

HyperNova: Recursive arguments for customizable constraint systems

Abhiram Kothapalli[†]

Srinath Setty^{*}

[†]Carnegie Mellon University

^{*}Microsoft Research

Abstract.

This paper introduces *HyperNova*, a recursive argument for proving incremental computations whose steps are expressed with CCS (Setty et al. ePrint 2023), a customizable constraint system that simultaneously generalizes Plonkish, R1CS, and AIR without overheads. A distinguishing aspect of HyperNova is that the prover’s cost at each step is dominated by a *single* multi-scalar multiplication (MSM) of size equal to the number of variables in the constraint system, which is optimal when using an MSM-based commitment scheme.

To construct HyperNova, we generalize folding schemes (CRYPTO 22), to allow instances from two (potentially) different NP relations, that share the *same* structure, to be folded; we refer to this generalization as *multi-folding schemes*. Furthermore, we devise a public-coin, multi-folding scheme for instances in CCS and *linearized CCS* (a variant of CCS that we introduce). This construction can be viewed as an “early stopping” version of Spartan (CRYPTO 20), applied to a carefully-crafted polynomial that includes claims about prior linearized CCS instances. The prover’s work in the multi-folding scheme is a linear number of finite field operations and the verifier’s work is a logarithmic number of finite field operations and a single group scalar multiplication. We then construct incrementally verifiable computation (IVC) from non-interactive multi-folding schemes with the lowest prover costs and the lowest recursion overheads in the literature. We also provide an alternate realization of HyperNova with a black box use of Nova, which nearly eliminates the need for deferred arithmetic when instantiated with a cycle of elliptic curves.

As an additional contribution, we describe **nlookup**, a lookup argument, that is particularly suited for recursive arguments based on folding schemes. Specifically, at a particular step in an incremental computation, for m lookups into a table of size n ($m \ll n$), the prover’s work is dominated by $O(n)$ finite field operations and it requires only $O(m \cdot \log n)$ degree-2 constraints and $O(\log n)$ hash evaluations in the incremental computation. **nlookup** is currently not suitable for efficiently encoding bitwise operations, but it provides a powerful tool for efficiently encoding (large) finite state machines and proving their transitions with recursive SNARKs.

1 Introduction

This paper continues a recent line of work that began with Nova [KST22] (and its generalization SuperNova [KS22]) to build prover-efficient succinct arguments. There are two somewhat different ways of interpreting results from this line of work. On the one hand, they provide the fastest succinct argument for semi-structured, non-deterministic computations where the computation to be proven can be decomposed into multiple steps (different steps do not need to have the same circuit representation) [Dra22]. On the other hand, they also provide the most efficient recursive argument scheme, with the lowest recursion overheads.

(Super)Nova’s computational model allows incremental computations where a particular step invokes one of the pre-defined “instructions”. This model captures many structured and semi-structured statements that arise frequently in practice. For example, in VDFs, there is only one instruction and it verifies that a certain iterations of a delay function (e.g., MinRoot [KMT22]) was executed correctly. As a more complex example, in the so-called “ZKVMs”, at each step, the incremental computation invokes one of the pre-defined instructions. Concretely, in the case of Ethereum’s virtual machine (EVM), an instruction might perform operations like add/mul, or an elliptic curve scalar multiplication over the BN-254 curve.

To prove such incremental computations, Nova provides an inexpensive mechanism, called a folding scheme, to (recursively) transform the task of proving N steps of a program execution into the task of proving a single step of the program execution. It then applies a general-purpose zkSNARK to prove that single step. This approach builds on a recent line of work to achieve efficient recursive arguments [BCTV14, BGH19, BCMS20, BDFG21, BCL⁺21], but it differs from them in that Nova does not rely on zkSNARKs or even (non-succinct) NARKs.

Compared to directly employing a general-purpose zkSNARK, built from from polynomial IOPs and polynomial commitment schemes (e.g., Spartan [Set20], Plonk [GWC19], Marlin [CHM⁺20], HyperPlonk [CBBZ23]), to prove the entire incremental computation, Nova’s approach is substantially cheaper (as long as each step is sufficiently large, to offset recursion overheads). Specifically, at each incremental step, Nova’s prover incurs only *two* MSMs of size proportional to the size of the circuit proven. Whereas, general-purpose zkSNARKs need many more MSMs (e.g., Marlin [CHM⁺20, Fig 1] reports 22 MSMs and many more FFTs of size proportional to the circuit size).

In addition, by design, Nova’s proof generation is incremental (i.e., it produces a proof for each step and then uses recursion capabilities to produce a single proof), so it can be more easily distributed and parallelized than with a non-recursive zkSNARK where one must unroll program executions into monolithic circuits. The latter rules out applications where one cannot statically unroll program executions (e.g., VDF) or makes it inconvenient (e.g., program executions on machines such as EVM or RISC-V). As presented, Nova does not immediately support parallel proof generation. There exists a generic compiler [BCCT13] to transform constructions such as Nova to support parallel proving.

Problem statement. In Nova, each step of an incremental computation is expressed with R1CS, an NP-complete problem that generalizes arithmetic circuits and is implicit in the QAPs formalism [GGPR13]. In practice, when monolithic zkSNARKs (e.g., Plonk) are applied to large program executions, practitioners often use custom constraint systems (e.g., Plonkish) that are tailored to particular classes of applications. In particular, constraints in Plonkish are multivariate high-degree polynomials, so they allow gates that can compute custom high-degree polynomials. In contrast, R1CS is restricted to checking quadratic constraints in a specific form.

A key question is whether one can build a recursive argument for Plonkish, with Nova-like performance characteristics. In particular, our goal is to prove CCS [STW23], a customizable constraint system that simultaneously generalizes Plonkish, R1CS, and AIR without overheads.

In more detail, for any solution that handles high-degree constraints, we require the prover’s cryptographic work to be independent of the degree of constraints supported. That is, the number of MSMs (or their sizes) performed by the prover must not depend on the degree of the supported constraints. This requirement is crucial because otherwise one can simply lower the degree of constraints to two and use Nova rather than attempting to use high-degree constraints.

Prior approaches and challenges.

Halo2 [BGH20] replaces the polynomial IOP in Halo [BGH19] from Sonic [MBKM19] to Plonk [GWC19], and its widely used implementation supports Plonkish. Unfortunately, it incurs substantial prover costs as the prover must necessarily produce a SNARK using Plonk at each step of the program execution. Furthermore, the prover incurs $O(n \cdot d)$ cryptographic operations, where n is the size of the circuit at each step and d is the maximum degree of constraints proven. Switching from Plonk to HyperPlonk [CBBZ23] would reduce the prover’s cryptographic operations to $O(n)$, but it does not avoid the need to produce a SNARK.

Split accumulation [BCL⁺21] reduces the recursion overheads of Halo-type approaches. It also avoids the need for succinct arguments to construct recursive SNARKs. Unfortunately, the existing construction of recursive SNARKs from split accumulation [BCL⁺21] is limited to R1CS. It is not clear how to extend their approach to handle Plonkish without making the prover incur $O(n \cdot d)$ cryptographic operations, which, as noted above, is undesirable.

Sangria [Moh23] shows that Nova’s folding scheme for R1CS can be easily adapted to handle Plonkish. However, the number of cross-terms that the prover must commit to increases linearly with d . Furthermore, the prover must incur $O(n \cdot d)$ cryptographic operations to commit to $O(d)$ cross-terms. As a result, in general, it is cheaper to lower the constraints to degree-2 and prove them with Nova than use higher-degree constraints with Sangria.

Concretely, a single iteration of the MinRoot delay function [KMT22] can be represented with a single degree-5 constraint in Plonkish [Zha23]. Whereas, with

R1CS, one would need three constraints. However, if Sangria is applied to prove MinRoot in degree-5 Plonkish, it incurs more work for the prover than when Nova is applied to prove MinRoot in R1CS.

1.1 Our approach: HyperNova

This section provides a high-level overview of HyperNova’s ingredients. As noted earlier, HyperNova’s target is to prove incremental computations where each step of the incremental computation is expressed with CCS [STW23], a customizable constraint system that simultaneously generalizes Plonkish, R1CS, and AIR without overheads. However, if we naively build a folding scheme for CCS, perhaps for a “relaxed” variant of CCS (analogous to relaxed R1CS in Nova), it will have the efficiency issues noted above for Sangria.

To avoid those issues, HyperNova takes a different approach that involves leveraging the power of the sum-check protocol [LFKN90], as done in Spartan [Set20] and its recent extension to handle CCS called SuperSpartan [STW23]. In particular, we provide two instantiations, one that directly builds HyperNova, and another that uses Nova as a black box. Both variants can be instantiated efficiently using cycles of elliptic curves and they both have the same prover efficiency. The second variant however has minimal “deferred arithmetic” (i.e., arithmetic that must be passed to the other curve in the curve cycle for efficiency) that shows up when implementing recursive arguments using cycles of elliptic curves.

Remark 1. Given that CCS is a strict generalization of Plonkish, we focus on describing HyperNova for CCS. HyperNova for CCS can be easily translated to HyperNova for Plonkish, with the same efficiency guarantees.

(1) Multi-folding schemes. To construct HyperNova, we introduce a generalization of folding schemes, and we refer to it as *multi-folding schemes*. Recall that a folding scheme for a relation \mathcal{R} is a protocol between a *prover* and *verifier* in which the prover and the verifier reduce the task of checking two instances in \mathcal{R} with the same structure \mathbf{s} into the task of checking a single instance in \mathcal{R} with structure \mathbf{s} . A multi-folding scheme is defined with respect to a pair of relations $(\mathcal{R}_1, \mathcal{R}_2)$. It is an interactive protocol between a prover and a verifier in which the prover and the verifier reduce the task of checking an instance in \mathcal{R}_1 and an instance in \mathcal{R}_2 that both share the *same* structure \mathbf{s} into the task of checking a single instance in \mathcal{R}_1 with structure \mathbf{s} . The paragraphs ahead will clarify how this generalization unlocks additional power for constructing IVC.

Remark 2. Multi-folding schemes are related to split-accumulation schemes [BCL⁺21]. However, multi-folding schemes use the language of NP relations and are specified in the standard interactive model. For example, our definitions do not require an auxiliary decider algorithm and the associated complexity. Furthermore, by using the language of NP relations, the split between a witness and instance is immediate in the context of folding schemes and multi-folding schemes. Overall, our definitions, security notions, and proofs are conceptually simpler.

(2) A multi-folding scheme for CCS. Our starting point to construct a multi-folding scheme for CCS is the following observation. We observe that the Spartan proof system [Set20] (more specifically its generalization to handle CCS called SuperSpartan [STW23]) transforms the task of checking the satisfiability of a CCS instance into the task of checking if a multivariate polynomial g of total degree $d + 1$, where d is the degree of the CCS constraints, sums to zero over a suitable Boolean hypercube. Spartan then invokes the sum-check protocol [LFKN90] to prove that claim about g . At the end of the sum-check invocation, the prover and the verifier are left with checking certain claims. Fortunately, these claims concern a restricted form of CCS (we formalize this restricted form of CCS and refer to it as *linearized CCS*). Note Spartan then proves those claims about the restricted form of CCS with an additional invocation of the sum-check protocol.

While an “early stopping” version of Spartan (the one with a single invocation of the sum-check protocol) provides a reduction of knowledge [KP22] from CCS to linearized CCS, it is not a folding (or a multi-folding) scheme. So, our second idea is to redefine the polynomial g to additionally include claims from a running linearized CCS instance using a random challenge from the verifier. This is possible as long as the running instance and the CCS instance that is being folded share the *same* structure (e.g., the same CCS matrices).

As a result, we obtain a public-coin multi-folding scheme that can fold a CCS instance with structure s into a linearized CCS instance with structure s to get a new linearized CCS instance with structure s . The prover’s work in the folding scheme is a linear number of finite field operations (the prover does not commit to any cross-terms!), and the verifier’s work is a logarithmic number of finite field operations. Since the multi-folding scheme is public coin, we make it non-interactive in the random oracle model using the Fiat-Shamir transform [FS86] and heuristically instantiate it in the plain model using a concrete cryptographic hash function.

Remark 3. Our multi-folding scheme for CCS is more powerful than described it thus far: it can take μ running linearized CCS instances and ν CCS instances that all share the same structure and output a single linearized CCS instance. So, our definition of multi-folding schemes is additionally parameterized by μ and ν .

(3) IVC and its generalizations from multi-folding schemes. Recall that Nova constructs IVC [Val08] from non-interactive folding schemes. Similarly, we show how to construct an IVC scheme from non-interactive multi-folding schemes such as the non-interactive version of the multi-folding scheme for CCS discussed above. Our construction and security proofs follow the same blueprint. The prover initializes a “default” running instance. Then, at each step of the incremental computation, the prover commits to the witness for an augmented CCS instance (the augmented CCS not only runs a step of the incremental computation, but it also runs the verifier of the non-interactive multi-folding scheme). The prover then runs the non-interactive multi-folding scheme to fold that CCS instance into a running instance. The prover then sends the CCS instance and the output

of the multi-folding instance as input to the next invocation of the augmented circuit, to prove the next step of the incremental computation.

As in Nova, in the “end” (or at any point in the incremental computation), the prover is left with a running instance U (which in the case of HyperNova is a linearized CCS instance) and a CCS instance u that represents the last step of the incremental computation augmented with a verifier circuit. The prover folds U and u to get U' with an invocation of the multi-folding scheme for CCS, and suppose that the prover’s output in the folding scheme is π_{fs} . HyperNova’s recursive proof is then $\pi = (U, u, \pi_{fs}, w)$ where w is a purported witness to U' . HyperNova’s prover can then send π to a verifier, but the size of the proof is linear in the size of the step of an incremental computation. Thus, as in Nova, HyperNova’s prover can use a general zkSNARK to prove the knowledge of a valid π . In particular, Spartan excluding the first sum-check invocation provides a necessary zkSNARK. Of course, other proof systems can also be used as long as they can prove the knowledge of a valid π .

Remark 4. A multi-folding scheme folds μ running instances and ν incoming instances into a single running instance, so it easily provides a natural generalization of IVC [Val08] called proof-carrying data [BCCT13]. However, we focus on IVC for its conceptual simplicity.

(4) A construction of IVC with a black box use of Nova In addition to the direct construction of IVC from a non-interactive multi-folding scheme for CCS, we describe an approach that leverages Nova as a black box. In particular, the high-level idea is to express the verifier of the non-interactive multi-folding scheme for CCS as a step circuit in Nova (i.e., with R1CS). The step circuit, as part of its public IO, stores a running instance (which in our case is a linearized CCS instance). At each step, the prover first represents a step of the incremental computation as a CCS instance u . It folds that into a running instance U to obtain a new running instance U' (and the associated witnesses) and a proof π_{fs} . The prover then runs Nova where the step circuit takes as non-deterministic input (u, π_{fs}) . The step circuit then folds u into a running instance that it tracks via its public IO using π_{fs} . In the “end” or at any point in the incremental computation, a verifier can verify the Nova proof associated with the incremental computation and check that the running instance in the public IO of the step circuit is indeed satisfying using a purported witness from the prover. As with the direct construction discussed above, the prover can compress such a recursive proof into a succinct proof using a general zkSNARK.

The primary advantage of this version of HyperNova is that it only needs to implement a step circuit. Furthermore, much of the work in representing the verifier of the non-interactive folding scheme is a logarithmic number of finite field operations and hash evaluations, both of which can be encoded natively and efficiently with R1CS (e.g., with a SNARK-friendly hash function). The step circuit will have to perform a *single* group scalar multiplication, but when using

a cycle of elliptic curves, this task can be easily (and with minimal public IO) be deferred to the other curve in the curve cycle.

2 Preliminaries

We use λ to denote the security parameter and \mathbb{F} to denote a finite field (e.g., the prime field \mathbb{F}_p for a large prime p). We use $\text{negl}(\lambda)$ to denote a negligible function in λ . Throughout the paper, the depicted asymptotics depend on λ , but we elide this for brevity. We use “PPT algorithms” to refer to probabilistic polynomial time algorithms.

2.1 Polynomials and low-degree extensions

We adapt this subsection from prior work [Set20].

We recall a few basic facts about polynomials.

- A polynomial g over \mathbb{F} is an expression consisting of a sum of *monomials* where each monomial is the product of a constant from \mathbb{F} and powers of one or more variables; all arithmetic is performed over \mathbb{F} .
- The degree of a monomial in a polynomial is the sum of the exponents of variables in the monomial; the degree of a polynomial g is the maximum degree of any monomial in g .
- The degree of a polynomial g in a particular variable x_i is the maximum exponent that x_i takes in any of the monomials in g .
- A *multivariate* polynomial is a polynomial with more than one variable; otherwise it is called a *univariate* polynomial.

Definition 1 (Multilinear polynomial). *A multivariate polynomial is called a multilinear polynomial if the degree of the polynomial in each variable is at most one.*

Definition 2 (Low-degree polynomial). *A multivariate polynomial g over a finite field \mathbb{F} is called low-degree polynomial if the degree of g in each variable is exponentially smaller than $|\mathbb{F}|$.*

Low-degree extensions (LDEs). Suppose $g : \{0, 1\}^\ell \rightarrow \mathbb{F}$ is a function that maps ℓ -bit elements into an element of \mathbb{F} . A *polynomial extension* of g is a low-degree ℓ -variate polynomial \tilde{g} such that $\tilde{g}(x) = g(x)$ for all $x \in \{0, 1\}^\ell$.

A *multilinear* polynomial extension (or simply, a multilinear extension, or MLE) is a low-degree polynomial extension where the extension is a multilinear polynomial (i.e., the degree of each variable in \tilde{g} is at most one). Given a function $Z : \{0, 1\}^\ell \rightarrow \mathbb{F}$, the multilinear extension of Z is the unique multilinear polynomial $\tilde{Z} : \mathbb{F}^\ell \rightarrow \mathbb{F}$.

It can be computed as follows.

$$\begin{aligned}
\tilde{Z}(x_1, \dots, x_\ell) &= \sum_{e \in \{0,1\}^\ell} Z(e) \cdot \prod_{i=1}^{\ell} (x_i \cdot e_i + (1 - x_i) \cdot (1 - e_i)) \\
&= \sum_{e \in \{0,1\}^\ell} Z(e) \cdot \tilde{e}q(x, e) \\
&= \langle (Z(0), \dots, Z(2^\ell - 1)), (\tilde{e}q(x, 0), \dots, \tilde{e}q(x, 2^\ell - 1)) \rangle
\end{aligned}$$

Note that $\tilde{e}q(x, e) = \prod_{i=1}^{\ell} (e_i \cdot x_i + (1 - e_i) \cdot (1 - x_i))$, which is the MLE of the following function:

$$eq(x, e) = \begin{cases} 1 & \text{if } x = e \\ 0 & \text{otherwise} \end{cases}$$

For any $r \in \mathbb{F}^\ell$, $\tilde{Z}(r)$ can be computed in $O(2^\ell)$ operations in \mathbb{F} [VSBW13,Tha13].

Dense representation for multilinear polynomials. Since the MLE of a function is unique, it offers the following method to represent any multilinear polynomial. Given a multilinear polynomial $g : \mathbb{F}^\ell \rightarrow \mathbb{F}$, it can be represented uniquely by the list of tuples L such that for all $i \in \{0, 1\}^\ell$, $(\text{to-field}(i), g(i)) \in L$ if and only if $g(i) \neq 0$, where to-field is the canonical injection from $\{0, 1\}^\ell$ to \mathbb{F} . We denote such a representation of g as $\text{DenseRepr}(g)$.

Definition 3. *A multilinear polynomial g in ℓ variables is a sparse multilinear polynomial if $|\text{DenseRepr}(g)|$ is sub-linear in $O(2^\ell)$. Otherwise, it is a dense multilinear polynomial.*

As an example, suppose $g : \mathbb{F}^{2^s} \rightarrow \mathbb{F}$. Suppose $|\text{DenseRepr}(g)| = O(2^s)$, then $g(\cdot)$ is a sparse multilinear polynomial because $O(2^s)$ is sublinear in $O(2^{2^s})$.

Lemma 1 (Sums over evaluations). *Consider size $\ell \in \mathbb{N}$. For multilinear polynomial $P \in \mathbb{F}[X_1, \dots, X_\ell]$ we have that*

$$P(X) = \sum_{x \in \{0,1\}^\ell} \tilde{e}q(X, x) \cdot P(x).$$

Proof. Let $Q(X) = \sum_{x \in \{0,1\}^\ell} \tilde{e}q(X, x) \cdot P(x)$ By the definition of $\tilde{e}q$, we have that

$$P(x) = Q(x)$$

for all $x \in \{0, 1\}^\ell$. However, because $P \in \mathbb{F}[X_1, \dots, X_\ell]$ is multilinear it is completely determined by 2^ℓ evaluation points. The same is holds for Q . Because P and Q agree on 2^ℓ points, they must be the same polynomial. \square

2.2 Schwartz-Zippel lemma

Lemma 2. *let $g : \mathbb{F}^\ell \rightarrow \mathbb{F}$ be an ℓ -variate polynomial of total degree at most d . Then, on any finite set $S \subseteq \mathbb{F}$,*

$$\Pr_{x \leftarrow S^\ell} [g(x) = 0] \leq d/|S|$$

2.3 The sum-check protocol

Suppose there is an ℓ -variate low-degree polynomial, g , where the degree of each variable in g is at most d . Suppose that a verifier \mathcal{V} is interested in checking a claim of the following form by an untrusted prover \mathcal{P} :

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

Of course, given g , \mathcal{V} can deterministically evaluate the above sum and verify whether the sum is T . But, this computation takes time exponential in ℓ . Lund et al. [LFKN90] describe the sum-check protocol that requires far less computation on \mathcal{V} 's behalf, but provides a probabilistic guarantee. 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 $c \leftarrow \langle \mathcal{P}, \mathcal{V}(r) \rangle(g, \ell, d, T)$.

For any ℓ -variate polynomial g with degree at most d in each variable, the sum-check protocol satisfies the following properties.

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

2.4 Arguments of knowledge and SNARKs

We adapt the definition provided in [KST22].

Definition 4. *Consider a relation \mathcal{R} over public parameters, structure, instance, and witness tuples. A non-interactive argument of knowledge for \mathcal{R} consists of PPT algorithms $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ and deterministic \mathcal{K} , denoting the generator, the prover, the verifier and the encoder respectively with the following interface.*

- $\mathcal{G}(1^\lambda) \rightarrow \text{pp}$: On input security parameter λ , samples public parameters pp .
- $\mathcal{K}(\text{pp}, \mathbf{s}) \rightarrow (\text{pk}, \text{vk})$: On input structure \mathbf{s} , representing common structure among instances, outputs the prover key pk and verifier key vk .

- $\mathcal{P}(\mathbf{pk}, u, w) \rightarrow \pi$: On input instance u and witness w , outputs a proof π proving that $(\mathbf{pp}, \mathbf{s}, u, w) \in \mathcal{R}$.
- $\mathcal{V}(\mathbf{vk}, u, \pi) \rightarrow \{0, 1\}$: On input the verifier key \mathbf{vk} , instance u , and a proof π , outputs 1 if the instance is accepting and 0 otherwise.

A non-interactive argument of knowledge satisfies completeness if for any PPT adversary \mathcal{A}

$$\Pr \left[\mathcal{V}(\mathbf{vk}, u, \pi) = 1 \mid \begin{array}{l} \mathbf{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\mathbf{s}, (u, w)) \leftarrow \mathcal{A}(\mathbf{pp}), \\ (\mathbf{pp}, \mathbf{s}, u, w) \in \mathcal{R}, \\ (\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, \mathbf{s}), \\ \pi \leftarrow \mathcal{P}(\mathbf{pk}, u, w) \end{array} \right] = 1.$$

A non-interactive argument of knowledge satisfies knowledge soundness if for all PPT adversaries \mathcal{A} there exists a PPT extractor \mathcal{E} such that for all randomness ρ

$$\Pr \left[\mathcal{V}(\mathbf{vk}, u, \pi) = 1, \begin{array}{l} (\mathbf{pp}, \mathbf{s}, u, w) \notin \mathcal{R} \\ \mathbf{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\mathbf{s}, u, \pi) \leftarrow \mathcal{A}(\mathbf{pp}; \rho), \\ (\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, \mathbf{s}), \\ w \leftarrow \mathcal{E}(\mathbf{pp}, \rho) \end{array} \right] = \text{negl}(\lambda).$$

A non-interactive argument of knowledge is succinct if the size of the proof π and the time to verify it are at most polylogarithmic in the size of the statement proven.

2.5 Incrementally Verifiable Computation

Definition 5 (Incrementally verifiable computation (IVC)). An incrementally verifiable computation (IVC) scheme is defined by PPT algorithms $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ and deterministic \mathcal{K} denoting the generator, the prover, the verifier, and the encoder respectively, with the following interface

- $\mathcal{G}(1^\lambda) \rightarrow \mathbf{pp}$: on input security parameter λ , samples public parameters \mathbf{pp} .
- $\mathcal{K}(\mathbf{pp}, F) \rightarrow (\mathbf{pk}, \mathbf{vk})$: on input public parameters \mathbf{pp} , and polynomial-time function F , deterministically produces a prover key \mathbf{pk} and a verifier key \mathbf{vk} .
- $\mathcal{P}(\mathbf{pk}, (i, z_0, z_i), \omega_i, \Pi_i) \rightarrow \Pi_{i+1}$: on input a prover key \mathbf{pk} , a counter i , an initial input z_0 , a claimed output after i iterations z_i , a non-deterministic advice ω_i , and an IVC proof Π_i attesting to z_i , produces a new proof Π_{i+1} attesting to $z_{i+1} = F(z_i, \omega_i)$.
- $\mathcal{V}(\mathbf{vk}, (i, z_0, z_i), \Pi_i) \rightarrow \{0, 1\}$: on input a verifier key \mathbf{vk} , a counter i , an initial input z_0 , a claimed output after i iterations z_i , and an IVC proof Π_i attesting to z_i , outputs 1 if Π_i is accepting, and 0 otherwise.

An IVC scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ satisfies the following requirements.

1. *Perfect Completeness:* For any PPT adversary \mathcal{A}

$$\Pr \left[\mathcal{V}(\mathbf{vk}, (i+1, z_0, z_{i+1}), \Pi_{i+1}) = 1 \mid \begin{array}{l} \mathbf{pp} \leftarrow \mathcal{G}(1^\lambda), \\ F, (i, z_0, z_i, \Pi_i) \leftarrow \mathcal{A}(\mathbf{pp}), \\ (\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, F), \\ z_{i+1} \leftarrow F(z_i, \omega_i), \\ \mathcal{V}(\mathbf{vk}, i, z_0, z_i, \Pi_i) = 1, \\ \Pi_{i+1} \leftarrow \mathcal{P}(\mathbf{pk}, (i, z_0, z_i), \omega_i, \Pi_i) \end{array} \right]$$

where F is a polynomial-time computable function.

2. *Knowledge Soundness:* Consider constant $n \in \mathbb{N}$. For all expected polynomial-time adversaries \mathcal{P}^* there exists an expected polynomial-time extractor \mathcal{E} such that over all randomness ρ

$$\Pr \left[\begin{array}{l} z_n = z \text{ where} \\ z_{i+1} \leftarrow F(z_i, \omega_i) \\ \forall i \in \{0, \dots, n-1\} \end{array} \mid \begin{array}{l} \mathbf{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (F, (z_0, z_i), \Pi) \leftarrow \mathcal{P}^*(\mathbf{pp}, \rho), \\ (\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, F), \\ \mathcal{V}(\mathbf{vk}, (n, z_0, z), \Pi) = 1, \\ (\omega_0, \dots, \omega_{n-1}) \leftarrow \mathcal{E}(\mathbf{pp}, \rho) \end{array} \right] \approx 1.$$

3. *Succinctness:* The size of an IVC proof Π is independent of the number of iterations n .

2.6 Polynomial commitment scheme

We adapt the definition from [BFS20]. A polynomial commitment scheme for multilinear polynomials is a tuple of four protocols $\text{PC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Eval})$:

- $\mathbf{pp} \leftarrow \text{Gen}(1^\lambda, \ell)$: takes as input ℓ (the number of variables in a multilinear polynomial); produces public parameters \mathbf{pp} .
- $\mathcal{C} \leftarrow \text{Commit}(\mathbf{pp}, g)$: takes as input a ℓ -variate multilinear polynomial over a finite field $g \in \mathbb{F}[\ell]$; produces a commitment \mathcal{C} .
- $b \leftarrow \text{Open}(\mathbf{pp}, \mathcal{C}, g)$: verifies the opening of commitment \mathcal{C} to the ℓ -variate multilinear polynomial $g \in \mathbb{F}[\ell]$; outputs $b \in \{0, 1\}$.
- $b \leftarrow \text{Eval}(\mathbf{pp}, \mathcal{C}, r, v, \ell, g)$ is a protocol between a PPT prover \mathcal{P} and verifier \mathcal{V} . Both \mathcal{V} and \mathcal{P} hold a commitment \mathcal{C} , the number of variables ℓ , a scalar $v \in \mathbb{F}$, and $r \in \mathbb{F}^\ell$. \mathcal{P} additionally knows an ℓ -variate multilinear polynomial $g \in \mathbb{F}[\ell]$. \mathcal{P} attempts to convince \mathcal{V} that $g(r) = v$. At the end of the protocol, \mathcal{V} outputs $b \in \{0, 1\}$.

Definition 6. A tuple of four protocols $(\text{Gen}, \text{Commit}, \text{Open}, \text{Eval})$ is an extractable polynomial commitment scheme for multilinear polynomials over a finite field \mathbb{F} if the following conditions hold.

- **Completeness.** For any ℓ -variate multilinear polynomial $g \in \mathbb{F}[\ell]$,

$$\Pr \left\{ \begin{array}{l} \mathbf{pp} \leftarrow \text{Gen}(1^\lambda, \ell); \mathcal{C} \leftarrow \text{Commit}(\mathbf{pp}, g); \\ \text{Eval}(\mathbf{pp}, \mathcal{C}, r, v, \ell, g) = 1 \wedge v = g(r) \end{array} \right\} \geq 1 - \text{negl}(\lambda)$$

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

$$\Pr \left\{ \begin{array}{l} \text{pp} \leftarrow \text{Gen}(1^\lambda, m); (\mathcal{C}, g_0, g_1) = \mathcal{A}(\text{pp}); \\ b_0 \leftarrow \text{Open}(\text{pp}, \mathcal{C}, g_0); b_1 \leftarrow \text{Open}(\text{pp}, \mathcal{C}, g_1): \\ b_0 = b_1 \neq 0 \wedge g_0 \neq g_1 \end{array} \right\} \leq \text{negl}(\lambda)$$

- **Knowledge soundness.** Eval is a succinct argument of knowledge for the following NP relation given $\text{pp} \leftarrow \text{Gen}(1^\lambda, \ell)$.

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

Definition 7. A polynomial commitment scheme for multilinear polynomials $\text{PC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Eval})$ is additively homomorphic if for all ℓ and public parameters pp produced from $\text{Gen}(1^\lambda, \ell)$, and for any $g_1, g_2 \in \mathbb{F}[\ell]$, $\text{Commit}(\text{pp}, g_1) + \text{Commit}(\text{pp}, g_2) = \text{Commit}(\text{pp}, g_1 + g_2)$.

3 Multi-folding schemes

Recall that a folding scheme [KST22] for a relation \mathcal{R} is a protocol between a *prover* and *verifier* in which the prover and the verifier reduce the task of checking two instances in \mathcal{R} with the same structure \mathbf{s} into the task of checking a single instance in \mathcal{R} with structure \mathbf{s} .

We introduce a generalization of folding schemes, which we refer to as *multi-folding schemes*. A multi-folding scheme is defined with respect to a pair of relations $(\mathcal{R}_1, \mathcal{R}_2)$ and size parameters μ and ν . It is an interactive protocol between a prover and a verifier in which the prover and the verifier reduce the task of checking a collection of μ instances in \mathcal{R}_1 and a collection of ν instances in \mathcal{R}_2 that all share the *same* structure \mathbf{s} into the task of checking a single instance in \mathcal{R}_1 with structure \mathbf{s} . Below, we formally define multi-folding schemes.

Definition 8 (Multi-folding schemes). Consider relations \mathcal{R}_1 and \mathcal{R}_2 over public parameters, structure, instance, and witness tuples such that \mathcal{R}_1 and \mathcal{R}_2 share the same structure, and size parameters $\mu, \nu \in \mathbb{N}$. A multi-folding scheme for $(\mathcal{R}_1, \mathcal{R}_2, \mu, \nu)$ consists of a PPT generator algorithm \mathcal{G} , a deterministic encoder algorithm \mathcal{K} , and a pair of PPT algorithms \mathcal{P} and \mathcal{V} denoting the prover and the verifier respectively, with the following interface:

- $\mathcal{G}(1^\lambda) \rightarrow \text{pp}$: on input security parameter λ , samples public parameters pp .
- $\mathcal{K}(\text{pp}, \mathbf{s}) \rightarrow (\text{pk}, \text{vk})$: on input pp , and a common structure \mathbf{s} among the instances to be folded, outputs a prover key pk and a verifier key vk .
- $\mathcal{P}(\text{pk}, (\bar{\mathbf{u}}_1, \bar{\mathbf{w}}_1), (\bar{\mathbf{u}}_2, \bar{\mathbf{w}}_2)) \rightarrow (\mathbf{u}, \mathbf{w})$: on input a vector of instances $\bar{\mathbf{u}}_1$ in \mathcal{R}_1 of size μ and a vector of instances $\bar{\mathbf{u}}_2$ in \mathcal{R}_2 of size ν all with the same structure \mathbf{s} , and corresponding witness vectors $\bar{\mathbf{w}}_1$ and $\bar{\mathbf{w}}_2$ outputs a folded instance-witness pair (\mathbf{u}, \mathbf{w}) in \mathcal{R}_1 .

- $\mathcal{V}(\text{vk}, \vec{u}_1, \vec{u}_2) \rightarrow u$: on input a vector of instances \vec{u}_1 and a vector of instances \vec{u}_2 outputs a new instance u .

Let $\langle \mathcal{P}, \mathcal{V} \rangle$ denote the interaction between \mathcal{P} and \mathcal{V} . We treat $\langle \mathcal{P}, \mathcal{V} \rangle$ as a function that takes as input $((\text{pk}, \text{vk}), (\vec{u}_1, \vec{u}_2), (\vec{w}_1, \vec{w}_2))$ and runs the interaction on prover input $(\text{pk}, (\vec{u}_1, \vec{u}_2), (\vec{w}_1, \vec{w}_2))$ and verifier input $(\text{vk}, (\vec{u}_1, \vec{u}_2))$. At the end of interaction $\langle \mathcal{P}, \mathcal{V} \rangle$ outputs (u, w) where u is the verifier's output folded instance, and w is the prover's output folded witness.

Let $\mathcal{R}^{(n)}$ be the relation such that $(\text{pp}, \text{s}, \vec{u}, \vec{w}) \in \mathcal{R}^{(n)}$ if and only if $(\text{pp}, \text{s}, \vec{u}_i, \vec{w}_i) \in \mathcal{R}$ for all $i \in [n]$. A multi-folding scheme for $(\mathcal{R}_1, \mathcal{R}_2, \mu, \nu)$ satisfies the following requirements.

1. *Perfect Completeness*: For all PPT adversaries \mathcal{A} , we have that

$$\Pr \left[(\text{pp}, \text{s}, u, w) \in \mathcal{R}_1 \left| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\text{s}, (\vec{u}_1, \vec{u}_2), (\vec{w}_1, \vec{w}_2)) \leftarrow \mathcal{A}(\text{pp}), \\ (\text{pp}, \text{s}, \vec{u}_1, \vec{w}_1) \in \mathcal{R}_1^{(\mu)}, (\text{pp}, \text{s}, \vec{u}_2, \vec{w}_2) \in \mathcal{R}_2^{(\nu)}, \\ (\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \text{s}), \\ (u, w) \leftarrow \langle \mathcal{P}, \mathcal{V} \rangle((\text{pk}, \text{vk}), (\vec{u}_1, \vec{u}_2), (\vec{w}_1, \vec{w}_2)) \end{array} \right. \right] = 1.$$

2. *Knowledge Soundness*: For any expected polynomial-time adversaries \mathcal{A} and \mathcal{P}^* there is an expected polynomial-time extractor \mathcal{E} such that over all randomness ρ

$$\Pr \left[\begin{array}{l} (\text{pp}, \text{s}, \vec{u}_1, \vec{w}_1) \in \mathcal{R}_1^{(\mu)}, \\ (\text{pp}, \text{s}, \vec{u}_2, \vec{w}_2) \in \mathcal{R}_2^{(\nu)} \end{array} \left| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\text{s}, (\vec{u}_1, \vec{u}_2), \text{st}) \leftarrow \mathcal{A}(\text{pp}, \rho), \\ (\vec{w}_1, \vec{w}_2) \leftarrow \mathcal{E}(\text{pp}, \rho) \end{array} \right. \right] \approx$$

$$\Pr \left[(\text{pp}, \text{s}, u, w) \in \mathcal{R}_1 \left| \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\text{s}, (\vec{u}_1, \vec{u}_2), \text{st}) \leftarrow \mathcal{A}(\text{pp}, \rho), \\ (\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \text{s}), \\ (u, w) \leftarrow \langle \mathcal{P}^*, \mathcal{V} \rangle((\text{pk}, \text{vk}), (\vec{u}_1, \vec{u}_2), \text{st}) \end{array} \right. \right]$$

3. *Efficiency*: 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.

A multi-folding scheme is secure in the random oracle model if the above requirements hold when all parties are provided access to a random oracle.

Definition 9 (Zero-knowledge). Let $\text{trace}(\langle \mathcal{P}, \mathcal{V} \rangle, \text{input})$ denote the non-deterministic function which takes as input an interaction function $\langle \mathcal{P}, \mathcal{V} \rangle$ and a prescribed input input , and produces an interaction transcript between \mathcal{P} and \mathcal{V} on input input . A multi-folding scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ for $(\mathcal{R}_1, \mathcal{R}_2, \mu, \nu)$ satisfies zero-knowledge if there exists a PPT simulator \mathcal{S} such that for all PPT adversaries \mathcal{A} and \mathcal{V}^* , and

input randomness ρ

$$\left\{ \begin{array}{l} (\text{pp}, \text{s}, (\vec{u}_1, \vec{u}_2), \text{tr}, \rho) \\ \left. \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\text{s}, (\vec{u}_1, \vec{w}_1), (\vec{u}_2, \vec{w}_2)) \leftarrow \mathcal{A}(\text{pp}), \\ (\text{pp}, \text{s}, \vec{u}_1, \vec{w}_1) \in \mathcal{R}_1^{(\mu)}, (\text{pp}, \text{s}, \vec{u}_2, \vec{w}_2) \in \mathcal{R}_2^{(\nu)}, \\ (\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \text{s}), \\ \text{tr} \leftarrow \text{trace}(\langle \mathcal{P}, \mathcal{V} \rangle, ((\text{pk}, \text{vk}), (\vec{u}_1, \vec{u}_2), (\vec{w}_1, \vec{w}_2))) \end{array} \right\} \cong \left\{ \begin{array}{l} (\text{pp}, \text{s}, (\vec{u}_1, \vec{u}_2), \text{tr}, \rho) \\ \left. \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\text{s}, (\vec{u}_1, \vec{w}_1), (\vec{u}_2, \vec{w}_2)) \leftarrow \mathcal{A}(\text{pp}), \\ (\text{pp}, \text{s}, \vec{u}_1, \vec{w}_1) \in \mathcal{R}_1, (\text{pp}, \text{s}, \vec{u}_2, \vec{w}_2) \in \mathcal{R}_2, \\ (\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \text{s}), \\ \text{tr} \leftarrow \mathcal{S}(\text{pp}, \text{s}, (\vec{u}_1, \vec{u}_2), \rho) \end{array} \right\}$$

Definition 10 (Non-interactive). A multi-folding scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ is non-interactive if the interaction between \mathcal{P} and \mathcal{V} consists of a single message from \mathcal{P} to \mathcal{V} . This single message is denoted as \mathcal{P} 's output and as \mathcal{V} 's input.

Definition 11 (Public coin). A multi-folding scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ is called public coin if all the messages sent from \mathcal{V} to \mathcal{P} are sampled from a uniform distribution.

3.1 Achieving non-interactivity for multi-folding schemes

Construction 1 (Fiat-Shamir transformation for multi-folding schemes).

Consider a public-coin multi-folding scheme $\Pi = (\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ for $(\mathcal{R}_1, \mathcal{R}_2, \mu, \nu)$ with ℓ rounds. Let ρ denote a random oracle. We construct a non-interactive multi-folding scheme $\Pi' = (\mathcal{G}', \mathcal{K}', \mathcal{P}', \mathcal{V}')$ for $(\mathcal{R}_1, \mathcal{R}_2, \mu, \nu)$ in the random oracle model as follows.

- $\mathcal{G}'(1^\lambda) \rightarrow \text{pp}$: Compute and output $\text{pp} \leftarrow \mathcal{G}(1^\lambda)$.
- $\mathcal{K}'(\text{pp}, \text{s}) \rightarrow \text{pp}$:
 1. Compute $(\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \text{s})$.
 2. Compute $\text{hs} \leftarrow \rho(\text{pp}, \text{s})$.
 3. Output $(\text{pk}', \text{vk}') \leftarrow ((\text{pk}, \text{hs}), (\text{vk}, \text{hs}))$.
- $\mathcal{P}'(\text{pk}', (\vec{u}_1, \vec{u}_2), (\vec{w}_1, \vec{w}_2))$:
 1. Parse pk' as (pk, hs)
 2. Run $\mathcal{P}(\text{pk}, (\vec{u}_1, \vec{u}_2), (\vec{w}_1, \vec{w}_2))$. On the i th message m_i , respond with verifier randomness $r_{i+1} \leftarrow \rho(m_i, r_i)$ where $r_1 = \text{hs}$. Let (u, w) be the output of \mathcal{P} and let $\pi = (m_1, \dots, m_\ell)$ consist of messages from \mathcal{P} .
 3. Send π to the verifier.
 4. Output (u, w) .

- $\mathcal{V}'(\text{vk}', (\vec{u}_1, \vec{u}_2))$:
 1. Parse vk' as (vk, hs)
 2. Recieve $\pi = (m_1, \dots, m_\ell)$ from the prover. Compute $r_{i+1} \leftarrow \rho(m_i, r_i)$ for $r_1 = \text{hs}$.
 3. Run $\mathcal{V}(\text{vk}, (\vec{u}_1, \vec{u}_2))$ with randomness $(r_1, \dots, r_{\ell+1})$. In round i , send the prover message m_i . Let u be the output of \mathcal{V} .
 4. Output u .

Lemma 3 (Fiat-Shamir transformation for multi-folding schemes). *Construction 1 transforms a public-coin multi-folding scheme for $(\mathcal{R}_1, \mathcal{R}_2, \mu, \nu)$ into a non-interactive multi-folding scheme for $(\mathcal{R}_1, \mathcal{R}_2, \mu, \nu)$ in the random oracle model.*

4 Customizable constraint systems

Setty et al. [STW23] recently introduced customizable constraint systems (CCS), a constraint system that simultaneously generalizes R1CS, Plonkish, and AIR without overheads. We first recall its definition and then describe variants that sections ahead will show are amenable to constructing multi-folding schemes. The definitions below are characterized by a finite field \mathbb{F} , but we leave this implicit.

Definition 12 (CCS [STW23]). *We define the customizable constraint system (CCS) relation \mathcal{R}_{CCS} as follows. Let the public parameter consists of size bounds $m, n, N, \ell, t, q, d \in \mathbb{N}$ where $n > \ell$.*

An \mathcal{R}_{CCS} structure \mathfrak{s} consists of:

- *a sequence of matrices $M_1, \dots, M_t \in \mathbb{F}^{m \times n}$ with at most $N = \Omega(\max(m, n))$ non-zero entries in total;*
- *a sequence of q multisets $[S_1, \dots, S_q]$, where an element in each multiset is from the domain $\{1, \dots, t\}$ 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} .*

An \mathcal{R}_{CCS} instance consists of public input $\mathbf{x} \in \mathbb{F}^\ell$. An \mathcal{R}_{CCS} witness consists of a vector $w \in \mathbb{F}^{n-\ell-1}$. An \mathcal{R}_{CCS} structure-instance tuple $(\mathfrak{s}, \mathbf{x})$ is satisfied by an \mathcal{R}_{CCS} witness w if

$$\sum_{i=1}^q c_i \cdot \bigcirc_{j \in S_i} M_j \cdot z = \mathbf{0},$$

where $z = (w, 1, \mathbf{x}) \in \mathbb{F}^n$, $M_j \cdot z$ denotes matrix-vector multiplication, \bigcirc denotes the Hadamard product between vectors, and $\mathbf{0}$ is an m -sized vector with entries equal to the the additive identity in \mathbb{F} .

4.1 Committed CCS

Consider a CCS structure $\mathfrak{s}_{\text{CCS}} = (m, n, N, \ell, t, q, d, [M_1, \dots, M_t], [S_1, \dots, S_t], [c_1, \dots, c_t])$.

Let $s = \log m$ and $s' = \log n$. We interpret each M_i (for $i \in [t]$) as functions with the following signature: $\{0, 1\}^s \times \{0, 1\}^{s'} \rightarrow \mathbb{F}$. For $i \in [t]$, let \widetilde{M}_i denote the MLE of M_i i.e., \widetilde{M}_i is the unique multilinear polynomial in $\log m + \log n$ variables such that $\forall x \in \{0, 1\}^s, y \in \{0, 1\}^{s'}, \widetilde{M}_i(x, y) = M_i(x, y)$. Similarly, for a purported witness $w \in \mathbb{F}^{n-\ell-1}$ let \widetilde{w} denote the unique MLE of w viewed as a function. WLOG, below, we let $|w| = \ell + 1$.

Definition 13 (Committed CCS). Let $PC = (\text{Gen}, \text{Commit}, \text{Open}, \text{Eval})$ denote an additively-homomorphic polynomial commitment scheme for multilinear polynomials over a finite field \mathbb{F} .

We define the committed customizable constraint system (CCCS) relation $\mathcal{R}_{\text{CCCS}}$ as follows. Let the public parameter consists of size bounds $m, n, N, \ell, t, q, d \in \mathbb{N}$ where $n = 2 \cdot (\ell + 1)$ and $\text{pp} \leftarrow \text{Gen}(1^\lambda, s' - 1)$. Let $s = \log m$ and $s' = \log n$.

An $\mathcal{R}_{\text{CCCS}}$ structure \mathfrak{s} 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 $N = \Omega(m)$ locations over the Boolean hypercube $\{0, 1\}^s \times \{0, 1\}^{s'}$;
- a sequence of q multisets $[S_1, \dots, S_q]$, where an element in each multiset is from the domain $\{1, \dots, t\}$ 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} .

An $\mathcal{R}_{\text{CCCS}}$ instance is (C, \mathfrak{x}) , where C is a commitment to a multilinear polynomial in $s' - 1$ variables and $\mathfrak{x} \in \mathbb{F}^\ell$. An $\mathcal{R}_{\text{CCCS}}$ witness consists of a multilinear polynomial \widetilde{w} in $s' - 1$ variables. An $\mathcal{R}_{\text{CCCS}}$ structure-instance tuple is satisfied by an $\mathcal{R}_{\text{CCCS}}$ witness if $\text{Commit}(\text{pp}, \widetilde{w}) = C$ and if $\forall x \in \{0, 1\}^s$,

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

where \widetilde{z} is an s' -variate multilinear polynomial such that $\widetilde{z}(x) = (\widetilde{w}, 1, \mathfrak{x})(x)$ for all $x \in \{0, 1\}^{s'}$.

4.2 Linearized committed CCS

Definition 14 (Linearized committed CCS). Let $PC = (\text{Gen}, \text{Commit}, \text{Open}, \text{Eval})$ denote an additively-homomorphic polynomial commitment scheme for multilinear polynomials over a finite field \mathbb{F} .

We define the linearized committed customizable constraint system (LCCS) relation $\mathcal{R}_{\text{LCCS}}$ as follows. Let the public parameter consists of size bounds

$m, n, N, \ell, t, q, d \in \mathbb{N}$ where $n = 2 \cdot (\ell + 1)$ and $\text{pp} \leftarrow \text{Gen}(1^\lambda, s' - 1)$. Let $s = \log m$ and $s' = \log n$.

An $\mathcal{R}_{\text{LCCCS}}$ structure \mathfrak{s} 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 $N = \Omega(m)$ locations over the Boolean hypercube $\{0, 1\}^s \times \{0, 1\}^{s'}$;
- a sequence of q multisets $[S_1, \dots, S_q]$, where an element in each multiset is from the domain $\{1, \dots, t\}$ 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} .

An $\mathcal{R}_{\text{LCCCS}}$ 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. An $\mathcal{R}_{\text{LCCCS}}$ witness is a multilinear polynomial \widetilde{w} in $s' - 1$ variables.

An $\mathcal{R}_{\text{LCCCS}}$ structure-instance tuple is satisfied by an $\mathcal{R}_{\text{LCCCS}}$ witness if $\text{Commit}(\text{pp}, \widetilde{w}) = C$ and if 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 an s' -variate multilinear polynomial such that $z(x) = (\widetilde{w}, u, \mathbf{x})(x)$ for all $x \in \{0, 1\}^{s'}$.

5 A multi-folding scheme for CCS

This section describes a multi-folding scheme for CCS. In particular, we provide a multi-folding scheme for $\mathcal{R}_1 = \mathcal{R}_{\text{LCCCS}}$ and $\mathcal{R}_2 = \mathcal{R}_{\text{CCCS}}$. To keep the description simple, we describe and prove the case of $\mu = \nu = 1$. However, we stress that both the construction and proofs easily generalize to the case of arbitrary values of μ and ν . In particular, the generalized version simply uses more powers of a random challenge when combining claims.

Construction 2 (A multi-folding scheme for CCS). We construct a multi-folding scheme for $(\mathcal{R}_{\text{LCCCS}}, \mathcal{R}_{\text{CCCS}}, \mu = 1, \nu = 1)$. Let $\text{PC} = (\text{Gen}, \text{Commit}, \text{Open}, \text{Eval})$ denote an additively-homomorphic polynomial commitment scheme for multilinear polynomials.

We define the generator and the encoder as follows.

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

1. Sample size bounds $m, n, N, \ell, t, q, d \in \mathbb{N}$ with $n = 2 \cdot (\ell + 1)$
2. $\text{pp}_{\text{PC}} \leftarrow \text{Gen}(1^\lambda, \log n - 1)$
3. Output $(m, n, N, \ell, t, q, d, \text{pp}_{\text{PC}})$

$\mathcal{K}(\text{pp}, ([\widetilde{M}_1, \dots, \widetilde{M}_t], [S_1, \dots, S_q], [c_1, \dots, c_q])) \rightarrow (\text{pk}, \text{vk})$:

1. $\text{pk} \leftarrow (\text{pp}, ([\widetilde{M}_1, \dots, \widetilde{M}_t], [S_1, \dots, S_q], [c_1, \dots, c_q]))$
2. $\text{vk} \leftarrow \perp$

3. Output (pk, vk)

The verifier \mathcal{V} takes a linearized committed CCS instance $(C_1, u, \mathbf{x}_1, r_x, v_1, \dots, v_t)$ and a committed CCS instance (C_2, \mathbf{x}_2) and the prover \mathcal{P} , in addition to the two instances, takes witnesses to both instances, \widetilde{w}_1 and \widetilde{w}_2 .

Let $s = \log m$ and $s' = \log n$. Let $\widetilde{z}_1 = (\widetilde{w}_1, u, \mathbf{x}_1)$ and $\widetilde{z}_2 = (\widetilde{w}_2, \mathbf{1}, \mathbf{x}_2)$.

The prover and the verifier proceed as follows.

1. $\mathcal{V} \rightarrow \mathcal{P}$: \mathcal{V} samples $\gamma \xleftarrow{\$} \mathbb{F}, \beta \xleftarrow{\$} \mathbb{F}^s$, and sends them to \mathcal{P} .
2. \mathcal{V} : Sample $r'_x \xleftarrow{\$} \mathbb{F}^s$.
3. $\mathcal{V} \leftrightarrow \mathcal{P}$: Run the sum-check protocol $c \leftarrow (\mathcal{P}, \mathcal{V}(r'_x))(g, s, d+1, \sum_{j \in [t]} \gamma^j \cdot v_j)$, where:

$$g(x) := \left(\sum_{j \in [t]} \gamma^j \cdot L_j(x) \right) + \gamma^{t+1} \cdot Q(x)$$

$$L_j(x) := \widetilde{e}q(r_x, x) \cdot \left(\sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(x, y) \cdot \widetilde{z}_1(y) \right)$$

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

4. $\mathcal{P} \rightarrow \mathcal{V}$: $((\sigma_1, \dots, \sigma_t), (\theta_1, \dots, \theta_t))$, where for all $i \in [t]$:

$$\sigma_i = \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \widetilde{z}_1(y)$$

$$\theta_i = \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \widetilde{z}_2(y)$$

5. \mathcal{V} : Compute $e_1 \leftarrow \widetilde{e}q(r_x, r'_x)$ and $e_2 \leftarrow \widetilde{e}q(\beta, r'_x)$, and abort if:

$$c \neq \left(\sum_{j \in [t]} \gamma^j \cdot e_1 \cdot \sigma_j + \gamma^{t+1} \cdot e_2 \cdot \left(\sum_{i=1}^q c_i \cdot \prod_{j \in S_i} \theta_j \right) \right)$$

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

7. \mathcal{V}, \mathcal{P} : Output the folded linearized committed CCS instance $(C', u', x', r'_x, v'_1, \dots, v'_t)$, where for all $i \in [t]$:

$$\begin{aligned} C' &\leftarrow C_1 + \rho \cdot C_2 \\ u' &\leftarrow u + \rho \cdot 1 \\ x' &\leftarrow x_1 + \rho \cdot x_2 \\ v'_i &\leftarrow \sigma_i + \rho \cdot \theta_i \end{aligned}$$

8. \mathcal{P} : Output the folded witness $\tilde{w}' \leftarrow \tilde{w}_1 + \rho \cdot \tilde{w}_2$.

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

Lemma 4 (Perfect Completeness). *Construction 2 satisfies perfect completeness.*

Proof. Consider public parameters $\text{pp} = (m, n, N, \ell, t, q, d, \text{pp}_{\text{PC}}) \leftarrow \mathcal{G}(1^\lambda)$ and let $s = \log m$ and $s' = \log n$. Consider arbitrary structure

$$\mathcal{S} = (\tilde{M}_1, \dots, \tilde{M}_t), (S_1, \dots, S_q), (c_1, \dots, c_q) \leftarrow \mathcal{A}(\text{pp}).$$

Consider prover and verifier key $(\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \mathcal{S})$. Suppose the prover and verifier are provided a linearized committed CCS instance

$$(C_1, u, x_1, r_x, v_1, \dots, v_t)$$

and committed CCS instance

$$(C_2, x_2).$$

Suppose the prover additionally is provided the the corresponding satisfying witnesses \tilde{w}_1 and \tilde{w}_2 .

Because the input linearized committed CCS instance-witness pair is satisfying, we have, for $\tilde{z}_1 = (\tilde{w}_1, u, x_1)$, that

$$\begin{aligned} v_j &= \sum_{y \in \{0,1\}^{s'}} \tilde{M}_j(r_x, y) \cdot \tilde{z}_1(y) && \text{By precondition.} \\ &= \sum_{x \in \{0,1\}^s} \tilde{e}q(r_x, x) \cdot \left(\sum_{y \in \{0,1\}^{s'}} \tilde{M}_j(x, y) \right) && \text{By Lemma 1} \\ &= \sum_{x \in \{0,1\}^s} L_j(x) && \text{By construction.} \end{aligned}$$

Moreover, because the input committed CCS instance-witness pair is satisfying, we have, for all $x \in \{0, 1\}^s$ and for $\widetilde{z}_2(x) = (\widetilde{w}, 1, x)(x)$, that

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

Treating the right-hand side of the above equation as a polynomial in x , because it is multilinear and vanishes on all $x \in \{0, 1\}^s$, we have that it must be the zero polynomial. Therefore, we have, for β sampled by the verifier, that

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

Therefore, for γ sampled by the verifier, by linearity, we have that

$$\begin{aligned} \sum_{j \in [t]} \gamma^j \cdot v_j &= \sum_{x \in \{0, 1\}^s} \left(\left(\sum_{j \in [t]} \gamma^j \cdot L_j(x) \right) + \gamma^{t+1} \cdot Q(x) \right) \\ &= \sum_{x \in \{0, 1\}^s} g(x) \quad \text{By construction.} \end{aligned}$$

Therefore, by the perfect completeness of the sum-check protocol, we have for $e_1 = \widetilde{e}q(r_x, r'_x)$, $e_2 = \widetilde{e}q(\beta, r'_x)$ and

$$\sigma_i = \sum_{y \in \{0, 1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \widetilde{z}_1(y) \quad \text{and} \quad \theta_i = \sum_{y \in \{0, 1\}^{s'}} \widetilde{M}_i(r'_x, y) \cdot \widetilde{z}_2(y)$$

that

$$\begin{aligned} c &= g(r'_x) \\ &= \left(\left(\sum_{j \in [t]} \gamma^j \cdot L_j(r'_x) \right) + \gamma^{t+1} \cdot Q(r'_x) \right) \\ &= \left(\left(\sum_{j \in [t]} \gamma^j \cdot e_1 \cdot \sigma_j \right) + \gamma^{t+1} \cdot e_2 \sum_{i \in [q]} c_i \cdot \prod_{j \in S_i} \theta_j \right). \end{aligned}$$

This implies that the verifier will not abort.

Now, consider the linearized CCS instance

$$(C_2, 1, x_2, r'_x, \theta_1, \dots, \theta_t).$$

By the precondition that the committed CCS instance (C_2, x_2) is satisfied by \tilde{w}_2 and by the definition of $\theta_1, \dots, \theta_t$ we have that this linearized CCS instance is satisfied by the witness \tilde{w}_2 .

Therefore, for random ρ sampled by the verifier, and for $C' = C_1 + \rho \cdot C_2$, $u' = u + \rho \cdot 1$, $x' = x_1 + \rho \cdot x_2$, $v'_i = \sigma_i + \rho \cdot \theta_i$, we have that the output linearized CCS instance

$$(C', u', x', r'_x, v'_1, \dots, v'_t)$$

is satisfied by the witness $\tilde{w}' = \tilde{w}_1 + \rho \cdot \tilde{w}_2$ by the linearity and the additive homomorphism property of the commitment scheme. \square

Some of our probabilistic analysis below is adapted from the proof of forking lemma for folding schemes [KST22], which itself builds on the proof of the forking lemma for interactive arguments [BCC⁺16].

Lemma 5 (Knowledge Soundness). *Construction 2 satisfies knowledge soundness.*

Proof. Consider an adversary \mathcal{A} that adaptively picks the structure and instances, and a malicious prover \mathcal{P}^* that succeeds with probability ϵ . Let $\mathbf{pp} \leftarrow \mathcal{G}(1^\lambda)$. Suppose on input \mathbf{pp} and random tape ρ , the adversary \mathcal{A} picks a CCS structure

$$s = ([M_1, \dots, M_t], [S_1, \dots, S_q], [c_1, \dots, c_q]),$$

a linearized committed CCS instance

$$\varphi_1 = (C_1, u, x_1, r_x, v_1, \dots, v_t),$$

a committed CCS instance

$$\varphi_2 = (C_2, x_2),$$

and some auxiliary state \mathbf{st} .

We construct an expected-polynomial time extractor \mathcal{E} that succeeds with probability $\epsilon - \text{negl}(\lambda)$ in obtaining satisfying witnesses for the original instances as follows.

$\mathcal{E}(\mathbf{pp}, \rho)$:

1. Obtain the output tuple from \mathcal{A} :

$$(s, \varphi_1, \varphi_2, \mathbf{st}) \leftarrow \mathcal{A}(\mathbf{pp}, \rho).$$

2. Compute $(\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, s)$.

3. Run the interaction

$$(\varphi^{(1)}, \tilde{w}^{(1)}) \leftarrow \langle \mathcal{P}^*, \mathcal{V} \rangle((\mathbf{pk}, \mathbf{vk}), \varphi_1, \varphi_2, \mathbf{st})$$

once with the final verifier challenge $\rho^{(1)} \xleftarrow{\$} \mathbb{F}$.

4. Abort if $(\mathbf{pp}, \mathbf{s}, \varphi^{(1)}, \tilde{w}^{(1)}) \notin \mathcal{R}_{\text{LCCCS}}$.

5. Rerun the interaction

$$(\varphi^{(2)}, \tilde{w}^{(2)}) \leftarrow \langle \mathcal{P}^*, \mathcal{V} \rangle((\mathbf{pk}, \mathbf{vk}), \varphi_1, \varphi_2, \mathbf{st})$$

with a different verifier's final challenge $\rho^{(2)} \xleftarrow{\$} \mathbb{F}$ while maintaining the same prior randomness. Keep doing so until $(\mathbf{pp}, \mathbf{s}, \varphi^{(2)}, \tilde{w}^{(2)}) \in \mathcal{R}_{\text{LCCCS}}$.

6. Abort if $\rho^{(1)} = \rho^{(2)}$.

7. Interpolating points $(\rho^{(1)}, \tilde{w}^{(1)})$ and $(\rho^{(2)}, \tilde{w}^{(2)})$, retrieve the witness polynomials \tilde{w}_1 and \tilde{w}_2 such that for $i \in \{1, 2\}$

$$\tilde{w}_1 + \rho^{(i)} \cdot \tilde{w}_2 = \tilde{w}^{(i)}. \quad (1)$$

8. Output $(\tilde{w}_1, \tilde{w}_2)$.

We first demonstrate that the extractor \mathcal{E} runs in expected polynomial time. Observe that \mathcal{E} runs the interaction once, and if it does not abort, keeps rerunning the interaction until \mathcal{P}^* succeeds. Thus, the expected number of times \mathcal{E} runs the interaction is

$$1 + \Pr[\text{First call to } \langle \mathcal{P}^*, \mathcal{V} \rangle \text{ succeeds}] \cdot \frac{1}{\Pr[\langle \mathcal{P}^*, \mathcal{V} \rangle \text{ succeeds}]} = 1 + \epsilon \cdot \frac{1}{\epsilon} = 2.$$

Therefore, we have that the extractor runs in expected polynomial-time.

We now analyze \mathcal{E} 's success probability. We must demonstrate that \mathcal{E} succeeds in producing \tilde{w}_1 and \tilde{w}_2 such that

$$(\mathbf{pp}, \mathbf{s}, \varphi_1, \tilde{w}_1) \in \mathcal{R}_{\text{LCCCS}} \quad \text{and} \quad (\mathbf{pp}, \mathbf{s}, \varphi_2, \tilde{w}_2) \in \mathcal{R}_{\text{CCCS}}$$

with probability $\epsilon - \text{negl}(\lambda)$.

To do so, we first show that the extractor successfully produces *some* output (i.e., does not abort) in under $\sqrt[3]{|\mathbb{F}|}$ rewinding steps with probability $\epsilon - \text{negl}(\lambda)$. Indeed, by the malicious prover's success probability, we have that the extractor does not abort in step (4) with probability ϵ . Given that the extractor does not abort in step (4), by Markov's inequality, we have that the extractor rewinds more than $\sqrt[3]{|\mathbb{F}|}$ times with probability $2/\sqrt[3]{|\mathbb{F}|}$. Thus, the probability that the extractor does not abort in step (4) and requires less than $\sqrt[3]{|\mathbb{F}|}$ rewinds is $(1 - 2/\sqrt[3]{|\mathbb{F}|}) \cdot \epsilon$.

Now suppose that that the extractor does not abort in step (4) and requires less than $\sqrt[3]{|\mathbb{F}|}$ rewinds. This ensures that the extractor tests at most $\sqrt[3]{|\mathbb{F}|}$ values

for ρ . Thus, the probability that $\rho^{(1)} \neq \rho^{(2)}$ is $1 - \sqrt[3]{|\mathbb{F}|^2}/|\mathbb{F}|$. Thus, assuming $\sqrt[3]{|\mathbb{F}|^2} \geq 2$, we have that the probability the extractor successfully produces some output in under $\sqrt[3]{|\mathbb{F}|}$ rewinding steps is

$$\left(1 - \frac{2}{\sqrt[3]{|\mathbb{F}|}}\right) \cdot \epsilon \cdot \left(1 - \frac{\sqrt[3]{|\mathbb{F}|^2}}{|\mathbb{F}|}\right) = \epsilon - \text{negl}(\lambda).$$

Next, if the extractor does not abort, we show that the extractor succeeds in producing satisfying witnesses with probability $1 - \text{negl}(\lambda)$. This brings the overall extractor success probability to $\epsilon - \text{negl}(\lambda)$.

Indeed, for $i \in \{1, 2\}$, let $\varphi^{(i)} = (C^{(i)}, u^{(i)}, x^{(i)}, r_x^{(i)}, v_1^{(i)}, \dots, v_t^{(i)})$. We first show that the retrieved polynomials are valid openings to the corresponding commitments in the instance. For $i \in \{1, 2\}$, because $\tilde{w}^{(i)}$ is a satisfying witness, by construction,

$$\begin{aligned} & \text{Commit}(\text{pp}, \tilde{w}_1) + \rho^{(i)} \cdot \text{Commit}(\text{pp}, \tilde{w}_2) \\ &= \text{Commit}(\text{pp}, \tilde{w}_1 + \rho^{(i)} \cdot \tilde{w}_2) && \text{By additive homomorphism.} \\ &= \text{Commit}(\text{pp}, \tilde{w}^{(i)}) && \text{By Equation (1).} \\ &= C^{(i)} && \text{Witness } \tilde{w}^{(i)} \text{ is a satisfying opening.} \\ &= C_1 + \rho^{(i)} \cdot C_2 && \text{By the verifier's computation.} \end{aligned}$$

Interpolating, we have that

$$\text{Commit}(\text{pp}, \tilde{w}_1) = C_1 \tag{2}$$

$$\text{Commit}(\text{pp}, \tilde{w}_2) = C_2. \tag{3}$$

Next, we must argue that \tilde{w}_1 and \tilde{w}_2 satisfy the remainder of the instances φ_1 and φ_2 respectively under the structure \mathbf{s} .

Indeed, consider $(\sigma_1, \dots, \sigma_t)$ and $(\theta_1, \dots, \theta_t)$ sent by the prover which by the extractor's construction are identical across all executions of the interaction. By the verifier's computation we have that for $i \in \{1, 2\}$ and all $j \in [t]$

$$\sigma_j + \rho^{(i)} \cdot \theta_j = v_j^{(i)} \tag{4}$$

Now, because $\tilde{w}^{(i)}$ is a satisfying witness, for $i \in \{1, 2\}$ we have for all $j \in [t]$ that

$$v_j^{(i)} = \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(r'_x, y) \cdot \tilde{z}^{(i)}(y),$$

where $\tilde{z}^{(i)} = (w^{(i)}, \widetilde{u^{(i)}}, x^{(i)})$.

However, by Equations (1) and (4), for $i \in \{1, 2\}$ and $j \in [t]$, this implies that

$$\sigma_j + \rho^{(i)} \cdot \theta_j = \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(r'_x, y) \cdot \widetilde{z}_1(y) + \rho^{(i)} \cdot \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(r'_x, y) \cdot \widetilde{z}_2(y),$$

where $\widetilde{z}_1 = (\widetilde{w}_1, u, x_1)$ and $\widetilde{z}_2 = (\widetilde{w}_2, 1, x_2)$. Interpolating, we have that, for all $j \in [t]$

$$\begin{aligned} \sigma_j &= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(r'_x, y) \cdot \widetilde{z}_1(y) \\ \theta_j &= \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(r'_x, y) \cdot \widetilde{z}_2(y) \end{aligned}$$

Thus, because that the verifier does not abort, we have that

$$\begin{aligned} c &= \left(\sum_{j \in [t]} \gamma^j \cdot e_1 \cdot \sigma_j \right) + \left(\gamma^{t+1} \cdot e_2 \cdot \sum_{i \in [q]} c_i \cdot \prod_{j \in S_i} \theta_j \right) \\ &= \left(\sum_{j \in [t]} \gamma^j \cdot \widetilde{e}q(r_x, r'_x) \cdot \sigma_j \right) + \left(\gamma^{t+1} \cdot \widetilde{e}q(\beta, r'_x) \cdot \sum_{i \in [q]} c_i \cdot \prod_{j \in S_i} \theta_j \right) \\ &= \left(\sum_{j \in [t]} \gamma^j \cdot \widetilde{e}q(r_x, r'_x) \cdot \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(r'_x, y) \cdot \widetilde{z}_1(y) \right) + \\ &\quad \left(\gamma^{t+1} \cdot \widetilde{e}q(\beta, r'_x) \cdot \sum_{i \in [q]} c_i \cdot \prod_{j \in S_i} \sum_{y \in \{0,1\}^{s'}} \widetilde{M}_j(r'_x, y) \cdot \widetilde{z}_2(y) \right) \\ &= \sum_{j \in [t]} \gamma_j \cdot L_j(r'_x) + \gamma^{t+1} \cdot Q(r'_x) \\ &= g(r'_x) \end{aligned}$$

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

$$\begin{aligned} \sum_{j \in [t]} \gamma^j \cdot v_j + \gamma^{t+1} \cdot 0 &= \sum_{x \in \{0,1\}^s} g(x) \\ &= \sum_{x \in \{0,1\}^s} \left(\sum_{j \in [t]} \gamma^j \cdot L_j(x) + \gamma^{t+1} \cdot Q(x) \right) \\ &= \sum_{j \in [t]} \gamma^j \cdot \left(\sum_{x \in \{0,1\}^s} L_j(x) \right) + \gamma^{t+1} \cdot \sum_{x \in \{0,1\}^s} Q(x) \end{aligned}$$

By the Schwartz-Zippel lemma [Sch80], this implies that with probability $1 - O(t)/|\mathbb{F}| = 1 - \text{negl}(\lambda)$ over the choice of γ , we have

$$v_j = \sum_{x \in \{0,1\}^s} L_j(x)$$

for all $j \in [t]$ and

$$0 = \sum_{x \in \{0,1\}^s} Q(x).$$

Now, for all $j \in [t]$, we have

$$\begin{aligned} v_j &= \sum_{x \in \{0,1\}^s} L_j(x) \\ &= \sum_{x \in \{0,1\}^s} \tilde{e}q(r_x, x) \cdot \left(\sum_{y \in \{0,1\}^{s'}} M_j(x, y) \cdot \tilde{z}_1(y) \right) \\ &= \sum_{y \in \{0,1\}^{s'}} M_j(r_x, y) \cdot \tilde{z}_1(y) \end{aligned}$$

This implies that \tilde{w}_1 is a satisfying witness to φ_1 .

Finally, we have that

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

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

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

This implies that \tilde{w}_2 is a satisfying witness to φ_2 .

Thus, if the extractor does not abort, it succeeds in producing satisfying witness $(\tilde{w}_1, \tilde{w}_2)$ with probability $1 - \text{negl}(\lambda)$. \square

Assumption 1 (Non-Interactive Multi-Folding Scheme). There exists a non-interactive multi-folding scheme for $(\mathcal{R}_{\text{LCCCS}}, \mathcal{R}_{\text{CCCS}}, 1, 1)$ in the plain model.

Justification. By applying the Fiat-Shamir transformation in the Construction 1 to the multi-folding scheme in Construction 2, we obtain a non-interactive multi-folding scheme for $(\mathcal{R}_{\text{LCCCS}}, \mathcal{R}_{\text{CCCS}}, 1, 1)$ in the random oracle model. By instantiating the random oracle with an appropriate cryptographic hash function, we heuristically obtain a non-interactive multi-folding scheme for $(\mathcal{R}_{\text{LCCCS}}, \mathcal{R}_{\text{CCCS}}, 1, 1)$ in the plain model. \square

6 HyperNova: IVC from multi-folding schemes

We now describe HyperNova, a prover-efficient and recursive zkSNARK for proving incremental computations where each step is expressed with committed CCS. For simplicity, we focus on constructing an IVC scheme, but our construction extends easily to provide a natural extension of IVC to distributed computations called proof-carrying data (PCD) [BCCT13]. We leave a full specification and an analysis of HyperNova’s PCD scheme to the near-term future work.

As noted earlier, we provide two constructions, one that directly builds on our non-interactive multi-folding scheme for CCS, and another that uses the multi-folding scheme in conjunction with Nova as a black box.

6.1 A direct approach

This section provides a direct construction of IVC using the non-interactive multi-folding scheme. Our construction below is an adaptation of Nova’s IVC scheme to the case of multi-folding schemes.

Construction 3 (A direct construction of HyperNova’s IVC). Let NIFS be the non-interactive multi-folding scheme for $(\mathcal{R}_{\text{LCCCS}}, \mathcal{R}_{\text{CCCS}}, 1, 1)$. Let $(\mathbf{u}_\perp, \mathbf{w}_\perp)$ be a default trivially satisfying $\mathcal{R}_{\text{LCCCS}}$ instance-witness pair for any structure and public parameters. We construct an IVC scheme as follows.

Consider a polynomial-time function F that takes non-deterministic input and a cryptographic hash function hash . We define our augmented function F' , where all input arguments are taken as non-deterministic advice, as follows.

$F'(\mathbf{vk}, \mathbf{U}_i, \mathbf{u}_i, (i, z_0, z_i), \omega_i, \pi) \rightarrow \mathbf{x}$:

1. If $i = 0$, output $\text{hash}(\mathbf{vk}, 1, z_0, F(z_0, \omega_0), \mathbf{u}_\perp)$
2. Otherwise:
 - (a) Check that $\mathbf{u}_i.\mathbf{x} = \text{hash}(\mathbf{vk}, i, z_0, z_i, \mathbf{U}_i)$, where $\mathbf{u}_i.\mathbf{x}$ is the public IO of \mathbf{u}_i
 - (b) Compute $\mathbf{U}_{i+1} \leftarrow \text{NIFS.V}(\mathbf{vk}, \mathbf{U}_i, \mathbf{u}_i, \pi)$
 - (c) Output $\text{hash}(\mathbf{vk}, i + 1, z_0, F(z_i, \omega_i), \mathbf{U}_{i+1})$

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

$$(\mathbf{u}_{i+1}, \mathbf{w}_{i+1}) \leftarrow \text{trace}(F', (\mathbf{vk}, \mathbf{U}_i, \mathbf{u}_i, (i, z_0, z_i), \omega_i, \pi))$$

denote the satisfying $\mathcal{R}_{\text{CCCS}}$ instance-witness pair $(\mathbf{u}_{i+1}, \mathbf{w}_{i+1})$ for the execution of F' on non-deterministic advice $(\mathbf{vk}, \mathbf{U}_i, \mathbf{u}_i, (i, z_0, z_i), \omega_i, \pi)$.

We define the IVC scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ as follows.

$\mathcal{G}(1^\lambda) \rightarrow \text{pp}$: Output $\text{NIFS.G}(1^\lambda)$.

$\mathcal{K}(\text{pp}, F) \rightarrow (\mathbf{pk}, \mathbf{vk})$:

1. Compute $(\mathbf{pk}_{\text{fs}}, \mathbf{vk}_{\text{fs}}) \leftarrow \text{NIFS.K}(\text{pp}, F')$
2. Output $(\mathbf{pk}, \mathbf{vk}) \leftarrow ((F, \mathbf{pk}_{\text{fs}}), (F, \mathbf{vk}_{\text{fs}}))$.

$\mathcal{P}(\mathbf{pk}, (i, z_0, z_i), \omega_i, \Pi_i) \rightarrow \Pi_{i+1}$:

1. Parse Π_i as $((\mathbf{U}_i, \mathbf{W}_i), (\mathbf{u}_i, \mathbf{w}_i))$
2. $(\mathbf{U}_{i+1}, \mathbf{W}_{i+1}, \pi) \leftarrow$ if $i = 0$ $\{(\mathbf{u}_\perp, \mathbf{w}_\perp, \perp)\}$ else $\{\text{NIFS.P}(\mathbf{pk}, (\mathbf{U}_i, \mathbf{W}_i), (\mathbf{u}_i, \mathbf{w}_i))\}$
3. Compute $(\mathbf{u}_{i+1}, \mathbf{w}_{i+1}) \leftarrow \text{trace}(F', (\mathbf{vk}, \mathbf{U}_i, \mathbf{u}_i, (i, z_0, z_i), \omega_i, \pi))$
4. Output $\Pi_{i+1} \leftarrow ((\mathbf{U}_{i+1}, \mathbf{W}_{i+1}), (\mathbf{u}_{i+1}, \mathbf{w}_{i+1}))$

$\mathcal{V}(\mathbf{vk}, (i, z_0, z_i), \Pi_i) \rightarrow \{0, 1\}$:

1. If $i = 0$, check that $z_i = z_0$
2. Otherwise:
 - (a) Parse Π_i as $((\mathbf{U}_i, \mathbf{W}_i), (\mathbf{u}_i, \mathbf{w}_i))$
 - (b) Check that $\mathbf{u}_i.x = \text{hash}(\mathbf{vk}, i, z_0, z_i, \mathbf{U}_i)$
 - (c) Check that \mathbf{W}_i and \mathbf{w}_i are satisfying witnesses to \mathbf{U}_i and \mathbf{u}_i respectively

Theorem 2 (A direct construction of HyperNova's IVC). *Construction 3 is an IVC scheme.*

Lemma 6 (Completeness). *Construction 3 is an IVC scheme that satisfies completeness.*

Proof. Consider $(F, i+1, z_0, z_{i+1})$ and the corresponding inputs (z_i, ω_i) such that

$$z_{i+1} = F(z_i, \omega_i)$$

Let $\text{pp} \leftarrow \mathcal{G}(1^\lambda)$, and let $(\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\text{pp}, F)$. Now consider a proof Π_i such that

$$\mathcal{V}(\mathbf{vk}, i, z_0, z_i, \Pi_i) = 1$$

We must show that given

$$\Pi_{i+1} \leftarrow \mathcal{P}(\text{pk}, (i, z_0, z_i), \omega_i, \Pi_i)$$

that

$$\mathcal{V}(\text{vk}, i + 1, z_0, z_{i+1}, \Pi_{i+1}) = 1$$

with probability 1. We show this by induction on i .

Base Case ($i = 0$): Suppose the prover is provided Π_0 such that

$$\mathcal{V}(\text{vk}, 0, z_0, z_0, \Pi_0) = 1.$$

By the base case of \mathcal{P} and F' , we have

$$\Pi_1 = ((\mathbf{u}_\perp, \mathbf{w}_\perp), (\mathbf{u}_1, \mathbf{w}_1))$$

for some $(\mathbf{u}_1, \mathbf{w}_1)$. By definition, the instance-witness pair $(\mathbf{u}_\perp, \mathbf{w}_\perp)$ is satisfying. Moreover, by construction, $(\mathbf{u}_1, \mathbf{w}_1)$ must also be satisfying. Additionally, by the construction of F' , we have

$$\mathbf{u}_1.x = \text{hash}(\text{vk}, 1, z_0, F(z_0, w_0), \mathbf{u}_\perp).$$

Therefore, we have

$$\mathcal{V}(\text{pp}, 1, z_0, z_1, \Pi_1) = 1.$$

Inductive Step ($i \geq 1$): Assume that for

$$\Pi_i = ((\mathbf{U}_i, \mathbf{W}_i), (\mathbf{u}_i, \mathbf{w}_i))$$

we have that

$$\mathcal{V}(\text{vk}, i, z_0, z_i, \Pi_i) = 1$$

and suppose that

$$\Pi_{i+1} = ((\mathbf{U}_{i+1}, \mathbf{W}_{i+1}), (\mathbf{u}_{i+1}, \mathbf{w}_{i+1})) \leftarrow \mathcal{P}(\text{pk}, (i, z_0, z_i), \omega_i, \Pi_i).$$

By the construction of \mathcal{P} , we have that

$$(\mathbf{U}_{i+1}, \mathbf{W}_{i+1}, \pi) \leftarrow \text{NIFS.P}(\text{pk}, (\mathbf{U}_i, \mathbf{W}_i), (\mathbf{u}_i, \mathbf{w}_i)).$$

Thus, by the completeness of the underlying folding scheme, and the premise that $(\mathbf{u}_i, \mathbf{w}_i)$ and $(\mathbf{U}_i, \mathbf{W}_i)$ are satisfying instance-witness pairs, we have that $(\mathbf{U}_{i+1}, \mathbf{W}_{i+1})$ is a satisfying instance-witness pair. Additionally, by the premise, we have that $\mathbf{u}_i.x = \text{hash}(\text{vk}, i, z_0, z_i, \mathbf{U}_i)$. Therefore, \mathcal{P} can construct a satisfying

instance-witness pair $(\mathbf{u}_{i+1}, \mathbf{w}_{i+1})$ that represents the correct execution of F' on input $(\mathbf{U}, \mathbf{u}, (i, z_0, z_i), \omega_i, \pi)$. By construction, this particular input implies that

$$\begin{aligned} \mathbf{u}_{i+1}.x &= \text{hash}(\text{vk}, \mathbf{u}_i.i + 1, \mathbf{u}_i.z_0, F(z_i, w_i), \text{NIFS.V}(\text{vk}, \mathbf{U}_i, \mathbf{u}_i, \pi)) \\ &= \text{hash}(\text{vk}, \mathbf{u}_i.i + 1, \mathbf{u}_i.z_0, z_{i+1}, \mathbf{U}_{i+1}) \end{aligned} \quad (5)$$

by the correctness of the underlying folding scheme. Thus, by Equation (5) we have

$$\mathcal{V}(\text{vk}, i + 1, z_0, z_{i+1}, \Pi_{i+1}) = 1.$$

□

Lemma 7 (Knowledge Soundness). *Construction 3 is an IVC scheme that satisfies knowledge soundness.*

Proof. Our approach is inspired by a recursive extraction technique described by Bünz et al [BCL⁺21]. Let $\mathbf{pp} \leftarrow \mathcal{G}(1^\lambda)$. Consider an expected polynomial-time adversary \mathcal{P}^* that outputs a function F on input \mathbf{pp} , and let $(\mathbf{pk}, \text{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, F)$. Suppose that, for a constant n , \mathcal{P}^* additionally outputs (z_0, z, Π) such that

$$\mathcal{V}(\text{vk}, n, z_0, z, \Pi) = 1$$

with probability ϵ . We must construct an expected polynomial-time extractor \mathcal{E} that with input (\mathbf{pp}, z_0, z) , outputs $(\omega_0, \dots, \omega_{n-1})$ such that by computing

$$z_i \leftarrow F(z_{i-1}, \omega_{i-1})$$

we have that $z_n = z$ with probability $\epsilon - \text{negl}(\lambda)$.

We show inductively that \mathcal{E} can construct an expected polynomial-time extractor $\mathcal{E}_i(\mathbf{pp})$ that outputs $((z_i, \dots, z_{n-1}), (\omega_i, \dots, \omega_{n-1}), \Pi_i)$ such that for all $j \in \{i + 1, \dots, n\}$,

$$z_j = F(z_{j-1}, \omega_{j-1})$$

and

$$\mathcal{V}(\text{vk}, i, z_0, z_i, \Pi_i) = 1 \quad (6)$$

for $z_n = z$ with probability $\epsilon - \text{negl}(\lambda)$. Then, because in the base case when $i = 0$, \mathcal{V} checks that $z_0 = z_i$, the values $(\omega_0, \dots, \omega_{n-1})$ retrieved by $\mathcal{E}_0(\mathbf{pp})$ are such that computing $z_{i+1} = F(z_i, \omega_i)$ for all $i \geq 1$ gives $z_n = z$. At a high level, to construct an extractor \mathcal{E}_{i-1} , we first assume the existence of \mathcal{E}_i that satisfies the inductive hypothesis. We then use $\mathcal{E}_i(\mathbf{pp})$ to construct an adversary for the non-interactive folding scheme (which we denote as $\tilde{\mathcal{P}}_{i-1}$). This in turn guarantees an extractor for the non-interactive folding scheme, which we denote as $\tilde{\mathcal{E}}_{i-1}$. We then use $\tilde{\mathcal{E}}_{i-1}$ to construct \mathcal{E}_{i-1} that satisfies the inductive hypothesis.

In the base case, for $i = n$, let $\mathcal{E}_n(\mathbf{pp}; \rho)$ output (\perp, \perp, Π_n) where Π_n is the output of $\mathcal{P}^*(\mathbf{pp}; \rho)$. By the premise, \mathcal{E}_n succeeds with probability ϵ in expected polynomial-time.

For $i \geq 1$, suppose \mathcal{E} can construct an expected polynomial-time extractor \mathcal{E}_i that outputs $((z_i, \dots, z_{n-1}), (\omega_i, \dots, \omega_{n-1}))$, and Π_i that satisfies the inductive hypothesis. To construct an extractor \mathcal{E}_{i-1} , \mathcal{E} first constructs an adversary $\tilde{\mathcal{P}}_{i-1}$ for the non-interactive folding scheme as follows:

$\tilde{\mathcal{P}}_{i-1}(\mathbf{pp}; \rho)$:

1. Let $((z_i, \dots, z_{n-1}), (\omega_i, \dots, \omega_{n-1}), \Pi_i) \leftarrow \mathcal{E}_i(\mathbf{pp}; \rho)$.
2. Parse Π_i as $((\mathbf{U}_i, \mathbf{W}_i), (\mathbf{u}_i, \mathbf{w}_i))$.
3. Parse \mathbf{w}_i to retrieve $(\mathbf{U}_{i-1}, \mathbf{u}_{i-1}, \pi_{i-1})$.
4. Output $(\mathbf{U}_{i-1}, \mathbf{u}_{i-1})$ and $((\mathbf{U}_i, \mathbf{W}_i), \pi_{i-1})$.

By the inductive hypothesis, we have that $\mathcal{V}(\mathbf{vk}, i, z_0, z_i, \Pi_i) = 1$, where $\Pi_i \leftarrow \mathcal{E}_i(\mathbf{pp})$ with probability $\epsilon - \text{negl}(\lambda)$. Therefore, by the the verifier's checks we have that $(\mathbf{u}_i, \mathbf{w}_i)$ and $(\mathbf{U}_i, \mathbf{W}_i)$ are satisfying instance-witness pairs, and that

$$\mathbf{u}_i \cdot \mathbf{x} = \text{hash}(\mathbf{vk}, i, z_0, z_i, \mathbf{U}_i).$$

Therefore, we have that \mathbf{w}_i is indeed a satisfying assignment for F' (and not just a trivially satisfying witness). Then, by the construction of F' and the binding property of the hash function, we have that

$$\mathbf{U}_i = \text{NIFS.V}(\mathbf{vk}, \mathbf{U}_{i-1}, \mathbf{u}_{i-1}, \pi_{i-1})$$

with probability $\epsilon - \text{negl}(\lambda)$. Thus, $\tilde{\mathcal{P}}_{i-1}$ succeeds in producing an accepting folded instance-witness pair $(\mathbf{U}_i, \mathbf{W}_i)$, for instances $(\mathbf{U}_{i-1}, \mathbf{u}_{i-1})$, with probability $\epsilon - \text{negl}(\lambda)$ in expected polynomial-time.

Then, by the knowledge soundness of the underlying non-interactive multi-folding scheme (Assumption 1) there exists an extractor $\tilde{\mathcal{E}}_{i-1}$ that outputs $(\mathbf{w}_{i-1}, \mathbf{W}_{i-1})$ such that $(\mathbf{u}_{i-1}, \mathbf{w}_{i-1})$ and $(\mathbf{U}_{i-1}, \mathbf{W}_{i-1})$ satisfy F' with probability $\epsilon - \text{negl}(\lambda)$ in expected polynomial-time.

Given an expected polynomial-time $\tilde{\mathcal{P}}_{i-1}$ and an expected polynomial-time $\tilde{\mathcal{E}}_{i-1}$, \mathcal{E} constructs an expected polynomial time \mathcal{E}_{i-1} as follows

$\mathcal{E}_{i-1}(\mathbf{pp}; \rho)$:

1. $((\mathbf{U}_{i-1}, \mathbf{u}_{i-1}), (\mathbf{U}_i, \mathbf{W}_i), \pi_{i-1}) \leftarrow \tilde{\mathcal{P}}_{i-1}(\mathbf{pp}; \rho)$
2. Retrieve $((z_i, \dots, z_{n-1}), (\omega_i, \dots, \omega_{n-1}), \Pi_i)$ from the internal state of $\tilde{\mathcal{P}}_{i-1}$.
3. Parse $\Pi_i \cdot \mathbf{w}_i$ to retrieve z_{i-1} and ω_{i-1}
4. Let $(\mathbf{w}_{i-1}, \mathbf{W}_{i-1}) \leftarrow \tilde{\mathcal{E}}_{i-1}(\mathbf{pp})$.

5. Let $\Pi_{i-1} \leftarrow ((U_{i-1}, W_{i-1}), (u_{i-1}, w_{i-1}))$.
6. Output $((z_{i-1}, \dots, z_{n-1}), (\omega_{i-1}, \dots, \omega_{n-1}), \Pi_{i-1})$.

We first reason that the output $(z_{i-1}, \dots, z_{n-1})$, and $(\omega_{i-1}, \dots, \omega_{n-1})$ are valid. By the inductive hypothesis, we already have that for all $j \in \{i+1, \dots, n\}$,

$$z_j = F(z_{j-1}, \omega_{j-1}),$$

and that $\mathcal{V}(\text{vk}, i, z_0, z_i, (U_i, W_i), (u_i, w_i)) = 1$ with probability $\epsilon - \text{negl}(\lambda)$. Because \mathcal{V} additionally checks that

$$u_i.x = \text{hash}(\text{vk}, i, z_0, z_i, U_i) \tag{7}$$

by the construction of F' and the binding property of the hash function, we have

$$F(z_{i-1}, \omega_{i-1}) = z_i$$

with probability $\epsilon - \text{negl}(\lambda)$. Next, we argue that Π_{i-1} is valid. Because (u_i, w_i) satisfies F' , and (U_{i-1}, u_{i-1}) were retrieved from w_i , by the binding property of the hash function, and by Equation (7), we have that

$$u_{i-1}.x = \text{hash}(\text{vk}, i-1, z_0, z_{i-1}, U_{i-1})$$

Additionally, in the case where $i = 1$, by the base case check of F' , we have that $z_{i-1} = z_0$. Because $\tilde{\mathcal{E}}_{i-1}$ succeeds with probability $\epsilon - \text{negl}(\lambda)$, we have that

$$\mathcal{V}(\text{vk}, i-1, z_0, z_{i-1}, \Pi_{i-1}) = 1$$

with probability $\epsilon - \text{negl}(\lambda)$. □

Lemma 8 (Efficiency). *When instantiated with the Pedersen commitment scheme, we have that $|F'| = |F| + o(\mathbf{G} + 2 \cdot \mathbf{H} + d \cdot \log m \cdot \mathbf{F} + \log m \cdot \mathbf{R})$, where $|F|$ denotes the number of CCS constraints to encode a function F , \mathbf{G} is the number of constraints required to encode a group scalar multiplication, \mathbf{H} is the number of constraints required to encode hash, \mathbf{F} is the number of constraints to encode field operations, and \mathbf{R} is the number of constraints to encode the RO ρ .*

Proof. On input instances U and u , NIFS.V computes $U.C \leftarrow U.C + \rho \cdot u.C$, which costs a single group scalar multiplication. Verifying the non-interactive sum-check proof in the non-interactive multi-folding scheme proof requires the verifier to perform $O(d \cdot \log m)$ finite field operations and $O(\log m)$ calls to the RO to obtain challenges in the sum-check protocol. Finally, F' makes two additional calls to hash (details are in the description of F'). □

6.2 A simple approach to build HyperNova: Use Nova as a black box

We design a step circuit for Nova that runs the verifier's logic in the non-interactive multi-folding scheme for $\mathcal{R}_{\text{LCCCS}}$ and $\mathcal{R}_{\text{CCCS}}$. The step circuit is encoded with

R1CS (a popular NP-complete constraint system [GGPR13]) and proven incrementally with Nova, but the step circuit is only in charge of running the verifier of the non-interactive multi-folding scheme, in addition to simple bookkeeping. As a result, this provides an IVC scheme, where each step of the incremental computation is expressed with committed CCS. Furthermore, we achieve this with a black box use of an IVC scheme for R1CS.

In Nova, each step circuit takes as input the output of the previous step and produces the output for the current step. In HyperNova, besides the application's IO, we augment them with the latest running instance. At each recursive step, the step circuit gets as non-deterministic input a purported instance u in $\mathcal{R}_{\text{CCCS}}$ and π , where π is the prover's output in the non-interactive multi-folding scheme. The step circuit checks that the public input of u matches the application's input provided to the step circuit. If so, it runs the verifier of the non-interactive folding scheme on (vk, U, u, π) , where vk is the verifier's key and U is the latest running instance passed from the prior step. It then provides uses the public output of u and the output of the folding scheme verifier to construct the step's output.

Construction 4 (A step circuit for Nova). Let NIFS be the non-interactive multi-folding scheme for $(\mathcal{R}_{\text{LCCCS}}, \mathcal{R}_{\text{CCCS}}, 1, 1)$. Let IVC denote the Nova's IVC scheme for functions expressed as R1CS constraints.

We first define a non-deterministic polynomial-time function **step**, represented as an R1CS structure, that iteratively folds instances expressed in $\mathcal{R}_{\text{CCCS}}$.

$\text{step}(vk, U_i, z_i; (u, \pi)) \rightarrow (vk, U_{i+1}, z_{i+1})$

1. Parse $u.x$ as (in, out)
2. Check that $in = z_i$
3. Compute $U_{i+1} \leftarrow \text{NIFS.V}(vk, U_i, u, \pi)$
4. Output (vk, U_{i+1}, out)

Given F , expressed as an $\mathcal{R}_{\text{CCCS}}$ structure, we define the corresponding IVC scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$, which uses Nova in a black-box manner.

$\mathcal{G}(1^\lambda) \rightarrow pp$: Output $(\text{NIFS.G}(1^\lambda), \text{IVC.G}(1^\lambda))$

$\mathcal{K}((pp_{\text{NIFS}}, pp_{\text{IVC}}), F) \rightarrow (pk, vk)$:

1. Compute $(pk_{\text{NIFS}}, vk_{\text{NIFS}}) \leftarrow \text{NIFS.K}(pp_{\text{NIFS}}, F)$
2. Compute $(pk_{\text{IVC}}, vk_{\text{IVC}}) \leftarrow \text{IVC.K}(pp_{\text{IVC}}, \text{step})$
3. Output $(pk, vk) \leftarrow ((F, pk_{\text{NIFS}}, vk_{\text{NIFS}}, pk_{\text{IVC}}), (\text{step}, pp_{\text{NIFS}}, vk_{\text{NIFS}}, vk_{\text{IVC}}))$.

$\mathcal{P}(pk, (i, z_0, z_i), \omega_i, \Pi_i) \rightarrow \Pi_{i+1}$:

1. Parse Π_i as (Π'_i, U_i, W_i)

2. Compute $(\mathbf{u}_i, \mathbf{w}_i) \leftarrow \text{trace}(F, (z_i, \omega_i))$ (where trace is defined as in Construction 3)
3. Compute $(\mathbf{U}_{i+1}, \mathbf{W}_{i+1}, \pi_{i+1}) \leftarrow \text{NIFS.P}(\text{pk}_{\text{NIFS}}, (\mathbf{U}_i, \mathbf{W}_i), (\mathbf{u}_i, \mathbf{w}_i))$
4. Compute $\Pi'_{i+1} \leftarrow \text{IVC.P}(\text{pk}_{\text{IVC}}, i, (\text{vk}_{\text{NIFS}}, \mathbf{u}_\perp, z_0), (\text{vk}_{\text{NIFS}}, \mathbf{U}_i, z_i), (\mathbf{u}_i, \pi_{i+1}), \Pi'_i)$
5. Output $\Pi_{i+1} = (\Pi'_{i+1}, \mathbf{U}_{i+1}, \mathbf{W}_{i+1})$

$\mathcal{V}(\text{vk}, (i, z_0, z_i), \Pi_i) \rightarrow \{0, 1\}$:

1. Parse Π_i as $(\Pi'_i, \mathbf{U}_i, \mathbf{W}_i)$.
2. Check that $\text{IVC.V}(\text{vk}_{\text{IVC}}, i, (\text{vk}_{\text{NIFS}}, \mathbf{u}_\perp, z_0), (\text{vk}_{\text{NIFS}}, \mathbf{U}_i, z_i), \Pi'_i) = 1$
3. Check that $(\text{pp}_{\text{NIFS}}, \text{step}, \mathbf{U}_i, \mathbf{W}_i) \in \mathcal{R}_{\text{LCCCS}}$

Theorem 3 (A simple construction for IVC). *Construction 4 is an IVC scheme.*

7 nlookup: A lookup argument for Nova

This section describes a lookup argument, which we refer to as **nlookup**, that is suitable for use in recursive arguments such as Nova, HyperNova, and others.

Suppose that there is a table T of size n . Now consider m variables v_1, \dots, v_m in a CCS instance and we wish to enforce that those values are contained in T .

A classic approach is to store T as a Merkle tree for which the circuit gets as public input a commitment. Then to prove that a certain value is in T , the prover could supply as non-deterministic advice to the circuit a Merkle proof of inclusion, and the circuit verifies the Merkle proof of inclusion. This unfortunately requires $O(m \cdot \log n)$ hash evaluations inside the circuit, which is prohibitive. Plookup [GW20] provides an approach where the number of constraints is $O(\max(m, n))$, which is acceptable when $m \approx n$. It is unsuitable in the context of recursive SNARKs such as Nova where a particular recursive step may perform $m \ll n$ lookup operations. A recent flurry of works (e.g., see cq [EFG22] for the latest in this line of work) consider the case where $m \ll n$, but it is unclear how to adapt them to the setting of recursive SNARKs without incurring high recursion overheads.

We provide a conceptually simple and yet efficient lookup argument, that we refer to as **nlookup**. For m lookups on a table of size n entries, **nlookup** requires $O(m \log n)$ multiplications and $O(\log n)$ hash operations inside a circuit (with small constants) and the prover performs $O(n)$ finite field operations. In particular, the prover does not commit to any additional polynomials. This lookup argument is *not* suitable for accelerating bitwise operations in the circuit model of computation, but it is a perfect tool for expressing finite state machines efficiently with Nova and HyperNova (e.g., see [Sol23, §2.4]).

nlookup in a nutshell. Without loss of generality, assume that $n = 2^\ell$. We can view T as a function from $\{0, 1\}^\ell \rightarrow \mathbb{F}$. Furthermore, let \tilde{T} denote the

unique multilinear extension of the function T . In other words, \tilde{T} is a multilinear polynomial in ℓ variables where the entries in the table are evaluations of \tilde{T} over the Boolean hypercube $\{0, 1\}^\ell$. To prove m lookup operations, the prover specifies m evaluation points q_1, \dots, q_m over the Boolean hypercube such that $\tilde{T}(q_i) = v_i$ for all $i \in [m]$. This requires $O(m \log n)$ Booleanity checks in the circuit to ensure that $q_i \in \{0, 1\}^\ell$ for all $i \in [m]$.

We now devise a multi-folding scheme where the prover and the verifier fold the task of checking the correctness of m lookup operations into task of checking an evaluation of \tilde{T} at a single point in its domain. Furthermore, in our context, the circuit maintains a running claim about an evaluation of \tilde{T} and the folding scheme folds incoming lookup claims into this running claim.

Suppose that the running claim is $\tilde{T}(q_r) \stackrel{?}{=} v_r$ for some $q_r \in \mathbb{F}^{\log n}$ and $v_r \in \mathbb{F}$. At initialization, q_r can be arbitrary and $v_r \leftarrow \tilde{T}(q_r)$. Now, the folding scheme reduces the following claim to an evaluation of \tilde{T} , where $\rho \in \mathbb{F}$ is picked by the verifier at random.

$$v_r + \sum_{i=\{1, \dots, m\}} \rho^i \cdot v_i \stackrel{?}{=} \sum_{j \in \{0, 1\}^{\log n}} \tilde{e}q(q_r, j) \cdot \tilde{T}(j) + \sum_{i=\{1, \dots, m\}} \rho^i \cdot \sum_{j \in \{0, 1\}^{\log n}} \tilde{e}q(q_i, j) \cdot \tilde{T}(j)$$

The folding scheme applies the sum-check protocol and outputs a new claim about $\tilde{T}(q'_r) \stackrel{?}{=} v'_r$. The prover's work in the folding scheme is $O(n)$ finite field operations. The verifier's work in the non-interactive folding scheme is $O(\log n)$ hash and field operations. Furthermore, at the end of the sum-check protocol (i.e., inside the folding scheme), the verifier computes evaluations of m $\tilde{e}q$ polynomials at a random point, this takes $O(m \cdot \log n)$ multiplications.

7.1 Details and security proofs

Definition 15 (Polynomial Evaluation Relation). *We define the polynomial evaluation relation $\mathcal{R}_{\text{poly}}$ as follows. Let the public parameters consist of size parameter $\ell \in \mathbb{N}$. An $\mathcal{R}_{\text{poly}}$ structure consists of \tilde{T} , a multilinear polynomial in ℓ variables. An $\mathcal{R}_{\text{poly}}$ instance is $(r, v) \in (\mathbb{F}^\ell, \mathbb{F})$ where r is an evaluation point and v is a claimed evaluation. An $\mathcal{R}_{\text{poly}}$ witness is \perp . We define $\mathcal{R}_{\text{poly}}$ as follows.*

$$\mathcal{R}_{\text{poly}} = \left\{ ((\ell, \tilde{T}), (r, v), \perp) \mid \ell \in \mathbb{N}, \tilde{T} \in \mathbb{F}[X_1, \dots, X_\ell], (r, v) \in (\mathbb{F}^\ell, \mathbb{F}) \right\}.$$

Definition 16 (Lookup Relation). *We define the lookup relation $\mathcal{R}_{\text{lookup}}$ as follows. Let the public parameters consist of size parameter $\ell \in \mathbb{N}$. For vector $T \in \mathbb{F}^n$ (where $n = 2^\ell$), an $\mathcal{R}_{\text{lookup}}$ structure consists of the corresponding multilinear extension in ℓ variables, \tilde{T} . An $\mathcal{R}_{\text{lookup}}$ instance consists of value*

$v \in \mathbb{F}$. An $\mathcal{R}_{\text{lookup}}$ witness consists of index $q \in \{0, 1\}^\ell$. We define $\mathcal{R}_{\text{lookup}}$ as follows.

$$\mathcal{R}_{\text{lookup}} = \left\{ ((\ell, \tilde{T}), v, q) \mid \begin{array}{l} \ell \in \mathbb{N}, \tilde{T} \in \mathbb{F}[X_1, \dots, X_\ell], v \in \mathbb{F}, q \in \{0, 1\}^\ell \\ \tilde{T}(q) = v \end{array} \right\}.$$

We now provide a multi-folding between two relations, a polynomial evaluation instance and a collection of lookup instances.

Construction 5 (A multi-folding scheme for lookup instances). We construct a multi-folding scheme for $(\mathcal{R}_{\text{poly}}, \mathcal{R}_{\text{lookup}}, \mu = 1, \nu)$ for arbitrary $\nu \in \mathbb{N}$.

We define the generator and the encoder as follows.

- $\mathcal{G}(1^\lambda) \rightarrow \text{pp}$:
 - Sample size bound $\ell \in \mathbb{N}$.
 - Output $\text{pp} = \ell$.
- $\mathcal{K}(\text{pp}, \tilde{T} \in \mathbb{F}[\text{pp}]) \rightarrow (\text{pk}, \text{vk})$: Output $(\text{pk}, \text{vk}) = ((\text{pp}, \tilde{T}), \text{pp})$.

The prover takes as input $\text{pk} = (\text{pp} = \ell, \tilde{T})$ and the verifier take as input $\text{vk} = \text{pp} = \ell$. The verifier \mathcal{V} takes a polynomial evaluation instance (q_r, v_r) and a vector of lookup instances (v_1, \dots, v_m) . The prover \mathcal{P} , in addition to the instances, takes witnesses to the lookup instances (q_1, \dots, q_m) .

The prover and the verifier proceed as follows.

1. $\mathcal{P} \rightarrow \mathcal{V}$: (q_1, \dots, q_m) .
2. \mathcal{V} : Check that for all $i \in [m]$, $q_i \in \{0, 1\}^\ell$.
3. $\mathcal{V} \rightarrow \mathcal{P}$: \mathcal{V} samples $\rho \xleftarrow{\$} \mathbb{F}$ and send it to \mathcal{P} .
4. \mathcal{V} : Sample $q'_r \xleftarrow{\$} \mathbb{F}^s$.
5. $\mathcal{V} \leftrightarrow \mathcal{P}$: Run the sum-check protocol $c \leftarrow \langle \mathcal{P}, \mathcal{V}(q'_r) \rangle(g, \ell, 2, v_r + \sum_{i \in [m]} \rho^i \cdot v_i)$, where:

$$g(x) := \tilde{e}q(q_r, x) \cdot \tilde{T}(x) + \sum_{i \in [m]} \rho^i \cdot \tilde{e}q(q_i, x) \cdot \tilde{T}(x)$$

6. $\mathcal{P} \rightarrow \mathcal{V}$: v'_r , where $v'_r = \tilde{T}(q'_r)$.
7. \mathcal{V} : Compute $e \leftarrow \tilde{e}q(q_r, q'_r)$ and $e_i \leftarrow \tilde{e}q(q_i, q'_r)$ for all $i \in [m]$. Abort if

$$c \neq e \cdot v'_r + \sum_{i \in [m]} \rho^i \cdot e_i \cdot v'_r.$$

8. \mathcal{V}, \mathcal{P} : Output the folded polynomial evaluation instance (q'_r, v'_r) .

Theorem 4. *Construction 5 is a public-coin multi-folding scheme for $(\mathcal{R}_{\text{poly}}, \mathcal{R}_{\text{lookup}}, \mu = 1, \nu)$ for arbitrary $\nu \in \mathbb{N}$.*

Lemma 9 (Perfect Completeness). *Construction 5 satisfies perfect completeness.*

Proof. Consider public parameters $\text{pp} = \ell \leftarrow \mathcal{G}(1^\lambda)$. Consider arbitrary structure $\tilde{T} \in \mathbb{F}[X_1, \dots, X_\ell]$. Consider prover and verifier key $(\text{pk}, \text{vk}) = (\tilde{T}, \perp) \leftarrow \mathcal{K}(\text{pp}, \tilde{T})$. Suppose the prover and verifier are provided $\mathcal{R}_{\text{poly}}$ instance

$$(q_r, v_r)$$

and a vector of $\mathcal{R}_{\text{lookup}}$ instances

$$(v_1, \dots, v_m).$$

Suppose the prover is additionally provided the corresponding satisfying witnesses for the $\mathcal{R}_{\text{lookup}}$ instances

$$(q_1, \dots, q_m).$$

By the satisfiability of the input instances we have that $v_r = \tilde{T}(q_r)$ and $v_i = \tilde{T}(q_i)$ for all $i \in [m]$

Therefore, for $\rho \in \mathbb{F}$ we have that

$$\begin{aligned} v_r + \sum_{i \in [m]} \rho^i \cdot v_i &= \tilde{T}(q_r) + \sum_{i \in [m]} \rho^i \cdot \tilde{T}(q_i) && \text{By precondition.} \\ &= \sum_{x \in \{0,1\}^\ell} \left(\tilde{e}q(q_r, x) \cdot \tilde{T}(x) + \sum_{i \in [m]} \rho^i \cdot \tilde{e}q(q_i, x) \cdot \tilde{T}(x) \right) && \text{By Lemma 1} \\ &= g(x) && \text{By definition.} \end{aligned}$$

Therefore, by the completeness of the sum-check protocol, we have that $c = g(q'_r)$. Thus, for $v'_r = \tilde{T}(q'_r)$, $e = \tilde{e}q(q_r, q'_r)$, and $e_i = \tilde{e}q(q_i, q'_r)$ for all $i \in [m]$, we have that

$$\begin{aligned} c &= g(q'_r) \\ &= \tilde{e}q(q_r, q'_r) \cdot \tilde{T}(q'_r) + \sum_{i \in [m]} \rho^i \cdot \tilde{e}q(q_i, q'_r) \cdot \tilde{T}(q'_r) \\ &= e \cdot v'_r + \sum_{i \in [m]} \rho^i \cdot e_i \cdot v'_r \end{aligned}$$

Therefore, we have that the verifier does not abort.

By construction, we have that $v'_r = \tilde{T}(q'_r)$. Therefore, the folded polynomial evaluation instance is satisfying. \square

Lemma 10 (Knowledge Soundness). *Construction 5 satisfies knowledge soundness assuming that $|\mathbb{F}| = \Theta(2^\lambda)$.*

Proof. Consider an adversary \mathcal{A} that adaptively picks the structure and instances, and a malicious prover \mathcal{P}^* that succeeds with probability ϵ . Let $\mathbf{pp} = \ell \leftarrow \mathcal{G}(1^\lambda)$. Suppose on input \mathbf{pp} and random tape ρ , the adversary \mathcal{A} picks a structure $\tilde{T} \in \mathbb{F}[X_1, \dots, X_\ell]$, an $\mathcal{R}_{\text{poly}}$ instance (q_r, v_r) , a vector of $\mathcal{R}_{\text{lookup}}$ instances (v_1, \dots, v_m) , and some auxiliary state \mathbf{st} .

We construct an extractor \mathcal{E} that succeeds with probability $\epsilon - \text{negl}(\lambda)$ in obtaining satisfying witnesses for the original instances. It works as follows. On input \mathbf{pp} and ρ , \mathcal{E} first obtains the following tuple from \mathcal{A} :

$$(\tilde{T}, (q_r, v_r), (v_1, \dots, v_m), \mathbf{st}) \leftarrow \mathcal{A}(\mathbf{pp}, \rho).$$

\mathcal{E} then computes $(\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, \tilde{T})$. Next, \mathcal{E} runs

$$((q'_r, v'_r), \perp) \leftarrow \langle \mathcal{P}^*, \mathcal{V} \rangle((\mathbf{pk}, \mathbf{vk}), (q_r, v_r), (v_1, \dots, v_m), \mathbf{st})$$

and obtains the first message (q_1, \dots, q_m) from \mathcal{P}^* by parsing the corresponding transcript. The extractor \mathcal{E} outputs $(\perp, (q_1, \dots, q_m))$ as the witness. Because the extractor only runs \mathcal{P}^* once, it runs in expected polynomial-time.

We must argue that \perp is a satisfying $\mathcal{R}_{\text{poly}}$ witness for (q_r, v_r) , and that (q_1, \dots, q_m) are satisfying $\mathcal{R}_{\text{lookup}}$ witnesses for the input instances (v_1, \dots, v_m) with probability $\epsilon - \text{negl}(\lambda)$.

Suppose we have that the witness \perp output by \mathcal{P}^* is satisfying for the corresponding verifier's output (q'_r, v'_r) with probability ϵ . By definition, this means that

$$v'_r = \tilde{T}(q'_r) \tag{8}$$

with probability ϵ . Moreover, this means that the verifier does not abort with probability at least ϵ , and thus we have the following:

$$\begin{aligned} c &= e \cdot v'_r + \sum_{i \in [m]} \rho^i \cdot e_i \cdot v'_r \\ &= \tilde{e}q(q_r, q'_r) \cdot v'_r + \sum_{i \in [m]} \rho^i \cdot \tilde{e}q(q_i, q'_r) \cdot v'_r && \text{By the verifiers computation.} \\ &= \tilde{e}q(q_r, q'_r) \cdot \tilde{T}(q'_r) + \sum_{i \in [m]} \rho^i \cdot \tilde{e}q(q_i, q'_r) \cdot \tilde{T}(q'_r) && \text{By Equation 8.} \\ &= g(q'_r) && \text{By definition.} \end{aligned}$$

with probability ϵ .

Then, by the soundness of the sum-check protocol, we must have that

$$\begin{aligned}
v_r + \sum_{i \in [m]} \rho^i \cdot v_i &= \sum_{x \in \{0,1\}^\ell} g(x) \\
&= \sum_{x \in \{0,1\}^\ell} \left(\tilde{e}q(q_r, x) \cdot \tilde{T}(x) + \sum_{i \in [m]} \rho^i \cdot \tilde{e}q(q_i, x) \cdot \tilde{T}(x) \right) && \text{By definition.} \\
&= \tilde{T}(q_r) + \sum_{i \in [m]} \rho^i \cdot \tilde{T}(q_i) && \text{By Lemma 1}
\end{aligned}$$

with probability $\epsilon - \text{negl}(\lambda)$. By the Schwartz-Zippel lemma over ρ , this implies that $v_r = \tilde{T}(q_r)$ and $v_i = \tilde{T}(q_i)$ for all $i \in [m]$ with probability $\epsilon - \text{negl}(\lambda)$.

Moreover, by the verifier's initial check we have that $q_i \in \{0, 1\}^\ell$ for all $i \in [m]$. Therefore, we have that

$$(\mathbf{pp}, \tilde{T}, (q_r, v_r), \perp) \in \mathcal{R}_{\text{poly}}$$

and

$$(\mathbf{pp}, \tilde{T}, (v_1, \dots, v_m), (q_1, \dots, q_m)) \in \mathcal{R}_{\text{lookup}}^{(m)}.$$

with probability $\epsilon - \text{negl}(\lambda)$. Because the extractor \mathcal{E} outputs the initial message from the prover, we have that the extractor succeeds with probability $\epsilon - \text{negl}(\lambda)$. \square

Acknowledgments

We thank Justin Drake, Ariel Gabizon, Justin Thaler, and Ioanna Tzialla for helpful conversations and comments on a prior version of this paper.

References

- BCC⁺16. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In *EUROCRYPT*, 2016.
- BCCT13. Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKs and proof-carrying data. In *STOC*, 2013.
- BCL⁺21. Benedikt Bünz, Alessandro Chiesa, William Lin, Pratyush Mishra, and Nicholas Spooner. Proof-carrying data without succinct arguments. In *CRYPTO*, 2021.
- BCMS20. Benedikt Bünz, Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. Proof-carrying data from accumulation schemes. In *TCC*, 2020.
- BCTV14. Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. In *CRYPTO*, 2014.
- BDFG21. Dan Boneh, Justin Drake, Ben Fisch, and Ariel Gabizon. Halo Infinite: Recursive zk-SNARKs from any Additive Polynomial Commitment Scheme. In *CRYPTO*, 2021.
- BFS20. Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. In *EUROCRYPT*, 2020.
- BGH19. Sean Bowe, Jack Grigg, and Daira Hopwood. Halo: Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Report 2019/1021, 2019.
- BGH20. Sean Bowe, Jack Grigg, and Daira Hopwood. Halo2, 2020. <https://github.com/zcash/halo2>.
- CBBZ23. Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang. Hyperplonk: Plonk with linear-time prover and high-degree custom gates. In *EUROCRYPT*, 2023.
- CHM⁺20. Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In *EUROCRYPT*, 2020.
- Dra22. Justin Drake. ZK Whiteboard Sessions – Module Fourteen: Nova Crash Course with Justin Drake. <https://www.youtube.com/watch?v=SwonTt0QzAk>, 2022.
- EFG22. Liam Eagen, Dario Fiore, and Ariel Gabizon. cq: Cached quotients for fast lookups. *Cryptology ePrint Archive*, 2022.
- FS86. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1986.
- GGPR13. Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In *EUROCRYPT*, 2013.
- GW20. Ariel Gabizon and Zachary J Williamson. plookup: A simplified polynomial protocol for lookup tables. *Cryptology ePrint Archive*, 2020.
- GWC19. Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. ePrint Report 2019/953, 2019.
- KMT22. Dmitry Khovratovich, Mary Maller, and Pratyush Ranjan Tiwari. MinRoot: candidate sequential function for Ethereum VDF. Cryptology ePrint Archive, Paper 2022/1626, 2022.
- KP22. Abhiram Kothapalli and Bryan Parno. Algebraic reductions of knowledge. Cryptology ePrint Archive, Paper 2022/009, 2022.

- KS22. Abhiram Kothapalli and Srinath Setty. SuperNova: Proving universal machine executions without universal circuits. Cryptology ePrint Archive, 2022.
- KST22. Abhiram Kothapalli, Srinath Setty, and Ioanna Tzialla. Nova: Recursive Zero-Knowledge Arguments from Folding Schemes. In *CRYPTO*, 2022.
- LFKN90. Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. In *FOCS*, October 1990.
- MBKM19. Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updateable structured reference strings. In *CCS*, 2019.
- Moh23. Nicolas Mohnblatt. Sangria: a folding scheme for PLONK. <https://geometry.xyz/notebook/sangria-a-folding-scheme-for-plonk>, 2023.
- Sch80. Jacob T Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4), 1980.
- Set20. Srinath Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In *CRYPTO*, 2020.
- Sol23. Tomer Solberg. A brief history of lookup arguments. <https://github.com/ingonyama-zk/papers/blob/main/lookups.pdf>, 2023.
- STW23. Srinath Setty, Justin Thaler, and Riad Wahby. Customizable constraint systems for succinct arguments. Cryptology ePrint Archive, 2023.
- Tha13. Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In *CRYPTO*, 2013.
- Val08. Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In *TCC*, pages 552–576, 2008.
- VSBW13. Victor Vu, Srinath Setty, Andrew J. Blumberg, and Michael Walfish. A hybrid architecture for verifiable computation. In *SECP*, 2013.
- Zha23. Zhenfei Zhang. Origami: Fold a Plonk for Ethereum’s VDF. Cryptology ePrint Archive, 2023.