Algebraic Reductions of Knowledge

Abhiram Kothapalli[1] and Bryan Parno

Carnegie Mellon University
{akothapalli,parno}@cmu.edu

Abstract. We introduce *reductions of knowledge*, a generalization of arguments of knowledge, which reduce checking knowledge of a witness in one relation to checking knowledge of a witness in another (simpler) relation. Reductions of knowledge unify a growing class of modern techniques as well as provide a compositional framework to modularly reason about individual steps in complex arguments of knowledge. As a demonstration, we simplify and unify recursive arguments over linear algebraic statements by decomposing them as a sequence of reductions of knowledge. To do so, we develop the *tensor reduction of knowledge*, which generalizes the central reductive step common to many recursive arguments. Underlying the tensor reduction of knowledge is a new information-theoretic reduction, which, for any modules U, U_1, and U_2 such that $U \cong U_1 \otimes U_2$, reduces the task of evaluating a homomorphism in U to evaluating a homomorphism in U_1 and evaluating a homomorphism in U_2.

Keywords: Arguments of knowledge · Composition

1 Introduction

Arguments of knowledge [27] are powerful cryptographic primitives that allow a verifier to efficiently check that a prover *knows* a satisfying witness for a claimed statement. Such arguments provide strong integrity and privacy guarantees that enable a large class of cryptographic applications [21, 28, 37, 40, 43].

However, a growing body of work challenges the traditional paradigm by describing interactions in which the verifier does not fully resolve the prover's statement to true or false, but rather reduces it to a simpler statement to be checked:

– The well-studied *inner-product argument* [10] (along with subsequent optimizations [13] and generalizations [11, 16]) relies on recursively applying an interactive reduction from the task of checking knowledge of size n vectors to the task of checking knowledge of size $n/2$ vectors.

[1] Abhiram Kothapalli was supported by a fellowship from Protocol Labs, a gift from Bosch, NSF Grant No. 1801369, and the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation program sponsored by DARPA. An extended version of this work is available on the Cryptology ePrint Archive [30].

- *Aggregation schemes* for polynomial commitments [9, 12] and *unbounded aggregation schemes* for linear-map vector commitments [18] can both be viewed as interactive reductions from checking proofs of several openings of a commitment to checking a proof of a single opening of a commitment.
- *Split-accumulation schemes* [14] can be viewed as interactive reductions from checking several proofs of knowledge and several accumulators to checking a single accumulator. *Folding schemes* [31] can be viewed as interactive reductions from checking knowledge of two instances in a relation to checking knowledge of a single instance in the relation.
- As observed by Ràfols and Zapico [35], most argument systems with universal and updatable trusted setups [17,19,29,39] construct an interactive reduction from the task of checking knowledge of a preimage of a matrix evaluation to the task of checking knowledge of a preimage of a vector evaluation.

Such interactive reductions, although central to modern arguments, lack a unifying theoretical foundation. As evidenced above, these reductions typically have case-by-case security definitions (if any at all) that are tailored towards the larger systems that rely on them. The lack of a common language makes it difficult to relate comparable techniques hidden under incomparable abstractions. Moreover, stitching together various techniques requires remarkably delicate (and often tedious) reasoning for how the soundness of the larger protocol reduces to the soundness of each subprotocol.

Contributions Towards a unifying language, we formalize the notion of an interactive reduction over statements of knowledge, in which the verifier reduces the task of checking the original statement to the task of checking a new (simpler) statement. We refer to such a protocol as a *reduction of knowledge*. We prove that reductions of knowledge can be composed sequentially and in parallel. As such, reductions of knowledge serve as both a crisp abstraction and a theory of composition. In particular, they can be stitched together to modularly construct complex arguments of knowledge. Under this treatment, each step of an argument is instilled with a meaningful (and composable) soundness guarantee. This enables significantly simpler soundness proofs and allows each subcomponent to be reused independently in other protocols.

As a technical contribution, we employ reductions of knowledge to unify recursive algebraic arguments and simplify the corresponding analysis. In particular, we develop the *tensor reduction of knowledge* as a generalization of the central recursive step for arguments in this class. By instantiating and recursively composing the tensor reduction of knowledge over appropriate spaces, we derive both new and existing arguments of knowledge for various linear algebraic structures. Most notably, we derive a new argument of knowledge for *bilinear forms* which are expressive enough to encode weighted (and permuted) inner-products and more generally any degree-two gate over vectors of inputs.

Throughout our development, we provide various examples to demonstrate how reductions of knowledge offer a promising route towards taming the complexity of modern arguments. Most notably, we compose our linear algebraic

reductions to construct an argument of knowledge for NP with logarithmic communication with minimal additional reasoning.

1.1 Reductions of Knowledge

Recall that arguments of knowledge are defined over a relation $\mathcal{R}$ and allow a prover to show for some statement u that it knows witness w such that $(u, w) \in \mathcal{R}$. In contrast, a reduction of knowledge is defined over a pair of relations $\mathcal{R}_1$ and $\mathcal{R}_2$, and enables a verifier to reduce the task of checking knowledge of a satisfying witness for a statement in $\mathcal{R}_1$ to the task of checking knowledge of a satisfying witness for a new statement in $\mathcal{R}_2$.

Definition 1 (Reduction of Knowledge, Informal). *A reduction of knowledge from $\mathcal{R}_1$ to $\mathcal{R}_2$ is an interactive protocol between a prover and a verifier. Both parties take as input a claimed statement u_1 to be checked, and the prover additionally takes as input a corresponding witness w_1 such that $(u_1, w_1) \in \mathcal{R}_1$. After interaction, the prover and verifier together output a new statement u_2 to be checked in place of the original statement, and the prover additionally outputs a corresponding witness w_2 such that $(u_2, w_2) \in \mathcal{R}_2$. A reduction of knowledge satisfies the following properties.*

(i) Completeness: If the prover is provided a satisfying witness w_1 for the verifier's input statement u_1, then the prover outputs a satisfying witness w_2 for the verifier's output statement u_2.
(ii) Knowledge Soundness: If an arbitrary prover provides a satisfying witness w_2 for the verifier's output statement u_2, then the prover almost certainly knows a satisfying witness w_1 for the verifier's input statement u_1.

We write $\Pi : \mathcal{R}_1 \to \mathcal{R}_2$ to denote that protocol Π is a reduction of knowledge from $\mathcal{R}_1$ to $\mathcal{R}_2$.

There are two ways to conceptually reconcile reductions of knowledge with arguments of knowledge. First, arguments of knowledge can be viewed as a special case of reductions of knowledge where the second relation $\mathcal{R}_2$ is fixed to encode true or false. This interpretation helps naturally translate existing tooling used to study arguments of knowledge to study reductions of knowledge. For instance, we can expect reductions of knowledge to be compatible with idealized soundness models such as the random oracle model [5] and the algebraic group model [23], idealized communication models such as interactive oracle proofs [6] and variants [15,17,19], and heuristic transformations such as Fiat-Shamir [22].

Second, reductions of knowledge can be interpreted as arguments for *conditional* statements in which a prover shows for some u_1 that it knows w_1 such that $(u_1, w_1) \in \mathcal{R}_1$ *contingent* on the fact that for u_2 output by the verifier it knows w_2 such that $(u_2, w_2) \in \mathcal{R}_2$. Put more plainly, reductions of knowledge are arguments for statements of the form "If you believe that I know a witness for statement u_2 in $\mathcal{R}_2$, then you should believe that I know a witness for statement

u_1 in $\mathcal{R}_1$". This interpretation helps characterize statements that reductions of knowledge can handle more naturally than arguments of knowledge.

Reductions of knowledge can also be viewed as a probabilistic variant of Levin reductions [2] (i.e., Karp reductions [2] that map witnesses as well as statements) that verifiably proceed through interaction. Under this interpretation, Levin reductions can be understood as deterministic reductions of knowledge with no interaction. Just as standard reductions are used for principled algorithm design, reductions of knowledge are intended for principled argument design.

Under any interpretation, we are interested in proving that reductions of knowledge can be composed sequentially and in parallel. Such a requirement holds immediately for standard notions of reductions, but requires subtle reasoning when considering knowledge soundness: To ensure that sequential composability holds, we additionally require that reductions of knowledge are *publicly reducible*. That is, given the input statement u_1 and the interaction transcript, any party should be able to reconstruct the output statement u_2. As we detail in Section 4, this seemingly innocuous requirement becomes the linchpin in arguing sequential composability. With public reducibility, we have the following.

Theorem 1 (Sequential Composition, Informal). *Consider relations $\mathcal{R}_1$, $\mathcal{R}_2$, and $\mathcal{R}_3$. For reductions of knowledge $\Pi_1 : \mathcal{R}_1 \to \mathcal{R}_2$ and $\Pi_2 : \mathcal{R}_2 \to \mathcal{R}_3$ we have that $\Pi_2 \circ \Pi_1$ is a reduction of knowledge from $\mathcal{R}_1$ to $\mathcal{R}_3$ where $\Pi_2 \circ \Pi_1$ denotes the protocol that first runs Π_1, and then runs Π_2 on the statement and witness output by Π_1.*

By parallel composition, we do not mean running both protocols at the same time, but rather that the composed protocol takes as input instance-witness pairs in parallel and and outputs instance-witness pairs in parallel. For relations $\mathcal{R}_1$ and $\mathcal{R}_2$, let relation $\mathcal{R}_1 \times \mathcal{R}_2$ be such that $((u_1, u_2), (w_1, w_2)) \in \mathcal{R}_1 \times \mathcal{R}_2$ if and only if $(u_1, w_1) \in \mathcal{R}_1$ and $(u_2, w_2) \in \mathcal{R}_2$. Then, we have the following.

Theorem 2 (Parallel Composition, Informal). *Consider relations $\mathcal{R}_1$, $\mathcal{R}_2$, $\mathcal{R}_3$, and $\mathcal{R}_4$. For reductions of knowledge $\Pi_1 : \mathcal{R}_1 \to \mathcal{R}_2$ and $\Pi_2 : \mathcal{R}_3 \to \mathcal{R}_4$ we have that $\Pi_1 \times \Pi_2$ is a reduction of knowledge from $\mathcal{R}_1 \times \mathcal{R}_3$ to $\mathcal{R}_2 \times \mathcal{R}_4$ where $\Pi_1 \times \Pi_2$ denotes the protocol that runs Π_1 on the statement-witness pair in $\mathcal{R}_1$, runs Π_2 on the statement-witness pair in $\mathcal{R}_3$, and outputs the pair of results.*

1.2 A Theory of Composition for Arguments of Knowledge

Reductions of knowledge can be viewed as a minimal compositional framework that can feasibly capture and tame the growing complexity of modern arguments. Regardless of how reductions are stitched together, our composition results abstract out the pedantic reasoning for how exactly to use the soundness of each subcomponent to prove the soundness of the composed reduction. We develop several examples to concretely demonstrate how the reductions of knowledge framework opens up new possibilities.

In more detail, the requirement that the prover *knows* a witness is formally stated as an extractability property: Given an expected polynomial-time prover

that can produce a satisfying interaction, there must exist a corresponding expected polynomial-time extractor that can extract the alleged witness (e.g., by running and rewinding the prover internally). This definition, while undoubtedly natural, requires subtle reasoning when constructing large arguments which rely on several sub-arguments: In general, the soundness analysis must meticulously detail how to use the successful prover to construct successful provers for each sub-argument and then use the corresponding extractors to derive an extractor for the overall argument.

In the public-coin setting (where the verifier only sends random challenges), Bootle et al. [10] abstract away some low-level reasoning by proving that *tree special soundness* implies the standard notion of knowledge soundness. Tree special soundness holds when a *tree of accepting transcripts* contains sufficient information to reconstruct the witness, with each path representing a unique transcript and each branch representing diverging verifier randomness. Both Lee [32] and Attema and Cramer [3] show that tree special soundness implies modularity by observing that tree special sound protocols can be sequentially composed to produce a tree special sound protocol.

As demonstrated by these works, tree special soundness is a remarkably useful abstraction for simplifying sequentially composed, uniformly structured arguments (e.g., arguments that recursively invoke themselves). However, when dealing with larger arguments that invoke various *independent* sub-arguments, such as modern arguments for NP, tree special soundness is no longer an appropriate abstraction: having a single transcript that weaves through all such sub-arguments and globally forks with each local challenge undermines the intended semantics and unnecessarily blows up the knowledge error (i.e., the extractor's failure probability).

Reductions of knowledge are designed precisely to reason about such arguments. Unlike prior work, our parallel composition operator enables us to capture arguments with arbitrary dependence topologies. For instance, most argument systems for NP, such as Spartan [39], Poppins [29], and Marlin [19], reduce a statement in an NP-complete relation such as R1CS [25] to several simpler linear algebraic statements (such as inner-product and polynomial evaluation claims), each of which is then checked using a tailored argument [35]. As a concrete example, we show that an argument for NP can be captured modularly in our framework by utilizing both sequential and parallel composition.

Moreover, because we demonstrate that *any* two publicly verifiable reductions can be composed, this opens up the ability to modularly reason about knowledge-assumption-based succinct non-interactive arguments of knowledge (SNARKs [7,26]) and incrementally verifiable computation [41], which currently fall back on composing extractors in intricate ways [14, 29, 31]. As a concrete example, we demonstrate how to succinctly express non-interactive *ℓ-folding schemes* [31,36] (i.e, folding schemes reducing ℓ initial instances) by utilizing a tree-like dependence topology in our reductions of knowledge framework.

In the public-coin setting, we incorporate prior progress into our framework by proving that tree special soundness implies our notion of knowledge soundness. As such, public-coin reductions can be analyzed using standard techniques.

1.3 A Unified Theory for Recursive Algebraic Arguments

Reductions of knowledge provide the necessary abstraction to view various techniques under a unifying lens. As a demonstration, we consolidate recursive arguments over homomorphic structures by recasting their central recursive step as instantiations of the tensor reduction of knowledge.

In more detail, modern arguments are designed around leveraging homomorphic structure to achieve better asymptotics and concrete efficiency. An influential line of work [3, 4, 10, 13, 32] studies the consequences of arguments over structurally nested homomorphic objects such as vectors, matrices and hypercubes. A key insight is that such objects contain sufficient algebraic structure for *recursive* arguments in which larger composed statements can be reduced to smaller constituent statements of the same form. For instance, Bootle et al. [10] show that the task of checking an inner-product over committed size n vectors can be split into the task of checking two inner-products over committed size $n/2$ vectors which can then be "folded" into the task of checking a single inner-product over committed size $n/2$ vectors. Homomorphic structures that enable recursive techniques have become a staple in constructing efficient argument systems for NP [13, 29, 39, 42]. However, while arguments over recursive homomorphic structures have become an essential tool in practice, the literature detailing such techniques is becoming increasingly dissonant with sparse progress on unifying the disparate approaches.

Bünz et al. [16] initiate the study of a unified theory by observing that existing inner-product arguments [10,13] only require a commitment scheme that is homomorphic over both the commitment keys and messages. Thus, such inner-product arguments can be viewed as instantiations of a generic inner-product argument that only leverages these properties. Bootle, et al. [11] further relax this requirement by observing that split-and-fold style techniques in general [3, 13,15,16] only require a commitment scheme that can be computed by summing over a hypercube. Leveraging this insight, Bootle et al. show that such techniques can be interpreted as instantiations of the familiar sum-check protocol [33].

We considerably sharpen the sufficient conditions with the following observation: Protocols such as the sumcheck protocol and the inner-product argument only require that the underlying linear-algebraic objects (e.g., polynomials, vectors, and matrices) form a module (i.e., have a notion of addition and scalar multiplication). Abstracting away the specific details of the associated modules, all such protocols reduce a claim in a "tensored" module to claims in constituent modules. Leveraging this insight, we design an information-theoretic protocol, the *tensor reduction*, as a sweeping generalization of protocols in this class. Conceptually, the tensor reduction explains why such a broad class of protocols look different but feel the same.

Theorem 3 (Tensor Reduction, Informal). *For modules U, U_1, and U_2 such that $U \cong U_1 \otimes U_2$, there exists an interactive reduction that reduces the task of evaluating a homomorphism in U to the task of evaluating a homomorphism in U_1 and evaluating a homomorphism in U_2.*

We explain in detail how the tensor reduction generalizes familiar patterns in Section 5. Essentially, the versatility of the tensor reduction stems from its ability to work over any pair of modules and any valid notion of a tensor product between these modules. In particular, the tensor product can be defined as *any* operator that satisfies the prescribed universality property: the tensor product of any two modules U_1 and U_2 must result in a new module, denoted $U_1 \otimes U_2$, such that any bilinear mapping $\varphi : U_1 \times U_2 \to V$ induces a unique homomorphism $\widetilde{\varphi} : U_1 \otimes U_2 \to V$ such that $\widetilde{\varphi}(u_1 \otimes u_2) = \varphi(u_1, u_2)$.

For instance, for field $\mathbb{F}$, let the tensor product denote the outer product and consider an arbitrary vector in $\mathbb{F}^n$. This vector can be interpreted as a matrix in $\mathbb{F}^{(n/2)\times 2}$ or equivalently as an element of $\mathbb{F}^{n/2} \otimes \mathbb{F}^2$ which consists of sums of outer products of vectors in $\mathbb{F}^{n/2}$ and $\mathbb{F}^2$. Thus, the tensor reduction can reduce a claim over a vector in $\mathbb{F}^n$ to a claim over a vector in $\mathbb{F}^{n/2}$ and a vector in $\mathbb{F}^2$. Similarly, by taking the tensor product to be polynomial multiplication, the tensor reduction can reduce a claim over a degree (m, n) bivariate polynomial in $\mathbb{F}[X, Y] \cong \mathbb{F}[X] \otimes \mathbb{F}[Y]$ to a claim over a degree m univariate polynomial in $\mathbb{F}[X]$ and a degree n univariate polynomial in $\mathbb{F}[Y]$. By taking the tensor product to be the Kronecker product, the tensor reduction can reduce a claim over a matrix in $\mathbb{F}^{mp\times nq}$ to a claim over a matrix in $\mathbb{F}^{m\times n}$ and a matrix in $\mathbb{F}^{p\times q}$. By taking the tensor product to be a pairing operation mapping groups $\mathbb{G}_1$ and $\mathbb{G}_2$ to $\mathbb{G}_\mathsf{T}$, the tensor reduction can reduce a claim over $\mathbb{G}_\mathsf{T}$ to claims over $\mathbb{G}_1$ and $\mathbb{G}_2$.

Just as the sum-check protocol can be used to design arguments of knowledge, the tensor reduction can be used to design reductions of knowledge. By instantiating the tensor reduction over vector spaces, we derive the *tensor reduction of knowledge*, an unconditionally secure protocol that generalizes the central reductive step common to most recursive algebraic arguments.

Theorem 4 (Tensor Reduction of Knowledge, Informal). *For vector space $\hom(W, V)$, denoting homomorphisms from vector space W to vector space V, and length n, there exists a reduction of knowledge that reduces the task of checking knowledge of $w \in W^n$ such that $u(w) = v$ for $u \in \hom(W^n, V)$ and $v \in V$ to the task of checking knowledge of $w' \in W$ such that $u'(w') = v'$ for $u' \in \hom(W, V)$ and $v' \in V$.*

Leveraging our composition result, we show that tensor reductions of knowledge can be recursively composed to recover various recursive arguments. In particular, we appropriately instantiate the vector spaces to recover a family of reductions of knowledge for vector commitments [9–11], and linear forms [3, 4]. Table 1 summarizes the concrete protocols synthesized under the various instantiations of the tensor reduction of knowledge.

We also develop a new family of arguments for bilinear forms which falls out naturally from our prior generalizations. In particular, consider prime order

group $\mathbb{G}$ and corresponding scalar field $\mathbb{F}$. For public key $G \in \mathbb{G}^m$, public matrix $M \in \mathbb{F}^{m \times m}$, commitments $\overline{A}, \overline{B} \in \mathbb{G}$, and scalar $\sigma \in \mathbb{F}$, a bilinear forms argument allows a verifier to check that a prover knows $A, B \in \mathbb{F}^m$ such that $A^\top M B = \sigma$, $\langle G, A \rangle = \overline{A}$ (i.e., the inner-product of G and A is $\overline{A}$), and $\langle G, B \rangle = \overline{B}$.

In practice, the matrix M in the bilinear forms relation can encode a variety of constraints. For instance, if M is the identity matrix then the verifier can check the inner-product of A and B (and more generally the inner product of any rearrangement of A and B). If instead M assigns weights to the diagonal, then the verifier can check a weighted inner-product [16,20]. More generally, M can encode any degree-two custom-gate [24], enabling an expressive constraint system for NP as we show in Section 7.

Structure	Module	Decomposition	
		$k = 2$	$k = \sqrt[4]{n}$
Vector Commitment	PO Groups	[10]	✓
	Bil Groups	[11]	✓
Linear Forms	PO Groups	[3]	[3]
	Bil Groups	[4]	✓
Bilinear Forms	Bil Groups	✓	✓

Table 1: Protocols synthesized by instantiating the tensor reduction of knowledge. We denote previously unexplored protocols with ✓. PO Group indicates prime order groups and Bil Group indicates symmetric bilinear groups. The parameter k denotes the number of chunks tensors are decomposed into in the tensor reduction of knowledge. For vectors of size n, $k = 2$ results in protocols with $O(\log n)$ rounds of communication and $O(1)$ messages per round. Alternatively, $k = \sqrt[4]{n}$ results in protocols with $O(1)$ rounds of communication and $O(\sqrt{n})$ messages per round.

1.4 Overview of the Upcoming Sections

The remainder of this work formally treats all of the introduced concepts. In Section 2, we study two concrete examples, the vector commitment argument [10] and ℓ-folding schemes [31,36], to both preface the tensor reduction of knowledge and demonstrate how our framework simplifies the corresponding analysis. In Section 4, we formally treat reductions of knowledge and the corresponding composition results. In Section 5, we formally introduce the tensor reduction, followed by the tensor reduction of knowledge as a generalization of the core reductive step common to most recursive algebraic arguments. In Section 6, we instantiate the tensor reduction of knowledge to derive arguments for vector commitments, linear forms, and bilinear forms. In Section 7, we show that the

linear algebraic reductions derived from the tensor reduction of knowledge can be composed to derive an argument of knowledge for NP with minimal effort.

2 Technical Overview

In this section, we demonstrate how reductions of knowledge can be used to modularly reason about the vector commitment argument of Bootle et al. [10] and folding schemes [31] in the non-interactive setting [36]. The former example, being a special case of the tensor reduction of knowledge, provides an introductory overview of its mechanics. The latter example demonstrates how the reductions of knowledge framework can significantly simplify arguments with non-linear dependence topologies. We additionally demonstrate how reductions of knowledge provide a unifying language by formally defining arguments of knowledge and folding schemes as particular types of reductions.

2.1 First Example: A Vector Commitment Argument

The vector commitment argument allows a prover to show that it knows the opening to a Pedersen vector commitment [34]. In particular, consider group $\mathbb{G}$ of prime order p and corresponding scalar field $\mathbb{F} = \mathbb{Z}_p$. Consider some public key $G \in \mathbb{G}^n$ where $n = 2^i$ for some $i \in \mathbb{N}$. Suppose a prover would like to *succinctly* demonstrate to a verifier that it knows $A \in \mathbb{F}^n$ such that $\langle G, A \rangle = \overline{A}$ (i.e., the inner-product of G and A is $\overline{A}$). That is, we would like to design an argument of knowledge for the following relation.

Definition 2 (Vector Commitment Relation). *The vector commitment relation is defined as* $\mathcal{R}_{\mathsf{VC}}(n) = \{((G, \overline{A}), A) \in ((\mathbb{G}^n, \mathbb{G}), \mathbb{F}^n) \mid \langle G, A \rangle = \overline{A}\}$.

Bootle et al. [10] provide an argument system with sublinear communication cost for the above relation. At a high level, the verifier splits the task of checking knowledge of vector A into the task of checking knowledge of the first and second half of A. Instead of checking each separately, the verifier "folds" the two checks into a single check using a random linear combination. The prover computes the corresponding random linear combination of the first and second half of A to produce a folded witness vector that is half the original size. This folding procedure is recursively run until the length of the vector to be checked is 1. At this point the prover directly sends the vector to the verifier.

While the vector commitment argument can be described in a straightforward manner, proving its soundness is considerably more involved. Recursive arguments typically require recursive extractors, and the vector commitment argument is no exception. To prove knowledge soundness, the malicious prover, which produces a length one witness vector as its final message, is used to build an extractor that can produce a length two folded witness vector (which is allegedly the result of folding the original witness vector $\log n - 1$ times). Such an extractor is recursively used to produce an extractor that can produce a

length four vector, and so on. Ensuring that the extractor can successively unfold in each recursive step while also ensuring that its runtime remains expected polynomial-time requires tedious low-level reasoning. Bootle et al. [10] and following works [3,13,32] use tree special soundness precisely to avoid such reasoning.

We show that the reductions of knowledge framework is equally as effective in simplifying the analysis for the vector commitment argument. In particular, we start with the simpler goal of designing a reduction of knowledge that reduces the task of checking knowledge of a size n vector to checking knowledge of a size $n/2$ vector. We can recursively compose such a reduction to design an argument of knowledge for the vector commitment relation.

Construction 1 (Vector Commitment Reduction of Knowledge). We construct a reduction of knowledge from $\mathcal{R}_{\mathsf{VC}}(n)$ to $\mathcal{R}_{\mathsf{VC}}(n/2)$ for $n = 2^i$ where $i \geq 1$. Suppose that the prover $\mathcal{P}$ and verifier $\mathcal{V}$ take as input statement $(G, \overline{A}) \in (\mathbb{G}^n, \mathbb{G})$ and that the prover additionally takes as input alleged witness vector $A \in \mathbb{F}^n$ such that

$$((G, \overline{A}), A) \in \mathcal{R}_{\mathsf{VC}}(n).$$

The reduction proceeds as follows.

1. $\mathcal{P}$: Let G_1 and G_2 (respectively A_1 and A_2) denote the first and second half of vector G (respectively A). The prover begins by sending $\overline{A}_{ij} \leftarrow \langle G_i, A_j \rangle$ for $i, j \in \{1, 2\}$. Here, $\overline{A}_{11}$ and $\overline{A}_{22}$ represent the first and second "half" of the original commitment $\overline{A}$, and $\overline{A}_{12}$ and $\overline{A}_{21}$ represent cross terms which will assist the verifier in folding the original statement.
2. $\mathcal{V}$: The verifier first checks the consistency of $\overline{A}_{11}$ and $\overline{A}_{22}$ with $\overline{A}$ by checking that $\overline{A}_{11} + \overline{A}_{22} = \overline{A}$. The verifier must still check that the prover knows A_1 and A_2 such that $\overline{A}_{11} = \langle G_1, A_1 \rangle$ and $\overline{A}_{22} = \langle G_2, A_2 \rangle$. Instead of checking each individually, the verifier folds them into a single check by using a random linear combination. In particular, the verifier sends random $r \in \mathbb{F}$ to $\mathcal{P}$.
3. $\mathcal{P}, \mathcal{V}$: Together, the prover and verifier output the folded key and corresponding commitment $(G', \overline{A}') \in (\mathbb{G}^{n/2}, \mathbb{G})$ where $G' \leftarrow G_1 + r \cdot G_2$ and $\overline{A}' \leftarrow \overline{A}_{11} + r \cdot (\overline{A}_{12} + \overline{A}_{21}) + r^2 \cdot \overline{A}_{22}$.
4. $\mathcal{P}$: The prover outputs the folded witness $A' \in \mathbb{F}^{n/2}$ where $A' \leftarrow A_1 + r \cdot A_2$.

Now, to check the original statement, it is sufficient for the verifier to check that the prover knows A' such that

$$((G', \overline{A}'), A') \in \mathcal{R}_{\mathsf{VC}}(n/2).$$

To prove knowledge soundness, we must show that given a prover that produces a witness for the output statement with non-negligible probability, we can derive an extractor that can use this prover to derive a witness for the input statement with nearly the same probability. Because the above reduction is public-coin, it suffices to show that there exists an extractor that can derive a satisfying input witness given a tree of transcripts and corresponding satisfying outputs (Lemma 6). Intuitively, the original extractor can generate such a tree

by repeatedly rewinding the prover and collecting transcripts in which the prover outputs a satisfying witness.

Lemma 1 (Vector Commitment Reduction of Knowledge). *For $n = 2^i$ where $i \geq 1$, Construction 1 is a reduction from $\mathcal{R}_{\mathsf{VC}}(n)$ to $\mathcal{R}_{\mathsf{VC}}(n/2)$.*

Proof. We reason via tree extractability (Lemma 6). Suppose an extractor is provided with a tree of transcripts which consists of three transcripts, where the kth transcript has the same initial message $\overline{A}_{ij}$ for $i, j \in \{1, 2\}$, random challenge r_k, and satisfying output instance-witness pairs $((G'_k, \overline{A}'_k), A'_k) \in \mathcal{R}_{\mathsf{VC}}(n/2)$. The extractor first solves for a_k for $k \in \{1, 2, 3\}$ such that

$$\begin{pmatrix} 1 & 1 & 1 \\ r_1 & r_2 & r_3 \\ r_1^2 & r_2^2 & r_3^3 \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$$

using an inverse Vandermonde matrix. The extractor then computes and outputs the unfolded witness $A = \left(\sum_k a_k \cdot A'_k, \sum_k a_k r_k \cdot A'_k\right)$. Indeed, by textbook algebra, we have that $\langle G, A\rangle = \overline{A}$ (we explicitly demonstrate this in the extended version [30]). Thus, we have that $((G, \overline{A}), A) \in \mathcal{R}_{\mathsf{VC}}(n)$. □

Later, in Section 6, we show that the vector commitment reduction of knowledge is precisely the tensor reduction of knowledge from homomorphisms in $\mathbb{G}^n \cong (\mathbb{G}^{n/2})^2$ to homomorphisms in $\mathbb{G}^{n/2}$.

We are still tasked with isolating the base case of the original vector commitment argument. Below we specify an *argument of knowledge* for $\mathcal{R}_{\mathsf{VC}}(1)$. An argument of knowledge can be succinctly formalized as a reduction of knowledge that reduces to the relation $\mathcal{R}_\top$ encoding true. A verifier reducing to $\mathcal{R}_\top$ can output true if it accepts and any other string (e.g., false) otherwise.

Definition 3 (Argument of Knowledge). *Let $\mathcal{R}_\top = \{(\mathsf{true}, \bot)\}$. An argument of knowledge for relation $\mathcal{R}$ is a reduction of knowledge from $\mathcal{R}$ to $\mathcal{R}_\top$.*

Construction 2 (Base Case). We construct an argument of knowledge for $\mathcal{R}_{\mathsf{VC}}(1)$. Given statement $(G, \overline{A})$ and corresponding witness A, the prover sends A directly to the verifier. The verifier outputs true if $\langle G, A\rangle = \overline{A}$.

We can compose the above reductions to modularly recover the original argument of knowledge for the vector commitment relation. By formalizing each step as a reduction of knowledge, our composition result abstracts away the brunt of the proof effort. In particular, the following corollary holds immediately.

Corollary 1 (Vector Commitment Argument of Knowledge). *Let Π_{VC} denote a reduction of knowledge from $\mathcal{R}_{\mathsf{VC}}(n)$ to $\mathcal{R}_{\mathsf{VC}}(n/2)$ and let Π_{base} denote an argument of knowledge for $\mathcal{R}_{\mathsf{VC}}(1)$. Then*

$$\Pi_{\mathsf{base}} \circ \underbrace{\Pi_{\mathsf{VC}} \circ \ldots \circ \Pi_{\mathsf{VC}}}_{\log n \text{ times}}$$

is an argument of knowledge for $\mathcal{R}_{\mathsf{VC}}(n)$ where $n = 2^i$ for $i \in \mathbb{N}$.

2.2 Second Example: Folding Schemes

The vector commitment reduction of knowledge can be further decomposed into two reductions of knowledge: The first reduction of knowledge splits the original instance into two half-sized instances (i.e., a reduction from $\mathcal{R}_{\mathsf{VC}}(n)$ to $\mathcal{R}_{\mathsf{VC}}(n/2) \times \mathcal{R}_{\mathsf{VC}}(n/2)$). The second folds the two instances into a single instance of the same size (i.e., a reduction from $\mathcal{R}_{\mathsf{VC}}(n/2) \times \mathcal{R}_{\mathsf{VC}}(n/2)$ to $\mathcal{R}_{\mathsf{VC}}(n/2)$).

Kothapalli, Setty, and Tzialla [31] abstract the latter pattern to arbitrary relations and refer to such protocols as *folding schemes*. In particular, a folding scheme is an interactive protocol that reduces the task of checking two instances in a relation to the task of checking a single instance in the relation. Folding schemes provide a minimal abstraction for various protocols in the literature. For instance, Kothapalli et al. show that there exists a folding scheme for NP instances with some fixed size and show that such a construction implies incrementally verifiable computation [41].

More recently, Ràfols and Zacharakis [36] provide non-interactive ℓ-folding schemes (i.e., folding schemes for ℓ initial statements) for the vector commitment relation, inner-product relation, and polynomial commitment relation. Such folding schemes help amortize the verifier's work over multiple instances in larger non-interactive arguments of knowledge which typically involve checking multiple instances of the same form.

As these folding schemes rely on knowledge assumptions rather than interaction, prior techniques cannot help modularize the corresponding soundness analysis. As promised, we can still achieve modularity by decomposing them as a sequence of *non-interactive* reductions of knowledge. Formally, a non-interactive reduction of knowledge is one in which the interaction only consists of messages from the prover. Non-interactive ℓ-folding schemes can be succinctly formalized as a particular class of non-interactive reductions of knowledge. Letting $\mathcal{R}^\ell$ denote $\mathcal{R} \times \ldots \times \mathcal{R}$ for ℓ times, we define the following.

Definition 4 (ℓ-Folding Schemes). *A (non-interactive) ℓ-folding scheme for relation $\mathcal{R}$ is a (non-interactive) reduction of knowledge from $\mathcal{R}^\ell$ to $\mathcal{R}$.*

Ràfols and Zacharakis achieve ℓ-folding schemes for various relations by recursively composing 2-folding schemes in a tree-like fashion. In particular, ℓ instances are treated as leaves in a tree. A 2-folding scheme is then used to fold each pair of adjacent instances to produce a total of $\ell/2$ instances. These $\ell/2$ instances are once more folded in a pairwise fashion to produce $\ell/4$ instances and so on until a single instance remains.

Once again, as demonstrated by Ràfols and Zacharakis, while the tree-folding protocol can be stated in a straightforward manner, the corresponding knowledge soundness analysis requires careful attention to detail. In particular, the corresponding proof involves demonstrating that the malicious prover induces a corresponding expected polynomial-time extractor that unfolds once. Such an extractor is then shown to induce a pair of expected polynomial-time malicious provers for the previous layer of the tree, and so on. Alternatively, by working in

the reductions of knowledge framework, nearly all of this reasoning is abstracted away. Indeed, we condense the original three-page proof into several lines.

Lemma 2 (ℓ-Folding Scheme). *Consider a (non-interactive) 2-folding scheme Π_{TF} for relation $\mathcal{R}$ and $\ell = 2^i$ for $i \in \mathbb{N}$ where $i \geq 1$. Then, Π_ℓ, inductively defined as follows, is a (non-interactive) ℓ-folding scheme for $\mathcal{R}$.*

$$\Pi_\ell = \Pi_{\mathsf{TF}} \circ (\Pi_{\ell/2} \times \Pi_{\ell/2})$$
$$\Pi_2 = \Pi_{\mathsf{TF}}$$

Proof. We reason inductively over i. In the base case, suppose $i = 1$. Then, by construction, Π_2 is a 2-folding scheme. Suppose instead $i \geq 2$. Suppose that for $\ell = 2^i$ we have that $\Pi_{\ell/2}$ is a $(\ell/2)$-folding scheme. Then, $\Pi_{\ell/2} \times \Pi_{\ell/2}$ is a reduction of knowledge from $\mathcal{R}^{\ell/2} \times \mathcal{R}^{\ell/2} = \mathcal{R}^\ell$ to $\mathcal{R}^2$. Thus, $\Pi_{\mathsf{TF}} \circ (\Pi_{\ell/2} \times \Pi_{\ell/2})$ is a reduction of knowledge from $\mathcal{R}^\ell$ to $\mathcal{R}$. □

3 Preliminaries

3.1 Module Theory

In this section, we introduce our notation, intuit the direct sum and tensor product, and recall several useful properties. In the extended version [30], we present formal definitions for rings, modules, direct sums, and tensor products.

Notation (Module Theory). We assume finite, unital, commutative rings and modules with a finite basis throughout. We use $\cong$ to denote that two modules are isomorphic. For ring R and R-modules W and V, let $\hom(W, V)$ denote the R-module of homomorphisms from W to V. For $n \in \mathbb{N}$, we let W^n denote $W \otimes \mathsf{R}^n$ (equivalently $W \oplus \ldots \oplus W$ for n times). We use $\{\delta_i\}$ to denote an orthonormal basis. We refer to elements of modules as tensors. As we use tensors to represent both homomorphisms and objects, for tensors g and a, we use $g(a)$ to denote evaluating the homomorphism tensor g on the object tensor a. For $n \in \mathbb{N}$, let $[n]$ denote $\{1, 2, \ldots, n\}$ and let $[i, n]$ for $i \leq n$ denote $\{i, i+1, \ldots, n\}$. When summing over a variable, we will omit the bounds when clear from context. We write $\langle a, b \rangle$ to denote the inner-product of a and b.

Modules Intuitively, modules are vector spaces over rings. That is, they support a notion of addition, can be scaled by ring elements, and have an identity element. We say a module is an R-module if it is scaled by ring R. Vectors, polynomials, matrices, tensors and scalars all form modules.

The Direct Sum Intuitively, a *direct sum* of two R-modules U and V, forms a new R-module denoted $U \oplus V$, which is essentially a Cartesian product of the original modules. Elements of $U \oplus V$ consist of pairs of elements in U and V which are denoted as $u \oplus v$ for $u \in U$ and $v \in V$. For example, for field $\mathbb{F}$, if $U \cong \mathbb{F}^n$ and $V \cong \mathbb{F}^m$ we have that $U \oplus V \cong \mathbb{F}^{n+m}$. We have that $U \oplus V$ forms a module, because we can naturally compute $u_1 \oplus v_1 + u_2 \oplus v_2 = (u_1 + u_2) \oplus (v_1 + v_2)$ and $r \cdot (u \oplus v) = (r \cdot u) \oplus (r \cdot v)$ for $r \in \mathsf{R}$.

The Tensor Product Intuitively, the tensor product, denoted $\otimes$, can be considered a generalized outer-product that distributes with respect to the direct sum. The tensor product of two modules U and V, forms a new module denoted $U \otimes V$. Elements of $U \otimes V$ include *simple tensors* which are outer products of elements in U and V and are denoted as $u \otimes v$ for $u \in U$ and $v \in V$. The module $U \otimes V$ also contains arbitrary sums of these simple tensors, which are denoted as $\sum_{i \in [\ell]} u_i \otimes v_i$ for $u_1, \ldots, u_\ell \in U$ and $v_1, \ldots, v_\ell \in V$. If $U \cong \mathbb{F}^n$ and $V \cong \mathbb{F}^m$ we have that $U \otimes V \cong \mathbb{F}^{n \times m}$ (i.e., $n \times m$ matrices over $\mathbb{F}$). Simple tensors in $\mathbb{F}^n \otimes \mathbb{F}^m$ consist of outer products of vectors in $\mathbb{F}^n$ and $\mathbb{F}^m$; however, the entire space is generated by sums over such outer products. We have that $U \otimes V$ forms a module because we can naturally add two sums and compute $r \cdot \sum_i u_i \otimes v_i = \sum_i (r \cdot u_i) \otimes v_i = \sum_i u_i \otimes (r \cdot v_i)$.

Abstracting the Direct Sum and Tensor Product Formally, the particular definitions of the direct sum and tensor product depend on the particular modules they are working over. For instance, the tensor product could mean the outer product when working over vectors and the Kronecker product when working over matrices. Even for a fixed pair of modules, there could exist multiple valid definitions. For instance, for vectors $v_1, v_2 \in \mathbb{F}^n$, we can define $v_1 \oplus v_2$ to be a vector in $\mathbb{F}^{2n}$ or a matrix in $\mathbb{F}^{2 \times n}$. To account for these considerations, we treat the direct sum and tensor product as *abstract* operations that can be implemented by any concrete operations that satisfy certain axioms (which we detail in the extended version [30]). In practice, much like how abstract groups and rings must be instantiated with concrete objects such as elliptic curves and polynomials, the direct sum and tensor product must be instantiated with concrete operations that respect the prescribed properties.

For the majority of our development, we are interested in taking the direct sum and tensor product of homomorphisms (represented as tensors). In this situation, we do not need to invoke the abstract definitions of these operations, but rather the identities that follow from their axioms.

Lemma 3 (Direct Sum of Homomorphisms). *Consider homomorphisms $r \in \hom(U_1, V)$ and $s \in \hom(U_2, V)$ over R-modules (where R is a commutative ring). Then $r \oplus s \in \hom(U_1 \oplus U_2, V)$ is a homomorphism where $(r \oplus s)(u_1 \oplus u_2) = r(u_1) + s(u_2)$. Symmetrically, homomorphisms $r \in \hom(U, V_1)$ and $s \in \hom(U, V_2)$ over R-modules induce a homomorphism $r \oplus s \in \hom(U, V_1 \oplus V_2)$ where $(r \oplus s)(u) = r(u) \oplus s(u)$.*

Example 1 (Direct Sum of Homomorphisms). Consider group $\mathbb{G}$ of prime order p and corresponding scalar field $\mathbb{F} \cong \mathbb{Z}_p$. We can interpret $\mathbb{G}^n$ as the module of homomorphisms from $\mathbb{F}^n$ to $\mathbb{G}$. In particular, for $g \in \mathbb{G}^n$ we can define $g(a) = \langle g, a \rangle$ for $a \in \mathbb{F}^n$. Then, for $g \in \mathbb{G}^n$ and $h \in \mathbb{G}^m$ we have that $g \oplus h \in \mathbb{G}^n \oplus \mathbb{G}^m \cong \mathbb{G}^{n+m}$ can be interpreted as a map from $\mathbb{F}^{n+m} \cong \mathbb{F}^n \oplus \mathbb{F}^m$ to $\mathbb{G}$. By definition, for $u \in \mathbb{F}^n$ and $v \in \mathbb{F}^m$, we have $(g \oplus h)(u \oplus v) = \langle g \oplus h, u \oplus v \rangle = \langle g, u \rangle + \langle h, v \rangle = g(u) + h(v)$.

Lemma 4 (Tensor Product of Homomorphisms). *Homomorphisms $r \in \hom(U, X)$ and $s \in \hom(V, Y)$ over R-modules (where R is a commutative ring) induce a homomorphism $r \otimes s \in \hom(U \otimes V, X \otimes Y)$, such that $(r \otimes s)(u \otimes v) = r(u) \otimes s(v)$. By linearity, we have that*

$$\Big(\sum_{i\in[I]} r_i \otimes s_i\Big)\Big(\sum_{j\in[J]} u_i \otimes v_i\Big) = \sum_{i\in[I],j\in[J]} r_i(u_j) \otimes s_i(v_j).$$

Example 2 (Tensor Product of Homomorphisms). Let $\otimes$ denote the outer product. For prime p and field $\mathbb{F} \cong \mathbb{Z}_p$ we can interpret $\mathbb{F}^n$ as the module of homomorphisms from $\mathbb{F}^n$ to $\mathbb{F}$. In particular, for $f \in \mathbb{F}^n$ we can define $f(a) = \langle f, a\rangle$ for $a \in \mathbb{F}^n$. Then, $f \in \mathbb{F}^n$ and $g \in \mathbb{F}^m$ induce a new map $f \otimes g \in \mathbb{F}^n \otimes \mathbb{F}^m \cong \mathbb{F}^{nm}$ from $\mathbb{F}^{nm}$ to $\mathbb{F} \otimes \mathbb{F} \cong \mathbb{F}$. By definition, for $u \in \mathbb{F}^n$ and $v \in \mathbb{F}^m$, we have $(f \otimes g)(u \otimes v) = (f \cdot g_1 \oplus \ldots \oplus f \cdot g_m)(u \cdot v_1 \oplus \ldots \oplus u \cdot v_m) = \sum_{j\in[m]} f(u) \cdot g_i(v_i) = f(u) \otimes g(v)$.

Lemma 5 (Useful Identities). *For commutative ring R and R-modules U, V, and W, we have that $(U \otimes V) \otimes W \cong U \otimes (V \otimes W)$, $U \otimes V \cong V \otimes U$, $U \otimes (V \oplus W) \cong (U \otimes V) \oplus (U \otimes W)$, and $\mathsf{R} \otimes U \cong U \otimes \mathsf{R} \cong U$.*

3.2 Cryptographic Preliminaries

Notation (Cryptography). We use λ globally to denote the security parameter, and negl to denote negligible functions. For events A and B, we let $\Pr[A] \approx \Pr[B]$ denote that $|\Pr[A] - \Pr[B]| = \mathsf{negl}(\lambda)$. We let PPT denote probabilistic polynomial-time. We write _ to denote unused terms.

For soundness to hold when randomly sampling over rings, the set of admissible values must be constrained. We define a valid sampling set over rings.

Definition 5 (Sampling Set [11]). *For ring R and R-module M, subset $Q \subseteq \mathsf{R}$ is a sampling set for M if for every $q_1, q_2 \in Q$, the map $\varphi_{q_1,q_2}(m) = (q_1 - q_2) \cdot m$ for $m \in M$ is injective.*

For certain relations, to be able to prove knowledge soundness, we will need to rely on computational hardness assumptions. We adapt the bilinear relation assumption [11], which can be viewed as a generalization of the discrete logarithm assumption, and the double pairing assumption [1].

Definition 6 (Bilinear Relation Assumption). *For ring R, length parameter n, and security parameter λ, consider R-modules U and V such that $|U| = O(2^\lambda)$ and $|V| = O(2^\lambda)$. The bilinear relation assumption holds for (U, V) (w.r.t. tensor product $\otimes$) if given random $u_1, \ldots, u_n \in U$, there exists no polynomial-time algorithm to find non-trivial $v_1, \ldots, v_n \in V$ such that $\sum_{i\in[n]} u_i \otimes v_i = 0$.*

Symmetrically, we can consider composite spaces such that given elements from both of the constituent spaces, it is *easy* to check that they satisfy the above relation. This ensures that that the verifier is able to perform its requisite checks efficiently. Throughout our development, we assume the coset equality assumption holds as necessary.

Definition 7 (Coset Equality Assumption). *For ring* R *and length parameter* n*, consider* R*-modules* U *and* V*. The coset equality assumption holds for* (U, V) *(w.r.t. tensor product* $\otimes$*) if for any* $u_1, \ldots, u_n \in U$ *and* $v_1, \ldots, v_n \in V$*, there exists a polynomial-time algorithm to check* $\sum_{i \in [n]} u_i \otimes v_i = 0$*.*

Example 3 (Bilinear Relation Assumption). Suppose U is a group of prime order p and V is the corresponding scalar field $\mathbb{Z}_p$. Let the tensor product between these two modules be defined as scalar multiplication. In this setting, the bilinear relation assumption is equivalent to the discrete logarithm assumption. Alternatively, suppose U and V are prime order groups such that there exists a corresponding pairing operation e from $U \times V$ into some target group. Let the tensor product be defined as this pairing operation. In this setting, the bilinear relation assumption is equivalent to the double pairing assumption.

4 Reductions of Knowledge

Recall that in contrast to arguments of knowledge, reductions of knowledge are defined over a pair of relations $\mathcal{R}_1$ and $\mathcal{R}_2$. A prover can use a reduction of knowledge to show for some u_1 that it knows w_1 such that $(u_1, w_1) \in \mathcal{R}_1$ contingent on the fact that it knows w_2 for some statement u_2 (derived from its interaction with the verifier) such that $(u_2, w_2) \in \mathcal{R}_2$. We start by intuiting the desired notion of knowledge soundness needed to capture such an interaction, before presenting a formal definition (Definition 8). We show that any two reductions of knowledge that respect this definition can be composed sequentially and in parallel (Theorems 5 and 6). We then observe that a more restricted — but simpler — notion of soundness, known as tree extractability, implies our definition of knowledge soundness (Lemma 6). In the following sections, we leverage this observation to prove that our reductions of knowledge for linear-algebraic statements are secure.

4.1 Defining Reductions of Knowledge

Intuitively, we would like that if a prover is able to convince a verifier on input u_1 to output some derived statement u_2 such that it knows a corresponding satisfying witness w_2, then it must have known a corresponding satisfying witness w_1 for u_1. We can capture this notion formally by stating that if a malicious prover can output a satisfying witness w_2 for the verifier's output statement u_2, then there must exist a corresponding extractor that can output a satisfying witness w_1 for the verifier's input statement u_1.

While this presents a stand-alone notion of knowledge soundness, we require a more nuanced definition to capture technical difficulties that arise when reasoning about sequential composability. In particular, existing definitions implicitly assume that the environment is provided access to the inputs and outputs of the prover and the verifier, and that some of this material (such as an adversarially chosen statement) is forwarded to the extractor. Unfortunately, when composing

such arguments, we end up in situations where intermediate inputs expected by the extractor are never exposed to the environment.

Concretely, consider a reduction of knowledge Π_1 with prover $\mathcal{P}_1$, verifier $\mathcal{V}_1$, and extractor $\mathcal{E}_1$, and a second reduction of knowledge Π_2 with corresponding $\mathcal{P}_2$, $\mathcal{V}_2$, and $\mathcal{E}_2$. Ideally, we would want to use $\mathcal{E}_1$ and $\mathcal{E}_2$ in a black-box manner to construct an extractor $\mathcal{E}$ for $\Pi_2 \circ \Pi_1$. A typical knowledge soundness definition would dictate that the statement provided to the verifier is forwarded to the extractor as well. Unfortunately, in the composed setting, the statement u_2 output by $\mathcal{V}_1$ as input to $\mathcal{V}_2$ is never exposed to the environment, and thus it is unclear how $\mathcal{E}$ can simulate the intermediate statement u_2 expected by $\mathcal{E}_2$.

To alleviate this issue, we stipulate an additional requirement that the verifier's output statement can be deterministically recovered from the mutual view of both the prover and verifier. Specifically, the mutual view consists of the public parameters, initial input statement, and interaction transcript. We refer to this property as *public reducibility*, which can be viewed as analogous to the public verifiability property common to most modern arguments. With public reducibility, we are afforded sequential composability.

We formally define reductions of knowledge as interactive protocols in the global common reference string model.

Definition 8 (Reduction of Knowledge). *Consider ternary relations $\mathcal{R}_1$ and $\mathcal{R}_2$ consisting of public parameters, statement, witness tuples. A reduction of knowledge from $\mathcal{R}_1$ to $\mathcal{R}_2$ is defined by PPT algorithms $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ denoting the generator, the prover, and the verifier respectively with the following interface.*

- $\mathcal{G}(\lambda) \to \mathsf{pp}$*: Takes security parameter λ. Outputs public parameters pp.*
- $\mathcal{P}(\mathsf{pp}, u_1, w_1) \to (u_2, w_2)$*: Takes as input public parameters pp, and statement-witness pair (u_1, w_1). Interactively reduces the statement $(\mathsf{pp}, u_1, w_1) \in \mathcal{R}_1$ to a new statement $(\mathsf{pp}, u_2, w_2) \in \mathcal{R}_2$.*
- $\mathcal{V}(\mathsf{pp}, u_1) \to u_2$*: Takes as input public parameters pp, and statement u_1 associated with $\mathcal{R}_1$. Interactively reduces the task of checking u_1 to the task of checking a new statement u_2 associated with $\mathcal{R}_2$.*

Let $\langle \mathcal{P}, \mathcal{V} \rangle$ denote the interaction between $\mathcal{P}$ and $\mathcal{V}$. We treat $\langle \mathcal{P}, \mathcal{V} \rangle$ as a function that takes as input (pp, u_1, w_1) and runs the interaction on prover input (pp, u_1, w_1) and verifier input (pp, u_1). At the end of the interaction, $\langle \mathcal{P}, \mathcal{V} \rangle$ outputs the verifier's statement u_2 and the prover's witness w_2. A reduction of knowledge $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ satisfies the following conditions.

(i) Completeness: For any PPT adversary $\mathcal{A}$, given $\mathsf{pp} \leftarrow \mathcal{G}(\lambda)$ and $(u_1, w_1) \leftarrow \mathcal{A}(\mathsf{pp})$ such that $(\mathsf{pp}, u_1, w_1) \in \mathcal{R}_1$, we have that the prover's output statement is equal to the verifier's output statement and that

$$(\mathsf{pp}, \langle \mathcal{P}, \mathcal{V} \rangle (\mathsf{pp}, u_1, w_1)) \in \mathcal{R}_2.$$

(ii) Knowledge Soundness: For any expected polynomial-time adversaries $\mathcal{A}$ and $\mathcal{P}^*$, there exists an expected polynomial-time extractor $\mathcal{E}$ such that given $\mathsf{pp} \leftarrow \mathcal{G}(\lambda)$ and $(u_1, \mathsf{st}) \leftarrow \mathcal{A}(\mathsf{pp})$, we have that

$$\Pr[(\mathsf{pp}, u_1, \mathcal{E}(\mathsf{pp}, u_1, \mathsf{st})) \in \mathcal{R}_1] \approx \Pr[(\mathsf{pp}, \langle \mathcal{P}^*, \mathcal{V} \rangle (\mathsf{pp}, u_1, \mathsf{st})) \in \mathcal{R}_2].$$

(iii) Public Reducibility: There exists a deterministic polynomial-time function φ such that for any PPT *adversary $\mathcal{A}$ and expected polynomial-time adversary $\mathcal{P}^*$, given $\mathsf{pp} \leftarrow \mathcal{G}(\lambda)$, $(u_1, \mathsf{st}) \leftarrow \mathcal{A}(\mathsf{pp})$, and $(u_2, w_2) \leftarrow \langle \mathcal{P}^*, \mathcal{V} \rangle (\mathsf{pp}, u_1, \mathsf{st})$ with interaction transcript tr, we have that $\varphi(\mathsf{pp}, u_1, \mathsf{tr}) = u_2$.*

We write $\Pi : \mathcal{R}_1 \to \mathcal{R}_2$ to denote that protocol Π is a reduction of knowledge from relation $\mathcal{R}_1$ to relation $\mathcal{R}_2$.

Definition 9 (Public-Coin). *A reduction of knowledge is public-coin if the verifier only sends uniformly random challenges to the prover.*

4.2 Composing Reductions of Knowledge

We now prove sequential and parallel composition theorems for reductions of knowledge. This allows us to construct complex arguments by stitching together simpler reductions sequentially and in parallel. In the case of sequential composition, much like recursive composition techniques [8, 14, 31, 41], each composition step induces a polynomial blowup in the corresponding extractor. Thus, sequential composition cannot be used more than a constant number of times without additional computational assumptions.[1] Our parallel composition operator is not parallel in the sense that both protocols are being run at the same time, but rather parallel in the sense that the composed protocol takes incoming instance-witness pairs in parallel and produces outgoing instance-witness pairs in parallel.

Theorem 5 (Sequential Composition). *Consider ternary relations $\mathcal{R}_1$, $\mathcal{R}_2$, and $\mathcal{R}_3$. For reductions of knowledge $\Pi_1 = (\mathcal{G}, \mathcal{P}_1, \mathcal{V}_1) : \mathcal{R}_1 \to \mathcal{R}_2$ and $\Pi_2 = (\mathcal{G}, \mathcal{P}_2, \mathcal{V}_2) : \mathcal{R}_2 \to \mathcal{R}_3$, we have that $\Pi_2 \circ \Pi_1 = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ is a reduction of knowledge from $\mathcal{R}_1$ to $\mathcal{R}_3$ where*

$$\mathcal{P}(\mathsf{pp}, u_1, w_1) = \mathcal{P}_2(\mathsf{pp}, \mathcal{P}_1(\mathsf{pp}, u_1, w_1))$$
$$\mathcal{V}(\mathsf{pp}, u_1) = \mathcal{V}_2(\mathsf{pp}, \mathcal{V}_1(\mathsf{pp}, u_1, w_1)).$$

Proof Intuition. Completeness and public reducibility follow by observation. As for knowledge soundness, assume there exists an adversarial prover $\mathcal{P}^*$ for Π that succeeds in producing an accepting witness w_3 with non-negligible probability. Using the second half of $\mathcal{P}^*$ (i.e., the part that interacts with $\mathcal{V}_2$), we can construct an adversary $\mathcal{P}_2^{**}$ for Π_2 that succeeds in producing an accepting witness w_3 with the same probability. By the knowledge soundness of Π_2, this implies an extractor $\mathcal{E}_2$ that succeeds in producing an intermediate witness w_2 with nearly the same probability. We can then leverage $\mathcal{E}_2$ to construct an adversary $\mathcal{P}_1^{**}$ for Π_1 that succeeds in producing an accepting witness w_2 with nearly the same probability. In particular, $\mathcal{P}_1^{**}$ first runs the first half of $\mathcal{P}^*$ and then runs extractor $\mathcal{E}_2$ on the intermediate statement u_2 (derived by the public reducibility of Π_1) and the intermediate state of $\mathcal{P}^*$ to produce the output

[1] We recommend Bitansky et al. [8, Remark 6.3] for details on such assumptions.

w_2. By the knowledge soundness of Π_1, this implies the desired extractor $\mathcal{E}_1$ that succeeds in producing the witness w_1 with nearly the same probability. We present a formal proof in the extended version [30]. □

Definition 10 (Relation Pair). *Consider ternary relations $\mathcal{R}_1$ and $\mathcal{R}_2$ over public parameters, statement, witness tuples. We define the relation $\mathcal{R}_1 \times \mathcal{R}_2 = \{(\mathsf{pp}, (u_1, u_2), (w_1, w_2)) \mid (\mathsf{pp}, u_1, w_1) \in \mathcal{R}_1, (\mathsf{pp}, u_2, w_2) \in \mathcal{R}_2\}$. We let $\mathcal{R}^\ell$ denote $\mathcal{R} \times \ldots \times \mathcal{R}$ for ℓ times.*

Theorem 6 (Parallel Composition). *Consider ternary relations $\mathcal{R}_1$, $\mathcal{R}_2$, $\mathcal{R}_3$, and $\mathcal{R}_4$. For reductions of knowledge $\Pi_1 = (\mathcal{G}, \mathcal{P}_1, \mathcal{V}_1) : \mathcal{R}_1 \to \mathcal{R}_2$ and $\Pi_2 = (\mathcal{G}, \mathcal{P}_2, \mathcal{V}_2) : \mathcal{R}_3 \to \mathcal{R}_4$, we have that $\Pi_1 \times \Pi_2 = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ is a reduction of knowledge from $\mathcal{R}_1 \times \mathcal{R}_3$ to $\mathcal{R}_2 \times \mathcal{R}_4$ where*

$$\mathcal{P}(\mathsf{pp}, (u_1, u_3), (w_1, w_3)) = (\mathcal{P}_1(\mathsf{pp}, u_1, w_1), \mathcal{P}_2(\mathsf{pp}, u_3, w_3))$$
$$\mathcal{V}(\mathsf{pp}, (u_1, u_3)) = (\mathcal{V}_1(\mathsf{pp}, u_1), \mathcal{V}_2(\mathsf{pp}, u_3)).$$

Proof Intuition. For $i \in \{1, 2\}$, we leverage a malicious prover $\mathcal{P}^*$ for Π to construct a prover $\mathcal{P}_i^*$ for protocol Π_i that succeeds in producing a satisfying output witness with the same probability. By the knowledge soundness of Π_i, this implies a corresponding extractor $\mathcal{E}_i$ that succeeds in producing a satisfying input witness with nearly the same probability. These extractors imply the desired extractor $\mathcal{E}$. We present a formal proof in the extended version [30]. □

4.3 Knowledge Soundness from Tree Extraction

When proving constructions secure, reasoning about knowledge soundness directly is typically cumbersome. To alleviate this issue, prior work [10] observes that most protocols are algebraic: The corresponding extractor typically runs the malicious prover multiple times with refreshed verifier randomness to retrieve accepting transcripts, which can be interpolated to retrieve the witness. Leveraging this insight, Bootle et al. [10] provide a general extraction lemma, which states that to prove knowledge soundness for algebraic protocols, it is sufficient to show that there exists an extractor that can produce a satisfying witness when provided a *tree of accepting transcripts* with refreshed verifier randomness at each layer. This proof technique has been adapted to various settings [11, 13, 14, 31], and we similarly provide the corresponding lemma for reductions of knowledge.

Definition 11 (Tree of Transcripts). *Consider an m-round public-coin interactive protocol $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ that satisfies the interface described in Definition 8. A $(n_1, \ldots, n_m)$-tree of accepting transcripts for statement u_1 is a tree of depth m where each vertex at layer i has n_i outgoing edges such that (1) each vertex in layer $i \in [m]$ is labeled with a prover message for round i; (2) each outgoing edge from layer $i \in [m]$ is labeled with a different choice of verifier randomness for round i; (3) each leaf is labeled with an accepting statement-witness pair output by the prover and verifier corresponding to the interaction along the path.*

Lemma 6 (Tree Extraction [11]). *Consider an m-round public-coin interactive protocol $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ that satisfies the interface described in Definition 8 and satisfies completeness. Then $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ is a reduction of knowledge if there exists a PPT extractor χ that, for all instances u_1, outputs a satisfying witness w_1 with probability $1 - \mathsf{negl}(\lambda)$, given an $(n_1, \ldots, n_m)$-tree of accepting transcripts for u_1 where the verifier's randomness is sampled from space Q such that $|Q| = O(2^\lambda)$, and $\prod_i n_i = \mathsf{poly}(\lambda)$.*

Proof Intuition. Our proof closely follows that of Bootle et al. [10]. At a high level, we construct an expected polynomial-time extractor $\mathcal{E}$ that repeatedly runs the malicious prover $\mathcal{P}^*$ and collects corresponding accepting transcripts and associated output statement-witness pairs. The extractor then passes these collected transcripts to χ which retrieves the desired witness by assumption. We present a formal proof in the extended version [30]. □

5 The Tensor Reduction of Knowledge

We start by defining a general tensor-based language to capture a large class of linear algebraic statements. We then design a general reduction, the tensor reduction, for such statements, by extending the sum-check protocol [33]. Next, we leverage the tensor reduction to construct the tensor reduction of knowledge, which, for any length vector space of homomorphisms $\hom(W, V)$ and length n, reduces the task of checking knowledge of a preimage of a vector in $\hom(W^n, V)$ to checking knowledge of a preimage in $\hom(W, V)$.

5.1 Tensor Evaluation Statements

We observe that arguments of knowledge built around statements over linear algebraic objects — such as matrices, vectors, polynomials, and homomorphisms — typically share hints of symmetry. Our goal is to generalize such statements, and more interestingly generalize interactive reductions for such statements.

Regardless of the underlying linear-algebraic objects, arguments over them tend to only rely on the fact they support some notion of addition and that they can be scaled by elements in a field (and more generally rings). This seems to suggest that designing a reduction over the most general objects that support these operations, namely tensors, would give a single universal protocol for such objects. From an algebraic standpoint, tensors unify objects such as scalars, vectors, matrices, and polynomials. More generally, tensors provide a unifying algebraic object for describing both functions (when viewed as homomorphisms) and objects (when viewed as elements of a module).

Take for instance the vector commitment relation: Given a prime order group $\mathbb{G}$ and an underlying scalar field $\mathbb{F}^n$, a prover claims that for public commitment key $G \in \mathbb{G}^n$ and commitment $\overline{A}$, it knows a vector $A \in \mathbb{F}^n$ such that $\langle G, A \rangle = \overline{A}$. As the spaces $\mathbb{G}^n$, $\mathbb{F}^n$ and $\mathbb{G}$ are all modules, we can build a corresponding "tensor evaluation" statement

$$G(A) = \overline{A}$$

where G is a tensor in $\mathbb{G}^n$ that maps tensors in $\mathbb{F}^n$ to tensors in $\mathbb{G}$.

Alternatively, suppose in addition to claiming that it knows a vector A underlying a commitment $\overline{A}$ with respect to commitment key G, the prover additionally claims that taking the inner-product of A against some public vector $B \in \mathbb{F}^n$ results in a scalar $\sigma \in \mathbb{F}$. Following our prior reasoning, this can be represented as two tensor evaluation statements: A claim that $G(A) = \overline{A}$ and a claim that $B(A) = \sigma$. But, under the rules of the direct sum (which can be interpreted as a Cartesian product), this is equivalent to applying the tensor $G \oplus B \in \mathbb{G}^n \oplus \mathbb{F}^n$ to A and checking that this results in $\overline{A} \oplus \sigma \in \mathbb{G} \oplus \mathbb{F}$. Namely, we have that the composite statement can be encoded as the following tensor evaluation statement:

$$(G \oplus B)(A) = \overline{A} \oplus \sigma.$$

The flexibility of tensor evaluation statements becomes more salient with the sum-check protocol [33]. In the sum-check protocol, the prover claims for multivariate polynomial $P : \mathbb{F}^n \to \mathbb{F}$ with degree d in each variable that

$$\sum_{x_1,\ldots,x_n \in \{0,1\}} P(x_1, \ldots, x_n) = \sigma \tag{1}$$

for some claimed sum $\sigma \in \mathbb{F}$. For $i \in [n]$, consider the tensor $\bigoplus_{j \in [0,d]} x_i^j$ which is just shorthand for the vector $(x_i^0, x_i^1, \ldots, x_i^d)$. Now, consider $\bigotimes_{i \in [n]} \bigoplus_{j \in [0,d]} x_i^j$, which is an n-dimensional matrix populated with all possible products of powers of $x_1, \ldots, x_n$. We can now define a tensor $\boldsymbol{X} = \sum_{x_1,\ldots,x_n \in \{0,1\}} \bigotimes_{i \in [n]} \bigoplus_{j \in [0,d]} x_i^j \in (\mathbb{F}^{d+1})^n$ which encodes all desired evaluation points. Additionally, let $\boldsymbol{P} \in (\mathbb{F}^{d+1})^n$ denote an n-dimensional tensor constituting of the coefficients of P. Specifically, let $\boldsymbol{P}$ contain at index $(j_1, \ldots, j_n)$ the coefficient of P associated with term $x_1^{j_1} x_2^{j_2} \ldots x_n^{j_n}$. Now, we have that checking the original sum-check statement is equivalent to checking the tensor evaluation statement

$$\boldsymbol{P}(\boldsymbol{X}) = \sigma.$$

The three examples above suggest that seemingly disparate linear-algebraic claims can be uniformly viewed as tensor evaluation claims. In light of this, we are interested in designing a reduction for statements of the form $u(w) = v$ for tensors u, w, and v.

5.2 The Tensor Reduction

To design a general reduction for tensor statements of the form $u(w) = v$, we start by generalizing the sum-check protocol for tensor evaluation statements. Recall that the sum-check protocol reduces the task of checking the claim in Equation (1) to the task of checking a sum-check claim over a polynomial with one less variable. In particular, the prover begins by sending

$$p(X) = \sum_{x_1,\ldots,x_{n-1} \in \{0,1\}} P(x_1, \ldots, x_{n-1}, X)$$

The verifier then checks that $p(0)+p(1) = \sigma$. The verifier must now check that p is consistent with P. To do so, the verifier samples a random $r \leftarrow \mathbb{F}$, and reduces to checking

$$\sum_{x_1,\ldots,x_{n-1}} P(x_1, \ldots, x_{n-1}, r) = p(r).$$

In essence, the sum-check protocol leverages the nested structure of polynomials to reduce the task of checking n-variate polynomials to checking $(n-1)$-variate polynomials. This intuition can be more lucidly expressed with the corresponding tensor evaluation statements: the sum-check protocol reduces the task of checking the evaluation of $\boldsymbol{P} \in (\mathbb{F}^{d+1})^n \cong (\mathbb{F}^{d+1})^{n-1} \otimes \mathbb{F}^{d+1}$ (representing P) to the task of checking the evaluation of $\boldsymbol{P}_r \in (\mathbb{F}^{d+1})^{n-1}$ (representing P evaluated on r) and $\boldsymbol{p} \in \mathbb{F}^{d+1}$ (representing p). That is, the sum-check protocol factors the original statement with respect to the tensor product.

The tensor reduction, which we detail below, follows from generalizing the involved spaces to handle arbitrary tensor evaluation statements: for any modules U, U_1, and U_2 such that $U \cong U_1 \otimes U_2$, we derive a mechanism to reduce an evaluation claim in U to an evaluation claim in U_1 and an evaluation claim in U_2. In the extended version [30], we show that we can recover the sum-check protocol when instantiating the tensor reduction over multivariate polynomials.

Construction 3 (Tensor Reduction). Suppose for tensors $u \in \hom(W_1, V_1) \otimes \hom(W_2, V_2)$ of rank I, $w \in W_1 \otimes W_2$ of rank J, and $v \in V_1 \otimes V_2$ over ring R, a verifier would like to check

$$u(w) = v \tag{2}$$

where $u = \sum_{i\in[I]} u_{1,i} \otimes u_{2,i}$, and $w = \sum_{j\in[J]} w_{1,j} \otimes w_{2,j}$. By definition, the verifier can check (2) by checking $\sum_{i,j} u_{1,i}(w_{1,j}) \otimes u_{2,i}(w_{2,j}) = v$. Therefore, the prover begins by computing and sending $v_{1,ij} \leftarrow u_{1,i}(w_{1,j})$ and $v_{2,ij} \leftarrow u_{2,i}(w_{2,j})$ for all $i \in [I], j \in [J]$. The verifier directly checks

$$\sum_{i\in[I],j\in[J]} v_{1,ij} \otimes v_{2,ij} = v.$$

The verifier must still check that $v_{1,ij} = u_{1,i}(w_{1,j})$ and $v_{2,ij} = u_{2,i}(w_{2,j})$ for all i, j. To do so, the verifier takes a random linear combination of these checks by sending random α, β from a valid sampling set $Q \subseteq \mathsf{R}$, and computing $v_1 = \sum_{i,j} \alpha^i \beta^j v_{1,ij}$ and $v_2 = \sum_{i,j} \alpha^i \beta^j v_{2,ij}$. The verifier then outputs $(\alpha, \beta, v_1, v_2)$, reducing the original check to the task of checking

$$\Big(\sum_i \alpha^i u_{1,i}\Big)\Big(\sum_j \beta^j w_{1,j}\Big) = v_1 \quad \text{and} \quad \Big(\sum_i \alpha^i u_{2,i}\Big)\Big(\sum_j \beta^j w_{2,j}\Big) = v_2.$$

Theorem 7 (Tensor Reduction). *For tensors $u = \sum_i u_{1,i} \otimes u_{2,i} \in \hom(W_1, V_1) \otimes \hom(W_2, V_2)$ of rank I, $w = \sum_j w_{1,j} \otimes w_{2,j} \in W_1 \otimes W_2$ of rank J, and $v \in V_1 \otimes V_2$ over ring R, the tensor reduction reduces the task of checking*

$$u(w) = v$$

to the task of checking

$$\Big(\sum_i \alpha^i u_{1,i}\Big)\Big(\sum_j \beta^j w_{1,j}\Big) = v_1 \quad \textit{and} \quad \Big(\sum_i \alpha^i u_{2,i}\Big)\Big(\sum_j \beta^j w_{2,j}\Big) = v_2$$

for verifier output $(\alpha, \beta, v_1, v_2)$*. Formally, if the former is true, then the latter is true with probability* 1*, and if the former is false, then the latter is false with probability at least* $1 - \frac{IJ}{|Q|}$*. The prover complexity, verifier complexity, and communication complexity are all proportional to* IJ*.*

Proof. This follows from the Schwartz-Zippel Lemma [38] extended to modules [11]. □

At first glance, it may seem that the communication cost of the tensor reduction is *greater* than the size of the witness: the witness only consists of J elements in $W_1 \otimes W_2$, but the prover sends IJ elements in V_1 and V_2. This is reconciled by the fact that elements of V_1 and V_2 are intended to be significantly smaller than elements in $W_1 \otimes W_2$. For instance, elements in $W_1 \otimes W_2$ may be long vectors that are mapped to short commitments in V_1 and V_2.

To build intuition for where tensor reductions are useful, we explain how to instantiate the tensor reduction to reconstruct the vector commitment reduction of knowledge presented in Section 2.

Example 4 (Vector Commitment Reduction of Knowledge). We construct a reduction of knowledge from $\mathcal{R}_{\mathsf{VC}}(n)$ to $\mathcal{R}_{\mathsf{VC}}(n/2)$ for $n = 2^i$ where $i \geq 1$. Consider group $\mathbb{G}$ of prime order p, and corresponding scalar field $\mathbb{F} \cong \mathbb{Z}_p$. Consider some public key $G \in \mathbb{G}^n$. Suppose a verifier would like to check for some commitment $\overline{A} \in \mathbb{G}$, that the prover knows vector $A \in \mathbb{F}^n$ such that $G(A) = \overline{A}$ where $G(A)$ is defined to be $\langle G, A \rangle$.

We observe that $\mathbb{G}^n \cong \mathbb{G}^{n/2} \otimes \mathbb{F}^2$ and $\mathbb{F}^n \cong \mathbb{F}^{n/2} \otimes \mathbb{F}^2$. Let $\{\delta_1, \delta_2\}$ be an orthonormal basis for $\mathbb{F}^2$ (i.e., we have that $\delta_i(\delta_j) = 1$ when $i = j$ and 0 otherwise). Then, we have that $G = G_1 \otimes \delta_1 + G_2 \otimes \delta_2$ and $A = A_1 \otimes \delta_1 + A_2 \otimes \delta_2$ for some $G_1, G_2 \in \mathbb{G}^{n/2}$ and $A_1, A_2 \in \mathbb{F}^{n/2}$. These terms can be interpreted as the first and second half of vectors G and A. Therefore, the verifier can equivalently check

$$\Big(\sum_i G_i \otimes \delta_i\Big)\Big(\sum_j A_j \otimes \delta_j\Big) = \overline{A}.$$

Applying the tensor reduction with respect to this decomposition, we have that the prover sends to the verifier $G_i(A_j), \delta_i(\delta_j)$ for $i, j \in \{1, 2\}$. Explicitly, letting $\overline{A}_{ij} = G_i(A_j)$, the prover sends the terms $(\overline{A}_{11}, 1)$, $(\overline{A}_{12}, 0)$, $(\overline{A}_{21}, 0)$, and $(\overline{A}_{22}, 1)$. We recognize that the first and last terms correspond with the first and second half of commitment $\overline{A}$, and the middle two terms are cross terms.

Upon receiving these terms, the verifier checks that

$$\overline{A}_{11} \otimes 1 + \overline{A}_{12} \otimes 0 + \overline{A}_{21} \otimes 0 + \overline{A}_{22} \otimes 1 = \overline{A}.$$

The verifier then samples and sends random $\alpha, \beta \leftarrow \mathbb{F}$, and sets the new statements to be checked to be $(G_1 + \alpha G_2)(A_1 + \beta A_2) = \sum_{i,j \in \{1,2\}} \overline{A}_{ij} \cdot \alpha^i \beta^j$ and $(\delta_1 + \alpha\delta_2)(\delta_1 + \beta\delta_2) = 1 + (\beta + \alpha) \cdot 0 + \alpha\beta \cdot 1$. The latter check holds immediately. As for the former check, the prover and verifier compute and output the new statement $G' \leftarrow G_1 + \alpha \cdot G_2 \in \mathbb{G}^{n/2}$ and $\overline{A} \leftarrow \sum_{i,j \in \{1,2\}} \overline{A}_{ij} \cdot \alpha^i \beta^j$. The prover privately computes and outputs the new witness vector $A' \leftarrow A_1 + \beta A_2 \in \mathbb{F}^{n/2}$. Now, it is sufficient for the verifier to check that the prover knows $A' \in \mathbb{F}^{n/2}$ such that $G'(A') = \overline{A}'$.

5.3 The Tensor Reduction of Knowledge

By generalizing Example 4 for arbitrary tensor statements, we arrive at the tensor reduction of knowledge, which is unconditionally secure. We start by defining the tensor relation which fixes the homomorphism and image as a statement and the preimage as the witness.[2] We then construct the tensor reduction of knowledge, which for a vector space of homomorphisms U and length n, reduces the task of checking knowledge of a preimage of a homomorphism in U^n to the task of checking knowledge of a preimage of a homomorphism in U. In the upcoming section, we show that the tensor reduction of knowledge can be instantiated to derive reductions of knowledge for various linear algebraic statements.

Definition 12 (Tensor Relation). *For* R*-modules* U*,* W *and* V*, such that* $U \cong \hom(W, V)$ *we define the tensor relation for* U *as follows*

$$\mathcal{R}(U) = \left\{ ((u, v), w) \;\middle|\; \begin{array}{l} u \in U, v \in V, w \in W, \\ u(w) = v \end{array} \right\}$$

Construction 4 (Tensor Reduction of Knowledge). Consider field $\mathbb{F}$, length parameter n, and $\mathbb{F}$-modules W and V. We construct a reduction of knowledge from $\mathcal{R}(\hom(W^n, V))$ to $\mathcal{R}(\hom(W, V))$. Let $\{\delta_i\}$ be an orthonormal basis for $\mathbb{F}^n$. Suppose the prover and verifier are provided statement $u = \sum_i u_i \otimes \delta_i \in \hom(W^n, V)$, and $v \in V$. Additionally, suppose the prover is provided an alleged witness $w = \sum_j w_j \otimes \delta_j \in W^n$ such that

$$((u, v), w) \in \mathcal{R}(\hom(W^n, V)).$$

The prover and verifier run a single tensor reduction on the equivalent statement

$$\Big(\sum_{i \in [n]} u_i \otimes \delta_i \Big) \Big(\sum_{j \in [n]} w_j \otimes \delta_j \Big) = v.$$

At the end of tensor reduction, the verifier outputs $(\alpha, \beta, v', \_)$. The prover and verifier compute $u' = \sum_i \alpha^i \cdot u_i$ and set the output statement to be (u', v'). The

[2] The tensor relation can be formally understood as a ternary relation where any public parameters are ignored. This makes it compatible with the reductions of knowledge framework which works over ternary relations defined over public parameter, statement, and witness tuples.

prover additionally computes the output witness $w' = \sum_j \beta^j \cdot w_j$ as dictated by the tensor reduction. Now, to check the original statement, it is sufficient for the verifier to check that the prover knows w' such that

$$((u', v'), w') \in \mathcal{R}(\hom(W, V)).$$

Theorem 8 (Tensor Reduction of Knowledge). *For field $\mathbb{F}$, length parameter n, and $\mathbb{F}$-modules W and V, Construction 4 is a reduction of knowledge from $\mathcal{R}(\hom(W^n, V))$ to $\mathcal{R}(\hom(W, V))$.*

We present a formal proof of Theorem 8 in the extended version [30].

6 Instantiating the Tensor Reduction of Knowledge

In this section, we demonstrate a unifying view of existing recursive algebraic arguments by deriving them by instantiating the tensor reduction of knowledge over the appropriate structures. We additionally derive new reductions of knowledge for bilinear forms by extending our techniques. In the extended version [30], we additionally discuss concrete modules each of these reductions can be instantiated over. In Section 7, we show how to stitch together these reductions to derive an argument for NP.

6.1 Vector Commitments and Linear Forms

We start by generalizing the vector commitment relation from Section 2 and then discuss how succinctly derive the vector commitment reduction of knowledge via the tensor reduction of knowledge. We then adapt the vector commitment reduction for linear forms. The high level approach is to first split all checks over size n vectors into k checks over size n/k vectors. These checks are then folded using a random linear combination. How exactly the vectors are split and folded is abstracted away by the tensor reduction of knowledge.

Consider size parameter $n \in \mathbb{N}$, and consider $\mathbb{F}$-modules $\mathbb{G}$ and $\mathbb{H}$ for field $\mathbb{F}$ such that $\mathbb{G} \cong \hom(\mathbb{H}, \mathbb{G} \otimes \mathbb{H})$. For public key $G \in \mathbb{G}^n$, and commitment $\overline{H} \in \mathbb{G} \otimes \mathbb{H}$, suppose a verifier would like to check that a prover knows $H \in \mathbb{H}^n$ such that $\sum_i G_i \otimes H_i = \overline{H}$. For example, suppose $\mathbb{G}$ is a group of prime order p where the discrete logarithm is hard, $\mathbb{H}$ and $\mathbb{F}$ are $\mathbb{Z}_p$, and $\otimes$ represents scalar multiplication. Then, this amounts to checking knowledge of the opening for a Pedersen commitment. Recall that the prover's claim can be expressed as a tensor statement $G(H) = \overline{H}$. Therefore, because $G \in \mathbb{G}^n$, we define the generalized vector commitment relation as the tensor relation over homomorphisms in $\mathbb{G}^n$.

Definition 13 (Generalized Vector Commitment Relation). *For length $n \in \mathbb{N}$ and group $\mathbb{G}$, the vector commitment relation is defined to be $\mathcal{R}(\mathbb{G}^n)$.*

Construction 5 (Vector Commitment Reduction of Knowledge). Because $\mathbb{G}^n \cong (\mathbb{G}^{n/k})^k$, we can directly apply the tensor reduction of knowledge to get a reduction from $\mathcal{R}(\mathbb{G}^n)$ to $\mathcal{R}(\mathbb{G}^{n/k})$.

Suppose that in addition to checking that the prover knows a vector opening to a commitment, the verifier would like to additionally check some public linear combination of the prover's opening. In particular, for public vector $A \in \mathbb{F}^n$, and $\sigma \in \mathbb{H}$, suppose the verifier would like to additionally check that $A(H) = \sigma$ where $A(H)$ is defined to be $\sum_{i \in [n]} A_i \otimes H_i$. For example, if $\otimes$ represents scalar multiplication, then this amounts to checking an inner-product. Recall, from Section 5, that this is equivalent to checking $(G \oplus A)(H) = \overline{H} \oplus \sigma$. Because $G \oplus A \in \mathbb{G}^n \oplus \mathbb{F}^n$, we define the linear forms relation as follows.

Definition 14 (Linear Forms Relation). *For length n and $\mathbb{F}$-module $\mathbb{G}$ for field $\mathbb{F}$, let $\mathsf{LF}_n = \mathbb{G}^n \oplus \mathbb{F}^n$. The linear forms relation is defined to be $\mathcal{R}(\mathsf{LF}_n)$.*

Construction 6 (Linear Forms Reduction of Knowledge). Consider $n, k \in \mathbb{N}$ such that k divides n. We construct a reduction of knowledge from $\mathcal{R}(\mathsf{LF}_n)$ to $\mathcal{R}(\mathsf{LF}_{n/k})$. In particular, we have that $\mathsf{LF}_n = (\mathbb{G} \oplus \mathbb{F})^n \cong (\mathbb{G} \oplus \mathbb{F})^{(n/k) \cdot k} = (\mathsf{LF}_{n/k})^k$. Therefore, the prover and verifier can apply the tensor reduction of knowledge with respect to this decomposition to reduce the task of checking a statement in $\mathcal{R}(\mathsf{LF}_n)$ to the task of checking a statement in $\mathcal{R}(\mathsf{LF}_{n/k})$.

Lemma 7 (Linear Forms Reduction of Knowledge). *Construction 6 is a reduction of knowledge from LF_n to $\mathsf{LF}_{n/k}$ with $O(n)$ prover and verifier time complexity and $O(k^2)$ communication complexity.*

As discussed in Section 2, we can construct a base case argument for LF_1 where the prover directly reveals the witness. Thus, we have the following.

Corollary 2 (Linear Forms Argument of Knowledge). *Consider $n, k \in \mathbb{N}$ such that k divides n. Let Π_{LF} be a reduction of knowledge from $\mathcal{R}(\mathsf{LF}_n)$ to $\mathcal{R}(\mathsf{LF}_{n/k})$. Let Π_{base} be an argument of knowledge for $\mathcal{R}(\mathsf{LF}_1)$. Then*

$$\Pi_{\mathsf{base}} \circ \underbrace{\Pi_{\mathsf{LF}} \circ \ldots \circ \Pi_{\mathsf{LF}}}_{\log_k n \text{ times}}$$

is an argument of knowledge for LF_n with $O(n)$ prover and verifier time complexity and $O(k^2 \cdot \log_k n)$ communication complexity.

6.2 Bilinear Forms

We extend the above methodology to develop a new reduction for bilinear forms. Recall that the public parameters consist of public key $G \in \mathbb{G}^m$, and the statement consists of matrix $M \in \mathbb{F}^{m \times m}$, commitments $\overline{A}, \overline{B} \in \mathbb{G}$, and scalar $\sigma \in \mathbb{F}$. A witness $(A, B) \in \mathbb{F}^m$ is satisfying if $A^\top M B = \sigma$, $\langle G, A \rangle = \overline{A}$, and $\langle G, B \rangle = \overline{B}$.

Below, we define a slight generalization where the length n of the vector B is some fraction of the length m. The key G is first (partially) compressed with respect to some public random vector $r \in \mathbb{F}^{m/n}$ to produce a new key $H \in \mathbb{G}^n$. This key is instead used to commit to the vector B. Our bilinear forms reduction will recursively compress G and B until $n = 1$. At this point the bilinear forms statement can be reduced to a linear forms statement.

Definition 15 (Bilinear Forms, Original). *Consider $\mathbb{F}$-module $\mathbb{G}$ for field $\mathbb{F}$. We define the bilinear forms relation, $\mathcal{R}_{\mathsf{Bil}}$, characterized by m rows and n columns as follows. The public parameters consist of key $G \in \mathbb{G}^m$. The statement consists of matrix $M \in \mathbb{F}^{m \times n}$, public randomness $r \in \mathbb{F}^{m/n}$, commitments $(\overline{A}, \overline{B}) \in \mathbb{G}$, and $\sigma \in \mathbb{F}$. A witness (A, B) is satisfying if $A^\top MB = \sigma$, $G(A) = \overline{A}$, and $G(r \otimes B) = \overline{B}$.*

Unlike vector commitments and linear forms, the bilinear forms relation cannot be encoded directly as a tensor evaluation statement. Our approach is to encode the original statement as the *related* statement,

$$(G \otimes H \oplus \boldsymbol{M})(A \otimes B) = (\overline{A} \otimes \overline{B} \oplus \sigma), \tag{3}$$

where $\boldsymbol{M} \in \mathbb{F}^m \otimes \mathbb{F}^n$ is a tensor such that $\boldsymbol{M}(A \otimes B) = A^\top MB$ and $H = G(r) \in \mathbb{G}^n$. The tensor-based statement implies checking the original statement so long as we additionally stipulate that the bilinear relation assumption holds for $(\mathbb{G}, \mathbb{F})$, and $(\mathbb{G}, \mathbb{G})$. Then, we can utilize the tensor reduction of knowledge to reduce the corresponding tensor relation $\mathcal{R}(\mathbb{G}^m \otimes \mathbb{G}^n \oplus \mathbb{F}^m \otimes \mathbb{F}^n)$.

In practice, $\mathbb{G}$ can be a symmetric bilinear group with the pairing operation acting as the tensor product and $\mathbb{G} \otimes \mathbb{G}$ denoting the target group. In this setting, the bilinear relation assumptions are equivalent to the discrete logarithm assumption over $\mathbb{G}$ and the double pairing assumption [1] over $(\mathbb{G}, \mathbb{G})$.

The computational hardness assumptions are a critical detail for arguing that checking Equation (3) is sufficient to check the original relation: the unconditional knowledge soundness property of the tensor reduction of knowledge only guarantees that the prover knows *some* satisfying witness in $\mathbb{F}^m \otimes \mathbb{F}^n$ which may be of the form $\sum_i A_i \otimes B_i$ (i.e., not a simple tensor). While this is a valid witness for the corresponding tensor statement, it is *not* a valid witness for the original statement. However, by assuming that the commitment scheme is computationally binding, we can argue that all A_i values must be the same. Leveraging this, we can show that the prover must know a single A and B vector that satisfies the statement. Formally, we define the bilinear forms relation as follows.

Definition 16 (Bilinear Forms, Tensor). *Consider $n, m \in \mathbb{N}$, and consider $\mathbb{F}$-module $\mathbb{G}$ for field $\mathbb{F}$ such that the bilinear relation assumption holds for $(\mathbb{G}, \mathbb{F})$, and $(\mathbb{G}, \mathbb{G})$. Let $\mathsf{BF}_{m,n} = (\mathbb{G}^m \otimes \mathbb{G}^n) \oplus (\mathbb{F}^m \otimes \mathbb{F}^n)$. We define the (tensor-based) bilinear form relation as the corresponding tensor relation $\mathcal{R}(\mathsf{BF}_{m,n})$.*

Next, we show how to recursively reduce $\mathcal{R}_{\mathsf{Bil}(m,n)}$ to $\mathcal{R}(\mathsf{LF}_m)$. To do so, we construct a reduction from $\mathcal{R}_{\mathsf{Bil}(m,n)}$ to $\mathcal{R}_{\mathsf{Bil}(m,n/k)}$, which internally uses the tensor reduction of knowledge from $\mathcal{R}(\mathsf{BF}_{m,n})$ to $\mathcal{R}(\mathsf{BF}_{m,n/k})$. We then construct a base case reduction from $\mathcal{R}_{\mathsf{Bil}(m,1)}$ to $\mathcal{R}(\mathsf{LF}_m)$.

Construction 7 (Bilinear Forms Reduction of Knowledge). Consider $n, k \in \mathbb{N}$ such that k divides n. We reduce from $\mathcal{R}_{\mathsf{Bil}(m,n)}$ to $\mathcal{R}_{\mathsf{Bil}(m,n/k)}$.

The generator samples public key $G \leftarrow \mathbb{G}^m$. Suppose that the prover and verifier take as input statement $(M, r, \overline{A}, \overline{B}, \sigma)$ and the prover additionally takes

as input and witness (A, B) such that

$$(G, (M, r, \overline{A}, \overline{B}, \sigma), (A, B)) \in \mathcal{R}_{\mathsf{Bil}(m,n)}$$

The prover and verifier begin by encoding the statement and witness as

$$((G \otimes H \oplus \boldsymbol{M}, \overline{A} \otimes \overline{B} \oplus \sigma), A \otimes B) \in \mathcal{R}(\mathsf{BF}_{m,n})$$

where $\boldsymbol{M} \in \mathbb{F}^m \otimes \mathbb{F}^n$ is such that $\boldsymbol{M}(A \otimes B) = A^\top M B$ and $H = G(r) \in \mathbb{G}^n$.

We observe that

$$\mathsf{BF}_{m,n} = \mathbb{G}^m \otimes \mathbb{G}^n \oplus \mathbb{F}^m \otimes \mathbb{F}^n \cong (\mathbb{G}^m \otimes \mathbb{G}^{n/k} \oplus \mathbb{F}^m \otimes \mathbb{F}^{n/k})^k = (\mathsf{BF}_{m,n/k})^k.$$

Therefore, the prover and verifier can apply the tensor reduction of knowledge with respect to this decomposition and reduce to the task of checking a statement in $\mathcal{R}(\mathsf{BF}_{m,n/k})$. At a high level, the tensor reduction prover and verifier partition $\boldsymbol{M}$ and H into k sets of columns and the prover partitions B into k corresponding sets of rows. The prover and verifier then take a random linear combination of these sets against weights $(s, s^2, \ldots, s^k)$ for some randomness $s \in \mathbb{F}$. By linearity, we have that the output statement is of the form

$$((G \otimes H' \oplus \boldsymbol{M}', \overline{A} \otimes \overline{B}' \oplus \sigma'), A \otimes B') \in \mathcal{R}(\mathsf{BF}_{m,n/k})$$

for some $H' = H((s, \ldots, s^k)) \in \mathbb{G}^{n/k}$, $\boldsymbol{M}' = \boldsymbol{M}((s, \ldots, s^k)) \in \mathbb{F}^m \otimes \mathbb{F}^{n/k}$, $\overline{B}' \in \mathbb{G}$, $\sigma' \in \mathbb{F}$, and $B' \in \mathbb{F}^{n/k}$. Together, the prover and verifier output the decoded statement $(M', (r \otimes (s, \ldots, s^k)), \overline{A}, \overline{B}', \sigma')$ and witness (A, B'). Now it is sufficient for the verifier to check that the prover knows (A, B') such that.

$$(G, (M', (r \otimes (s, \ldots, s^k)), \overline{A}, \overline{B}', \sigma'), (A, B')) \in \mathcal{R}_{\mathsf{Bil}(m,n/k)}.$$

Lemma 8 (Bilinear Forms Reduction of Knowledge). *Construction 7 is a reduction of knowledge from $\mathcal{R}_{\mathsf{Bil}(m,n)}$ to $\mathcal{R}_{\mathsf{Bil}(m,n/k)}$ with $O(mn)$ prover and verifier time complexity and $O(k^2)$ communication complexity.*

We present a formal proof of Lemma 8 in the extended version [30]. We now present the base case reduction.

Construction 8 (Bilinear Forms Base Case). We construct a reduction of knowledge from $\mathcal{R}_{\mathsf{Bil}(m,1)}$ to $\mathcal{R}(\mathsf{LF}_m)$. Once again the generator samples public key $G \leftarrow \mathbb{G}^m$. Consider statement $(M, r, \overline{A}, \overline{B}, \sigma)$ and alleged witness (A, B). The prover begins the reduction by directly sending B to the verifier. The verifier immediately checks that $H(B) = \overline{B}$ for $H = G(r)$. Additionally, as $M \in \mathbb{F}^m$ and $B \in \mathbb{F}$, the verifier computes the vector $V \leftarrow M \cdot B$. The verifier is left with checking that the prover knows $A \in \mathbb{F}^m$ such that $G(A) = \overline{A}$ and $V(A) = \sigma$. This is equivalent to checking that $((G \oplus V, \overline{A} \oplus \sigma), A) \in \mathcal{R}(\mathsf{LF}_m)$.

Lemma 9 (Bilinear Forms Base Case). *Construction 8 is a reduction of knowledge from $\mathcal{R}_{\mathsf{Bil}(m,1)}$ to $\mathcal{R}(\mathsf{LF}_m)$ with $O(m)$ prover and verifier time complexity and $O(1)$ communication complexity.*

Corollary 3 (Bilinear Forms to Linear Forms). *Consider $n, k \in \mathbb{N}$ such that k divides n. Let Π_{Bil} be the reduction of knowledge from $\mathcal{R}_{\mathsf{Bil}(m,n)}$ to $\mathcal{R}_{\mathsf{Bil}(m,n/k)}$. Let Π_{base} be the reduction of knowledge from $\mathcal{R}(\mathsf{Bil}(m,1))$ to $\mathcal{R}(\mathsf{LF}_m)$. Then*

$$\Pi_{\mathsf{base}} \circ \underbrace{\Pi_{\mathsf{Bil}} \circ \ldots \circ \Pi_{\mathsf{Bil}}}_{\log_k n \text{ times}}$$

is a reduction of knowledge from $\mathcal{R}_{\mathsf{Bil}(m,n)}$ to $\mathcal{R}(\mathsf{LF}_m)$ with $O(mn)$ prover and verifier time complexity and $O(k^2 \cdot \log_k n)$ communication complexity.

7 An Argument of Knowledge for NP

In this section, we develop an argument of knowledge for NP with logarithmic communication by leveraging our reductions of knowledge for linear algebraic statements. In particular, we first show that an NP-complete relation, $\mathcal{R}_{\mathsf{ACS}}$, can be encoded as a sequence of linear and bilinear forms constraints over the same commitment. We then develop helper reductions of knowledge that reduce the task of checking many linear and bilinear forms over the same commitment to a single linear and bilinear form. We then apply our reductions of knowledge for linear forms and bilinear forms.

Definition 17 (Algebraic Constraint System [29]). *Consider group $\mathbb{G}$ and corresponding field $\mathbb{F}$ such that the bilinear relation assumption holds for $(\mathbb{G}, \mathbb{F})$ and $(\mathbb{G}, \mathbb{G})$. We define the* NP*-complete algebraic constraint relation, $\mathcal{R}_{\mathsf{ACS}}$, characterized by n variables, $m = O(n)$ constraints, and ℓ inputs as follows. The public parameters consist of $G \in \mathbb{G}^n$. The statement consists of m sparse constraint matrices $M_1, \ldots, M_m \in \mathbb{F}^{n \times n}$ such that the total number of non-zero values in* all *matrices combined is $O(n)$, public inputs and outputs vector $X \in \mathbb{F}^\ell$, and witness commitment $\overline{Z} \in \mathbb{G}$. A witness vector $W \in \mathbb{F}^{n-\ell}$ is satisfying if for $Z = (X, W)$, $Z^\top M_i Z = 0$ for all $i \in [m]$, and $G(Z) = \overline{Z}$.*

We can encode $\mathcal{R}_{\mathsf{ACS}}$ to tensor relations as follows: First, the verifier can check that $((G \oplus \delta_i, \overline{Z} \oplus X_i), Z) \in \mathcal{R}(\mathsf{LF}_n)$ for all $i \in [\ell]$ to ensure that Z contains public vector X. To check the commitment and constraints, it is sufficient for the verifier to check that the prover knows $Z_1, Z_2 \in \mathbb{F}^n$ such that $(G, (M_i, 1, \overline{Z}, \overline{Z}, 0), (Z_1, Z_2)) \in \mathcal{R}_{\mathsf{Bil}(n,n)}$ for all $i \in [m]$. The bilinear relation assumptions ensure that Z, Z_1 and Z_2 are equal.

Next, we leverage the fact that all linear form checks and all bilinear form checks are over the same commitment to reduce these checks. We formally capture the set of linear and bilinear form checks over the same commitment as the multiple linear and bilinear forms relations.

Definition 18 (Multiple Linear Forms). *We define $\mathcal{R}_{\mathsf{MLF}(n,\ell)}$ such that $((G, (V_1, \ldots, V_\ell), (\sigma_1, \ldots, \sigma_\ell), \overline{Z}), Z) \in \mathcal{R}_{\mathsf{MLF}(n,\ell)}$ if and only if $((G \oplus V_i, \overline{Z} \oplus \sigma_i), Z) \in \mathcal{R}(\mathsf{LF}_n)$ for all i in $[\ell]$.*

Definition 19 (Multiple Bilinear Forms). *We define* $\mathcal{R}_{\mathsf{MBil}(m,n,\ell)}$ *such that* $(G, ((M_1, \ldots, M_\ell), r, (\sigma_1, \ldots, \sigma_\ell), \overline{Z}_1, \overline{Z}_2), (Z_1, Z_2)) \in \mathcal{R}_{\mathsf{MBil}(m,n,\ell)}$ *if and only if* $(G, (M_i, r, \overline{Z}_1, \overline{Z}_2, \sigma_i), (Z_1, Z_2)) \in \mathcal{R}_{\mathsf{Bil}(m,n)}$ *for all* i *in* $[\ell]$.

With these relations, the above encoding can be captured as a reduction of knowledge in which the prover and verifier do not interact but rather take as input an $\mathcal{R}_{\mathsf{ACS}}$ statement-witness pair and output the corresponding tensor-based statements and witnesses in the multiple linear forms and bilinear forms relations. This step can be interpreted as a Levin reduction.

Lemma 10 (Encoding NP as Tensor Relations). *There exists a reduction of knowledge from* $\mathcal{R}_{\mathsf{ACS}(m,n,\ell)}$ *to* $\mathcal{R}_{\mathsf{MBil}(n,n,m)} \times \mathcal{R}_{\mathsf{MLF}(n,\ell)}$ *with* $O(n)$ *prover and verifier complexity, and no communication.*

Because all ℓ checks for $\mathcal{R}_{\mathsf{MLF}(n,\ell)}$ concern the same committed value, we observe that they can be batched into a single check for $\mathcal{R}(\mathsf{LF}_n)$ using a random linear combination. In particular, the verifier can send a random challenge $r \in \mathbb{F}$. Together the prover and verifier can compute $V \leftarrow \sum_i V_i \cdot r^i$ and $\sigma \leftarrow \sum_i \sigma_i \cdot r^i$ and reduce to checking that the prover knows Z such that $((G \oplus V, \overline{Z} \oplus \sigma), Z) \in \mathcal{R}(\mathsf{LF}_n)$. Similarly, we can reduce multiple bilinear forms over the same commitment to a single bilinear form. Formally, we have the following reductions.

Lemma 11 (Linear Forms Batch Reduction). *For* $n, m, \ell \in \mathbb{N}$*, there exists a reduction of knowledge from* $\mathcal{R}_{\mathsf{MLF}(n,\ell)}$ *to* $\mathcal{R}_{\mathsf{LF}(n)}$ *with* $O(n\ell)$ *prover and verifier time complexity, and* $O(1)$ *communication complexity.*

Lemma 12 (Bilinear Forms Batch Reduction). *For* $n, m, \ell \in \mathbb{N}$*, there exists a reduction of knowledge from* $\mathcal{R}_{\mathsf{MBil}(m,n,\ell)}$ *to* $\mathcal{R}_{\mathsf{Bil}(m,n)}$ *with* $O(mn\ell)$ *prover and verifier time complexity, and* $O(1)$ *communication complexity.*

Putting everything together, we arrive at an argument of knowledge for NP.

Corollary 4 (An Argument of Knowledge for NP). *Let* Π_{encode} *be the reduction of knowledge from* $\mathcal{R}_{\mathsf{ACS}(n,m,\ell)}$ *to* $\mathcal{R}_{\mathsf{MBil}(n,n,m)} \times \mathcal{R}_{\mathsf{MLF}(n,\ell)}$ *(Lemma 10). Let* Π_{batchLF} *be the batching scheme for linear forms (Lemma 11). Let* Π_{batchBil} *be the batching scheme for bilinear forms (Lemma 12). Let* Π_{LF_n} *be the argument of knowledge for* $\mathcal{R}(\mathsf{LF}_n)$ *with decomposition parameter* k *(Construction 2). Let* $\Pi_{\mathsf{Bil}(n,n)}$ *be the reduction of knowledge from* $\mathcal{R}_{\mathsf{Bil}(n,n)}$ *to* $\mathcal{R}_{\mathsf{LF}(n)}$ *with decomposition parameter* k *(Corollary 3). Let* Π_{id} *be the identity reduction of knowledge (i.e., the prover and verifier output their inputs). Let* Π_{foldBool} *be a* 2*-folding scheme for* $\mathcal{R}_\top$ *(i.e., the verifier outputs* true *if both its inputs are* true*). Then*

$$\Pi_{\mathsf{foldBool}} \circ (\Pi_{\mathsf{id}} \times \Pi_{\mathsf{LF}_m}) \circ (\Pi_{\mathsf{LF}_n} \times \Pi_{\mathsf{Bil}(n,n)}) \circ (\Pi_{\mathsf{batchLF}} \times \Pi_{\mathsf{batchBil}}) \circ \Pi_{\mathsf{encode}}$$

is an argument of knowledge for $\mathcal{R}_{\mathsf{ACS}(n,m,\ell)}$ *with* $O(n)$ *prover and verifier time complexity, and* $O(k^2 \log_k n)$ *communication complexity.*

Acknowledgments We thank Jonathan Bootle, Quang Dao, Vipul Goyal, Yael Tauman Kalai, Jonathan Lee, Srinath Setty, Elaine Shi, and Zoe Wellner for comments on earlier versions of this work.

References

[1] Abe, M., Fuchsbauer, G., Groth, J., Haralambiev, K., Ohkubo, M.: Structure-preserving signatures and commitments to group elements. In: Annual Cryptology Conference. pp. 209–236. Springer (2010)

[2] Arora, S., Barak, B.: Computational complexity: a modern approach. Cambridge University Press (2009)

[3] Attema, T., Cramer, R.: Compressed-protocol theory and practical application to plug & play secure algorithmics. In: Advances in Cryptology–CRYPTO 2020: 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17–21, 2020, Proceedings, Part III. pp. 513–543. Springer (2020)

[4] Attema, T., Cramer, R., Rambaud, M.: Compressed Sigma-protocols for bilinear group arithmetic circuits and application to logarithmic transparent threshold signatures. In: Advances in Cryptology–ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part IV. pp. 526–556. Springer (2021)

[5] Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Proceedings of the 1st ACM Conference on Computer and Communications Security. pp. 62–73 (1993)

[6] Ben-Sasson, E., Chiesa, A., Spooner, N.: Interactive oracle proofs. In: Theory of Cryptography: 14th International Conference, TCC 2016-B, Beijing, China, October 31-November 3, 2016, Proceedings, Part II 14. pp. 31–60. Springer (2016)

[7] Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. pp. 326–349 (2012)

[8] Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: Recursive composition and bootstrapping for SNARKs and proof-carrying data. In: Proceedings of the forty-fifth annual ACM symposium on Theory of computing. pp. 111–120 (2013)

[9] Boneh, D., Drake, J., Fisch, B., Gabizon, A.: Halo infinite: Proof-carrying data from additive polynomial commitments. In: Malkin, T., Peikert, C. (eds.) Advances in Cryptology – CRYPTO 2021. pp. 649–680. Springer International Publishing, Cham (2021)

[10] Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In: Advances in Cryptology–EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II 35. pp. 327–357. Springer (2016)

[11] Bootle, J., Chiesa, A., Sotiraki, K.: Sumcheck arguments and their applications. In: Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41. pp. 742–773. Springer (2021)

[12] Bowe, S., Grigg, J., Hopwood, D.: Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Paper 2019/1021 (2019)

[13] Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: 2018 IEEE symposium on security and privacy (SP). pp. 315–334. IEEE (2018)

[14] Bünz, B., Chiesa, A., Lin, W., Mishra, P., Spooner, N.: Proof-carrying data without succinct arguments. In: Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41. pp. 681–710. Springer (2021)

[15] Bünz, B., Fisch, B., Szepieniec, A.: Transparent snarks from dark compilers. In: Advances in Cryptology–EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part I 39. pp. 677–706. Springer (2020)

[16] Bünz, B., Maller, M., Mishra, P., Tyagi, N., Vesely, P.: Proofs for inner pairing products and applications. In: Advances in Cryptology–ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part III 27. pp. 65–97. Springer (2021)

[17] Campanelli, M., Faonio, A., Fiore, D., Querol, A., Rodríguez, H.: Lunar: a toolbox for more efficient universal and updatable zkSNARKs and commit-and-prove extensions. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 3–33. Springer (2021)

[18] Campanelli, M., Nitulescu, A., Ràfols, C., Zacharakis, A., Zapico, A.: Linear-map vector commitments and their practical applications. In: Agrawal, S., Lin, D. (eds.) Advances in Cryptology – ASIACRYPT 2022. pp. 189–219. Springer Nature Switzerland, Cham (2022)

[19] Chiesa, A., Hu, Y., Maller, M., Mishra, P., Vesely, N., Ward, N.: Marlin: Preprocessing zkSNARKs with universal and updatable srs. In: Advances in Cryptology–EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part I 39. pp. 738–768. Springer (2020)

[20] Chung, H., Han, K., Ju, C., Kim, M., Seo, J.H.: Bulletproofs+: Shorter proofs for a privacy-enhanced distributed ledger. IEEE Access **10**, 42067–42082 (2022)

[21] Delignat-Lavaud, A., Fournet, C., Kohlweiss, M., Parno, B.: Cinderella: Turning shabby x. 509 certificates into elegant anonymous credentials with the magic of verifiable computation. In: 2016 IEEE Symposium on Security and Privacy (SP). pp. 235–254. IEEE (2016)

[22] Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Advances in Cryptology—CRYPTO'86: Proceedings 6. pp. 186–194. Springer (1987)
[23] Fuchsbauer, G., Kiltz, E., Loss, J.: The algebraic group model and its applications. In: Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part II 38. pp. 33–62. Springer (2018)
[24] Gabizon, A., Williamson, Z.J., Ciobotaru, O.: PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953 (2019)
[25] Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct NIZKs without PCPs. In: Advances in Cryptology–EUROCRYPT 2013: 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings 32. pp. 626–645. Springer (2013)
[26] Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: Proceedings of the forty-third annual ACM symposium on Theory of computing. pp. 99–108 (2011)
[27] Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof-systems. In: Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali, pp. 203–225 (2019)
[28] Kosba, A., Miller, A., Shi, E., Wen, Z., Papamanthou, C.: Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In: 2016 IEEE symposium on security and privacy (SP). pp. 839–858. IEEE (2016)
[29] Kothapalli, A., Masserova, E., Parno, B.: Poppins: A direct construction for asymptotically optimal zkSNARKs. Cryptology ePrint Archive, Report 2020/1318 (2020)
[30] Kothapalli, A., Parno, B.: Algebraic reductions of knowledge. Cryptology ePrint Archive, Paper 2022/009 (2022)
[31] Kothapalli, A., Setty, S., Tzialla, I.: Nova: Recursive zero-knowledge arguments from folding schemes. In: Advances in Cryptology–CRYPTO 2022: 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15–18, 2022, Proceedings, Part IV. pp. 359–388. Springer (2022)
[32] Lee, J.: Dory: Efficient, transparent arguments for generalised inner products and polynomial commitments. In: Theory of Cryptography: 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8–11, 2021, Proceedings, Part II. pp. 1–34. Springer (2021)
[33] Lund, C., Fortnow, L., Karloff, H., Nisan, N.: Algebraic methods for interactive proof systems. Journal of the ACM (JACM) **39**(4), 859–868 (1992)
[34] Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Advances in Cryptology—CRYPTO'91: Proceedings. pp. 129–140. Springer (2001)
[35] Ràfols, C., Zapico, A.: An algebraic framework for universal and updatable snarks. In: Advances in Cryptology–CRYPTO 2021: 41st Annual In-

ternational Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I. pp. 774–804. Springer (2021)
[36] Ràfols, C., Zacharakis, A.: Folding schemes with selective verification. Cryptology ePrint Archive, Paper 2022/1576 (2022)
[37] Sasson, E.B., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., Virza, M.: Zerocash: Decentralized anonymous payments from bitcoin. In: 2014 IEEE symposium on security and privacy. pp. 459–474. IEEE (2014)
[38] Schwartz, J.T.: Fast probabilistic algorithms for verification of polynomial identities. Journal of the ACM (JACM) **27**(4), 701–717 (1980)
[39] Setty, S.: Spartan: Efficient and general-purpose zksnarks without trusted setup. In: Advances in Cryptology–CRYPTO 2020: 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17–21, 2020, Proceedings, Part III. pp. 704–737. Springer (2020)
[40] Tzialla, I., Kothapalli, A., Parno, B., Setty, S.: Transparency dictionaries with succinct proofs of correct operation. In: Network and Distributed System Security (NDSS) 2022 (April 2022)
[41] Valiant, P.: Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In: Theory of Cryptography: Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008. Proceedings 5. pp. 1–18. Springer (2008)
[42] Wahby, R.S., Tzialla, I., Shelat, A., Thaler, J., Walfish, M.: Doubly-efficient zksnarks without trusted setup. In: 2018 IEEE Symposium on Security and Privacy (SP). pp. 926–943. IEEE (2018)
[43] Zhang, Y., Katz, J., Papamanthou, C.: Integridb: Verifiable sql for outsourced databases. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. pp. 1480–1491 (2015)