability $1 - \text{negl}(\lambda)$. Hence, we have that $\Pr[E_2] = 1 - \text{negl}(\lambda)$.

Therefore, we have that

$$\Pr[\mathcal{E} \text{ succeeds}] = \Pr[E_1] \cdot \Pr[E_2] = (\epsilon - \text{negl}(\lambda)) \cdot (1 - \text{negl}(\lambda)) = \epsilon - \text{negl}(\lambda).$$

We conclude that Construction 2 has knowledge soundness. \square

A Non-interactive Multi-folding Scheme. Since our multi-folding scheme is public-coin, we could transform it into a non-interactive multi-folding scheme $\text{NIMFS} = (\mathcal{G}', \mathcal{K}', \mathcal{P}', \mathcal{V}')$ for the tuple $(\mathcal{R}_{\text{LCCCS}}, \mathcal{R}_{\text{CCCS}}, \text{compat}, \mu, \nu)$ using the Fiat-Shamir transformation, according to the Lemma 3 of Kothapalli and Setty [23].

Efficiency. For the prover's cost, the sum-check protocol requires the prover to generate some proof, which could be completed with $O(\mu(N + tm) + \nu(N + tm + qmd \log^2 d))$ field operations according to [29]. For $\{\sigma_{k,i}, \theta_{k',i}\}_{k \in [\mu], k' \in [\nu], i \in [t]}$, we have that

$$\begin{aligned} \sigma_{k,i} &:= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \phi_k \cdot \widetilde{z}_1(y) \\ &= \sum_{y \in \{0,1\}^{s'}} \left(\sum_{a \in \{0,1\}^s} \sum_{b \in \{0,1\}^{s'}} M_i(a, b) \cdot \widetilde{e}q((r'_x, y), (a, b)) \right) \cdot \phi_k \cdot \widetilde{z}_1(y) \quad \text{By multilinear extension} \\ &= \sum_{y \in \{0,1\}^{s'}} \left(\sum_{a \in \{0,1\}^s} \widetilde{e}q(r'_x, a) \sum_{b \in \{0,1\}^{s'}} M_i(a, b) \cdot \widetilde{e}q(y, b) \right) \cdot \phi_k \cdot \widetilde{z}_1(y) \\ &= \sum_{a \in \{0,1\}^s} \widetilde{e}q(r'_x, a) \cdot \sum_{b \in \{0,1\}^{s'}} M_i(a, b) \cdot \phi_k \cdot \widetilde{z}_1(b) \quad \text{By the definition of } \widetilde{e}q \end{aligned}$$

According to [33], $\{\tilde{e}q(r'_x, a)\}_{a \in \{0,1\}^s}$ could be computed with $O(m)$ field operations. Based on these values, the computation of $\{\sigma_{k,i}, \theta_{k',i}\}_{k \in [\mu], k' \in [\nu], i \in [t]}$ could be accomplished with $O(m + (\mu + \nu) \cdot N)$ field operations. When instantiating the polynomial commitment scheme with Bulletproofs [6], the computation of C' and \tilde{w}' requires one MSM of size $\mu + \nu$ and $O((\mu + \nu) \cdot n)$ field operations, respectively. Thus, the total cost is dominated by $O(\mu(N + tm + n) + \nu(N + tm + n + qmd \log^2 d))$ field operations and one MSM of size $\mu + \nu$.

For the verifier's cost, verification in the sum-check protocol requires $O(d \log m)$ field operations [30]. The computation of $\{e_{k,1} := \tilde{e}q(\phi_k \cdot r_x, r'_x)\}_{k \in [\mu]}, e_2 := \tilde{e}q(\beta, r'_x)$ and checking c require $O(\mu \log m + \mu t + \nu dq)$ field operations. The computation of C' requires one MSM of size $\mu + \nu$. Thus, the total cost is dominated by $O(d \log m + \mu \log m + \mu t + \nu dq)$ field operations and one MSM of size $\mu + \nu$.

The communication cost consists of elements sent from \mathcal{P} to \mathcal{V} in the sum-check protocol and $\{\sigma_{k,i}, \theta_{k',i}\}_{k \in [\mu], k' \in [\nu], i \in [t]}$, which sums up to $O(d \log m) + t \cdot (\mu + \nu)$ field elements.

4 PCD from Non-interactive Multi-folding Schemes

Recall that PCD enables a set of parties to carry out an indefinitely long distributed computation where every step along the way is accompanied by a proof of correctness. We rely on our non-interactive multi-folding scheme to construct PCD scheme where the computation and the circuit expressing verifier at each step are together expressed as CCCS, a NP-complete constraint system.

4.1 Overview

At each step, the prover \mathcal{P} receives r previous outputs $\{z_i\}_{i \in [r]}$ each with a proof Π_i that consists of a satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pair (U_i, W_i) and a satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pair (u_i, w_i) in our case. In addition, \mathcal{P} has some local input z_{loc} . With $z_{\text{loc}}, z_1, \dots, z_r$, \mathcal{P} computes z such that $\varphi(z, z_{\text{loc}}, z_1, \dots, z_r)$ accepts. Then he should provide a proof to the next party that attests not only to the correctness of his local computation, but also to the correctness of all his inputs $\{z_i\}_{i \in [r]}$. To this end, \mathcal{P} first invokes the non-interactive multi-folding scheme to fold $\{(U_i, W_i), (u_i, w_i)\}_{i \in [r]}$ into a single $\mathcal{R}_{\text{LCCCS}}$ instance-witness pair (U, W) . Define a circuit R_φ that represents the compliance predicate φ as well as the verifier of the non-interactive multi-folding scheme, which could be expressed as a $\mathcal{R}_{\text{CCCS}}$ structure. Then \mathcal{P} generates a satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pair (u, w) by computing R_φ on input appropriate values from his prior computations.

Let the new PCD proof Π consist of $(U, W), (u, w)$. Now, by checking (u, w) , we have that $\varphi(z, z_{\text{loc}}, z_1, \dots, z_r)$ accepts and $\{U_i, u_i\}_{i \in [r]}$ is corrected folded into U . Further by checking (U, W) and the knowledge soundness of the multi-folding scheme, we have that there exists satisfied witnesses $\{W_i, w_i\}_{i \in [r]}$ for instances $\{U_i, u_i\}_{i \in [r]}$, which attests to the correctness of $\{z_i\}_{i \in [r]}$. Therefore, we maintain the invariant that if (U, W) is a satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pair and (u, w) is a satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pair, then the distributed computation is correct up to this step.

However, the above description elides some subtlety. Since the proof Π passed to the next party includes U , the public IO of the circuit R_φ , i.e., $u.x$ should include U . The next party will fold u, U into a new $\mathcal{R}_{\text{LCCCS}}$ instance as described above. But now U is part of u , they can not be folded. To address this issue, inspired by Nova [24], we modify R_φ to additionally hash z, U and let the output be the circuit's public IO, which ensures that $u.x$ does not contain U but still binds with it. To achieve recursion, we further modify R_φ to include the verifier's work of checking that the inputs satisfy the hash relation, i.e., $u_i.x$ is the hash of z_i, U_i for $i \in [r]$.

Note that since the PCD proof includes the entire witnesses W, w , its size is linear in the size of the circuit. However, as in HyperNova [23], we could use a general SNARK such as SuperSpartan [29] to prove the knowledge of a valid PCD proof, which could reduce the proof size exponentially.

4.2 Formal Description

Construction 3 (A PCD Scheme from Non-interactive Multi-folding Schemes). We formally present our PCD scheme as follows. Let (u_\perp, w_\perp) be a default trivially satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pair. Given a compliance predicate φ , the circuit R_φ that realizes the recursion is as follows.

$0/1 \leftarrow R_\varphi(h; (z, z_{\text{loc}}, \{z_i, U_i, u_i\}_{i=1}^r, \text{fvk}', U, \pi))$:

1. Check that the compliance predicate $\varphi(z, z_{\text{loc}}, z_1, \dots, z_r)$ accepts.
2. If $z_i = \perp$ for all $i \in [r]$, then check that $h = \text{hash}(\text{fvk}', z, u_\perp)$.

Else, check that

- (a) for $i \in [r]$, $u_i.x = \text{hash}(\text{fvk}', z_i, U_i)$, where $u_i.x$ is the public IO of u_i .
- (b) $U = \text{NIMFS.V}'(\text{fvk}', \{U_i\}_{i \in [r]}, \{u_i\}_{i \in [r]}, \pi)$.
- (c) $h = \text{hash}(\text{fvk}', z, U)$.

3. If the above checks hold, output 1; otherwise, output 0.

Since R_φ can be computed in polynomial time, it can be represented as a $\mathcal{R}_{\text{CCCS}}$ structure. Let

$$(u, w) \leftarrow \text{trace}\left(R_\varphi, (h, (z, z_{\text{loc}}, \{z_i, U_i, u_i\}_{i=1}^r, \text{fvk}', U, \pi))\right)$$

denote the satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pair for the execution of R_φ on input $(h, (z, z_{\text{loc}}, \{z_i, U_i, u_i\}_{i=1}^r, \text{fvk}', U, \pi))$.

We define the PCD scheme (G, K, P, V) as follows.

$\text{pp} \leftarrow G(1^\lambda)$:

1. Compute $\text{fpp}' \leftarrow \text{NIMFS.G}'(1^\lambda)$, and output $\text{pp} := \text{fpp}'$.

$(\text{pk}, \text{vk}) \leftarrow K(\text{pp}, \varphi)$:

1. Compute $(\text{fpk}', \text{fvk}') \leftarrow \text{NIMFS.K}'(\text{fpp}', R_\varphi)$, and output $(\text{pk}, \text{vk}) := ((\text{fpk}', \text{fvk}'), \text{fvk}')$.

$\text{II} \leftarrow P(\text{pk}, z, z_{\text{loc}}, \{z_i, \text{II}_i\}_{i=1}^r)$:

1. For $i \in [r]$, parse II_i as $((U_i, W_i), (u_i, w_i))$, where (U_i, W_i) is a satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pair and (u_i, w_i) is a satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pair.
2. If $z_i = \perp$ for all $i \in [r]$, then set $(U, W, \pi) := (u_\perp, w_\perp, \perp)$.
Else, compute $(U, W, \pi) \leftarrow \text{NIMFS.P}'(\text{fpk}', \{(U_i, W_i)\}_{i \in [r]}, \{(u_i, w_i)\}_{i \in [r]})$
3. Compute $h \leftarrow \text{hash}(\text{fvk}', z, U)$.

4. $(\mathbf{u}, \mathbf{w}) \leftarrow \text{trace}\left(R_\varphi, (h, (z, z_{\text{loc}}, \{z_i, \mathbf{U}_i, \mathbf{u}_i\}_{i=1}^r, \text{fvk}', \mathbf{U}, \pi))\right)$.

5. Output $\Pi := ((\mathbf{U}, \mathbf{W}), (\mathbf{u}, \mathbf{w}))$.

$0/1 \leftarrow \mathbf{V}(\text{vk}, z, \Pi)$:

1. Parse Π as $((\mathbf{U}, \mathbf{W}), (\mathbf{u}, \mathbf{w}))$.

2. Check that $\mathbf{u}.x = \text{hash}(\text{fvk}', z, \mathbf{U})$.

3. Check that \mathbf{W} is a satisfied $\mathcal{R}_{\text{LCCCS}}$ witness to \mathbf{U} and \mathbf{w} is a satisfied $\mathcal{R}_{\text{CCCS}}$ witness to \mathbf{u} .

4. If the above checks hold, output 1; otherwise, output 0.

Theorem 2. *Construction 3 is a PCD scheme with perfect completeness and knowledge soundness.*

Proof. We now describe the proof of Theorem 2.

Perfect Completeness. For public parameter \mathbf{pp} , consider arbitrary adversarially chosen $(\varphi, z, z_{\text{loc}}, \{z_i, \Pi_i\}_{i=1}^r)$ such that the perfect completeness precondition is satisfied. We show that given $\Pi \leftarrow \mathbf{P}(\mathbf{pk}, z, z_{\text{loc}}, \{z_i, \Pi_i\}_{i=1}^r)$, $\mathbf{V}(\text{vk}, z, \Pi) = 1$ with probability 1.

If $z_i = \perp$ for all $i \in [r]$, by the construction of \mathbf{P} , we have that $(\mathbf{U}, \mathbf{W}) = (\mathbf{u}_\perp, \mathbf{w}_\perp)$ is a trivially satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pair, and $h = \text{hash}(\text{fvk}', z, \mathbf{U})$. By the perfect completeness precondition, we have that $\varphi(z, z_{\text{loc}}, z_1, \dots, z_r) = 1$. Therefore, \mathbf{P} could construct a satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pair (\mathbf{u}, \mathbf{w}) that represents the correct computation of R_φ . Moreover, by the construction of R_φ , we have $\mathbf{u}.x = \text{hash}(\text{fvk}', z, \mathbf{U})$. Therefore, $\mathbf{V}(\text{vk}, z, \Pi) = 1$ with probability 1.

If $\exists i \in [r]$ such that $z_i \neq \perp$, by the perfect completeness precondition, $\{(\mathbf{U}_i, \mathbf{W}_i)\}_{i \in [r]}$ are satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pairs, $\{(\mathbf{u}_i, \mathbf{w}_i)\}_{i \in [r]}$ are satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pairs, and $\mathbf{u}_i.x = \text{hash}(\text{fvk}', z_i, \mathbf{U}_i)$. Then by the perfect completeness of the multi-folding scheme, we have that (\mathbf{U}, \mathbf{W}) is a satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pair. Therefore, \mathbf{P} could construct a satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pair (\mathbf{u}, \mathbf{w}) that represents the correct computation of R_φ . Additionally, by the construction of R_φ , we have that $\mathbf{u}.x = \text{hash}(\text{fvk}', z, \mathbf{U})$. Therefore, $\mathbf{V}(\text{vk}, z, \Pi) = 1$ with probability 1.

In conclusion, we conclude that Construction 3 has perfect completeness.

Knowledge Soundness. We mainly refer to the proof of the PCD scheme of Bünz et al [7]. Fix a set Z , and for $\mathbf{pp} \leftarrow \mathbf{G}(1^\lambda)$, $\mathbf{ai} \leftarrow \mathcal{D}(\mathbf{pp})$, consider an expected polynomial time adversary \mathbf{P}^* that succeeds with probability ϵ . We construct an expected polynomial time extractor $\mathcal{E}_{\mathbf{P}^*}$ that with input $(\mathbf{pp}, \mathbf{ai})$, outputs $(\varphi, \mathbf{T}, \mathbf{ao})$ such that $\varphi \in \mathbf{F}$, $(\mathbf{pp}, \mathbf{ai}, \varphi, \mathbf{o}(\mathbf{T}), \mathbf{ao}) \in Z$ and \mathbf{T} is φ -compliant with probability $\epsilon - \text{negl}(\lambda)$.

Referring to existing works [8,7], we assume that every node has a unique outgoing edge. Thus, the extracted transcript \mathbf{T} will be a tree. Let every node v be labeled with a local data $z_{\text{loc}}^{(v)}$, the label $z^{(v)}$ of its unique outgoing edge and a proof $\Pi^{(v)}$ that proves the correctness of $z^{(v)}$. We first present the construction of $\mathcal{E}_{\mathbf{P}^*}$ that extracts the labels of all the nodes in \mathbf{T} , and then explain its correctness.

We construct $\mathcal{E}_{\mathbf{P}^*}$ via an iterative process that constructs a sequence of extractors $\mathcal{E}_0, \mathcal{E}_1, \dots, \mathcal{E}_d$, where d is the depth of \mathbf{T} and for $i \in \{0, 1, \dots, d\}$, \mathcal{E}_i outputs a tree \mathbf{T}_i of depth $i + 1$. In particular, the nodes of \mathbf{T}_d at depth $d + 1$ are all empty nodes. We argue later that the extractor \mathcal{E}_d is then equal to $\mathcal{E}_{\mathbf{P}^*}$.

In the base case, we define $\mathcal{E}_0(\text{pp}, \text{ai})$ to compute $(\varphi, \mathbf{o}, \Pi, \text{ao}) \leftarrow \text{P}^*(\text{pp}, \text{ai})$ and output $(\varphi, \mathbb{T}_0, \text{ao})$, where \mathbb{T}_0 is a single node labeled with (\mathbf{o}, Π) .

Next, we construct the extractor \mathcal{E}_i inductively for each $i \in [d]$, given that we have already constructed \mathcal{E}_{i-1} . We denote $S_{\mathbb{T}}(i)$ as the set of nodes of \mathbb{T} at depth i . We first construct an adversary \mathcal{P}_{i-1}^* for the non-interactive multi-folding scheme using \mathcal{E}_{i-1} , which implies an extractor $\mathcal{E}_{\mathcal{P}_{i-1}^*}$ corresponding to \mathcal{P}_{i-1}^* by the knowledge soundness of the multi-folding scheme, and then construct \mathcal{E}_i using $\mathcal{P}_{i-1}^*, \mathcal{E}_{\mathcal{P}_{i-1}^*}$.

$\mathcal{P}_{i-1}^*(\text{pp}, \text{ai}, \rho)$:

1. Compute $(\varphi, \mathbb{T}_{i-1}, \text{ao}) \leftarrow \mathcal{E}_{i-1}(\text{pp}, \text{ai})$. If \mathbb{T}_{i-1} is not a tree of depth i , abort.
2. For each node $v \in S_{\mathbb{T}_{i-1}}(i)$, denote its label as $(z^{(v)}, \Pi^{(v)})$.
3. Parse $\Pi^{(v)}$ as $((\mathbf{U}^{(v)}, \mathbf{W}^{(v)}), (\mathbf{u}^{(v)}, \mathbf{w}^{(v)}))$.
4. Obtain $\{\mathbf{U}_j^{(v)}, \mathbf{u}_j^{(v)}, z_j^{(v)}\}_{j \in [r]}, \pi^{(v)}$ from $\mathbf{w}^{(v)}$.
5. Let $S_{i-1} := \{v \in S_{\mathbb{T}_{i-1}}(i) \mid \exists j \in [r], z_j^{(v)} \neq \perp\}$.
6. Output $\left(\left\{ \{\mathbf{U}_j^{(v)}, \mathbf{u}_j^{(v)}\}_{j \in [r]}, \mathbf{U}^{(v)}, \mathbf{W}^{(v)}, \pi^{(v)} \right\}_{v \in S_{i-1}}, (\varphi, \mathbb{T}_{i-1}, \text{ao}) \right)$.

By the knowledge soundness of the multi-folding scheme, there exists an extractor $\mathcal{E}_{\mathcal{P}_{i-1}^*}$ that for $v \in S_{i-1}$, outputs $\{\mathbf{W}_j^{(v)}, \mathbf{w}_j^{(v)}\}_{j \in [r]}$ such that $\{(\mathbf{U}_j^{(v)}, \mathbf{W}_j^{(v)})\}_{j \in [r]}$ are satisfied $\mathcal{R}_{\text{LCCCS}}$ instance-witness pairs and $\{(\mathbf{u}_j^{(v)}, \mathbf{w}_j^{(v)})\}_{j \in [r]}$ are satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pairs.

Given $\mathcal{P}_{i-1}^*, \mathcal{E}_{\mathcal{P}_{i-1}^*}$, we construct \mathcal{E}_i as follows.

$(\varphi, \mathbb{T}_i, \text{ao}) \leftarrow \mathcal{E}_i(\text{pp}, \text{ai})$:

1. Compute $\left(\left\{ (\mathbf{U}_j^{(v)}, \mathbf{W}_j^{(v)}), (\mathbf{u}_j^{(v)}, \mathbf{w}_j^{(v)}) \right\}_{j \in [r], v \in S_{i-1}}, (\varphi, \mathbb{T}_{i-1}, \text{ao}) \right) \leftarrow \mathcal{E}_{\mathcal{P}_{i-1}^*}(\text{pp}, \text{ai}, \rho)$. If \mathbb{T}_{i-1} is not a tree of depth i , abort.
2. Retrieve $\{\mathbf{w}^{(v)}\}_{v \in S_{\mathbb{T}_{i-1}}(i)}$ from the internal state of \mathcal{P}_{i-1}^* , and obtain $z_{\text{loc}}^{(v)}, \{z_j^{(v)}\}_{j \in [r]}$ from $\mathbf{w}^{(v)}$.
3. Append $z_{\text{loc}}^{(v)}$ to the label of $v \in S_{\mathbb{T}_{i-1}}(i)$.
4. For each node $v \in S_{i-1}$, let $S_v := \{j \in [r] \mid z_j^{(v)} \neq \perp\}$. Construct \mathbb{T}_i of depth $i + 1$ from \mathbb{T}_{i-1} by adding, for each node $v \in S_{i-1}$, $(z_j^{(v)}, \Pi_j^{(v)})$ to the label of its child $j \in S_v$, where $\Pi_j^{(v)} = ((\mathbf{U}_j^{(v)}, \mathbf{W}_j^{(v)}), (\mathbf{u}_j^{(v)}, \mathbf{w}_j^{(v)}))$.
5. Output $(\varphi, \mathbb{T}_i, \text{ao})$.

We now show inductively that the extractors are correct. We define the *inductive hypothesis* as that for $i \in \{0, 1, \dots, d\}$, $\mathcal{E}_i(\text{pp}, \text{ai})$ outputs $(\varphi, \mathbb{T}_i, \text{ao})$ in expected polynomial time such that with probability $\epsilon - \text{negl}(\lambda)$, 1) $\varphi \in \mathbb{F}$, $(\text{pp}, \text{ai}, \varphi, \mathbf{o}(\mathbb{T}_i), \text{ao}) \in \mathbb{Z}$, 2) \mathbb{T}_i is φ -compliant up to depth i , and 3) for all $v \in S_{\mathbb{T}_i}(i + 1)$, $\mathbf{V}(\mathbf{vk}, z^{(v)}, \Pi^{(v)}) = 1$.

In the base case, by the premise of P^* , \mathcal{E}_0 satisfies the inductive hypothesis.

Next, supposing that \mathcal{E}_{i-1} satisfies the inductive hypothesis, we show that \mathcal{E}_i also satisfies the inductive hypothesis. By the premise of \mathcal{E}_{i-1} , we have that with probability $\epsilon - \text{negl}(\lambda)$,

$\varphi \in \mathbf{F}, (\mathbf{pp}, \mathbf{ai}, \varphi, \mathbf{o}(\mathbb{T}_{i-1}), \mathbf{ao}) \in \mathbf{Z}, \mathbb{T}_{i-1}$ is φ -compliant up to depth $i - 1$, and for all $v \in S_{\mathbb{T}_{i-1}(i)}$, $\mathbf{V}(\mathbf{vk}, z^{(v)}, \mathbf{II}^{(v)}) = 1$. By the check of \mathbf{V} , we have that (1) $\{(\mathbf{U}^{(v)}, \mathbf{W}^{(v)}), (\mathbf{u}^{(v)}, \mathbf{w}^{(v)})\}_{v \in S_{\mathbb{T}_{i-1}(i)}}$ are satisfied instance-witness pairs. Additionally by the construction of R_φ and the collision-resistant property of the hash function, we have that (2) for $v \in S_{\mathbb{T}_{i-1}(i)}$, $\varphi(z^{(v)}, z_{\text{loc}}^{(v)}, z_1^{(v)}, \dots, z_r^{(v)})$ accepts, (3) for $v \in S_{i-1}$, $\mathbf{U}^{(v)} = \text{NIMFS.V}'(\mathbf{fvk}', \{\mathbf{U}_j^{(v)}\}_{j \in [r]}, \{\mathbf{u}_j^{(v)}\}_{j \in [r]}, \pi^{(v)})$, and (4) for $v \in S_{i-1}, j \in [r]$, $\mathbf{u}_j^{(v)}.x = \text{hash}(\mathbf{fvk}', z_j^{(v)}, \mathbf{U}_j^{(v)})$. Condition (2) implies that \mathbb{T}_i is φ -compliant up to depth i , and $\varphi \in \mathbf{F}, (\mathbf{pp}, \mathbf{ai}, \varphi, \mathbf{o}(\mathbb{T}_i), \mathbf{ao}) \in \mathbf{Z}$ with probability $\epsilon - \text{negl}(\lambda)$. Conditions (1)(3) imply that \mathcal{P}_{i-1}^* succeeds in producing satisfied folded instance-witness pairs $\{\mathbf{U}^{(v)}, \mathbf{W}^{(v)}\}_{v \in S_{i-1}}$ for instances $\{\mathbf{U}_j^{(v)}, \mathbf{u}_j^{(v)}\}_{j \in [r], v \in S_{i-1}}$ with probability $\epsilon - \text{negl}(\lambda)$. Then by the knowledge soundness of the multi-folding scheme, we have that $\mathcal{E}_{\mathcal{P}_{i-1}^*}$ succeeds in outputting satisfied witnesses $\{\mathbf{W}_j^{(v)}, \mathbf{w}_j^{(v)}\}_{j \in [r], v \in S_{i-1}}$ with probability $\epsilon - \text{negl}(\lambda)$. Additionally by Condition (4), we have that $\mathbf{V}(\mathbf{vk}, z^{(v)}, \mathbf{II}^{(v)}) = 1$ for all $v \in S_{\mathbb{T}_i}(i + 1)$ with probability $\epsilon - \text{negl}(\lambda)$. Since \mathcal{E}_{i-1} runs in expected polynomial time, $\mathcal{E}_{\mathcal{P}_{i-1}^*}$ also runs in expected polynomial time, and thereby so does \mathcal{E}_i . Therefore, we have that \mathcal{E}_i satisfies the inductive hypothesis.

In conclusion, we conclude that Construction 3 has knowledge soundness. \square

Efficiency. The recursion overhead, i.e., verifier's computations expressed as circuits is dominated by the computations in R_φ except for checking the compliance predicate φ , containing $r + 1$ calls to `hash` and one call to `NIMFS.V'`. The cost of `NIMFS.V'` is dominated by $O(d \log m + r \cdot (\log m + t + dq))$ field operations, one MSM of size $2r$ and $2 \log_2 m + 2$ calls to the random oracle `RO` to achieve non-interactivity, where `RO` could be instantiated with an appropriate cryptographic hash function.

The prover's work at each step is dominated by invoking `NIMFS.P'`, and computing the satisfied $\mathcal{R}_{\text{CCCS}}$ instance-witness pair (\mathbf{u}, \mathbf{w}) for the execution of R_φ . The cost of `NIMFS.P'` is dominated by $O(r \cdot (N + tm + n + qmd \log^2 d))$ field operations, one MSM of size $2r$, and $2 \log_2 m + 2$ calls to `RO` to achieve non-interactivity. The cost of computing (\mathbf{u}, \mathbf{w}) is dominated by computing the commitment C which requires one MSM of size $O(n)$ when instantiating the polynomial commitment scheme with Bulletproofs [6].

The proof \mathbf{II} consists of $(\mathbf{U}, \mathbf{W}, \mathbf{u}, \mathbf{w})$ whose size is linear in the size of R_φ . However, as in HyperNova [23], we could use a general SNARK to compress the proof. Specifically, the prover invokes $(\mathbf{U}', \mathbf{W}', \pi') \leftarrow \text{NIMFS.P}'(\mathbf{fpk}', (\mathbf{U}, \mathbf{W}), (\mathbf{u}, \mathbf{w}))$, and then uses a general SNARK to generate a proof $\pi_{\mathbf{U}'}$ that proves the knowledge of \mathbf{W}' . Now \mathbf{II} consists of $(\mathbf{U}, \mathbf{u}, \pi', \pi_{\mathbf{U}'})$. When instantiating the SNARK with SuperSpartan [29] excluding the first sum-check invocation and using the polynomial commitment scheme based on Bulletproofs [6], the proof size is then dominated by $O(d \log m + t + \log n)$ field elements and $O(\log n)$ group elements.

The verifier's work is dominated by checking $(\mathbf{U}, \mathbf{W}), (\mathbf{u}, \mathbf{w})$, whose cost is linear in the size of R_φ . However, by compressing the proof, the work is now dominated by performing $\mathbf{U}' \leftarrow \text{NIMFS.V}'(\mathbf{fvk}', \mathbf{U}, \mathbf{u}, \pi')$ and verifying $\pi_{\mathbf{U}'}$, whose cost is dominated by $O(d \log m + t + dq + \log n)$ field operations, $O(\log m + \log n)$ calls to `RO`, and one MSM of size $O(n)$.

5 Conclusion

In this paper, we first construct a multi-folding scheme for arbitrary number of instances, which could reduce the task of checking multiple instances into the task of checking one. Based on this

scheme, we construct a PCD scheme having the smallest prover’s cost at each step and recursion overhead in the literature. Supposing that there are r incoming edges of a node at certain step and n variables in the constraint system, the prover’s cost is dominated by one MSM of size $O(n)$, and the recursion overhead is dominated by one MSM of size $2r$. Besides, our PCD scheme supports more expressive constraint system, i.e., CCCS that allows gates to compute high-degree polynomials.

Recently, Eagen and Gabizon [17] proposed a folding scheme for multiple instances—ProtoGalaxy with novel efficiency characteristics. How to construct PCD schemes based on ProtoGalaxy and analyse the concrete efficiency is an interesting future study.

References

1. Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Scalable zero knowledge via cycles of elliptic curves. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO. LNCS, vol. 8617, pp. 276–294. Springer (2014)
2. Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: Recursive composition and bootstrapping for SNARKS and proof-carrying data. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013. pp. 111–120. ACM (2013)
3. Boneh, D., Drake, J., Fisch, B., Gabizon, A.: Halo infinite: Proof-carrying data from additive polynomial commitments. In: Malkin, T., Peikert, C. (eds.) CRYPTO. LNCS, vol. 12825, pp. 649–680. Springer (2021)
4. Bonneau, J., Meckler, I., Rao, V., Shapiro, E.: Coda: Decentralized cryptocurrency at scale. Cryptology ePrint Archive, Paper 2020/352 (2020), <https://eprint.iacr.org/2020/352>
5. Bowe, S., Grigg, J., Hopwood, D.: Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Paper 2019/1021 (2019), <https://eprint.iacr.org/2019/1021>
6. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: S&P. pp. 315–334. IEEE Computer Society (2018)
7. Bünz, B., Chiesa, A., Lin, W., Mishra, P., Spooner, N.: Proof-carrying data without succinct arguments. In: Malkin, T., Peikert, C. (eds.) CRYPTO. LNCS, vol. 12825, pp. 681–710. Springer (2021)
8. Bünz, B., Chiesa, A., Mishra, P., Spooner, N.: Recursive proof composition from accumulation schemes. In: Pass, R., Pietrzak, K. (eds.) TCC. LNCS, vol. 12551, pp. 1–18. Springer (2020)
9. Bünz, B., Fisch, B., Szepieniec, A.: Transparent snarks from DARK compilers. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT. LNCS, vol. 12105, pp. 677–706. Springer (2020)
10. Bünz, B., Chen, B.: Protostar: Generic efficient accumulation/folding for special sound protocols. Cryptology ePrint Archive, Paper 2023/620 (2023), <https://eprint.iacr.org/2023/620>
11. Chen, W., Chiesa, A., Dauterman, E., Ward, N.P.: Reducing participation costs via incremental verification for ledger systems. Cryptology ePrint Archive, Paper 2020/1522 (2020), <https://eprint.iacr.org/2020/1522>
12. Chiesa, A., Ojha, D., Spooner, N.: Fractal: Post-quantum and transparent recursive proofs from holography. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT. LNCS, vol. 12105, pp. 769–793. Springer (2020)
13. Chiesa, A., Tromer, E.: Proof-carrying data and hearsay arguments from signature cards. In: Yao, A.C. (ed.) ICS. pp. 310–331. Tsinghua University Press (2010)
14. Chiesa, A., Tromer, E., Virza, M.: Cluster computing in zero knowledge. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT. LNCS, vol. 9057, pp. 371–403. Springer (2015)
15. Chong, S., Tromer, E., Vaughan, J.A.: Enforcing language semantics using proof-carrying data. Cryptology ePrint Archive, Paper 2013/513 (2013), <https://eprint.iacr.org/2013/513>
16. Cormode, G., Mitzenmacher, M., Thaler, J.: Practical verified computation with streaming interactive proofs. In: Goldwasser, S. (ed.) ITCS. pp. 90–112. ACM (2012)
17. Eagen, L., Gabizon, A.: Protogalaxy: Efficient protostar-style folding of multiple instances. Cryptology ePrint Archive, Paper 2023/1106 (2023), <https://eprint.iacr.org/2023/1106>
18. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO. LNCS, vol. 263, pp. 186–194. Springer (1986)
19. Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct nizks without pcps. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT. LNCS, vol. 7881, pp. 626–645. Springer (2013)
20. Kattis, A., Bonneau, J.: Proof of necessary work: Succinct state verification with fairness guarantees. Cryptology ePrint Archive, Paper 2020/190 (2020), <https://eprint.iacr.org/2020/190>
21. Kothapalli, A., Setty, S.: Supernova: Proving universal machine executions without universal circuits. Cryptology ePrint Archive, Paper 2022/1758 (2022), <https://eprint.iacr.org/2022/1758>

22. Kothapalli, A., Setty, S.: Cyclefold: Folding-scheme-based recursive arguments over a cycle of elliptic curves. Cryptology ePrint Archive, Paper 2023/1192 (2023), <https://eprint.iacr.org/2023/1192>
23. Kothapalli, A., Setty, S.: Hypernova: Recursive arguments for customizable constraint systems. Cryptology ePrint Archive, Paper 2023/573 (2023), <https://eprint.iacr.org/2023/573>
24. Kothapalli, A., Setty, S.T.V., Tzialla, I.: Nova: Recursive zero-knowledge arguments from folding schemes. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO. LNCS, vol. 13510, pp. 359–388. Springer (2022)
25. Lund, C., Fortnow, L., Karloff, H.J., Nisan, N.: Algebraic methods for interactive proof systems. In: FOCS. pp. 2–10. IEEE Computer Society (1990)
26. Mohnblatt, N.: Sangria: a folding scheme for plonk (2023), https://github.com/geometryresearch/technical_notes/blob/main/sangria_folding_plonk.pdf
27. Naveh, A., Tromer, E.: Photoproof: Cryptographic image authentication for any set of permissible transformations. In: S&P. pp. 255–271. IEEE Computer Society (2016)
28. Schwartz, J.T.: Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM* **27**(4), 701–717 (1980)
29. Setty, S., Thaler, J., Wahby, R.: Customizable constraint systems for succinct arguments. Cryptology ePrint Archive, Paper 2023/552 (2023), <https://eprint.iacr.org/2023/552>
30. Thaler, J.: Proofs, arguments, and zero-knowledge. *Foundations and Trends in Privacy and Security* **4**(2-4), 117–660 (2022)
31. Tyagi, N., Fisch, B., Zitek, A., Bonneau, J., Tessaro, S.: Versa: Verifiable registries with efficient client audits from RSA authenticated dictionaries. In: Yin, H., Stavrou, A., Cremers, C., Shi, E. (eds.) CCS. pp. 2793–2807. ACM (2022)
32. Valiant, P.: Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In: Canetti, R. (ed.) TCC. LNCS, vol. 4948, pp. 1–18. Springer (2008)
33. Vu, V., Setty, S.T.V., Blumberg, A.J., Walfish, M.: A hybrid architecture for interactive verifiable computation. In: S&P. pp. 223–237. IEEE Computer Society (2013)