

Nova: Recursive Zero-Knowledge Arguments from Folding Schemes

Abhiram Kothapalli[†] Srinath Setty^{*} Ioanna Tzialla[‡]

[†]Carnegie Mellon University ^{*}Microsoft Research [‡]New York University

Abstract. We introduce a new approach to realize incrementally verifiable computation (IVC), in which the prover recursively proves the correct execution of incremental computations of the form $y = F^{(\ell)}(x)$, where F is a (potentially non-deterministic) computation, x is the input, y is the output, and $\ell > 0$. Unlike prior approaches to realize IVC, our approach avoids succinct non-interactive arguments of knowledge (SNARKs) entirely and arguments of knowledge in general. Instead, we introduce and employ *folding schemes*, a weaker, simpler, and more efficiently-realizable primitive, which reduces the task of checking two instances in some relation to the task of checking a single instance. We construct a folding scheme for a characterization of NP and show that it implies an IVC scheme with improved efficiency characteristics: (1) the “recursion overhead” (i.e., the number of steps that the prover proves in addition to proving the execution of F) is a constant and it is dominated by two group scalar multiplications expressed as a circuit (this is the smallest recursion overhead in the literature), and (2) the prover’s work at each step is dominated by two multiexponentiations of size $O(|F|)$, providing the fastest prover in the literature. The size of a proof is $O(|F|)$ group elements, but we show that using a variant of an existing zkSNARK, the prover can prove the knowledge of a valid proof succinctly and in zero-knowledge with $O(\log |F|)$ group elements. Finally, our approach neither requires a trusted setup nor FFTs, so it can be instantiated efficiently with any cycles of elliptic curves where DLOG is hard.

1 Introduction

We revisit the problem of realizing *incrementally-verifiable computation (IVC)* [46]: a cryptographic primitive that enables producing proofs of correct execution of “long running” computations such that a verifier can efficiently verify the correct execution of any prefix of the computation. IVC enables a wide variety of applications including verifiable delay functions [9, 49], succinct blockchains [13, 33], and incrementally-verifiable versions of verifiable state machines [35, 41].

A well-known approach to construct IVC is to use succinct non-interactive arguments of knowledge (SNARKs) for NP [25, 26, 32, 38]: at each incremental step i , the prover produces a SNARK proving that it has applied F correctly to the output of step $i - 1$ and that the SNARK verifier *represented as a circuit* has accepted the SNARK from step $i - 1$ [6, 8]. However, it is well-known that

this approach is impractical [6, 20]. Alternatively, one can use SNARKs without trusted setup [18, 22, 40, 43] but their verifiers are more expensive than those of SNARKs with trusted setup, both asymptotically and concretely. Recent works [10, 12, 16, 17] aim to address the inefficiency of SNARK-based IVC, with an innovative approach: at each step, the verifier circuit “defers” expensive steps in verifying a SNARK for NP instances (e.g., verifying polynomial evaluation proofs) by accumulating those steps into a single instance that is later checked efficiently. However, these works still require the prover to produce a SNARK at each step and the verifier circuit to partially verify that SNARK.

We introduce a new approach that avoids SNARKs (and more generally arguments of knowledge) entirely and relies purely on deferral to realize IVC. In a nutshell, instead of accumulating expensive steps of verifying a SNARK for NP instances, the verifier circuit in our approach accumulates the NP instances themselves. We formalize this technique as a new and minimal primitive, which we refer to as a *folding scheme*. A folding scheme is weaker, simpler, and far more efficient compared to arguments of knowledge including SNARKs. Indeed, realizing IVC via folding schemes results in improved efficiency over prior work (Figure 3): (1) the verifier circuit is constant-sized and its size is dominated by two group scalar multiplications; this is the smallest verifier circuit in the literature (in the context of recursive proof composition); and (2) the prover’s work at each step is dominated by two multiexponentiations of size $O(|F|)$, providing the fastest prover in the literature, both asymptotically and concretely. Section 1.4 provides a detailed comparison between our approach and prior work.

1.1 Folding Schemes

A folding scheme is defined with respect to an NP relation, and it is a protocol between an untrusted *prover* and a *verifier*. Both entities hold two N -sized NP instances, and the prover in addition holds purported witnesses for both instances. The protocol enables the prover and the verifier to output a single N -sized NP instance, which we refer to as a *folded instance*. Furthermore, the prover privately outputs a purported witness to the folded instance using purported witnesses for the original instances. Informally, a folding scheme guarantees that the folded instance is satisfiable only if the original instances are satisfiable. A folding scheme is said to be *non-trivial* if the verifier’s costs and the communication are lower in the case where the verifier participates in the folding scheme and then verifies a purported NP witness for the folded instance than the case where the verifier verifies purported NP witnesses for each of the original instances.

Several existing techniques exhibit the two-to-one reduction pattern of folding schemes. Examples include the sumcheck protocol [37] and the split-and-fold techniques in inner product arguments [11]. Appendix A provides further details.

Remark 1 (Folding Schemes vs. SNARKs). SNARKs for NP [7, 25, 26, 32, 38] trivially imply a folding scheme for NP: given two NP instances u_1 and u_2 and the corresponding witnesses, the prover proves u_1 by producing a SNARK. The verifier checks that SNARK and then sets u_2 to be the folded instance. However,

we construct a folding scheme for NP without relying on SNARKs (or more generally arguments of knowledge). Specifically, our folding scheme is weaker than any argument of knowledge (succinct or otherwise) because it merely *reduces* the satisfiability of two NP instances to the satisfiability of a single NP instance.¹

To design a folding scheme for NP, we start with a popular NP-complete language that generalizes arithmetic circuit satisfiability: R1CS (Definition 10). As we illustrate later, it is difficult to devise a folding scheme for R1CS. To address this, we introduce a variant of R1CS, called *relaxed R1CS*, which, like R1CS, not only characterizes NP, but, unlike R1CS, can support a folding scheme. The following theorem captures the cryptographic and efficiency characteristics of our folding scheme for relaxed R1CS.

Theorem 1. *There exists a constant-round, public-coin, zero-knowledge folding scheme for relaxed R1CS where for N -sized relaxed R1CS instances over a finite field \mathbb{F} with the same “structure” (i.e., R1CS coefficient matrices), the prover’s work is $O_\lambda(N)$, and the verifier’s work and the communication are both $O_\lambda(1)$, assuming the existence of any additively-homomorphic commitment scheme that provides $O_\lambda(1)$ -sized commitments to N -sized vectors over \mathbb{F} (e.g., Pedersen’s commitments), where λ is the security parameter.*

Because our folding scheme is public coin, it can be made non-interactive in the random oracle model using the Fiat-Shamir transform [24], and be instantiated (heuristically) in the standard model using a concrete hash function. We rely on such a non-interactive folding scheme to construct IVC.

1.2 IVC from Non-Interactive Folding Schemes

We show how to realize IVC using a non-interactive version of our folding scheme for relaxed R1CS. We refer to our construction as Nova.

Recall that an IVC is an argument of knowledge [32, 38]² for incremental computations of the form $y = F^{(\ell)}(x)$, where F is a (possibly non-deterministic) computation, $\ell > 0$, x is a public input, and y is the public output. At each incremental step, the IVC prover produces a proof that the step was computed correctly *and* it has verified a proof for the prior step. In other words, at each incremental step, the IVC prover produces a proof of satisfiability for an augmented circuit that augments the circuit for F with a “verifier circuit” that verifies the proof of the prior step. Recursively, the final proof proves the correctness of the entire incremental computation. A key aspect of IVC is that neither the IVC verifier’s work nor the IVC proof size depends on the number of

¹ This work realizes IVC using our folding scheme. As IVC implies SNARKs (e.g., see [6]), one might wonder whether folding schemes are in general weaker than SNARKs. However, existing constructions of IVC (including our own) rely on additional assumptions (§4.2), which the resulting IVC-based SNARK inherits.

² An *argument of knowledge* for circuit satisfiability enables an untrusted polynomial-time prover to prove to a verifier the knowledge of a witness w such that $\mathcal{C}(w, x) = y$, where \mathcal{C} is a circuit, x is some public input, and y is some public output.

steps in the incremental computation. In particular, the IVC verifier only verifies the proof produced at the last step of the incremental computation.

In Nova, we consider incremental computations, where each step of the incremental computation is expressed with R1CS (all the steps in the incremental computation share the same R1CS coefficient matrices). At step i of the incremental computation, as in other approaches to IVC, Nova’s prover proves that the step i was computed correctly. Furthermore, at step i , instead of verifying a proof for step $i - 1$ (as in traditional approaches to IVC), Nova’s approach treats the computation at step $i - 1$ as an R1CS instance and folds that into a running relaxed R1CS instance. Specifically, at each step, Nova’s prover proves that it has performed the step’s computation and has folded its prior step represented as an R1CS instance into a running relaxed R1CS instance. In other words, the circuit satisfiability instance that the prover proves at each incremental step computes a step of the incremental computation and includes a circuit for the computation of the verifier in the non-interactive folding scheme for relaxed R1CS.

A distinctive aspect of Nova’s approach to IVC is that it achieves the smallest “verifier circuit” in the literature. Since the verifier’s costs in the non-interactive version of the folding scheme for relaxed R1CS is $O_\lambda(1)$, the size of the computation that Nova’s prover proves at each incremental step is $\approx|F|$, assuming N -sized vectors are committed with an $O_\lambda(1)$ -sized commitments (e.g., Pedersen’s commitments). In particular, the verifier circuit in Nova is constant-sized and its size is dominated by two *group scalar multiplications*. Furthermore, Nova’s prover’s work at each step is dominated by two multiexponentiations of size $\approx|F|$. Note that Nova’s prover does not perform any FFTs, so it can be instantiated efficiently using *any* cycles of elliptic curves where DLOG is hard.

With the description thus far, the size of an IVC proof (which is a purported witness for the running relaxed R1CS instance) is $O_\lambda(|F|)$. Instead of sending such a proof to a verifier, at any point in the incremental computation, Nova’s prover can prove the knowledge of a satisfying witness to the running relaxed R1CS instance in zero-knowledge with an $O_\lambda(\log |F|)$ -sized succinct proof using a zkSNARK that we design by adapting Spartan [40]. The following theorem summarizes our key result.

Theorem 2. *For any incremental function where each step of the incremental function applies a (non-deterministic) function F , there exists an IVC scheme with the following efficiency characteristics, assuming N -sized vectors are committed with an $O_\lambda(1)$ -sized commitments.*

- *IVC proof sizes are $O(|F|)$ and the verifier’s work to verify them is $O_\lambda(|F|)$. The prover’s work at each incremental step is $\approx|F|$. Specifically, the prover’s work at each step is dominated by two multiexponentiations of size $\approx|F|$.*
- *Succinct zero-knowledge proofs of valid IVC proofs are size $O_\lambda(\log |F|)$, and the verifier’s work to verify them is either $O_\lambda(\log |F|)$ or $O_\lambda(|F|)$ depending on the commitment scheme for vectors. The prover’s work to produce this succinct zero-knowledge proof is $O_\lambda(|F|)$.*

1.3 Implementation and Performance Evaluation

We implement Nova as a library in about 6,000 lines of Rust [3]. The library is generic over a cycle of elliptic curves and a hash function (used internally as the random oracle). The library provides candidate implementations with the Pasta cycle of elliptic curves [4] and Poseidon [2, 28]. Finally, the library accepts F (i.e., a step of the incremental computation) as a bellperson gadget [1].

Recursion Overheads. We measure the size of Nova’s verifier circuit, as it determines the *recursion overhead*: the number of additional constraints that the prover must prove at each incremental step besides proving an invocation of F .

We find that Nova’s verifier circuit is $\approx 20,000$ R1CS constraints. This is the smallest verifier circuit in the literature and hence Nova incurs the lowest recursion overhead. Specifically, Nova’s recursion overhead is $> 10\times$ lower than in SNARK-based IVC [6] with state-of-the-art per-circuit trusted setup SNARK [29], and over $100\times$ smaller than with a SNARK without trusted setup [22]. Compared to recent works, Nova’s recursion overhead is over $7\times$ lower than Halo’s [12], and over $2\times$ lower than the scheme of Bunz et al. [16].

	Primary Curve	Secondary Curve
Scalar multiplications	12,362	12,362
Random oracle call	1,431	1,434
Collision-resistant hash	2,300	2,306
Non-native arithmetic	3,240	3,240
Glue code	1,251	1,782
Total	20,584	21,124

Fig. 1: A detailed breakdown of sub-routines in Nova’s verifier’s circuit and the associated number of R1CS constraints. The verifier circuit on each of the curves in the cycle are not identical as they have slightly different base cases. We find that a majority of constraints in the verifier circuit step from the group scalar multiplications.

Performance of Nova. We experiment with Nova on an Azure Standard F32s_v2 VM (16 physical CPUs, 2.70 GHz Intel(R) Xeon(R) Platinum 8168, and 64 GB memory). In our experiments, we vary the number of constraints in F . Our performance metrics are: the prover time, the verifier time, and proof sizes. We measure these for Nova’s IVC scheme as well as its Spartan-based zkSNARK to compress IVC proofs. Figure 2 depicts our results, and we find the following.

- The prover’s per-step cost to produce an IVC proof and compress it scale sub-linearly with the size of F (since the cost is dominated by two multiexponentiations, which scale sub-linearly due to the Pippenger algorithm and parallelize better at larger sizes). When $|F| \approx 2^{20}$ constraints, the prover’s per-step cost to produce an IVC proof is $\approx 1\mu\text{s}/\text{constraint}$. For the same F , the cost to produce a compressed IVC proof is $\approx 24\mu\text{s}/\text{constraint}$.³

³ If the prover produces a compressed IVC proof every ≈ 24 steps, the prover incurs at most $2\times$ overhead to compress IVC proofs. Similarly, if the prover compresses its IVC proof every ≈ 240 steps, the overhead drops to $\approx 20\%$.




Fig. 2: Performance of Nova as a function of $|F|$. See the text for details.

- Compressed IVC proofs are $\approx 8\text{--}9\text{KB}$ and are significantly shorter than IVC proofs (e.g., they are $\approx 7,400\times$ shorter when $|F| \approx 2^{20}$ constraints).
- Verifying a compressed proof is only $\approx 2\times$ higher costs than verifying a significantly longer IVC proof.

1.4 A More Detailed Comparison with Prior Work

Figure 3 compares Nova with prior approaches. Nova’s approach can be viewed as taking Halo’s approach to the extreme. Specifically:

- At each incremental step, Halo’s verifier circuit verifies a “partial” SNARK. This still requires Halo’s prover to perform $|F|$ -sized FFTs and $O(|F|)$ exponentiations (i.e., *not* an $|F|$ -sized multiexponentiation). Whereas, in Nova, the verifier circuit folds an entire NP instance representing computation at the prior step into a running relaxed R1CS instance. This only requires Nova’s prover to commit to a satisfying assignment of an $\approx |F|$ -sized circuit (which computes F and performs the verifier’s computation in a folding scheme for relaxed R1CS), so at each step, Nova’s prover only computes an $O(|F|)$ -sized multiexponentiation and does not compute any FFTs. So, Nova’s prover incurs lower costs than Halo’s prover, both asymptotically and concretely.
- The verifier circuit in Halo is of size $O_\lambda(\log |F|)$ whereas in Nova, it is $O_\lambda(1)$. Concretely, the dominant operations in Halo’s circuit is $O(\log |F|)$ group scalar multiplications, whereas in Nova, it is two group scalar multiplications.
- Halo and Nova have the same proof sizes $O_\lambda(\log |F|)$ and verifier time $O_\lambda(|F|)$.

Bünz et al. [17] apply Halo’s approach to other polynomial commitment schemes. Halo Infinite [10] generalizes the approach in Halo [12] to any homomorphic polynomial commitment scheme; they also obtain PCD (and hence IVC) even when polynomial commitment schemes do not satisfy succinctness.

Bünz et al. [16] propose a variant of the approach in Halo, where they realize PCD (and hence IVC) without relying on succinct arguments. Specifically, they first devise a non-interactive argument of knowledge (NARK) for R1CS with $O_\lambda(N)$ -sized proofs and $O_\lambda(N)$ verification times for N -sized R1CS instances. Then, they show that most of the NARK’s verifier’s computation can be deferred by performing $O_\lambda(1)$ work in the verifier circuit. For zero-knowledge, Nova relies on zero-knowledge arguments with succinct proofs, whereas their approach does

	“Verifier circuit” (dominant ops)	Prover (each step)	Proof size	Verifier	assumptions
BCTV14 [6] with [29] [†]	3 \mathbb{P}	$O(C)$ FFT $O(C)$ MSM	$O_\lambda(1)$	$O_\lambda(1)$	q-type
Spartan [40]-based IVC	$O(\sqrt{C})$ \mathbb{G}	$O(C)$ MSM	$O_\lambda(\sqrt{C})$	$O_\lambda(\sqrt{C})$	DLOG, RO
Fractal [22]	$O_\lambda(\log^2 C)$ \mathbb{F} $O(\log^2 C)$ \mathbb{H}	$O(C)$ FFT $O(C)$ MHT	$O_\lambda(\log^2 C)$	$O_\lambda(\log^2 C)$	RO
Halo [12]	$O(\log C)$ \mathbb{G}	$O(C)$ FFT $O(C)$ EXP	$O_\lambda(\log C)$	$O_\lambda(C)$	DLOG, RO
BCLMS [16]*	8 \mathbb{G}	$O(C)$ FFT $O(C)$ MSM	$O_\lambda(C)$	$O_\lambda(C)$	DLOG, RO
Nova (this work)	2 \mathbb{G}	$O(C)$ MSM	$O_\lambda(\log C)$	$O_\lambda(C)$	DLOG, RO
Nova (this work)	2 \mathbb{G}_T	$O(C)$ MSM	$O_\lambda(\log C)$	$O_\lambda(\log C)$	SXDH, RO

[†] Requires per-circuit trusted setup and is undesirable in practice
 $O(C)$ FFT: FFT over an $O(C)$ -sized vector costing $O(C \log C)$ operations over \mathbb{F}
 $O(C)$ MHT: Merkle tree over an $O(C)$ -sized vector costing $O(C)$ hash computations
 $O(C)$ EXP: $O(C)$ exponentiations in a cryptographic group
 $O(C)$ MSM: $O(C)$ -sized multi-exponentiation in a cryptographic group

Fig. 3: Asymptotic costs of Nova and its baselines to produce and verify a proof for an incremental computation where each incremental step applies a function F . C denotes the size of the computation at each incremental step, i.e., $|F| + |\mathcal{C}_V|$, where \mathcal{C}_V is the “verifier circuit” in IVC. The “verifier circuit” column depicts the number of dominant operations in \mathcal{C}_V , where \mathbb{P} denotes a pairing in a pairing-friendly group, \mathbb{F} denotes the number of finite field operations, \mathbb{H} denotes a hash computation, and \mathbb{G} denotes a scalar multiplication in a cryptographic group. The prover column depicts the cost to the prover for each step of the incremental computation, and proof sizes and verifier times refer respectively to the size of the proof of the incremental computation and the associated verification times. For Nova’s proof sizes and verification times, we depict the compressed proof sizes (otherwise, they are $O_\lambda(C)$) and the time to verify a compressed proof (otherwise, they are $O_\lambda(C)$). Rows with RO require heuristically instantiating the random oracle with a concrete hash function in the standard model.

not rely on succinct arguments. However, Nova’s approach has several conceptual and efficiency advantages over the work of Bünz et al [16]:

- Nova introduces a new primitive called a folding scheme, which is conceptually simpler and is easier to realize than prior notions such as (split) accumulation schemes used in prior work [16, 17]. Furthermore, a folding scheme for NP directly leads to IVC and is again easier to analyze than with prior notions.
- At each step, their prover performs an $O(|F|)$ -sized FFT (which costs $O(|F| \log |F|)$ operations over \mathbb{F}). Whereas, Nova does not perform any FFTs.
- Their prover’s work for multiexponentiations at each step and the size of their verifier circuit are both higher than in Nova by $\approx 4\times$.
- Proof sizes are $O_\lambda(|F|)$ in their work, whereas in Nova, they are $O_\lambda(\log |F|)$. We believe, in theory, they can also compress their proofs, using a succinct argument, but unlike Nova, they do not specify how to do so in a concretely efficient manner. Furthermore, using succinct arguments is inconsistent with their goal of not employing them.

Concurrent work. In an update concurrent with this work, Bünz et al. [16] provide an improved construction of their NARK for R1CS, which leads to an IVC that, like Nova, avoids FFTs. Furthermore, they improve the size of the verifier circuit by $\approx 2\times$, which is still larger than Nova’s verifier circuit by $\approx 2\times$. The per-step computation of the prover remains $4\times$ higher than Nova.

1.5 An Overview of the Rest of the Paper

Section 2 provides the necessary background. Section 3 formally defines folding schemes and their properties. In Section 4, we introduce a variant of R1CS called relaxed R1CS for which we provide a folding scheme satisfying Theorem 1. Then, in Section 5, we use a non-interactive version of the folding scheme (§4.2) to construct an IVC scheme and a scheme to compress IVC proofs satisfying Theorem 2 by assuming the existence of a zkSNARK for relaxed R1CS with logarithmic-sized proofs. Finally, in Section 6, we construct such a zkSNARK.

2 Preliminaries

Let \mathbb{F} denote a finite field with $|\mathbb{F}| = 2^{\Theta(\lambda)}$, where λ is the security parameter. Let \cong denote computational indistinguishability with respect to a PPT adversary. We globally assume that generator algorithms that produce public parameters are additionally provided appropriate size bounds.

2.1 A Commitment Scheme for Vectors over \mathbb{F}

We require a commitment scheme for vectors over \mathbb{F} that is additively homomorphic and succinct. We formally define these two properties and others noted below in Appendix F. Below, we define the syntax for commitment schemes.

Definition 1 (A Commitment Scheme for Vectors). *A commitment scheme for \mathbb{F}^m is a tuple of three protocols with the following interface.*

- $\text{Gen}(1^\lambda, m) \rightarrow \text{pp}$: takes length parameter m ; produces public parameters pp .
- $\text{Com}(\text{pp}, v, r) \rightarrow C$: takes vector $v \in \mathbb{F}^m$ and $r \in \mathbb{F}$; produces commitment C .
- $\text{Open}(\text{pp}, C, v, r) \rightarrow \{0, 1\}$: verifies the opening of commitment C to $v \in \mathbb{F}^m$.

A commitment scheme satisfies hiding (the commitment reveals no information), binding (a PPT adversary cannot open a commitment to two different values), and succinctness (the commitment size is logarithmic in the opening size).

2.2 Non-Interactive Arguments of Knowledge

Definition 2 (Non-Interactive Argument of Knowledge). *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 .
- $\mathcal{K}(\text{pp}, \text{s}) \rightarrow (\text{pk}, \text{vk})$: On input structure s , representing common structure among instances, outputs the prover key pk and verifier key vk .
- $\mathcal{P}(\text{pk}, u, w) \rightarrow \pi$: On input instance u and witness w , outputs a proof π proving that $(\text{pp}, \text{s}, u, w) \in \mathcal{R}$.
- $\mathcal{V}(\text{vk}, u, \pi) \rightarrow \{0, 1\}$: Checks instance u given proof π .

An argument of knowledge satisfies completeness if for any PPT adversary \mathcal{A}

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

An 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

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

Definition 3 (Zero-Knowledge). An argument of knowledge $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ for relation \mathcal{R} satisfies zero-knowledge if there exists PPT simulator \mathcal{S} such that for all PPT adversaries \mathcal{A}

$$\left\{ (\text{pp}, \text{s}, u, \pi) \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\text{s}, (u, w)) \leftarrow \mathcal{A}(\text{pp}), \\ (\text{pp}, \text{s}, u, w) \in \mathcal{R}, \\ (\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \text{s}), \\ \pi \leftarrow \mathcal{P}(\text{pk}, u, w) \end{array} \right\} \cong \left\{ (\text{pp}, \text{s}, u, \pi) \mid \begin{array}{l} (\text{pp}, \tau) \leftarrow \mathcal{S}(1^\lambda), \\ (\text{s}, (u, w)) \leftarrow \mathcal{A}(\text{pp}), \\ (\text{pp}, \text{s}, u, w) \in \mathcal{R}, \\ (\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \text{s}), \\ \pi \leftarrow \mathcal{S}(\text{pp}, u, \tau) \end{array} \right\}$$

Definition 4 (Succinctness). A non-interactive argument system is succinct if the size of the proof π is polylogarithmic in the size of the witness w .

2.3 Incrementally Verifiable Computation

Incrementally verifiable computation (IVC) [46] enables verifiable computation for repeated function application. Intuitively, for a function F , with initial input z_0 , an IVC scheme allows a prover to produce a proof Π_i for the statement $z_i = F^{(i)}(z_0)$ (i.e., i applications of F on input z_0) given a proof Π_{i-1} for the statement $z_{i-1} = F^{(i-1)}(z_0)$. Formally, IVC schemes additionally permit F to take auxiliary input ω . We recall the definition of IVC using notational conventions of modern argument systems.

Definition 5 (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. An IVC scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ satisfies perfect completeness if for any PPT adversary \mathcal{A}

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

where F is a polynomial time computable function. Likewise, an IVC scheme satisfies knowledge-soundness if for any constant $n \in \mathbb{N}$, and expected polynomial time adversaries \mathcal{P}^* there exists expected polynomial-time extractor \mathcal{E} such that for any input randomness ρ

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

An IVC scheme satisfies succinctness if the size of the IVC proof Π does not grow with the number of applications n .

We note that in the definition above, the number of steps n is treated as a fixed environment variable that characterizes the extractor. This model is required for all known general recursive techniques as they rely on recursive extractors that blowup polynomially for each additional recursive step [10, 13, 16, 17, 22]. Bitansky et al. [8] avoid such a restriction by making non-blackbox assumptions about the extractors runtime with respect to that of the malicious prover. In any case, there are no known attacks on arbitrary depth recursion.

3 Folding Schemes

This section formally defines folding schemes. Intuitively, a folding scheme for a relation \mathcal{R} is a protocol that reduces the task of checking two instances in \mathcal{R} to the task of checking a single instance in \mathcal{R} .

Definition 6 (Folding Scheme). Consider a relation \mathcal{R} over public parameters, structure, instance, and witness tuples. A folding scheme for \mathcal{R} 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 verifier respectively, with the following interface:

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

- $\mathcal{P}(\mathbf{pk}, (u_1, w_1), (u_2, w_2)) \rightarrow (u, w)$: On input instance-witness tuples (u_1, w_1) and (u_2, w_2) outputs a new instance-witness tuple (u, w) of the same size.
- $\mathcal{V}(\mathbf{vk}, u_1, u_2) \rightarrow u$: On input instances u_1 and u_2 , outputs a new instance u .

Let

$$(u, w) \leftarrow \langle \mathcal{P}(\mathbf{pk}, w_1, w_2), \mathcal{V}(\mathbf{vk}) \rangle(u_1, u_2)$$

denote the the verifier's output instance u and the prover's output witness w from the interaction of \mathcal{P} and \mathcal{V} on witnesses (w_1, w_2) , prover key \mathbf{pk} , verifier key \mathbf{vk} and instances (u_1, u_2) . Likewise, let

$$\text{tr} = \langle \mathcal{P}(\mathbf{pk}, w_1, w_2), \mathcal{V}(\mathbf{vk}) \rangle(u_1, u_2)$$

denote the corresponding interaction transcript. A folding scheme satisfies perfect completeness if for all PPT adversaries \mathcal{A}

$$\Pr \left[(\mathbf{pp}, \mathbf{s}, u, w) \in \mathcal{R} \left| \begin{array}{l} \mathbf{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\mathbf{s}, (u_1, w_1), (u_2, w_2)) \leftarrow \mathcal{A}(\mathbf{pp}), \\ (\mathbf{pp}, \mathbf{s}, u_1, w_1), (\mathbf{pp}, \mathbf{s}, u_2, w_2) \in \mathcal{R}, \\ (\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, \mathbf{s}), \\ (u, w) \leftarrow \langle \mathcal{P}(\mathbf{pk}, w_1, w_2), \mathcal{V}(\mathbf{vk}) \rangle(u_1, u_2) \end{array} \right. \right] = 1.$$

A folding scheme satisfies knowledge soundness if for any expected polynomial-time adversary \mathcal{P}^* there is an expected polynomial-time extractor \mathcal{E} such that

$$\Pr \left[\begin{array}{l} (\mathbf{pp}, \mathbf{s}, u_1, w_1) \in \mathcal{R}, \\ (\mathbf{pp}, \mathbf{s}, u_2, w_2) \in \mathcal{R} \end{array} \left| \begin{array}{l} \mathbf{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\mathbf{s}, (u_1, u_2)) \leftarrow \mathcal{P}^*(\mathbf{pp}, \rho), \\ (w_1, w_2) \leftarrow \mathcal{E}(\mathbf{pp}, \rho) \end{array} \right. \right] \geq$$

$$\Pr \left[(\mathbf{pp}, \mathbf{s}, u, w) \in \mathcal{R} \left| \begin{array}{l} \mathbf{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\mathbf{s}, (u_1, u_2)) \leftarrow \mathcal{P}^*(\mathbf{pp}, \rho), \\ (\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, \mathbf{s}), \\ (u, w) \leftarrow \langle \mathcal{P}^*(\mathbf{pk}, \rho), \mathcal{V}(\mathbf{vk}) \rangle(u_1, u_2) \end{array} \right. \right] - \text{negl}(\lambda)$$

where ρ denotes arbitrary input randomness for \mathcal{P}^* . We call a transcript an accepting transcript if \mathcal{P} outputs a satisfying folded witness w for the folded instance u . We consider a folding scheme non-trivial if the communication costs and \mathcal{V} 's computation are lower in the case where \mathcal{V} participates in the 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.

Definition 7 (Non-Interactive). A 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 an output of \mathcal{P} , and an input to \mathcal{V} .

Definition 8 (Zero-Knowledge). A folding scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ satisfies zero-knowledge for relation \mathcal{R} if there exists a PPT simulator \mathcal{S} such that for all PPT

adversaries \mathcal{A} , and \mathcal{V}^* , and input randomness ρ

$$\left\{ \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\mathbf{s}, (u_1, w_1), (u_2, w_2)) \leftarrow \mathcal{A}(\text{pp}), \\ (\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \mathbf{s}), \\ (\text{pp}, \mathbf{s}, u_1, w_1), (\text{pp}, \mathbf{s}, u_2, w_2) \in \mathcal{R}, \\ \text{tr} = \langle \mathcal{P}(\text{pk}, w_1, w_2), \mathcal{V}^*(\text{vk}, \rho) \rangle(u_1, u_2) \end{array} \right\} \cong \left\{ \begin{array}{l} \text{pp} \leftarrow \mathcal{G}(1^\lambda), \\ (\mathbf{s}, (u_1, w_1), (u_2, w_2)) \leftarrow \mathcal{A}(\text{pp}), \\ (\text{pp}, \mathbf{s}, u_1, w_1), (\text{pp}, \mathbf{s}, u_2, w_2) \in \mathcal{R}, \\ (\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, \mathbf{s}), \\ \text{tr} \leftarrow \mathcal{S}^{\mathcal{V}^*(\text{vk}, \rho)}(\text{pk}, u_1, u_2) \end{array} \right\}$$

Definition 9 (Public Coin). A 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.

Typically, knowledge soundness is difficult to prove directly. To assist these proofs, prior works employ the forking lemma [11], which abstracts away much of the probabilistic reasoning. The original forking lemma shows that to prove knowledge soundness it is sufficient to construct a PPT extractor that takes as input a “tree” of accepting transcripts and outputs a satisfying witness. However, in our setting, this extractor must *additionally* take as input the prover’s output (i.e., the folded instance and witness) for each of these transcripts, which contains information needed to reconstruct the original witness. So, we introduce a small variant of the forking lemma that captures this modification.

Lemma 1 (Forking Lemma for Folding Schemes). Consider a $(2\mu + 1)$ -move folding scheme $\Pi = (\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$. Π satisfies knowledge soundness if there exists a PPT \mathcal{X} such that for all input instance pairs (u_1, u_2) , outputs satisfying witnesses (w_1, w_2) with probability $1 - \text{negl}(\lambda)$, given public parameters pp , a structure \mathbf{s} , and an (n_1, \dots, n_μ) -tree of accepting transcripts and the corresponding folded instance-witness pairs (u, w) . This tree comprises of n_1 transcripts (and the corresponding instance-witness pairs) with fresh randomness in \mathcal{V} ’s first message; and for each such transcript, n_2 transcripts (and the corresponding instance-witness pairs) with fresh randomness in \mathcal{V} ’s second message; etc., for a total of $\prod_{i=1}^\mu n_i$ leaves bounded by $\text{poly}(\lambda)$.

Proof Intuition. A proof for our variant of the forking lemma is similar to that of Bootle et al. [11]. We present a formal proof in Appendix E. \square

4 A Folding Scheme for NP

In this section, we describe a public-coin, zero-knowledge interactive folding scheme for NP. We additionally discuss how to make it non-interactive. We leverage the non-interactivity property to realize IVC in the next section, and the zero-knowledge property to achieve zero-knowledge IVC proof compression.

4.1 A Public-Coin, Zero-Knowledge Folding Scheme

To design a folding scheme for NP, we need an NP-complete language. While theoretically any NP-complete language is a viable candidate, we focus on R1CS,⁴ a popular algebraic representation that generalizes arithmetic circuit satisfiability.

Definition 10 (R1CS). *Consider a finite field \mathbb{F} . Let the public parameters consist of size bounds $m, n, \ell \in \mathbb{N}$ where $m > \ell$. The R1CS structure consists of sparse matrices $A, B, C \in \mathbb{F}^{m \times m}$ with at most $n = \Omega(m)$ non-zero entries in each matrix. An instance $x \in \mathbb{F}^\ell$ consists of public inputs and outputs and is satisfied by a witness $W \in \mathbb{F}^{m-\ell-1}$ if $(A \cdot Z) \circ (B \cdot Z) = C \cdot Z$, where $Z = (W, x, 1)$.*

As we show in the next section, to realize IVC, we only need a folding scheme that can fold two R1CS instances with the same R1CS matrices (A, B, C) . Specifically, given R1CS matrices (A, B, C) , and two corresponding instance-witness pairs (x_1, W_1) and (x_2, W_2) , we would like to devise a scheme that reduces the task of checking both instances into the task of checking a single new instance-witness pair (x, W) against the same R1CS matrices (A, B, C) . Unfortunately, as we illustrate now, it is difficult to devise a folding scheme for R1CS such that it satisfies completeness, let alone knowledge soundness.

First Attempt. As R1CS is an algebraic system, the most direct approach would be to take a random linear combination. Ignoring efficiency concerns, suppose that the prover sends witnesses W_1 and W_2 in the first step. The verifier responds with a random $r \in \mathbb{F}$; the prover and the verifier both compute

$$\begin{aligned} x &\leftarrow x_1 + r \cdot x_2 \\ W &\leftarrow W_1 + r \cdot W_2, \end{aligned}$$

and set the new instance-witness pair to be (x, W) . However, for non-trivial $Z_1 = (W_1, x_1, 1)$ and $Z_2 = (W_2, x_2, 1)$, and $Z = (W, x, 1)$, we *roughly* have that

$$\begin{aligned} AZ \circ BZ &= A(Z_1 + r \cdot Z_2) \circ B(Z_1 + r \cdot Z_2) \\ &= AZ_1 \circ BZ_1 + r \cdot (AZ_1 \circ BZ_2 + AZ_2 \circ BZ_1) + r^2 \cdot (AZ_2 \circ BZ_2) \\ &\neq CZ. \end{aligned}$$

The failed attempt exposes three issues. First, we must account for an additional cross-term, $r \cdot (AZ_1 \circ BZ_2 + AZ_2 \circ BZ_1)$. Second, the terms excluding the cross-term combine to produce a term that does not equal CZ :

$$AZ_1 \circ BZ_1 + r^2 \cdot (AZ_2 \circ BZ_2) = CZ_1 + r^2 \cdot CZ_2 \neq CZ_1 + r \cdot CZ_2 = CZ.$$

Third, we do not even have that $Z = Z_1 + r \cdot Z_2$ because $Z_1 + r \cdot Z_2 = (W, x, 1 + r \cdot 1)$.

⁴ R1CS is implicit in the QAPs formalism of GGPR [25], but it was made explicit in subsequent work [42]; they refer to it as a “constraint system in quadratic form”.

Second Attempt. To handle the first issue, we introduce a “slack” (or error) vector $E \in \mathbb{F}^m$ which absorbs the cross terms generated by folding. To handle the second and third issues, we introduce a scalar u , which absorbs an extra factor of r in $CZ_1 + r^2 \cdot CZ_2$ and in $Z = (W, x, 1 + r \cdot 1)$. We refer to a variant of R1CS with these additional terms as *relaxed R1CS*.

Definition 11 (Relaxed R1CS). Consider a finite field \mathbb{F} . Let the public parameters consist of size bounds $m, n, \ell \in \mathbb{N}$ where $m > \ell$. The relaxed R1CS structure consists of sparse matrices $A, B, C \in \mathbb{F}^{m \times m}$ with at most $n = \Omega(m)$ non-zero entries in each matrix. A relaxed R1CS instance consists of an error vector $E \in \mathbb{F}^m$, a scalar $u \in \mathbb{F}$, and public inputs and outputs $x \in \mathbb{F}^\ell$. An instance (E, u, x) is satisfied by a witness $W \in \mathbb{F}^{m-\ell-1}$ if $(A \cdot Z) \circ (B \cdot Z) = u \cdot (C \cdot Z) + E$, where $Z = (W, x, u)$.

Note that any R1CS instance can be expressed as a relaxed R1CS instance by augmenting it with $u = 1$ and $E = 0$, so relaxed R1CS retains NP-completeness.

Building on the first attempt, the prover and verifier can now use E to accumulate the cross-terms. In particular, for $Z_i = (W_i, x_i, u_i)$, the prover and verifier additionally compute

$$\begin{aligned} u &\leftarrow u_1 + r \cdot u_2 \\ E &\leftarrow E_1 + r \cdot (AZ_1 \circ BZ_2 + AZ_2 \circ BZ_1 - u_1 CZ_2 - u_2 CZ_1) + r^2 \cdot E_2, \end{aligned}$$

and set the new instance-witness pair to be $((E, u, x), W)$. Conveniently, updating u in this manner also keeps track of how the constant term in Z should be updated, which motivates our choice to use u in $Z = (W, x, u)$ rather than introducing a new variable. Now, for $Z = (W, x, u)$, and for random $r \in \mathbb{F}$,

$$\begin{aligned} AZ \circ BZ &= AZ_1 \circ BZ_1 + r \cdot (AZ_1 \circ BZ_2 + AZ_2 \circ BZ_1) + r^2 \cdot (AZ_2 \circ BZ_2) \\ &= (u_1 CZ_1 + E_1) + r \cdot (AZ_1 \circ BZ_2 + AZ_2 \circ BZ_1) + r^2 \cdot (u_2 CZ_2 + E_2) \\ &= (u_1 + r \cdot u_2) \cdot C(Z_1 + rZ_2) + E \\ &= uCZ + E. \end{aligned}$$

This implies that, for R1CS matrices (A, B, C) , the folded witness W is a satisfying witness for the folded instance (E, u, x) as promised. A few issues remain: in the above scheme, the prover sends witnesses (W_1, W_2) for the verifier to compute E . As a result, the folding scheme is *not* non-trivial; it is also not zero-knowledge.

Final Protocol. To circumvent these issues, we use succinct and hiding additively homomorphic commitments to W and E in the instance, and treat both W and E as the witness. We refer to this variant of relaxed R1CS as *committed relaxed R1CS*. Below, we describe a folding scheme for committed relaxed R1CS, where the prover sends a single commitment to aid the verifier in computing commitments to the folded witness (W, E) .

Definition 12 (Committed Relaxed R1CS). Consider a finite field \mathbb{F} and a commitment scheme Com over \mathbb{F} . Let the public parameters consist of size bounds

$m, n, \ell \in \mathbb{N}$ where $m > \ell$, and commitment parameters pp_W and pp_E for vectors of size m and $m - \ell - 1$ respectively. The committed relaxed R1CS structure consists of sparse matrices $A, B, C \in \mathbb{F}^{m \times m}$ with at most $n = \Omega(m)$ non-zero entries in each matrix. A committed relaxed R1CS instance is a tuple $(\bar{E}, u, \bar{W}, \mathbf{x})$, where \bar{E} and \bar{W} are commitments, $u \in \mathbb{F}$, and $\mathbf{x} \in \mathbb{F}^\ell$ are public inputs and outputs. An instance $(\bar{E}, u, \bar{W}, \mathbf{x})$ is satisfied by a witness $(E, r_E, W, r_W) \in (\mathbb{F}^m, \mathbb{F}, \mathbb{F}^{m-\ell-1}, \mathbb{F})$ if $\bar{E} = \text{Com}(\text{pp}_E, E, r_E)$, $\bar{W} = \text{Com}(\text{pp}_W, W, r_W)$, and $(A \cdot Z) \circ (B \cdot Z) = u \cdot (C \cdot Z) + E$, where $Z = (W, \mathbf{x}, u)$.

Construction 1 (A Folding Scheme for Committed Relaxed R1CS).

Consider a finite field \mathbb{F} and a succinct, hiding, and homomorphic commitment scheme Com over \mathbb{F} . We define the generator and the encoder as follows.

- $\mathcal{G}(1^\lambda) \rightarrow \text{pp}$: output size bounds $m, n, \ell \in \mathbb{N}$, and commitment parameters pp_W and pp_E for vectors of size m and $m - \ell - 1$ respectively.
- $\mathcal{K}(\text{pp}, (A, B, C)) \rightarrow (\text{pk}, \text{vk})$: output $\text{pk} \leftarrow (\text{pp}, (A, B, C))$ and $\text{vk} \leftarrow \perp$.

The verifier \mathcal{V} takes two committed relaxed R1CS instances $(\bar{E}_1, u_1, \bar{W}_1, \mathbf{x}_1)$ and $(\bar{E}_2, u_2, \bar{W}_2, \mathbf{x}_2)$. The prover \mathcal{P} , in addition to the two instances, takes witnesses to both instances, $(E_1, r_{E_1}, W_1, r_{W_1})$ and $(E_2, r_{E_2}, W_2, r_{W_2})$. Let $Z_1 = (W_1, \mathbf{x}_1, u_1)$ and $Z_2 = (W_2, \mathbf{x}_2, u_2)$. The prover and the verifier proceed as follows.

1. \mathcal{P} : Send $\bar{T} := \text{Com}(\text{pp}_E, T, r_T)$, where $r_T \leftarrow_R \mathbb{F}$ and with cross term

$$T = AZ_1 \circ BZ_2 + AZ_2 \circ BZ_1 - u_1 \cdot CZ_2 - u_2 \cdot CZ_1.$$

2. \mathcal{V} : Sample and send challenge $r \leftarrow_R \mathbb{F}$.
3. \mathcal{V}, \mathcal{P} : Output the folded instance $(\bar{E}, u, \bar{W}, \mathbf{x})$ where

$$\begin{aligned} \bar{E} &\leftarrow \bar{E}_1 + r \cdot \bar{T} + r^2 \cdot \bar{E}_2 \\ u &\leftarrow u_1 + r \cdot u_2 \\ \bar{W} &\leftarrow \bar{W}_1 + r \cdot \bar{W}_2 \\ \mathbf{x} &\leftarrow \mathbf{x}_1 + r \cdot \mathbf{x}_2 \end{aligned}$$

4. \mathcal{P} : Output the folded witness (E, r_E, W, r_W) , where

$$\begin{aligned} E &\leftarrow E_1 + r \cdot T + r^2 \cdot E_2 \\ r_E &\leftarrow r_{E_1} + r \cdot r_T + r^2 \cdot r_{E_2} \\ W &\leftarrow W_1 + r \cdot W_2 \\ r_W &\leftarrow r_{W_1} + r \cdot r_{W_2} \end{aligned}$$

Theorem 3 (A Folding Scheme for Committed Relaxed R1CS). *Construction 1 is a public-coin folding scheme for committed relaxed R1CS with perfect completeness, knowledge soundness, and zero-knowledge.*

Proof Intuition. With textbook algebra, we can show that if witnesses $(E_1, r_{E_1}, W_1, r_{W_1})$ and $(E_2, r_{E_2}, W_2, r_{W_2})$ are satisfying witnesses, then the folded witness (E, r_E, W, r_W) must be a satisfying witness. We prove knowledge soundness via the forking lemma (Lemma 1) by showing that the extractor can produce the initial witnesses given three accepting transcripts and the corresponding folded witnesses. Specifically, the extractor uses all three transcripts to compute E_i and r_{E_i} , and any two transcripts to compute W_i and r_{W_i} for $i \in \{1, 2\}$. The choice of which two transcripts does not matter due to the binding property of the commitment scheme. We present a formal proof in Appendix B. \square

4.2 Achieving Non-Interactivity via the Fiat-Shamir Transform

To design Nova’s IVC scheme, we require our folding scheme for committed relaxed R1CS to be non-interactive in the standard model. To do so we first achieve non-interactivity in the random oracle model using the (strong) Fiat-Shamir transform [24]. Next, we heuristically instantiate the random oracle using a cryptographic hash function. As a result, we can only heuristically argue the security of the resulting non-interactive folding scheme. Note that all existing IVC constructions in the standard model require instantiating the random oracle with a cryptographic hash function [12, 16, 22, 46].

Construction 2 (A Non-Interactive Folding Scheme). We achieve non-interactivity in the random oracle model using the strong Fiat-Shamir transform [24]. Let ρ denote a random oracle sampled during parameter generation and provided to all parties. Let $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ represent our interactive folding scheme (Construction 1). We construct a non-interactive folding scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ as follows:

- $\mathcal{G}(1^\lambda)$: output $\text{pp} \leftarrow \mathcal{G}(1^\lambda)$.
- $\mathcal{K}(\text{pp}, (A, B, C))$: $\text{vk} \leftarrow \rho(\text{pp}, \text{s})$ and $\text{pk} \leftarrow (\text{pp}, (A, B, C), \text{vk})$; output (vk, pk) .
- $\mathcal{P}(\text{pk}, (u_1, w_1), (u_2, w_2))$: runs $\mathcal{P}(\text{pk}, \text{pp}, \text{pk}, (A, B, C))$ to retrieve its first message \bar{T} , and sends \bar{T} to \mathcal{V} ; computes $r \leftarrow \rho(\text{vk}, u_1, u_2, \bar{T})$, forwards this to \mathcal{P} , and outputs the resulting output.
- $\mathcal{V}(\text{vk}, u_1, u_2, \bar{T})$: runs \mathcal{V} with \bar{T} as the message from the prover and with randomness $r \leftarrow \rho(\text{vk}, u_1, u_2, \bar{T})$, and outputs the resulting output.

Assumption 1 (RO instantiation). Construction 2 is a non-interactive folding scheme that satisfies completeness, knowledge soundness, and zero-knowledge in the standard model when ρ is instantiated with a cryptographic hash function.

5 Nova: An IVC Scheme with Proof Compression

This section describes Nova, an IVC scheme designed from a non-interactive folding scheme, which when instantiated with any additively-homomorphic commitment scheme with succinct commitments achieves the claimed efficiency (Lemma 4). In addition, Nova incorporates an efficient zkSNARK to prove the knowledge

of valid IVC proofs succinctly and in zero-knowledge, providing a succinct, zero-knowledge proof of knowledge of a valid IVC proof.

In Nova, at each incremental step, the prover folds a particular step of the incremental computation (represented as a committed relaxed R1CS instance-witness pair) into a running committed relaxed R1CS instance-witness pair. At any step in the incremental computation, a valid “IVC proof”, in a nutshell, is a satisfying witness of the running committed relaxed R1CS instance (which an honest prover can compute by folding witnesses associated with each step of the incremental computation) along with the running committed relaxed R1CS instance. Furthermore, at any incremental step, Nova’s prover can prove in zero-knowledge and with a succinct proof—using a variant of an existing zkSNARK [40] (Section 6)—that it knows a valid IVC proof (i.e., a satisfying witness) to the running committed relaxed R1CS instance (Construction 4).

Note that Nova is *not* a zero-knowledge IVC scheme, as that would additionally require an IVC proof to be zero-knowledge (in Nova’s case, an IVC proof does *not* hide witnesses associated with steps of the incremental computation). This difference is immaterial in the context of a single prover since it can use Nova’s auxiliary zkSNARK to provide a zero-knowledge proof of knowledge of a valid IVC proof; we leave it to future work to achieve zero-knowledge IVC.

5.1 Constructing IVC from a Folding Scheme for NP

Recall that an IVC scheme allows a prover to show that $z_n = F^{(n)}(z_0)$ for some count n , initial input z_0 , and output z_n . We now show how to construct an IVC scheme for a non-deterministic, polynomial-time computable function F using our non-interactive folding scheme for committed relaxed R1CS (Construction 2).⁵

In our construction, as in a SNARK-based IVC, the prover uses an augmented function F' (Figure 4), which, in addition to invoking F , performs additional bookkeeping to fold proofs of prior invocations of itself.

We first describe a simplified version of F' , to provide intuition. F' takes as non-deterministic advice two committed relaxed R1CS instances u_i and U_i . Suppose that U_i represents the correct execution of invocations $1, \dots, i-1$ of F' so long as u_i represents the correct execution of invocation i of F' . F' performs two tasks. First, it executes a step of the incremental computation: instance u_i contains z_i which F' uses to output $z_{i+1} = F(z_i)$. Second, F' invokes the verifier of the non-interactive folding scheme to fold the task of checking u_i and U_i into the task of checking a single instance U_{i+1} . The IVC prover then computes a new instance u_{i+1} which attests to the correct execution of invocation $i+1$ of F' , thereby attesting that $z_{i+1} = F(z_i)$ and U_{i+1} is the result of folding u_i and U_i . Now, we have that U_{i+1} represents the correct execution of invocations $1, \dots, i$ of F' so long as u_{i+1} represents the correct execution of invocation $i+1$ of F' .

The above description glossed over a subtle discrepancy: Because F' must output the running instance U_{i+1} for the next invocation to use, it is contained

⁵ While, in theory, we can use any folding scheme for NP, we specifically invoke our construction for committed relaxed R1CS for a simpler presentation.

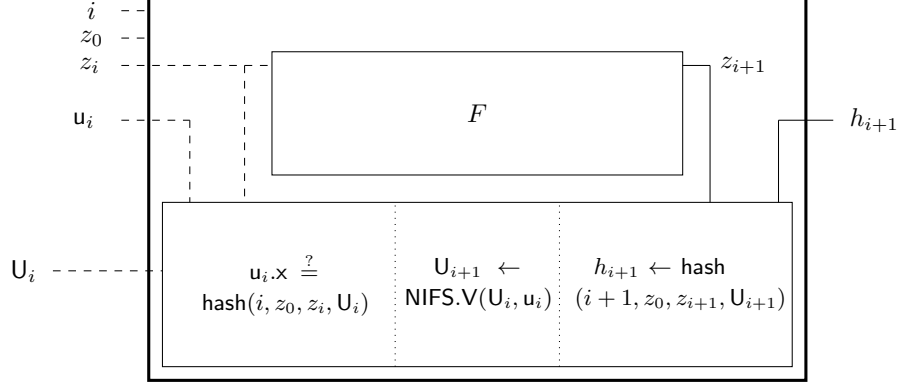


Fig. 4: Overview of F' . F' represented as a committed relaxed R1CS instance \mathbf{u}_{i+1} encodes the statement that there exists $((i, z_0, z_i, \mathbf{u}_i, \mathbf{U}_i), \mathbf{U}_{i+1}, \bar{T})$ such that $\mathbf{u}_{i+1}.x = \text{hash}(\text{vk}, i, z_0, z_i, \mathbf{U}_i)$, $h_{i+1} = \text{hash}(\text{vk}, i+1, z_0, F(z_i), \mathbf{U}_{i+1})$, $\mathbf{U}_{i+1} = \text{NIFS.V}(\text{vk}, \mathbf{U}_i, \mathbf{u}_i, \bar{T})$, and that F' outputs h_{i+1} . The diagram omits depicting vk , ω , and \bar{T} .

in $\mathbf{u}_{i+1}.x$ (i.e., the public IO of \mathbf{u}_{i+1}). But, in the next iteration, F' must fold $\mathbf{u}_{i+1}.x$ into $\mathbf{U}_{i+1}.x$, meaning that F' is stuck trying to squeeze \mathbf{U}_{i+1} into $\mathbf{u}_{i+1}.x$. To handle this inconsistency, we modify F' to output a *collision-resistant hash* of its public IO rather than producing it directly (this ensures that the public IO of F' is a constant number of finite field elements). The next invocation of F' then additionally takes the preimage of this hash as non-deterministic advice. We assume that the hash function takes an additional random input (which provides hiding) but for notational convenience we do not explicitly depict this.

Producing IVC Proofs. Let $(\mathbf{u}_\perp, \mathbf{w}_\perp)$ be the trivially satisfying instance-witness pair, where E, W , and x are appropriately-sized zero vectors, $r_E = 0$, $r_W = 0$, and \bar{E} and \bar{W} are commitments of E and W respectively.

Now, in iteration $i+1$, the IVC prover runs F' and computes \mathbf{u}_{i+1} and \mathbf{U}_{i+1} as well as the corresponding witnesses \mathbf{w}_{i+1} and \mathbf{W}_{i+1} . Because \mathbf{u}_{i+1} and \mathbf{U}_{i+1} together attest to the correctness of $i+1$ invocations of F' (which indirectly attests to $i+1$ invocations of F) the IVC proof Π_{i+1} is $((\mathbf{U}_{i+1}, \mathbf{W}_{i+1}), (\mathbf{u}_{i+1}, \mathbf{w}_{i+1}))$. Moreover, succinctness is maintained by the properties of the underlying folding scheme. We formally describe our construction below.

Construction 3 (IVC). Let $\text{NIFS} = (\text{G}, \text{K}, \text{P}, \text{V})$ be the non-interactive folding scheme for committed relaxed R1CS (Construction 2). Consider a polynomial-time function F that takes non-deterministic input, and a cryptographic hash function hash . We define our augmented function F' as follows (all arguments to F' are taken as non-deterministic advice):

$$\underline{F'(\text{vk}, \mathbf{U}_i, \mathbf{u}_i, (i, z_0, z_i), \omega_i, \bar{T}) \rightarrow x:}$$

If i is 0, output $\text{hash}(\text{vk}, 1, z_0, F(z_0, \omega_i), \mathbf{u}_\perp)$;

otherwise,

- (1) check that $u_i.x = \text{hash}(\text{vk}, i, z_0, z_i, U_i)$, where $u_i.x$ is the public IO of u_i ,
- (2) check that $(u_i.\overline{E}, u_i.u) = (u_\perp.\overline{E}, 1)$,
- (3) compute $U_{i+1} \leftarrow \text{NIFS.V}(\text{vk}, U, u, \overline{T})$, and
- (4) output $\text{hash}(\text{vk}, i+1, z_0, F(z_i, \omega_i), U_{i+1})$.

Because F' can be computed in polynomial time, it can be represented as a committed relaxed R1CS structure $s_{F'}$. Let

$$(u_{i+1}, w_{i+1}) \leftarrow \text{trace}(F', (\text{vk}, U_i, u_i, (i, z_0, z_i), \omega_i, \overline{T}))$$

denote the satisfying committed relaxed R1CS instance-witness pair (u_{i+1}, w_{i+1}) for the execution of F' on non-deterministic advice $(\text{vk}, U_i, u_i, (i, z_0, z_i), \omega_i, \overline{T})$.

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

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

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

Compute $(\text{pk}_{\text{fs}}, \text{vk}_{\text{fs}}) \leftarrow \text{NIFS.K}(\text{pp}, s_{F'})$ and output $(\text{pk}, \text{vk}) \leftarrow ((F, \text{pk}_{\text{fs}}), (F, \text{vk}_{\text{fs}}))$.

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

Parse Π_i as $((U_i, W_i), (u_i, w_i))$ and then

- (1) if i is 0, compute $(U_{i+1}, W_{i+1}, \overline{T}) \leftarrow (u_\perp, w_\perp, u_\perp.\overline{E})$;
otherwise, compute $(U_{i+1}, W_{i+1}, \overline{T}) \leftarrow \text{NIFS.P}(\text{pk}, (U_i, W_i), (u_i, w_i))$,
- (2) compute $(u_{i+1}, w_{i+1}) \leftarrow \text{trace}(F', (\text{vk}, U_i, u_i, (i, z_0, z_i), \omega_i, \overline{T}))$, and
- (3) output $\Pi_{i+1} \leftarrow ((U_{i+1}, W_{i+1}), (u_{i+1}, w_{i+1}))$.

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

If i is 0, check that $z_i = z_0$;

otherwise,

- (1) parse Π_i as $((U_i, W_i), (u_i, w_i))$,
- (2) check that $u_i.x = \text{hash}(\text{vk}, i, z_0, z_i, U_i)$,
- (3) check that $(u_i.\overline{E}, u_i.u) = (u_\perp.\overline{E}, 1)$, and
- (4) check that W_i and w_i are satisfying witnesses to U_i and u_i respectively.

Lemma 2 (Completeness). *Construction 3 is an IVC scheme that satisfies completeness.*

Proof Intuition. Given a satisfying IVC proof $\Pi_i = ((U_i, W_i), (u_i, w_i))$ suppose that \mathcal{P} outputs $\Pi_{i+1} = ((U_{i+1}, W_{i+1}), (u_{i+1}, w_{i+1}))$. Because Π_i is a valid IVC proof, (u_i, w_i) and (U_i, W_i) are satisfying instance-witness pairs. Because (U_{i+1}, W_{i+1}) is obtained by folding (u_i, w_i) and (U_i, W_i) , it must be satisfying by the folding scheme's completeness. By construction, (u_{i+1}, w_{i+1}) is satisfying instance-witness pair that satisfies the IVC verifier's auxiliary checks. Thus, Π_{i+1} is satisfying. Appendix C provides a formal proof. \square

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

Proof Intuition. For function F , constant n , $\mathbf{pp} \leftarrow \mathcal{G}(1^\lambda)$, and $(\mathbf{pk}, \mathbf{vk}) \leftarrow \mathcal{K}(\mathbf{pp}, F)$, consider an adversary \mathcal{P}^* that outputs (z_0, z, Π) such that $\mathcal{V}(\mathbf{vk}, (n, z_0, z), \Pi) = 1$ with probability ϵ . We construct an extractor \mathcal{E} that with input (\mathbf{pp}, z_0, z) , outputs $(\omega_0, \dots, \omega_{n-1})$ such that by computing $z_i \leftarrow F(z_{i-1}, \omega_{i-1})$ for all $i \in \{1, \dots, n\}$ we have that $z_n = z$ with probability $\epsilon - \text{negl}(\lambda)$. We show inductively that \mathcal{E} can construct an extractor \mathcal{E}_i that outputs (z_i, \dots, z_{n-1}) , $(\omega_i, \dots, \omega_{n-1})$, and Π_i such that for all $j \in \{i+1, \dots, n\}$, $z_j = F(z_{j-1}, \omega_{j-1})$, $\mathcal{V}(\mathbf{vk}, i, z_0, z_i, \Pi_i) = 1$, and $z_n = z$ with probability $\epsilon - \text{negl}(\lambda)$. Then, because in the base case when $i = 0$, \mathcal{V} checks that $z_0 = z_i$, it is sufficient for \mathcal{E} to run \mathcal{E}_0 to retrieve values $(\omega_0, \dots, \omega_{n-1})$. Initially, \mathcal{E}_n simply runs the assumed \mathcal{P}^* to get a satisfying Π_n . Given extractor \mathcal{E}_i that satisfies the inductive hypothesis, we can construct extractor \mathcal{E}_{i-1} . Appendix C provides a formal proof. \square

Lemma 4 (Efficiency). *When instantiated with the Pedersen commitment scheme (Construction 6), we have that $|F'| = |F| + o(2 \cdot \mathbf{G} + 2 \cdot \mathbf{H} + \mathbf{R})$, where $|F|$ denotes the number of R1CS 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, and \mathbf{R} is the number of constraints to encode the RO ρ .*

Proof. On input instances \mathbf{U} and \mathbf{u} , NIFS.V computes $\overline{E} \leftarrow \mathbf{U} \cdot \overline{E} + r \cdot \overline{T} + r^2 \cdot \mathbf{u} \cdot \overline{E}$ and $\overline{W} \leftarrow \mathbf{U} \cdot \overline{W} + r \cdot \mathbf{u} \cdot \overline{W}$. However, by construction, $\mathbf{u} \cdot \overline{E} = \mathbf{u}_\perp \cdot \overline{E} = \overline{0}$. So, NIFS.V computes two group scalar multiplications, as it does not need to compute $r^2 \cdot \mathbf{u} \cdot \overline{E}$. NIFS.V additionally invokes the RO once to obtain a random scalar. Finally, F' makes two additional calls to hash (details are in the description of F'). \square

5.2 Compressing IVC Proofs with zkSNARKs

To prove a statement about an incremental computation, the prover can produce an IVC proof using the construction in the prior section and send the IVC proof to the verifier. However, this does not satisfy zero-knowledge (as the IVC proof described in the prior section does not hide the prover's non-deterministic inputs) and succinctness (as the IVC proof size is linear in the size F). In theory, one can address this problem with any zkSNARK for NP. Specifically, $\overline{\mathcal{P}}$ can produce a zkSNARK proving that it knows Π_i such that IVC verifier \mathcal{V} accepts for statement (i, z_0, z_i) . Naturally, the proof sent to the verifier is succinct and zero-knowledge due to the corresponding properties of the zkSNARK.

Unfortunately, employing an off-the-shelf zkSNARK makes the overall solution impractical as the zkSNARK prover must prove, among other things, the knowledge of vectors whose commitments equal a particular value; this requires encoding a linear number of group scalar multiplications in the programming model of zkSNARKs (e.g., R1CS or circuits). To address this, we design a zkSNARK tailored for our particular purpose and we describe it in Section 6. Below, we describe how to use a zkSNARK to prove the knowledge of a valid IVC proof.

Construction 4 (A zkSNARK of a Valid IVC Proof). Let IVC denote the IVC scheme in Construction 3, let NIFS denote the non-interactive folding scheme in Construction 2, and let hash denote a randomized cryptographic hash function. Assume a zero-knowledge succinct non-interactive argument of knowledge (Definition 2), zkSNARK, for committed relaxed R1CS. That is, given public parameters pp , structure \mathbf{s} , and instance \mathbf{u} , zkSNARK.P can convince zkSNARK.V in zero-knowledge and with a succinct proof (e.g., $O_\lambda(\log N)$ -sized proof) that it knows a corresponding witness \mathbf{w} such that $(\text{pp}, \mathbf{s}, \mathbf{u}, \mathbf{w})$ is a satisfying committed relaxed R1CS tuple.

Consider a polynomial-time computable function F . Suppose $\text{pp} \leftarrow \text{IVC.G}(1^\lambda)$ and $(\text{pk}, \text{vk}) \leftarrow \text{IVC.K}(\text{pp}, F)$. Suppose the prover \mathcal{P} and verifier \mathcal{V} are provided an instance (i, z_0, z_i) . We construct a zkSNARK that allows the prover to show that it knows an IVC proof Π_i such that $\text{IVC.V}(\text{vk}, i, z_0, z_i, \Pi_i) = 1$.

In a nutshell, we leverage the fact that Π is two committed relaxed R1CS instance-witness pairs. So, \mathcal{P} first folds instance-witness pairs (\mathbf{u}, \mathbf{w}) and (\mathbf{U}, \mathbf{W}) to produce a folded instance-witness pair $(\mathbf{U}', \mathbf{W}')$, using NIFS.P. Next, \mathcal{P} runs zkSNARK.P to prove that it knows a valid witness for \mathbf{U}' . In more detail, for polynomial-time computable function F and corresponding function F' as defined in Construction 3 (and instantiated with hash), we define $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ as follows.

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

- (1) Compute $\text{pp}_{\text{NIFS}} \leftarrow \text{NIFS.G}(1^\lambda)$
- (2) Compute $\text{pp}_{\text{zkSNARK}} \leftarrow \text{zkSNARK.G}(1^\lambda)$
- (3) Output $(\text{pp}_{\text{NIFS}}, \text{pp}_{\text{zkSNARK}})$

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

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

$\mathcal{P}(\text{pk}, (i, z_0, z_i), \Pi) \rightarrow \pi$:

If i is 0, output \perp ;

otherwise,

- (1) parse Π as $((\mathbf{U}, \mathbf{W}), (\mathbf{u}, \mathbf{w}))$
- (2) compute $(\mathbf{U}', \mathbf{W}', \bar{T}) \leftarrow \text{NIFS.P}(\text{pk}_{\text{NIFS}}, (\mathbf{U}, \mathbf{W}), (\mathbf{u}, \mathbf{w}))$
- (3) compute $\pi_{\mathbf{U}'} \leftarrow \text{zkSNARK.P}(\text{pk}_{\text{zkSNARK}}, \mathbf{U}', \mathbf{W}')$
- (4) output $(\mathbf{U}, \mathbf{u}, \bar{T}, \pi_{\mathbf{U}'})$.

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

If i is 0, check that $z_0 = z_i$;

otherwise,

- (1) parse π as $(\mathbf{U}, \mathbf{u}, \bar{T}, \pi_{\mathbf{U}'})$,
- (2) check that $\mathbf{u}.x = \text{hash}(\text{vk}_{\text{NIFS}}, i, z_0, z_i, \mathbf{U})$,
- (3) check that $(\mathbf{u}.\bar{E}, \mathbf{u}.u) = (\mathbf{u}_\perp.\bar{E}, 1)$,
- (4) compute $\mathbf{U}' \leftarrow \text{NIFS.V}(\text{vk}_{\text{NIFS}}, \mathbf{U}, \mathbf{u}, \bar{T})$, and
- (5) check that $\text{zkSNARK.V}(\text{vk}_{\text{zkSNARK}}, \mathbf{U}', \pi_{\mathbf{U}'}) = 1$.

Theorem 4 (A zkSNARK of a Valid IVC Proof). *Construction 4 is a zkSNARK of a valid IVC proof produced by Construction 3.*

Proof Intuition. Completeness and knowledge soundness hold due to the completeness and knowledge soundness of the underlying zkSNARK and the non-interactive folding scheme. Assuming the non-interactive folding scheme satisfies succinctness (e.g., by using the Pedersen commitment scheme), succinctness holds due to the fact that \mathbf{u} , \mathbf{U} , and \bar{T} are succinct, and due to the succinctness of the underlying zkSNARK.

To prove zero-knowledge, we construct a simulator \mathcal{S} that first iteratively simulates $(\mathbf{U}_i, \mathbf{u}_i)$ for all $i \in \{1, \dots, n\}$. Specifically, given a simulated proof $(\mathbf{U}_i, \mathbf{u}_i)$, \mathcal{S} first uses the simulator of the non-interactive folding scheme to simulate \bar{T}_i . \mathcal{S} then folds \mathbf{U}_i and \mathbf{u}_i using \bar{T}_i to produce \mathbf{U}_{i+1} . \mathcal{S} simulates \mathbf{u}_i using the observation that all terms are randomized. In the final round, \mathcal{S} folds \mathbf{u}_n and \mathbf{U}_n (again using a simulated \bar{T}_n) to produce an instance \mathbf{U}' , and then uses the simulator of the zkSNARK to produce $\pi_{\mathbf{U}'}$. \mathcal{S} outputs $(\mathbf{U}_n, \mathbf{u}_n, \bar{T}_n, \pi_{\mathbf{U}'})$. We provide a formal proof in Appendix D. \square

6 A zkSNARK for Committed Relaxed R1CS

As described in Section 5.2, Nova needs a zkSNARK for committed relaxed R1CS to prove the knowledge of a valid IVC proof succinctly and in zero-knowledge. This section presents such a zkSNARK by adapting Spartan [40]. We build on Spartan [40] to avoid FFTs and a trusted setup.

6.1 Background

We assume familiarity with polynomials. We provide background in Appendix G.

Definition 13 (Polynomial Extension). *Suppose $f : \{0, 1\}^\ell \rightarrow \mathbb{F}$ is a function that maps ℓ -bit strings to an element of \mathbb{F} . A polynomial extension of f is a low-degree ℓ -variate polynomial $\tilde{f} : \mathbb{F}^\ell \rightarrow \mathbb{F}$ such that $\tilde{f}(x) = f(x)$ for all $x \in \{0, 1\}^\ell$. A multilinear extension (MLE) of a function $f : \{0, 1\}^\ell \rightarrow \mathbb{F}$ is a low-degree polynomial extension where the extension is a multilinear polynomial.*

Every function $f : \{0, 1\}^\ell \rightarrow \mathbb{F}$ has a unique MLE, and conversely every ℓ -variate multilinear polynomial over \mathbb{F} extends a unique function mapping $\{0, 1\}^\ell \rightarrow \mathbb{F}$. Below, we use \tilde{f} to denote the unique MLE of f .

Lemma 5 (The Sum-Check Protocol [37]). *For ℓ -variate polynomial G over \mathbb{F} with degree at most μ in each variable, there exists a public-coin interactive proof protocol (known as the sum-check protocol) to reduce the task of checking $\sum_{x \in \{0,1\}^\ell} G(x) = T$ to the task of checking $G(r) = e$ for $r \in \mathbb{F}^\ell$. The interaction consists of a total of ℓ rounds, where in each round the verifier sends a single element of \mathbb{F} and the prover responds with $\mu + 1$ elements of \mathbb{F} .*

6.2 A Polynomial IOP for Idealized Relaxed R1CS

Our exposition below is based on Spartan [40] and its recent recapitulation [36]. The theorem below and its proof is a verbatim adaptation of Spartan’s polynomial IOP for R1CS to relaxed R1CS.

Recall that an interactive proof (IP) [27] for a relation \mathcal{R} is an interactive protocol between a prover and a verifier where the prover proves the knowledge of a witness w for a prescribed instance u such that $(u, w) \in \mathcal{R}$. An interactive oracle proof (IOP) [5, 39] generalizes interactive proofs where in each round the prover may send an oracle (e.g., a string) and the verifier may query a previously-sent oracle during the remainder of the protocol. A polynomial IOP [18] is an IOP in which the oracle sent by the prover is a polynomial and the verifier may query for an evaluation of the polynomial at a point in its domain. We consider a (minor) variant of polynomial IOPs, where the verifier has oracle access to polynomials in the R1CS structure and instance.

We first construct a polynomial IOP for an idealized version of relaxed R1CS (Definition 14) where the instance contains a purported witness. We then compile it into a zkSNARK for committed relaxed R1CS (Definition 12).

Definition 14 (Idealized Relaxed R1CS). *Consider a finite field \mathbb{F} . Let the public parameters consist of size bounds $m, n, \ell \in \mathbb{N}$ where $m > \ell$. The idealized relaxed R1CS structure consists of sparse matrices $A, B, C \in \mathbb{F}^{m \times m}$ with at most $n = \Omega(m)$ non-zero entries in each matrix. A idealized relaxed R1CS instance consists of an error vector $E \in \mathbb{F}^m$, a scalar $u \in \mathbb{F}$, witness vector $W \in \mathbb{F}^m$, and public inputs and outputs $x \in \mathbb{F}^\ell$. An instance (E, u, W, x) is satisfying if $(A \cdot Z) \circ (B \cdot Z) = u \cdot (C \cdot Z) + E$, where $Z = (W, x, u)$.*

Construction 5 (Polynomial IOP for Idealized Relaxed R1CS). Consider an idealized relaxed R1CS statement φ consisting of public parameters (m, n, ℓ) , structure (A, B, C) , and instance (E, u, W, x) . Without loss of generality, we assume that m and n are powers of 2 and that $m = 2 \cdot (\ell + 1)$.

Let $s = \log m$. We interpret the matrices A, B, C as functions with signature $\{0, 1\}^{\log m} \times \{0, 1\}^{\log m} \rightarrow \mathbb{F}$ in a natural manner. In particular, an input in $\{0, 1\}^{\log m} \times \{0, 1\}^{\log m}$ is interpreted as the binary representation of an index $(i, j) \in [m] \times [m]$, where $[m] := \{1, \dots, m\}$ and the function outputs (i, j) th entry of the matrix. As such, let \tilde{A} , \tilde{B} , and \tilde{C} denote multilinear extensions of A , B , and C interpreted as functions, so they are $2 \log m$ -variate sparse multilinear polynomials of size n . Similarly, we interpret E and W as functions with respective signatures $\{0, 1\}^{\log m} \rightarrow \mathbb{F}$ and $\{0, 1\}^{\log m - 1} \rightarrow \mathbb{F}$. Furthermore, let \tilde{E} and \tilde{W}

denote the multilinear extensions of E and W interpreted as functions, so they are multilinear polynomials in $\log m$ and $\log m - 1$ variables respectively.

As noted earlier, the verifier has an oracle access to the following polynomials: \tilde{A} , \tilde{B} , \tilde{C} , \tilde{E} , and \tilde{W} . Additionally, the verifier reads u and x in entirety.

Let $Z = (W, x, u)$. Similar to how we interpret matrices as functions, we interpret Z and (x, u) as functions with the following respective signatures: $\{0, 1\}^s \rightarrow \mathbb{F}$ and $\{0, 1\}^{s-1} \rightarrow \mathbb{F}$. Observe that the MLE \tilde{Z} of Z satisfies

$$\tilde{Z}(X_1, \dots, X_s) = (1 - X_1) \cdot \tilde{W}(X_2, \dots, X_s) + X_1 \cdot \widetilde{(x, u)}(X_2, \dots, X_s) \quad (1)$$

Similar to [40, Theorem 4.1], checking if φ is satisfiable is equivalent, except for a soundness error of $\log m/|\mathbb{F}|$ over the choice of $\tau \in \mathbb{F}^s$, to checking if the following identity holds:

$$0 \stackrel{?}{=} \sum_{x \in \{0, 1\}^s} \tilde{\text{eq}}(\tau, x) \cdot F(x), \quad (2)$$

where

$$F(x) = \left(\sum_{y \in \{0, 1\}^s} \tilde{A}(x, y) \cdot \tilde{Z}(y) \right) \cdot \left(\sum_{y \in \{0, 1\}^s} \tilde{B}(x, y) \cdot \tilde{Z}(y) \right) - \left(u \cdot \sum_{y \in \{0, 1\}^s} \tilde{C}(x, y) \cdot \tilde{Z}(y) + \tilde{E}(x) \right),$$

and $\tilde{\text{eq}}$ is the multilinear extension of $\text{eq} : \{0, 1\}^s \times \{0, 1\}^s \rightarrow \mathbb{F}$ where $\text{eq}(x, e) = 1$ if $x = e$ and 0 otherwise.

That is, if φ is satisfiable, then Equation (2) holds with probability 1 over the choice of τ , and if not, then Equation (2) holds with probability at most $O(\log m/|\mathbb{F}|)$ over the random choice of τ .

To compute the right-hand side in Equation (2), the prover and the verifier apply the sum-check protocol to the following polynomial: $g(x) := \tilde{\text{eq}}(\tau, x) \cdot F(x)$. From the verifier's perspective, this reduces the task of computing the right-hand side of Equation (2) to the task of evaluating g at a random input $r_x \in \mathbb{F}^s$. Note that the verifier can locally evaluate $\tilde{\text{eq}}(\tau, r_x)$ in $O(\log m)$ field operations via $\tilde{\text{eq}}(\tau, r_x) = \prod_{i=1}^s (\tau_i r_{x,i} + (1 - \tau_i)(1 - r_{x,i}))$. With $\tilde{\text{eq}}(\tau, r_x)$ in hand, $g(r_x)$ can be computed in $O(1)$ time given the four quantities: $\sum_{y \in \{0, 1\}^s} \tilde{A}(r_x, y) \cdot \tilde{Z}(y)$, $\sum_{y \in \{0, 1\}^s} \tilde{B}(r_x, y) \cdot \tilde{Z}(y)$, $\sum_{y \in \{0, 1\}^s} \tilde{C}(r_x, y) \cdot \tilde{Z}(y)$, and $\tilde{E}(r_x)$.

The last quantity can be computed with a single query to polynomial \tilde{E} . Furthermore, the first three quantities can be computed by applying the sum-check protocol three more times in parallel, once to each of the following three polynomials (using the same random vector of field elements, $r_y \in \mathbb{F}^s$, in each of the three invocations): $\tilde{A}(r_x, y) \cdot \tilde{Z}(y)$, $\tilde{B}(r_x, y) \cdot \tilde{Z}(y)$, and $\tilde{C}(r_x, y) \cdot \tilde{Z}(y)$.

To perform the verifier's final check in each of these three invocations of the sum-check protocol, it suffices for the verifier to evaluate each of the above three

polynomials at the random vector r_y , which means it suffices for the verifier to evaluate $\tilde{A}(r_x, r_y)$, $\tilde{B}(r_x, r_y)$, $\tilde{C}(r_x, r_y)$, and $\tilde{Z}(r_y)$. The first three evaluations can be obtained via the verifier's assumed query access to $(\tilde{A}, \tilde{B}, \tilde{C})$. $\tilde{Z}(r_y)$ can be computed (via Equation (1)) from a query to \tilde{W} and from computing (\tilde{x}, \tilde{u}) .

In summary, we have the following polynomial IOP.

1. $\mathcal{V} \rightarrow \mathcal{P}$: $\tau \in_R \mathbb{F}^s$
2. $\mathcal{V} \leftrightarrow \mathcal{P}$: run the sum-check protocol to reduce the check in Equation (2) to checking if the following hold, where r_x, r_y are vectors in \mathbb{F}^s chosen at random by the verifier over the course of the sum-check protocol:
 - $\tilde{A}(r_x, r_y) \stackrel{?}{=} v_A$, $\tilde{B}(r_x, r_y) \stackrel{?}{=} v_B$, and $\tilde{C}(r_x, r_y) \stackrel{?}{=} v_C$;
 - $\tilde{E}(r_x) \stackrel{?}{=} v_E$; and
 - $\tilde{Z}(r_y) \stackrel{?}{=} v_Z$.
3. \mathcal{V} :
 - check if $\tilde{A}(r_x, r_y) \stackrel{?}{=} v_A$, $\tilde{B}(r_x, r_y) \stackrel{?}{=} v_B$, and $\tilde{C}(r_x, r_y) \stackrel{?}{=} v_C$, with a query to $\tilde{A}, \tilde{B}, \tilde{C}$ at (r_x, r_y) ;
 - check if $\tilde{E}(r_x) \stackrel{?}{=} v_E$ with an oracle query to \tilde{E} ; and
 - check if $\tilde{Z}(r_y) \stackrel{?}{=} v_Z$ by checking if: $v_Z = (1 - r_y[1]) \cdot v_W + r_y[1] \cdot (\tilde{x}, \tilde{u})(r_y[2..])$, where $r_y[2..]$ refers to a slice of r_y without the first element of r_y , and $v_W \leftarrow \tilde{W}(r_y[2..])$ via an oracle query (see Equation (1)).

Theorem 5. *Construction 5 is a polynomial IOP for idealized relaxed R1CS defined over a finite field \mathbb{F} , with the following parameters, where m denotes the dimension of the R1CS matrices, and n denotes the number of non-zero entries in the matrices: Soundness error is $O(\log m)/|\mathbb{F}|$; round complexity is $O(\log m)$; The verifier has query access to $2 \log m$ -variate multilinear polynomials $\tilde{A}, \tilde{B}, \tilde{C}$ in the structure, and $(\log m)$ -variate multilinear polynomial \tilde{E} , and $(\log m - 1)$ -variate multilinear polynomial \tilde{W} in the instance; the verifier issues a single query to polynomials $\tilde{A}, \tilde{B}, \tilde{C}$, and \tilde{W}, \tilde{E} , and otherwise performs $O(\log m)$ operations over \mathbb{F} ; the prover performs $O(n)$ operations over \mathbb{F} to compute its messages in the polynomial IOP and to respond to the verifier's queries to $(\tilde{W}, \tilde{E}, \tilde{A}, \tilde{B}, \tilde{C})$.*

Proof. Perfect completeness follows from perfect completeness of the sum-check protocol and the fact that Equation (2) holds with probability 1 over the choice of τ if φ is satisfiable. Applying a standard union bound to the soundness error introduced by probabilistic check in Equation (2) with the soundness error of the sum-check protocol [37], we conclude that the soundness error for the depicted polynomial IOP is at most $O(\log m)/|\mathbb{F}|$.

The sum-check protocol is applied four times (although three of the invocations occur in parallel and in practice combined into one [40]). In each invocation, the polynomial to which the sum-check protocol is applied has degree at most 3 in each variable, and the number of variables is $s = \log m$. Hence, the round complexity of the polynomial IOP is $O(\log m)$. Since each polynomial has degree at most 3 in each variable, the total communication cost is $O(\log m)$ field elements.

The claimed verifier runtime is immediate from the verifier’s runtime in the sum-check protocol, and the fact that $\tilde{\mathbf{e}}\mathbf{q}$ can be evaluated at any input $(\tau, r_x) \in \mathbb{F}^{2s}$ in $O(\log m)$ field operations. As in Spartan [40], the prover’s work in the polynomial IOP in $O(n)$ operations over \mathbb{F} using prior techniques [44, 50]. \square

6.3 Compiling Polynomial IOPs to zkSNARKs

As in prior works [18, 21, 40], we compile our polynomial IOP into a zkSNARK using a polynomial commitment scheme [31] and the Fiat-Shamir transform [24].

Interpreting commitments to vectors as polynomial commitments. It is well known that commitments to m -sized vectors over \mathbb{F} are commitments to $\log m$ -variate multilinear polynomials represented with evaluations over $\{0, 1\}^m$ [34, 40, 48, 51]. Furthermore, there is a polynomial commitment scheme for $\log m$ -variate multilinear polynomials if there exists an argument protocol to prove an inner product computation between a committed vector and an m -sized public vector $((r_1, 1 - r_1) \otimes \dots \otimes (r_{\log m}, 1 - r_{\log m}))$, where $r \in \mathbb{F}^{\log m}$ is an evaluation point. There are two candidate constructions in the literature. Note that the primary difference between two schemes is in the verifier’s time.

1. PC_{BP} . If the commitment scheme for vectors over \mathbb{F} is Pedersen’s commitments (Construction 6), as in prior work [48], Bulletproofs [15] provides a suitable inner product argument protocol. The polynomial commitment scheme here achieves the following efficiency characteristics, assuming the hardness of the discrete logarithm problem. For a $\log m$ -variate multilinear polynomial, committing takes $O_\lambda(m)$ time to produce an $O_\lambda(1)$ -sized commitment; the prover incurs $O_\lambda(m)$ costs to produce an evaluation proof of size $O_\lambda(\log m)$ that can be verified in $O_\lambda(m)$. Note that PC_{BP} is a special case of Hyrax’s polynomial commitment scheme [48].
2. PC_{Dory} . If vectors over \mathbb{F} are committed with a two-tiered “matrix” commitment (see for example, [19, 34]), which provides $O_\lambda(1)$ -sized commitments to m -sized vectors under the SXDH assumption. With this commitment scheme, Dory [34] provides the necessary inner product argument. The polynomial commitment here achieves the following efficiency characteristics, assuming the hardness of SXDH. For a $\log m$ -variate multilinear polynomial, committing takes $O_\lambda(m)$ time to produce an $O_\lambda(1)$ -sized commitment; the prover incurs $O_\lambda(m)$ costs to produce an evaluation proof of size $O_\lambda(\log m)$ that can be verified in $O_\lambda(\log m)$.

Polynomial commitments for sparse multilinear polynomials. In our constructions below, we require polynomial commitment schemes that can efficiently handle sparse multilinear polynomials. Spartan [40, §7] (and its optimization [43, §6]) provides a generic compiler to transform existing polynomial commitment schemes for multilinear polynomials into those that can efficiently handle sparse multilinear polynomials. Specifically, we apply [36, Theorem 5] (which captures Spartan’s compiler in a generic manner) to PC_{BP} and PC_{Dory} to

obtain their variants that can efficiently handle sparse multilinear polynomials; we refer to them as “Sparse-PC_{BP}” and “Sparse-PC_{Dory}” respectively.

Theorem 6 (A zkSNARK from PC_{BP}). *Assuming the hardness of the discrete logarithm problem, there exists a zkSNARK in the random oracle model for committed relaxed R1CS with the following efficiency characteristics, where m denotes the dimensions of R1CS matrices and n denotes the number of non-zero entries in the matrices: The encoder runs in time $O_\lambda(n)$; The prover runs in time $O_\lambda(n)$; The proof length is $O_\lambda(\log n)$; and the verifier runs in time $O_\lambda(n)$.*⁶

Proof. For R1CS structure (A, B, C) , we first have the encoder directly provide $(\tilde{A}, \tilde{B}, \tilde{C})$ in the prover key, and additionally provide sparse polynomial commitments to $\tilde{A}, \tilde{B}, \tilde{C}$ using Sparse-PC_{BP} in both the prover and verifier keys. Next, we apply the compiler of [18] using PC_{BP} to the polynomial IOP from Construction 5. At a high level, this replaces all of the oracles provided to the verifier with PC_{BP} commitments, which the prover and verifier then use to simulate ideal queries to a committed oracle. By [18, Theorem 6] this provides a public-coin honest-verifier zero-knowledge interactive argument of knowledge. In particular, we can treat the resulting protocol as an argument for committed relaxed R1CS because the verifier is now provided with (polynomial) commitments to E and W . Applying the Fiat-Shamir transform [24] achieves non-interactivity and zero-knowledge in the random oracle model.

The claimed efficiency follows from the efficiency of the polynomial IOP, PC_{BP}, and Sparse-PC_{BP}. In more detail, using Sparse-PC_{BP}, the encoder takes $O_\lambda(n)$ time to create commitments $2 \log m$ -variate sparse multilinear polynomials $\tilde{A}, \tilde{B}, \tilde{C}$. The prover’s costs in the polynomial IOP is $O(n)$. Furthermore, proving the evaluations of two $O(\log m)$ -variate multilinear polynomials using PC_{BP}, it takes $O_\lambda(m)$ time. And, to prove the evaluations of three $2 \log m$ -variate sparse multilinear polynomials of size n , using Sparse-PC_{BP}, it takes $O_\lambda(n)$ time. In total, the prover time is $O_\lambda(n)$. The proof length in the polynomial IOP is $O(\log m)$, and the proof sizes in the polynomial evaluation proofs is $O_\lambda(\log n)$, so the proof length is $O_\lambda(\log n)$. The verifier’s time in the polynomial IOP is $O(\log m)$. In addition, it verifies five polynomial evaluations, which costs $O_\lambda(n)$ time: the two polynomial in the instance take $O_\lambda(m)$ time using PC_{BP}, and the three polynomials in the structure takes $O_\lambda(n)$ time using Sparse-PC_{BP}. So, in total, the verifier time is $O_\lambda(n)$. \square

Corollary 1 (A zkSNARK from PC_{Dory}). *Assuming the hardness of the SXDH problem, there exists a zkSNARK in the random oracle model for committed relaxed R1CS with the following efficiency characteristics, where m denotes the dimensions of R1CS matrices and n denotes the number of non-zero entries in the matrices: The encoder runs in time $O_\lambda(n)$; The prover runs in time $O_\lambda(n)$; The proof length is $O_\lambda(\log n)$; and the verifier runs in time $O_\lambda(\log n)$.*

⁶ Appendix H describes a minor optimization and a corresponding Corollary.

Acknowledgments

We thank Justin Drake, Bryan Parno, and Justin Thaler for various conversations and helpful comments on a prior version of this paper. We thank Andy Polyakov and Kelly Olson for help with `pasta-msm`, a high-performance library used by Nova to perform multiscalar multiplications over the Pasta cycle of curves. We thank Chhi'mèd Künzang for helpful conversations and for PRs to the Nova repository, including an adapter that enables transforming circuits written with `bellperson` to a format that Nova uses. Abhiram Kothapalli (while at Carnegie Mellon University) was supported by a gift from Bosch, NSF Grant No. 1801369, and by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

References

- [1] bellperson. <https://github.com/filecoin-project/bellperson>
- [2] neptune. <https://github.com/filecoin-project/neptune>
- [3] Nova: Recursive SNARKs without trusted setup. <https://github.com/Microsoft/Nova>
- [4] Pasta curves. <https://github.com/zcash/pasta>
- [5] Ben-Sasson, E., Chiesa, A., Spooner, N.: Interactive Oracle Proofs. In: TCC (2016)
- [6] Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Scalable zero knowledge via cycles of elliptic curves. In: CRYPTO (2014)
- [7] Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In: ITCS (2012)
- [8] Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: Recursive composition and bootstrapping for SNARKs and proof-carrying data. In: STOC (2013)
- [9] Boneh, D., Bünz, B., Fisch, B.: A survey of two verifiable delay functions. Cryptology ePrint Archive, Report 2018/712 (2018)
- [10] Boneh, D., Drake, J., Fisch, B., Gabizon, A.: Halo Infinite: Recursive zk-SNARKs from any Additive Polynomial Commitment Scheme. Cryptology ePrint Archive, Report 2020/1536 (2020)
- [11] Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In: EUROCRYPT (2016)
- [12] Bowe, S., Grigg, J., Hopwood, D.: Halo: Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Report 2019/1021 (2019)
- [13] Bowe, S., Grigg, J., Hopwood, D.: Halo2 (2020), <https://github.com/zcash/halo2>
- [14] Brakerski, Z., Holmgren, J., Kalai, Y.T.: Non-interactive delegation and batch NP verification from standard computational assumptions. In: STOC. pp. 474–482 (2017)
- [15] Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: S&P (2018)
- [16] Bünz, B., Chiesa, A., Lin, W., Mishra, P., Spooner, N.: Proof-carrying data without succinct arguments. Cryptology ePrint Archive, Report 2020/1618 (2020)
- [17] Bünz, B., Chiesa, A., Mishra, P., Spooner, N.: Proof-carrying data from accumulation schemes. In: TCC (2020)
- [18] Bünz, B., Fisch, B., Szepieniec, A.: Transparent SNARKs from DARK compilers. In: EUROCRYPT (2020)
- [19] Bünz, B., Maller, M., Mishra, P., Vesely, N.: Proofs for inner pairing products and applications. Cryptology ePrint Archive, Report 2019/1177 (2019)

- [20] Chen, W., Chiesa, A., Dauterman, E., Ward, N.P.: Reducing participation costs via incremental verification for ledger systems. Cryptology ePrint Archive, Report 2020/1522 (2020)
- [21] Chiesa, A., Hu, Y., Maller, M., Mishra, P., Vesely, N., Ward, N.: Marlin: Pre-processing zkSNARKs with universal and updatable SRS. In: EUROCRYPT (2020)
- [22] Chiesa, A., Ojha, D., Spooner, N.: Fractal: Post-quantum and transparent recursive proofs from holography. In: EUROCRYPT (2020)
- [23] Drucker, A.: New limits to classical and quantum instance compression. *SIAM Journal of Computing* **44**(5), 144–3–1479 (2015)
- [24] Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: CRYPTO. pp. 186–194 (1986)
- [25] Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct NIZKs without PCPs. In: EUROCRYPT (2013)
- [26] Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: STOC. pp. 99–108 (2011)
- [27] Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof-systems. In: STOC (1985)
- [28] Grassi, L., Khovratovich, D., Rechberger, C., Roy, A., Schofnegger, M.: Poseidon: A new hash function for zero-knowledge proof systems. Cryptology ePrint Archive, Paper 2019/458 (2019)
- [29] Groth, J.: On the size of pairing-based non-interactive arguments. In: EUROCRYPT (2016)
- [30] Harnik, D., Naor, M.: On the compressibility of \mathcal{NP} instances and cryptographic applications. *SIAM Journal of Computing* **39**(5), 1667–1713 (2010)
- [31] Kate, A., Zaverucha, G.M., Goldberg, I.: Constant-size commitments to polynomials and their applications. In: ASIACRYPT. pp. 177–194 (2010)
- [32] Kilian, J.: A note on efficient zero-knowledge proofs and arguments (extended abstract). In: STOC (1992)
- [33] Labs, O.: Mina cryptocurrency (2020), <https://minaprotocol.com>
- [34] Lee, J.: Dory: Efficient, transparent arguments for generalised inner products and polynomial commitments. Cryptology ePrint Archive, Report 2020/1274 (2020)
- [35] Lee, J., Nikitin, K., Setty, S.: Replicated state machines without replicated execution. In: S&P (2020)
- [36] Lee, J., Setty, S., Thaler, J., Wahby, R.: Linear-time zero-knowledge SNARKs for R1CS. Cryptology ePrint Archive, Report 2021/030 (2021)
- [37] Lund, C., Fortnow, L., Karloff, H., Nisan, N.: Algebraic methods for interactive proof systems. In: FOCS (Oct 1990)
- [38] Micali, S.: CS proofs. In: FOCS (1994)
- [39] Reingold, O., Rothblum, G.N., Rothblum, R.D.: Constant-round interactive proofs for delegating computation. In: STOC. pp. 49–62 (2016)
- [40] Setty, S.: Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In: CRYPTO (2020)
- [41] Setty, S., Angel, S., Gupta, T., Lee, J.: Proving the correct execution of concurrent services in zero-knowledge. In: OSDI (Oct 2018)

- [42] Setty, S., Braun, B., Vu, V., Blumberg, A.J., Parno, B., Walfish, M.: Resolving the conflict between generality and plausibility in verified computation. In: EuroSys (Apr 2013)
- [43] Setty, S., Lee, J.: Quarks: Quadruple-efficient transparent zkSNARKs. Cryptology ePrint Archive, Report 2020/1275 (2020)
- [44] Thaler, J.: Time-optimal interactive proofs for circuit evaluation. In: CRYPTO (2013)
- [45] Thaler, J.: Proofs, arguments, and zero-knowledge. <http://people.cs.georgetown.edu/jthaler/ProofsArgsAndZK.html> (2020)
- [46] Valiant, P.: Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In: TCC. pp. 552–576 (2008)
- [47] Wahby, R.S., Ji, Y., Blumberg, A.J., Shelat, A., Thaler, J., Walfish, M., Wies, T.: Full accounting for verifiable outsourcing. In: CCS (2017)
- [48] Wahby, R.S., Tzialla, I., Shelat, A., Thaler, J., Walfish, M.: Doubly-efficient zkSNARKs without trusted setup. In: S&P (2018)
- [49] Wesolowski, B.: Efficient verifiable delay functions. In: EUROCRYPT. pp. 379–407 (2019)
- [50] Xie, T., Zhang, J., Zhang, Y., Papamanthou, C., Song, D.: Libra: Succinct zero-knowledge proofs with optimal prover computation. In: CRYPTO (2019)
- [51] Zhang, Y., Genkin, D., Katz, J., Papadopoulos, D., Papamanthou, C.: vSQL: Verifying arbitrary SQL queries over dynamic outsourced databases. In: S&P (2017)

Supplementary Materials

A Additional Related Work

A folding scheme is reminiscent of the technique used in protocols such as Bulletproofs [15], where the prover splits an N -sized inner product instance into two $N/2$ -sized inner product instances, and then the prover and the verifier interactively combine the two $N/2$ -sized instances into a single $N/2$ -sized instance. A folding scheme differs from this technique in that it targets folding NP relations. Furthermore, our folding scheme for relaxed R1CS folds arbitrary NP instances.

A folding scheme is also reminiscent of the sum-check protocol [37] applied to layered circuits [47, 48], where the prover and the verifier interactively combine $\beta \geq 2$ instances of depth-1 circuit satisfiability into a single depth-1 circuit satisfiability instance, and the combined instance is proven using additional invocations of the sum-check protocol. Once the protocol combines β instances of depth-1 circuit satisfiability into a single depth-1 circuit satisfiability instance, it is not entirely clear how to fold additional circuit satisfiability instances into a single running instance. Whereas, with a folding scheme, the prover and the verifier can continually fold NP instances into a single running instance.

Finally, a folding scheme is also related to other prior notions such as “batch verification” [14, 39] and “instance compression” [23, 30]. Given β NP instances, batch verification protocols aim to produce a succinct proof that all β instances are satisfiable. In contrast, a folding scheme’s goal is to transform the satisfiability of two NP instances into a single NP-instance. (AND)-instance compression schemes efficiently compress β SAT instances into a single short string that retains information about whether all instances are satisfiable. A folding scheme differs from this notion in that compression is aided by an (untrusted) prover who knows (purported) witnesses to the original β instances. Furthermore, unlike instance compression, a folding scheme can be interactive.

B Proof of Theorem 3 (A Folding Scheme for Committed Relaxed R1CS)

Lemma 6 (Completeness). *Construction 1 is a folding scheme for committed relaxed R1CS with perfect completeness.*

Proof. Consider a finite field \mathbb{F} , a succinct, hiding, and homomorphic commitment scheme Com over \mathbb{F} , and public parameters consisting of size bounds $m, n, \ell \in \mathbb{N}$ where $m > \ell$, and commitment parameters pp_W and pp_E for vectors of size m and $m - \ell - 1$ respectively sampled by the generator. Consider an adversarially chosen R1CS structure $(A, B, C) \in \mathbb{F}^{m \times m}$ and two committed relaxed R1CS instances

$$\varphi_1 = (\overline{E}_1, u_1, \overline{W}_1, x_1) \quad \text{and} \quad \varphi_2 = (\overline{E}_2, u_2, \overline{W}_2, x_2).$$

Suppose that the prover \mathcal{P} , in addition to the two instances, holds *satisfying* witnesses to both instances, $(E_1, r_{E_1}, W_1, r_{W_1})$ and $(E_2, r_{E_2}, W_2, r_{W_2})$. Let $Z_1 = (W_1, x_1, u_1)$ and $Z_2 = (W_2, x_2, u_2)$.

Now suppose that \mathcal{P} and \mathcal{V} compute a folded instance $\varphi = (\bar{E}, u, \bar{W}, x)$, and suppose that \mathcal{P} computes a folded witness (E, r_E, W, r_W) . To prove completeness, we must show that (E, W) is a satisfying witness for instance φ . Let $Z = (W, x, u)$.

For (E, W) to be a satisfying witness, we must have the following:

$$AZ \circ BZ = u \cdot CZ + E \quad (3)$$

and

$$\bar{E} = \text{Com}(\text{pp}_E, E, r_E) \quad (4)$$

$$\bar{W} = \text{Com}(\text{pp}_W, W, r_W) \quad (5)$$

It is easy to see that Equations (4) and (5) hold from the additive homomorphism of the commitment scheme.

Thus, we focus on proving that Equation (3) holds. By construction, for Equation (3) to hold, we must have for $r \in_R \mathbb{F}$

$$A(Z_1 + r \cdot Z_2) \circ B(Z_1 + r \cdot Z_2) = (u_1 + r \cdot u_2) \cdot C(Z_1 + r \cdot Z_2) + E.$$

Distributing, we must have

$$\begin{aligned} & AZ_1 \circ BZ_1 + r(AZ_1 \circ BZ_2 + AZ_2 \circ BZ_1) + r^2(AZ_2 \circ BZ_2) = \\ & u_1 \cdot CZ_1 + r(u_1 \cdot CZ_2 + u_2 CZ_1) + r^2 \cdot u_2 \cdot CZ_2 + E. \end{aligned}$$

Aggregating by powers of r , we must have

$$\begin{aligned} & (AZ_1 \circ BZ_1 - u_1 \cdot CZ_1) + \\ & r(AZ_1 \circ BZ_2 + AZ_2 \circ BZ_1 - u_1 \cdot CZ_2 - u_2 CZ_1) + \\ & r^2(AZ_2 \circ BZ_2 - u_2 \cdot CZ_2) \\ & = E. \end{aligned} \quad (6)$$

However, because W_1 and W_2 are satisfying witnesses, we have

$$\begin{aligned} & AZ_1 \circ BZ_1 - u_1 \cdot CZ_1 = E_1 \\ & AZ_2 \circ BZ_2 - u_2 \cdot CZ_2 = E_2. \end{aligned}$$

Additionally, by construction we have

$$AZ_1 \circ BZ_2 + AZ_2 \circ BZ_1 - u_1 \cdot CZ_2 - u_2 \cdot CZ_1 = T.$$

Thus, by substitution, for Equation (6) to hold we must have

$$E_1 + r \cdot T + r^2 \cdot E_2 = E,$$

which holds by construction. □

Lemma 7 (Knowledge Soundness). *Construction 1 is a knowledge sound folding scheme for committed relaxed R1CS.*

Proof. Consider a finite field \mathbb{F} , a succinct, hiding, and homomorphic commitment scheme Com over \mathbb{F} , and public parameters consisting of size bounds $m, n, \ell \in \mathbb{N}$ where $m > \ell$, and commitment parameters pp_W and pp_E for vectors of size m and $m - \ell - 1$ respectively sampled by the generator. Consider an adversarially chosen R1CS structure $(A, B, C) \in \mathbb{F}^{m \times m}$ and two committed relaxed R1CS instances

$$\varphi_1 = (\overline{E}_1, u_1, \overline{W}_1, x_1) \quad \text{and} \quad \varphi_2 = (\overline{E}_2, u_2, \overline{W}_2, x_2).$$

We prove knowledge soundness via the forking lemma (Lemma 1). That is, we prove that there exists a PPT algorithm \mathcal{X} such that when given public parameters pp , structure (A, B, C) , and a tree of accepting transcripts and the corresponding folded instance-witness pairs outputs a satisfying witness with probability $1 - \text{negl}(\lambda)$.

In more detail, suppose \mathcal{X} is provided three transcripts (τ_1, τ_2, τ_3) with the same initial commitment \overline{T} from the prover. Note that a transcript τ_i for $i \in \{1, 2, 3\}$ additionally comes attached with an accepting witness $\tau_i.(W, E, r_W, r_E)$ and the verifier's randomness $\tau_i.r$. Interpolating points $(\tau_1.r, \tau_1.W)$ and $(\tau_2.r, \tau_2.W)$, \mathcal{X} retrieves (W_1, W_2) such that

$$W_1 + \tau_i.r \cdot W_2 = \tau_i.W \tag{7}$$

for $i \in \{1, 2\}$. Similarly interpolating points $(\tau_1.r, \tau_1.E)$, $(\tau_2.r, \tau_2.E)$, $(\tau_3.r, \tau_3.E)$, \mathcal{X} retrieves (E_1, E_2) and a cross-term T such that

$$E_1 + \tau_i.r \cdot T + \tau_i.r^2 \cdot E_2 = \tau_i.E \tag{8}$$

for $i \in \{1, 2, 3\}$. Using the same approach, \mathcal{X} can interpolate for r_{W_1}, r_{W_2} and r_{E_1}, r_T, r_{E_2} . We must argue that $(W_1, E_1, r_{W_1}, r_{E_1})$ and $(W_2, E_2, r_{W_2}, r_{E_2})$ are indeed satisfying witnesses for φ_1 and φ_2 respectively.

We first show that the retrieved witness elements are valid openings to the corresponding commitments in the instance. For $i \in \{1, 2\}$, because $\tau_i.W, \tau_i.r_W$ is part of a satisfying witness, by construction,

$$\begin{aligned} & \text{Com}(\text{pp}_W, W_1, r_{W_1}) + \tau_i.r \cdot \text{Com}(\text{pp}_W, W_2, r_{W_2}) \\ &= \text{Com}(\text{pp}_W, W_1 + \tau_i.r \cdot W_2, r_{W_1} + \tau_i.r \cdot r_{W_2}) \\ &= \text{Com}(\text{pp}_W, \tau_i.W, \tau_i.r_W) \\ &= \overline{W}_1 + \tau_i.r \cdot \overline{W}_2. \end{aligned}$$

Interpolating, we must have that

$$\text{Com}(\text{pp}_W, W_1, r_{W_1}) = \overline{W}_1 \tag{9}$$

$$\text{Com}(\text{pp}_W, W_2, r_{W_2}) = \overline{W}_2. \tag{10}$$

Similarly, for $i \in \{1, 2, 3\}$, because $\tau_i.E, \tau_i.r_E$ is part of a satisfying witness, by construction, we must have

$$\begin{aligned}
& \text{Com}(\text{pp}_E, E_1, r_{E_1}) + \tau_i.r \cdot \text{Com}(\text{pp}_E, T, r_T) + \tau_i.r^2 \cdot \text{Com}(\text{pp}_E, E_2, r_{E_2}) \\
&= \text{Com}(\text{pp}_E, E_1 + \tau_i.r \cdot T + \tau_i.r^2 \cdot E_2, r_{E_1} + \tau_i.r \cdot r_T + \tau_i.r^2 \cdot r_{E_2}) \\
&= \text{Com}(\text{pp}_E, \tau_i.E, \tau_i.r_E) \\
&= \bar{E}_1 + \tau_i.r \cdot \bar{T} + \tau_i.r^2 \cdot \bar{E}_2.
\end{aligned}$$

Interpolating, we must have that

$$\begin{aligned}
\text{Com}(\text{pp}_E, E_1, r_{E_1}) &= \bar{E}_1 \\
\text{Com}(\text{pp}_E, E_2, r_{E_2}) &= \bar{E}_2
\end{aligned}$$

Next, we must show that (W_1, E_1) and (W_2, E_2) satisfy the relaxed R1CS relation. To show this, we must first argue that Equation (7) holds for $i = 3$ as well (i.e., $W_1 + \tau_3.r \cdot W_2 = \tau_3.W$). Indeed, Equations (9) and (10) imply that

$$\begin{aligned}
& \text{Com}(\text{pp}_W, W_1 + \tau_3.r \cdot W_2, r_{W_1} + \tau_3.r \cdot r_{W_2}) \\
&= \text{Com}(\text{pp}_W, W_1, r_{W_1}) + \tau_3.r \cdot \text{Com}(\text{pp}_W, W_2, r_{W_2}) \\
&= \bar{W}_1 + \tau_3.r \cdot \bar{W}_2 \\
&= \text{Com}(\text{pp}_W, \tau_3.W, \tau_3.W)
\end{aligned}$$

Thus, by the binding property of Com , we must additionally have that

$$W_1 + \tau_3.r \cdot W_2 = \tau_3.W \quad (11)$$

with probability $1 - \text{negl}(\lambda)$.

Because $(\tau_i.W, \tau_i.E)$ is part of a satisfying witness, for $i \in \{1, 2, 3\}$ we have

$$A(\tau_i.Z) \circ B(\tau_i.Z) = u \cdot C(\tau_i.Z) + \tau_i.E$$

where $\tau_i.Z = (\tau_i.W, \tau_i.x, \tau_i.u)$. However, by Equations (7), (8), and (11) for $i \in \{1, 2, 3\}$, this implies that with probability $1 - \text{negl}(\lambda)$

$$\begin{aligned}
& A(Z_1 + \tau_i.r \cdot Z_2) \circ B(Z_1 + \tau_i.r \cdot Z_2) = \\
& (u_1 + \tau_i.r \cdot u_2) \cdot C(Z_1 + \tau_i.r \cdot Z_2) + (E_1 + \tau_i.r \cdot T + \tau_i.r^2 \cdot E_2)
\end{aligned}$$

where $Z_1 = (W_1, x_1, u_1)$ and $Z_2 = (W_2, x_2, u_2)$. Expanding and interpolating, we have that

$$\begin{aligned}
AZ_1 \circ BZ_1 &= u_1 \cdot CZ_1 + E_1 \\
AZ_2 \circ BZ_2 &= u_2 \cdot CZ_2 + E_2
\end{aligned}$$

with probability $1 - \text{negl}(\lambda)$. Thus, $(W_1, E_1, r_{W_1}, r_{E_1})$ and $(W_2, E_2, r_{W_2}, r_{E_2})$ meet all the requirements to be satisfying witnesses for φ_1 and φ_2 respectively with probability $1 - \text{negl}(\lambda)$. \square

Lemma 8 (Zero-Knowledge). *Construction 1 is a zero-knowledge folding scheme for committed relaxed R1CS.*

Proof. Intuitively, zero-knowledge holds because the prover only sends a single hiding commitment. More formally, the simulator \mathcal{S} samples random $T \in \mathbb{F}^m$ and $r \in \mathbb{F}$ and computes $\bar{T} = \text{Com}(\text{pp}_E, T, r)$. Next, \mathcal{S} derives the verifier's challenge, r , using ρ and outputs $\text{tr} = (\bar{T}, r)$. (Perfect) zero-knowledge holds from the (perfect) hiding property of the underlying commitment scheme. \square

C Proofs for IVC from Folding Schemes (Construction 3)

Lemma 9 (IVC Completeness). *Construction 3 is an IVC scheme that satisfies perfect 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 $(\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, F)$. Now consider a proof Π_i such that

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

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 $(\mathbf{u}_\perp, \mathbf{w}_\perp)$ satisfies F' . Moreover, by construction, $(\mathbf{u}_1, \mathbf{w}_1)$ must also be satisfying. Additionally, by the construction of F' , we have

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

And, by the construction of \mathcal{P} , we have $\mathbf{w}_1 \cdot E = \mathbf{0}$ and $\mathbf{w}_1 \cdot r_E = 0$ (implying that $\mathbf{u}_1 \cdot \bar{E} = 0$), and $\mathbf{u}_1 \cdot u = 1$. Therefore, we have

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

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

$$\Pi_i = ((U_i, W_i), (u_i, w_i))$$

we have that

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

and suppose that

$$\Pi_{i+1} = ((U_{i+1}, W_{i+1}), (u_{i+1}, w_{i+1})) \leftarrow \mathcal{P}(\text{pk}, (i, z_0, z_i), \omega_i, \Pi_i).$$

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

$$(U_{i+1}, W_{i+1}, \bar{T}) \leftarrow \text{NIFS.P}(\text{pk}, (U_i, W_i), (u_i, w_i)).$$

Thus, by the completeness of the underlying folding scheme, and the premise that (u_i, w_i) and (U_i, W_i) are satisfying instance-witness pairs, we have that (U_{i+1}, W_{i+1}) is a satisfying instance-witness pair. Additionally, by the premise, we have that $u_i.x = \text{hash}(\text{vk}, i, z_0, z_i, U_i)$, $u_i.\bar{E} = 0$, and $u_i.u = 1$. Therefore, \mathcal{P} can construct a satisfying instance-witness pair (u_{i+1}, w_{i+1}) that represents the correct execution of F' on input $(U, u, (i, z_0, z_i), \omega_i, \bar{T})$. By construction, this particular input implies that

$$\begin{aligned} u_{i+1}.x &= \text{hash}(\text{vk}, u_i.i + 1, u_i.z_0, F(z_i, w_i), \text{NIFS.V}(\text{vk}, U_i, u_i, \bar{T})) \\ &= \text{hash}(\text{vk}, u_i.i + 1, u_i.z_0, z_{i+1}, U_{i+1}) \end{aligned} \quad (12)$$

by the correctness of the underlying folding scheme. As in the base case, by the construction of \mathcal{P} , we have $w_{i+1}.E = \mathbf{0}$, $w_{i+1}.r_E = 0$ (implying that $u_{i+1}.\bar{E} = 0$) and $u_{i+1}.u = 1$. Thus, by Equation (12) we have

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

□

Lemma 10 (IVC 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 [16]. Let $\text{pp} \leftarrow \mathcal{G}(1^\lambda)$. Consider an expected polynomial-time adversary \mathcal{P}^* that outputs a function F on input pp , and let $(\text{pk}, \text{vk}) \leftarrow \mathcal{K}(\text{pp}, F)$. 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 (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 $\mathcal{E}_i(\text{pp})$ 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

$$\mathcal{V}(\text{vk}, i, z_0, z_i, \Pi_i) = 1 \tag{13}$$

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 $\mathcal{E}_0(\text{pp})$ 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 $\mathcal{E}_i(\text{pp})$ 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 $\mathcal{E}_n(\text{pp}, \rho)$ output (\perp, \perp, Π_n) where Π_n is the output of $\mathcal{P}^*(\text{pp}, \rho)$. 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:

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

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

By the inductive hypothesis, we have that $\mathcal{V}(\text{vk}, i, z_0, z_i, \Pi_i) = 1$, where $\Pi_i \leftarrow \mathcal{E}_i(\text{pp})$ 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

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

Because \mathcal{V} ensures that $(\mathbf{u}_i \cdot \bar{E}, \mathbf{u}_i \cdot u) = (\bar{0}, 1)$ 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

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

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 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

$\mathcal{E}_{i-1}(\text{pp}, \rho)$:

- (1) $((\mathbf{U}_{i-1}, \mathbf{u}_{i-1}), (\mathbf{U}_i, \mathbf{W}_i), \bar{T}_{i-1}) \leftarrow \tilde{\mathcal{P}}_{i-1}(\text{pp}, \rho)$
- (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 $(\mathbf{w}_{i-1}, \mathbf{W}_{i-1}) \leftarrow \tilde{\mathcal{E}}_{i-1}(\text{pp})$.
- (5) Let $\Pi_{i-1} \leftarrow ((\mathbf{U}_{i-1}, \mathbf{W}_{i-1}), (\mathbf{u}_{i-1}, \mathbf{w}_{i-1}))$.
- (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 $\mathcal{V}(\text{vk}, i, z_0, z_i, (\mathbf{U}_i, \mathbf{W}_i), (\mathbf{u}_i, \mathbf{w}_i)) = 1$ with probability $\epsilon - \text{negl}(\lambda)$. Because \mathcal{V} additionally checks that

$$\mathbf{u}_i \cdot \mathbf{x} = \text{hash}(\text{vk}, i, z_0, z_i, \mathbf{U}_i) \quad (14)$$

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 $(\mathbf{u}_i, \mathbf{w}_i)$ satisfies F' , and $(\mathbf{U}_{i-1}, \mathbf{u}_{i-1})$ were retrieved from \mathbf{w}_i , by the binding property of the hash function, and by Equation (14), we have that

$$\begin{aligned} \mathbf{u}_{i-1} \cdot \mathbf{x} &= \text{hash}(\text{vk}, i-1, z_0, z_{i-1}, \mathbf{U}_{i-1}) \\ (\mathbf{u}_{i-1} \cdot \bar{E}, \mathbf{u}_{i-1} \cdot u) &= (\bar{0}, 1) \end{aligned}$$

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 $\epsilon - \text{negl}(\lambda)$. □

D Proof of Theorem 4 (A zkSNARK of a Valid IVC Proof)

Theorem. *Construction 4 is a zero-knowledge succinct non-interactive argument of knowledge (zkSNARK) of an IVC proof as described in Construction 3.*

Proof. Completeness and knowledge soundness hold due to the completeness and knowledge soundness of the underlying succinct non-interactive argument of knowledge and the non-interactive folding scheme. Succinctness holds due to the succinctness of the commitment scheme underlying the non-interactive folding scheme, the succinctness of the hash function, and the succinctness of the underlying non-interactive argument. To prove zero-knowledge, we first construct a simulator \mathcal{S} as follows

$\mathcal{S}(\text{pp}, (i, z_0, z_i), \tau) \rightarrow \pi$:

If n is 0,

output \perp ;

otherwise,

- (1) for $(r_x, r_W) \leftarrow_R \mathbb{F}$,
 $\mathbf{u}_1 \leftarrow (\text{hash}(\text{vk}, 1, z_0, \perp, \mathbf{u}_\perp; r_x), (0, 1, \text{Com}(0; r_W)))$
 $\mathbf{U}_1 \leftarrow \mathbf{u}_\perp$
- (2) for $i \in \{1, \dots, n-1\}$,
 $\bar{T}_i \leftarrow \text{NIFS.S}(\text{pp}_{\text{NIFS}}, \mathbf{U}_i, \mathbf{u}_i)$
 $\mathbf{U}_{i+1} \leftarrow \text{NIFS.V}(\text{vk}_{\text{NIFS}}, \mathbf{U}_i, \mathbf{u}_i, \bar{T}_i)$
 $(r_x, r_W) \leftarrow_R \mathbb{F}$
 $\mathbf{u}_{i+1} \leftarrow (\text{hash}(\text{vk}, i+1, z_0, \perp, \mathbf{U}_{i+1}; r_x), (0, 1, \text{Com}(0; r_W)))$
- (3) $\bar{T}_n \leftarrow \text{NIFS.S}(\text{pp}_{\text{NIFS}}, \mathbf{U}_n, \mathbf{u}_n)$
- (4) $\pi_{U'} \leftarrow \text{zkSNARK.S}(\text{pp}_{\text{zkSNARK}}, \text{NIFS.V}(\text{vk}_{\text{NIFS}}, \mathbf{U}_n, \mathbf{u}_n, \bar{T}_n), \tau)$
- (5) output $(\mathbf{U}_n, \mathbf{u}_n, \bar{T}_n, \pi_{U'})$.

where NIFS.S and NIZK.S are the simulators for the underlying non-interactive folding scheme and the zkSNARK respectively.

In the base case, both the simulator and honest prover output \perp . In the non-base case, we reason that (\mathbf{U}, \mathbf{u}) are indistinguishable from that produced by an honest prover. We do so by showing that if $(\mathbf{U}_i, \mathbf{u}_i)$ is indistinguishable then $(\mathbf{U}_{i+1}, \mathbf{u}_{i+1})$ is indistinguishable.

- Suppose $i = 1$. By the definition of IVC.P , $\mathbf{U}_1 = \mathbf{u}_\perp$ in both the real and simulated setting. Additionally we have that $\mathbf{u}_1 \cdot \bar{E} = 0$ and $\mathbf{u} \cdot \mathbf{u} = 1$ in both the real and simulated setting. Due to the hiding property of the hash function, we have that the simulated hash of the public IO, $\text{hash}(\text{vk}, i+1, z_0, \perp, \mathbf{U}_{i+1}, r_x)$, is indistinguishable from the one generated honestly. Due to the hiding property of the commitment scheme, we additionally have that the simulated commitment to the witness, $\text{Com}(0; r_W)$, is indistinguishable from the one generated honestly.
- Suppose now that $(\mathbf{U}_i, \mathbf{u}_i)$ is indistinguishable in the real and the simulated setting. Due to the zero-knowledge property of the underlying non-interactive folding scheme, we have that $\bar{T}_i \leftarrow \text{NIFS.S}(\text{pp}, \mathbf{U}_i, \mathbf{u}_i)$ is indistinguishable in the real and the simulated setting. Thus, we have that $\mathbf{U}_{i+1} \leftarrow \text{NIFS.V}(\text{vk}, \mathbf{U}_i, \mathbf{u}_i, \bar{T}_i)$ is indistinguishable in the real and the simulated setting. Then, we have that $\text{hash}(\text{vk}, i+1, z_0, \perp, \mathbf{U}_{i+1}, r_x)$ is indistin-

guishable in the real and the simulated setting. The remaining pieces of \mathbf{u}_{i+1} are indistinguishable in the real and the simulated setting for the same reasons as the base case. Thus, we have that $(\mathbf{U}_{i+1}, \mathbf{u}_{i+1})$ is indistinguishable in the real and the simulated setting.

Now, because $(\mathbf{u}_n, \mathbf{U}_n)$ are indistinguishable from that produced by honest IVC.P, we have that \bar{T}_n produced by $\text{NIFS.S}(\mathbf{pp}, \mathbf{U}_n, \mathbf{u}_n)$ is indistinguishable from that produced by \mathcal{P} due to the zero-knowledge property of the underlying non-interactive folding scheme. This means that \mathbf{U}' generated by $\text{NIFS.V}(\mathbf{vk}, \mathbf{U}_n, \mathbf{u}_n, \bar{T})$ is indistinguishable from that produced honestly. This means that $\pi_{\mathbf{U}'}$ is indistinguishable in the real and the simulated setting due to the zero-knowledge property of the underlying zkSNARK. Thus, the simulated zkSNARK $(\mathbf{U}_n, \mathbf{u}_n, \bar{T}_n, \pi_{\mathbf{U}'})$ is indistinguishable from the one generated honestly. \square

E Proof of Lemma 1 (Forking Lemma for Folding Schemes)

We provide a proof for our variant of the forking lemma by adapting the proof of the forking lemma from Bootle et al. [11].

Lemma. *Let \mathbb{F} denote a finite field. Consider a $(2\mu + 1)$ -move folding scheme $\Pi = (\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$. Π satisfies knowledge soundness if there exists a PPT \mathcal{X} such that for all input instance pairs (u_1, u_2) , outputs satisfying witnesses (w_1, w_2) with probability $1 - \text{negl}(\lambda)$, given public parameters \mathbf{pp} , a structure \mathbf{s} , and an (n_1, \dots, n_μ) -tree of accepting transcripts and the corresponding folded instance-witness pairs (u, w) . This tree comprises of n_1 transcripts (and the corresponding instance-witness pairs) with fresh randomness in \mathcal{V} 's first message; and for each such transcript, n_2 transcripts (and the corresponding instance-witness pairs) with fresh randomness in \mathcal{V} 's second message; etc., for a total of $\prod_{i=1}^\mu n_i$ leaves bounded by $\text{poly}(\lambda)$.*

Proof. Consider a $(2\mu + 1)$ -move folding scheme $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ for relation \mathcal{R} with randomness sampled uniformly from \mathbb{F} where $|\mathbb{F}| \geq 2^\lambda$. Suppose there exists a PPT \mathcal{X} such that for arbitrary input instances (u_1, u_2) , outputs the corresponding satisfying witnesses (w_1, w_2) with probability $1 - \text{negl}(\lambda)$, given public parameters \mathbf{pp} , a structure \mathbf{s} , and an (n_1, \dots, n_μ) -tree of accepting transcripts and the corresponding folded instance-witness pairs where $\prod_{i=1}^\mu n_i$. Consider an arbitrary expected polynomial-time adversary \mathcal{P}^* , and $\mathbf{pp} \leftarrow \mathcal{G}(1^\lambda)$. Let ϵ be the probability that, given $(u_1, u_2) \leftarrow \mathcal{P}^*(\mathbf{pp}, \rho)$, $\mathcal{P}^*(\mathbf{pp}, \rho)$ succeeds in producing a valid witness w for the folded instance u . To prove that $(\mathcal{G}, \mathcal{K}, \mathcal{P}, \mathcal{V})$ satisfies knowledge-soundness we will show that there exists a corresponding extractor \mathcal{E} that outputs a valid witness pair (w_1, w_2) for instances (u_1, u_2) given public parameters \mathbf{pp} and the prover randomness ρ with probability $\epsilon - \text{negl}(\lambda)$.

We first construct an extractor \mathcal{E} as follows:

$\underline{\mathcal{E}(\mathbf{pp}, \rho) \rightarrow (w_1, w_2)}$:

- (1) Compute $(s, (u_1, u_2)) \leftarrow \mathcal{P}^*(\text{pp}, \rho)$.
- (2) Compute $\text{tree} \leftarrow \mathcal{T}(1)$, where the function \mathcal{T} is defined below.
- (3) If tree is not a valid (n_1, \dots, n_μ) -tree (i.e. there are collisions in the verifier's randomness) return \perp .
- (4) Output $(w_1, w_2) \leftarrow \mathcal{X}(\text{pp}, s, u_1, u_2, \text{tree})$.

$\mathcal{T}(i) \rightarrow \text{tree}$:

- (1) Sample the verifier's randomness for round i ; denote it as r_i .
- (2) If $i = \mu + 1$, compute

$$(u, w) \leftarrow \langle \mathcal{P}^*(\rho), \mathcal{V}(u_1, u_2) \rangle$$

and let tr be the corresponding transcript. If $(u, w) \in \mathcal{R}$, output $\{(\text{tr}, (u, w))\}$; otherwise output \perp .

- (3) With fixed verifier randomness (r_1, \dots, r_i) , compute $\text{tree} \leftarrow \mathcal{T}(i + 1)$ *once*. If $\text{tree} = \perp$ output \perp .
- (4) With fixed verifier randomness (r_1, \dots, r_i) , repeatedly run $\mathcal{T}(i + 1)$ until $n_i - 1$ additional lists of accepting transcripts are acquired. Append all the results to tree .
- (5) Output tree .

Let E_1 denote the event that $\mathcal{T}(1)$ outputs $\text{tree} \neq \perp$ in less than T time steps (we specify T later). Given E_1 , let E_2 denote the event that the resulting tree is a valid (n_1, \dots, n_μ) -tree (i.e. there are no collisions in the verifier's randomness). Given E_1 and E_2 , let E_3 denote the event that \mathcal{X} succeeds with tree as input. Then, we have that \mathcal{E} succeeds with probability

$$P_{\mathcal{E}} = \Pr[E_3] \cdot \Pr[E_2] \cdot \Pr[E_1].$$

We now compute each of these probabilities. To compute $\Pr[E_1]$, we observe that $\mathcal{T}(1)$ fails (i.e. returns \perp) only when its *first* call to $\mathcal{T}(2)$ fails. Likewise, $\mathcal{T}(2)$ fails only when its first call to $\mathcal{T}(3)$ fails. Chaining these assertions, we have that $\mathcal{T}(1)$ fails with probability $(1 - \epsilon)$, which, by assumption, is the probability that $\mathcal{T}(\mu + 1)$ fails. Thus, the expected number of times $\mathcal{T}(i)$ calls $\mathcal{T}(i + 1)$ is

$$\begin{aligned} & 1 + \Pr[\text{First call to } \mathcal{T}(i + 1) \text{ succeeds}] \cdot \frac{(n_i - 1)}{\Pr[\mathcal{T}(i + 1) \text{ returns } \text{tree} \neq \perp]} \\ &= 1 + \epsilon \cdot \frac{(n_i - 1)}{\epsilon} \\ &= n_i. \end{aligned}$$

Hence, the total runtime is expected to be $t = O(\prod_{i=1}^{\mu} n_i)$ which is bounded above by $\text{poly}(\lambda)$ by assumption. Moreover, by Markov's inequality, we have that \mathcal{T} runs for time longer than $T > t$ with probability $\frac{t}{T}$. Thus, we have that

$$\Pr[E_1] = (1 - \frac{t}{T}) \cdot \epsilon.$$

Given E_1 we have that \mathcal{T} runs in at most T time-steps. This ensures that there are at most T random challenges produced for the verifier, and that the probability of collision is at most $\frac{T^2}{|\mathbb{F}|}$. Thus, we have that

$$\Pr[E_2] = 1 - \frac{T^2}{|\mathbb{F}|}.$$

Finally, given E_1 and E_2 , we have that $\Pr[E_3] = 1 - \text{negl}(\lambda)$ by assumption. Consolidating our calculations we have that \mathcal{E} succeeds with probability

$$P_{\mathcal{E}} = (1 - \text{negl}(\lambda)) \cdot \left(1 - \frac{T^2}{|\mathbb{F}|}\right) \cdot \left(1 - \frac{t}{T}\right) \cdot \epsilon$$

Setting $T = \sqrt[3]{|\mathbb{F}|}$ (and assuming that $T \geq t$) we have that

$$\begin{aligned} P_{\mathcal{E}} &= (1 - \text{negl}(\lambda)) \cdot \left(1 - \frac{1}{\sqrt[3]{|\mathbb{F}|}}\right) \cdot \left(1 - \frac{t}{\sqrt[3]{|\mathbb{F}|}}\right) \cdot \epsilon \\ &= \epsilon - \text{negl}(\lambda) \end{aligned}$$

□

F A Commitment Scheme for Vectors

We first recall the syntax for commitment schemes for vectors introduced in Section 2.1 and then define the required binding and hiding properties. Next, we define the additionally required additively homomorphic and succinctness properties. For the sake of concreteness, we reproduce Pedersen's commitment scheme, which satisfies all of the required properties.

Definition 15 (Commitment Scheme). *A commitment scheme for \mathbb{F}^m is a tuple of three protocols with the following syntax that satisfy the two properties listed below:*

- $\text{pp} \leftarrow \text{Gen}(1^\lambda, m)$: produces public parameters pp .
- $C \leftarrow \text{Com}(\text{pp}, x, r)$: takes as input $x \in \mathbb{F}^m$ and $r \in \mathbb{F}$; produces a public commitment C .
- $b \leftarrow \text{Open}(\text{pp}, C, x, r)$: verifies the opening of commitment C to $x \in \mathbb{F}^m$ and $r \in \mathbb{F}$; outputs $b \in \{0, 1\}$.

(1) **Binding.** For any PPT adversary \mathcal{A} , the following probability is $\text{negl}(\lambda)$:

$$\Pr \left[\begin{array}{l} b_0 = b_1 = 1, \\ x_0 \neq x_1 \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow \text{Gen}(1^\lambda, m), \\ (C, x_0 \in \mathbb{F}^m, x_1 \in \mathbb{F}^m, r_0 \in \mathbb{F}, r_1 \in \mathbb{F}) \leftarrow \mathcal{A}(\text{pp}), \\ b_0 \leftarrow \text{Open}(\text{pp}, C, x_0, r_0), \\ b_1 \leftarrow \text{Open}(\text{pp}, C, x_1, r_1) \end{array} \right]$$

(2) **Hiding.** For all PPT adversaries $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, the following probability is $\text{negl}(\lambda)$:

$$\left| \frac{1}{2} - \Pr \left[b = \bar{b} \left| \begin{array}{l} (x_0, x_1, \text{st}) \leftarrow \mathcal{A}_0(\text{pp}), \\ b \leftarrow_R \{0, 1\}, r \leftarrow_R \mathbb{F}, \\ C \leftarrow \text{Com}(\text{pp}, x_b, r), \\ \bar{b} \leftarrow \mathcal{A}_1(\text{st}, C) \end{array} \right. \right] \right|$$

If hiding holds for all adversaries, then the commitment is statistically hiding.

Definition 16 (Additively Homomorphic). A commitment scheme for vectors over \mathbb{F}^m , $(\text{Gen}, \text{Com}, \text{Open})$, is additively homomorphic if for all public parameters pp produced from $\text{Gen}(1^\lambda, m)$, and for any $x_1, x_2 \in \mathbb{F}^m$ and for any $r_1, r_2 \in \mathbb{F}$, $\text{Com}(\text{pp}, x_1, r_1) + \text{Com}(\text{pp}, x_2, r_2) = \text{Com}(\text{pp}, x_1 + x_2, r_1 + r_2)$.

Definition 17 (Succinctness). A commitment scheme for vectors over \mathbb{F}^m , $(\text{Gen}, \text{Com}, \text{Open})$, provides succinct commitments if for all public parameters pp produced from $\text{Gen}(1^\lambda, m)$, and any $x \in \mathbb{F}^m$ and $r \in \mathbb{F}$, $|\text{Com}(\text{pp}, x, r)| = O_\lambda(\text{polylog}(|x|))$.

Construction 6 (Pedersen Commitment). For group \mathbb{G} such that $|\mathbb{G}| \geq 2^\lambda$ with scalar field \mathbb{F} and where the discrete logarithm problem is hard, the Pedersen commitment scheme for vectors over \mathbb{F} is defined as follows:

- $\text{pp} \leftarrow \text{Gen}(1^\lambda, m)$: Sample $g \leftarrow_R \mathbb{G}^m$, $h \leftarrow_R \mathbb{G}$. Output (g, h) .
- $C \leftarrow \text{Com}(\text{pp}, x \in \mathbb{F}^m, r \in \mathbb{F})$: Output $h^r \cdot \prod_{i \in \{0, \dots, m\}} g_i^{x_i}$.

Lemma 11 (Pedersen Commitment). The Pedersen commitment scheme is hiding, binding, additively homomorphic, and succinct.

G Background on Polynomials

This section is adapted from prior work [36, 40, 45]. 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 and powers of one or more variables; all arithmetic is performed over \mathbb{F} .
- The degree of a monomial is the sum of the exponents of variables in the monomial; the (total) degree of a polynomial g is the maximum degree of any monomial in g . Furthermore, the degree of a polynomial g in a particular variable x_i is the maximum exponent that x_i takes in any g 's monomials.
- A *multivariate* polynomial is a polynomial with more than one variable. A multivariate polynomial is called a *multilinear* polynomial if the degree of the polynomial in each variable is at most one.
- A multivariate polynomial g over a finite field \mathbb{F} is called *low-degree* if the degree of g in each variable is bounded above by a constant.

H An Optimization to Theorem 6

In Theorem 6, given that the verifier's runtime to verify an evaluation is at least linear in the polynomial size, the verifier does not gain by querying commitments to structure polynomials $\tilde{A}, \tilde{B}, \tilde{C}$. Rather, the verifier can directly evaluate them, which not only simplifies the protocol (as one does not need Sparse-PCBP), but also improves costs, both asymptotically and concretely.

Corollary 2. *Assuming the hardness of the discrete logarithm problem, there exists a non-interactive zero-knowledge argument of knowledge for committed relaxed RICS with the following efficiency characteristics, where m denotes the dimensions of the RICS coefficient matrices and n denotes the number of non-zero entries in the matrices: The encoder runs in time $O_\lambda(m)$; The prover runs in time $O(n) + O_\lambda(m)$; The proof length is $O_\lambda(\log m)$; and the verifier runs in time $n + O_\lambda(m)$.*