

LegoSNARK: Modular Design and Composition of Succinct Zero-Knowledge Proofs

Matteo Campanelli¹, Dario Fiore¹, and Anaïs Querol^{1,2}

¹ IMDEA Software Institute

² Universidad Politécnica de Madrid

matteo.campanelli@imdea.org

dario.fiore@imdea.org

anais.querol@imdea.org

Full Version

Abstract. We study the problem of building non-interactive proof systems modularly by linking small specialized “gadget” SNARKs in a lightweight manner. Our motivation is both theoretical and practical. On the theoretical side, modular SNARK designs would be flexible and reusable. In practice, specialized SNARKs have the potential to be more efficient than general-purpose schemes, on which most existing works have focused. If a computation naturally presents different “components” (e.g. one arithmetic circuit and one boolean circuit), a general-purpose scheme would homogenize them to a single representation with a subsequent cost in performance. Through a modular approach one could instead exploit the nuances of a computation and choose the best gadget for each component.

Our contribution is LegoSNARK, a “toolbox” (or framework) for commit-and-prove zkSNARKs (CP-SNARKs) that includes:

1) General composition tools: build new CP-SNARKs from proof gadgets for basic relations *simply*.

2) A “lifting” tool: a compiler to add commit-and-prove capabilities to a broad class of existing zkSNARKs *efficiently*. This makes them interoperable (linkable) within the same computation. For example, one QAP-based scheme can be used to prove one component; another GKR-based scheme can be used to prove another.

3) A collection of succinct proof gadgets for a variety of relations. Additionally, through our framework and gadgets, we are able to obtain new succinct proof systems. Notably:

– **LegoGro16**, a commit-and-prove version of Groth16 zkSNARK, that operates over data committed with a classical Pedersen vector commitment, and that achieves a 5000× speedup in proving time.

– **LegoUAC**, a pairing-based SNARK for arithmetic circuits that has a universal, circuit-independent, CRS, and proving time linear in the number of circuit gates (vs. the recent scheme of Groth et al. (CRYPTO’18) with quadratic CRS and quasilinear proving time).

– CP-SNARKs for matrix multiplication that achieve optimal proving complexity.

4) A codebase written in C++ for highly composable zkSNARKs with commit-and-prove capabilities^a.

^a Available at <https://github.com/imdea-software/legosnark> .

Table of Contents

1	Introduction	4
1.1	Our Results	6
1.2	Related Work	7
1.3	Roadmap	9
2	Preliminaries	9
2.1	Relations	10
2.2	Commitment Schemes	10
2.3	Zero-Knowledge SNARKs	10
3	Building the LegoSNARK Framework	12
3.1	Commit and Prove SNARKs	12
3.2	Composition Properties of CP-SNARKs	13
3.3	Commit-Carrying SNARKs	14
3.4	Existing CP-SNARKs and cc-SNARKs	17
3.5	Bootstrapping our Framework	18
4	CP-SNARKs for Pedersen-like Commitments	19
4.1	CP-SNARK for Pedersen Verification	20
4.2	CP-SNARK for Linear Properties	22
5	Efficient CP-SNARKs for Polynomial Commitments	23
5.1	Preliminaries and Building Blocks	23
5.2	A CP-SNARK for Sum-Check	25
5.3	A CP-SNARK for Hadamard Products	27
5.4	A CP-SNARK for Self Permutation	28
5.5	A CP-SNARK for Linear Properties of Committed Vector	31
5.6	A CP-SNARK for Matrix Multiplication	32
6	LegoSNARK Applications and Evaluation	35
6.1	Preliminaries and Building Blocks	35
6.2	Arithmetic Circuit Satisfiability	36
6.3	Parallel Computation on Joint Inputs	38
7	Experimental Evaluation	40
7.1	Commit-and-Prove SNARKs	40
7.2	Matrix Multiplication	41
7.3	LegoAC1 for Arithmetic Circuits	41
7.4	Parallel Checks on Joint Inputs	43
8	Conclusions	44

A	Security proof of CP-SNARK composition	48
A.1	Proof of Knowledge Soundness	48
A.2	Proof of Zero-Knowledge	50
B	Proofs for the General Compiler	50
B.1	Proof of Knowledge Soundness	50
B.2	Proof of Zero-Knowledge	53
C	Supplementary Results on CP_{link}	55
C.1	Proof of CP_{link} Security	55
C.2	An extension of CP_{link} for Prefixes of a Committed Vector	56
D	A zkSNARK for Linear Subspaces	56
E	A Construction of PolyCom and CP_{poly} from zk-vSQL	58
F	Additional Material on CP-SNARKs for PolyCom	59
F.1	Proof of our CP_{sc}	59
F.2	Proof of Security of CP_{had}	60
F.3	Proof of CP_{sfprm}	61
F.4	Proof of CP_{lin}	62
F.5	A CP-SNARK for Data-Parallel Computations	63
G	A CP-SNARK for Internal Products from Thaler’s Protocol	64
G.1	CMT Protocol	64
G.2	Thaler’s Protocol for Trees of Multiplications	65
G.3	Adapting zk-vSQL to Thaler’s Protocol	65
H	Commit and Prove SNARKs from existing schemes	67
H.1	“Adaptive Pinocchio” [Vee17]	68
H.2	Lipmaa’s Hadamard Product Argument [Lip16]	68
H.3	zk-vSQL [ZGK ⁺ 17b]	68
H.4	Geppetto [CFH ⁺ 15]	69
H.5	cc-SNARKs based on Groth’s SNARK	70

1 Introduction

Zero-knowledge proofs (ZKPs), introduced by Goldwasser, Micali and Rackoff [GMR89], let a prover convince a verifier of a statement without revealing more information than its validity. This power of ZKPs—simultaneously providing *integrity* (the prover cannot cheat) and *privacy* (the verifier does not learn any of the prover’s secrets)—has found countless applications, including multiparty computation [GMW87], signature schemes [Sch91], public-key encryption [NY90], and, more recently, blockchain systems [BCG⁺14, AJ18].

Some zero-knowledge proof systems—called *succinct* or simply *zkSNARKs*, zero-knowledge Succinct Non-interactive Argument of Knowledge—have short and efficiently verifiable proofs [Mic00, GW11, BCCT12]. Succinctness is desirable in general but is especially critical in applications where verifiers would not invest significant computational resources (e.g. if they are unwilling to do it for reasons of scalability and cost, or if they are computationally weak).

Motivation. The last years have seen remarkable progress in the construction of zkSNARKs. Different lines of work (cf. Section 1.2 for a detailed review) have built a variety of schemes that are highly expressive, supporting general computations in the class NP. The general-purpose nature of these schemes makes them very attractive to practitioners. At the same time, this high expressivity comes at a cost in terms of performance. To achieve generality, these constructions abstract specific features of computation by assuming one *single unifying representation* (e.g., boolean or arithmetic circuits, state-machine transitions, RAM computations), and this abstraction is often a source of overhead, for two main reasons.

First, *general-purpose zk-SNARKs may miss opportunities for significant optimizations* by not exploiting the nuances of a computation. In contrast, specialized solutions can gain efficiency by exploiting specific structural properties. For example, recent works [CMT12, WTas⁺17] show how to highly optimize the GKR protocol [GKR08] for the case of parallel computations. A further example is the specialized protocol for the multiplication of $n \times n$ matrices we propose in Section 5.6. Here, our prover runs in $O(n^2)$ time as opposed to any circuit-based approach running in at least $O(n^3)$ time.

Second, *computation tends to be heterogeneous*, often consisting of several subroutines of different nature, e.g. both arithmetic and boolean components. If we design SNARKs assuming one single general representation then we will not be able to provide the best match for all the different subroutines. In this context specialized protocols are clearly not an answer either as they fail whenever faced with a non homogeneous computation. As a concrete example, the GKR-like protocols mentioned above are highly efficient when executed on parallel computations, but they fail to be succinct if a computation *also* includes heavily sequential subroutines (e.g. iterated block ciphers).

In contrast, specialized solutions can gain efficiency by exploiting specific structural properties. For example, recent works [CMT12, WTas⁺17] show how to highly optimize the GKR protocol [GKR08] for the case of parallel computations. A further example is the specialized protocol for the multiplication of $n \times n$ matrices we propose in Section 5.6. Here, our prover runs in $O(n^2)$ time as opposed to a circuit-based approach running in at least $O(n^3)$ time.

A Modular Approach for zk-SNARKs. In this paper we study an alternative approach to the design of zkSNARKs that would gain the advantages of specialized proof systems without inheriting their shortcomings when applied to heterogeneous computations. With this goal in mind we propose to build zkSNARKs by proceeding in a modular “bottom-up” fashion. Most current works use a “top-down” approach: they build general-purpose schemes adopting one single representation that *must*

be shared across all the different subroutines in the program. On the other hand, in this work we consider designing a “global” SNARK for a computation C through a (lightweight) linking of “smaller” specialized SNARKs for the different subroutines composing C . We call these interlinked specialized SNARKs *proof gadgets*, as they act as basic building blocks that one can compose and reuse according to the situation.

The modular approach has multiple benefits.³ First, it allows for *reducing complexity*: instead of focusing on handling arbitrary computation using a single representation, one can focus on a smaller, more specific problem (e.g., log-depth computation, membership proof, range proof, algebraic group relation etc.), and exploit its nuances to get a more efficient solution. This way, one could maximize efficiency by letting each subroutine of C be handled by a different proof system, specialized and efficient for that type of computation. Second, modularity allows for flexibility and costs reduction: a proof gadget can be reused in several systems and one can easily plug in a new solution, or replace an old one.

Modularity from Commit-and-Prove SNARKs. To realize this modular approach we rely on the well known *commit-and-prove* (CP) methodology [Kil89, CLOS02]. With a CP scheme one can prove statements of the form “ $c_{\text{ck}}(x)$ contains x such that $R(x, w)$ ” where $c_{\text{ck}}(x)$ is a commitment. To see how the CP capability can be used for modular composition consider the following example of sequential composition in which one wants to prove that $\exists w : z = h(x; w)$, where $h(x; w) := g(f(x; w); w)$. Such a proof can be built by combining two CP systems Π_f and Π_g for its two building blocks, i.e., respectively f and g : the prover creates a commitment $c_{\text{ck}}(y)$ of y , and then uses Π_f (resp. Π_g) to prove that “ $c_{\text{ck}}(y)$ contains $y = f(x; w)$ (resp. contains y such that $z = g(y; w)$)”.

Challenges of the CP modular composition. The composition idea sketched above implicitly assumes that Π_f and Π_g work on the same commitment $c_{\text{ck}}(y)$. Namely, *in order to be composed, different CP schemes must be compatible with the same commitment scheme* (and commitment key). Essentially we need a *sort of universal commitment scheme* that is as decoupled⁴ as possible from the specific argument systems that will operate on it.

We argue that achieving such universality with state-of-the-art zkSNARKs entails major challenges:

- (a) Most of the popular zkSNARKs, e.g., [PHGR13, Gro16], are not explicitly commit-and-prove. This limitation can be overcome using a (somewhat folklore) approach in which the SNARK Π additionally proves the correct opening of the commitment, i.e., $R(x, w) \wedge “c_{\text{ck}}(x)$ opens to $x”$. This approach has two main drawbacks: (i) Π must be expressive enough to include the commitment verification in its language, but in our vision Π is a SNARK for a specialized task and may not have this capability; (ii) even if Π were expressive enough (e.g., supports arbitrary circuits), encoding commitment verification incurs significant overheads.⁵
- (b) Some existing SNARKs have commit-and-prove capabilities [Gro10, CFH⁺15, Lip16, Vee17]. Yet, each of these schemes uses its own specific commitment scheme. In some cases [CFH⁺15] the commitment keys are *relation-dependent*, which means commitments cannot be generated before

³ Most of these benefits are the typical ones of modularity, a design approach that is successfully used in a variety of fields, such as architecture, manufacturing, software design, and programming.

⁴ We find it apt to describe this notion in terms of *coupling*, the common measure of how interconnected two components are in a software system.

⁵ For example, we experimentally found that, when handling a Pedersen commitment to a vector of length 2048 with [Gro16], the proving overhead is 428 secs (7 minutes).

fixing one or multiple relations.⁶ In the other cases, despite being relation-independent, commitment keys have a very specific structure that may not fit other proof systems. In summary, a main limitation of existing commit-and-prove SNARKs is their incompatibility, between them and with other potentially more efficient candidates to be developed.

1.1 Our Results

LegoSNARK Framework. We present LegoSNARK, a framework for commit-and-prove zk-SNARKs (CP-SNARKs) that includes:

- *Definitions* that formalize CP-SNARKs and their variants.
- *Composition recipes* that show how to use different CP-SNARKs in a generic and secure way for handling conjunction, disjunction and sequential composition of relations. This composition result enables the use of modularity in designing CP-SNARKs for complex relations out of schemes for simpler relations.
- *A generic construction* to efficiently turn a broad class of zkSNARKs into CP-SNARKs that can be composed together. This class includes several existing schemes such as ones based on quadratic arithmetic programs [PHGR13, CFH⁺15, Gro16], or zk-vSQL [ZGK⁺17a, ZGK⁺17b]. For this transformation we only need a “minimal” CP-SNARK, CP_{link} , for proving that two commitments (under different schemes) open to the same value.

LegoSNARK Gadgets. We populate our framework by constructing new CP-SNARKs for several basic relations, such as:

- CP_{link} for proving that two *different* Pedersen-like commitments open to the same vector.⁷ Plugging CP_{link} in our generic construction solves the challenges (a) and (b) mentioned above and gives us interoperable versions of several existing schemes.
- CP_{lin} for proving that a linear relation $\mathbf{F} \cdot \mathbf{u} = \mathbf{x}$ holds for a committed vector \mathbf{u} , a public matrix \mathbf{F} and public vector \mathbf{x} .
- CP_{had} for proving that a vector \mathbf{u}_0 is the Hadamard product of \mathbf{u}_1 and \mathbf{u}_2 , when all the three vectors are committed.
- CP_{sfrm} for proving a self-permutation, i.e., that $y_i = y_{\phi(i)}$ for a public permutation ϕ and a committed vector \mathbf{y} .
- CP_{mm} for proving that matrix \mathbf{X} is the product of committed matrices \mathbf{A} and \mathbf{B} .

All the aforementioned schemes have succinct proofs and work for Pedersen-like commitments in bilinear groups. This means that by using our generic construction with CP_{link} they can be turned to support the same commitment and then be composed.

LegoSNARK Applications and Evaluation. Using our initial set of specialized proof gadgets, our next step is to combine them in order to build new succinct proof systems for different use cases, mentioned below. Our results offer various improvements over the state of the art. We have also implemented some of our solutions to test their concrete performance.

⁶ This could be mitigated by using universal circuits, paying a (multiplicative) logarithmic overhead in parameters size and prover complexity.

⁷ By “Pedersen-like” we mean schemes where the verification algorithm is the same as in Pedersen scheme [Ped92] for vectors (but the bases can have a different distribution).

1) **EFFICIENT COMMIT-AHEAD-OF-TIME.** Through our generic construction instantiated with CP_{link} we also obtained commit-and-prove versions of popular efficient zkSNARKs, such as Groth’s [Gro16], that can prove statements about data committed using the Pedersen scheme for vectors [Ped92], in which bases are random group elements that can be generated without trusted setup. Such commit-and-prove schemes are useful in applications where one needs to commit *before* the SNARK keys for a relation are created, e.g., to post commitments on a blockchain so that one can later prove statements about the committed data. By applying our solution to [Gro16] we obtain a scheme that is $5000\times$ faster than Groth16, where the commitment is encoded in the circuit.

2) **CP-SNARKS FOR PARALLEL COMPUTATION.** Consider the problem of proving (in zero-knowledge) correctness of a computation that consists of the same subcircuit executed in parallel. The recent Hyrax system [WTs⁺18] is suitably designed for and shows good performances on this type of circuit. It requires, however, an additional verification cost whenever the repeated subcircuits share (non-deterministic) inputs, which is common. The verifier thus pays an additional factor linear in the total width of the circuit. Using our LegoSNARK framework we show how to build a new CP-SNARK based on Hyrax that avoids this problem. The idea is that parallel computation on joint inputs can be expressed as the combination of a fully parallel computation (after inputs were appropriately duplicated) and a permutation check to ensure that inputs have been duplicated correctly. We build this by combining our CP_{lin} gadget with a version of Hyrax modified to work with the polynomial commitment of zk-vSQL [ZGK⁺17b].

3) **CP-SNARKS FOR ARITHMETIC CIRCUITS.** We give two main constructions of CP-SNARKs for arithmetic circuit (AC) satisfiability. Table 1 summarizes a theoretical comparison with other schemes in the literature (selected among the ones with similar succinctness).

Our first scheme, **LegoAC**, relies on an encoding of AC based on Hadamard products and linear constraints from [BCC⁺16] and can be built from CP_{lin} and CP_{had} gadgets. We evaluate two instantiations:

- **LegoAC1**—from our CP_{lin} and a CP_{had} from [Lip16]—is secure in the generic group model (GGM), enjoys constant-size proofs, and has a $\log n$ factor in proving time (similar to [PHGR13, Gro16]);
- **LegoAC2**—from our CP_{lin} and CP_{had} gadgets—is secure in the GGM and random oracle model, it has $\log n$ -size proofs but only *linear* proving time.

The second CP-SNARK, **LegoUAC**, builds on an encoding of AC based on Hadamard products, additions and permutation from [Gro09, BCG⁺17] and can be built from our CP_{had} and CP_{sfprm} gadgets.⁸ The main novelty of **LegoUAC** is to admit a *universal, circuit-independent* CRS, in the “specialization” model of [GKM⁺18] where the universal CRS can be specialized to a circuit C with a deterministic algorithm. **LegoUAC**’s CRS has $O(N)$ size where N is an upper bound on the number of gates of the circuits; in contrast, the CRS has quadratic size in the recent scheme in [GKM⁺18]. Our **LegoUAC** also improves on the approach applying an efficient system, say [Gro16], on a universal circuit [Val76, GKS17], which would incur at least a logarithmic multiplicative factor in circuit size.

1.2 Related Work

The idea of combining two different NIZKs to improve efficiency when handling heterogeneous computations has been considered by Chase et al. [CGM16] and more recently by Agrawal et

⁸ Additions are handled for free if the commitment is linearly homomorphic.

Scheme	Uni	KG time	Prove time	Ver. time	crs	\pi
[PHGR13, Gro16]	-	$n + m$	$m + n \log n$	$ x $	$n + m$	$O(1)$
LegoAC1	-	$n + m$	$n \log n$	$ x $	n	$O(1)$
LegoAC2	-	$n + m$	n	$ x + \log n$	n	$\log n$
[GKM ⁺ 18]	✓	n^2	$m + n \log n$	$ x $	n^2	$O(1)$
LegoUAC	✓	N^*	N	$ x + \log^2 N$	N^*	$\log^2 N$

Table 1: Comparing pairing-based zkSNARKs for arithmetic circuits with m wires and N gates, of which n are multiplication gates, n_a (resp. n_c) are addition (resp. multiplication-by-constant) gates, and $N^* = \max(n, n_a, n_c)$. Numbers in the table are in $O(\cdot)$ notation.

al. [AGM18]. In [AGM18], they propose combining the Pinocchio scheme [PHGR13] with Sigma-protocol-based NIZKs and show an efficient construction for computations that combines algebraic relations in a cryptographic group and arbitrary computation. Their approach reveals beneficial and improves performances. The solution in [AGM18] is tailored to two specific proof systems and their combination methodology does not always preserve succinctness. In contrast, our techniques are *general*, apply to a variety of existing proof systems and preserve succinctness (they compose succinct schemes into succinct schemes).

Succinct ZK Proofs. In the past years several research lines have built a variety of zk-SNARKs for general NP statements. Here we provide an overview of each line, especially focusing on their differences in performance.

A major research line is the one based on the seminal paper of Gennaro et al. [GGPR13] who proposed a pairing-based SNARK based on the NP-complete language of quadratic span/arithmetic programs. This approach improves on previous approaches by Ishai et al. [IKO07], Groth [Gro10] and Lipmaa [Lip12], and is the basis of several works such as [PHGR13, BCG⁺13, BFR⁺13, BCTV14, KPP⁺14, CFH⁺15, BBFR15, WSR⁺15, Gro16, FFG⁺16, GKM⁺18]. The zkSNARKs in this family enjoy constant-size proofs and fast verification, the latter depending only linearly on the statement size; on the downside, they feature large overheads at proving time, costly (although amortizable) preprocessing and security properties based on non-standard non-falsifiable assumptions.

A second research line builds on the MPC-in-the-head approach of Ishai et al. [IKOS07] to construct a ZK argument from an MPC protocol. The first scheme that refined and experimented this approach is ZKBoo [GMO16], then improved in [CDG⁺17]; a more recent work in this line is Ligerio [AHIV17]. These schemes do not need trusted setup and show excellent proving performances on Boolean circuits, since they rely only on symmetric-key cryptographic primitives. On the downside their proofs are not fully succinct, being linear in the circuit size $|C|$ in [GMO16], and $\tilde{O}(\sqrt{|C|})$ in [AHIV17].

The works [ZGK⁺17a, ZGK⁺17b, WTs⁺18] stem from the interactive proof techniques for low-depth circuits pioneered in Goldwasser et al. [GKR08] and later refined in [CMT12, Tha13, WJB⁺17]. The resulting succinct ZK arguments are made non-interactive in the random oracle model. These schemes offer good proving performance and use asymptotically fewer cryptographic operations than those from the MPC-in-the-head family; they can be instantiated without [WTs⁺18] (or with a circuit-independent [ZGK⁺17b]) trusted setup. On the other hand their proof size and verification time depend on the structure of the circuit at hand, notably on the depth and in some cases on the width.

Building on the work of Groth [Gro09], two recent proposals [BCC⁺16, BBB⁺17] give ZK arguments for arithmetic circuit satisfiability that can be instantiated without trusted setup. The first

scheme of Bootle et al. [BCC⁺16] has proofs of size $O(\sqrt{M})$ where M is the number of multiplication gates in the circuit, while their second scheme (improved in [BBB⁺17]) has proofs of size $O(\log M)$ but has a linear time verifier.

Compared to the results from the latter three research lines we described, our instantiations have the disadvantage of needing a trusted setup⁹, although in some cases this is universal and thus reusable. In terms of performances, however, our results are more succinct, both in terms of proof size and verifier time.

A recent line of work [BSBHR18] builds on the seminal works of Kilian [Kil92] and Micali [Mic94], and generalizations of PCPs (IOPs) [BCS16, RRR16] in order to construct systems (dubbed zkSTARKs) that are general-purpose (capturing very general computations that can be expressed as state-machine transitions), do not require trusted setup and offer good timings for prover and verifier. On the downside, the memory costs for the prover are still high and their security relies on a non-standard conjecture about Reed-Solomon codes.

1.3 Roadmap

The paper is organized as follows. Section 2 introduces notation and preliminar definitions. Section 3 provides the basis for building our framework: composing CP-SNARKs, the notion of cc-SNARKs and our compiler to import existing schemes in the framework. Sections 4 and 5 present constructions both for Pedersen-like commitments and polynomial commitments. Section 6 explains how to apply LegoSNARK to build schemes for arithmetic circuits. Section 7 gives experimental details of our library. We conclude in Section 8.

This text is the full work of our shorter version published at CCS'19. Several results only appear in this long version. Namely: security proofs, formal definitions, more schemes and constructions, and further details.

2 Preliminaries

We use $\lambda \in \mathbb{N}$ to denote the security parameter, and 1^λ to denote its unary representation. Throughout the paper we assume that all the algorithms of the cryptographic schemes take as input 1^λ , and thus we omit it from the list of inputs. For a distribution D , we denote by $x \leftarrow D$ the fact that x is being sampled according to D . We remind the reader that an ensemble $\mathcal{X} = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ is a family of probability distributions over a family of domains $\mathcal{D} = \{D_\lambda\}_{\lambda \in \mathbb{N}}$. We say two ensembles $\mathcal{D} = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}' = \{D'_\lambda\}_{\lambda \in \mathbb{N}}$ are statistically indistinguishable (denoted by $\mathcal{D} \approx_s \mathcal{D}'$) if $\frac{1}{2} \sum_x |D_\lambda(x) - D'_\lambda(x)| < \text{negl}(\lambda)$. If $\mathcal{A} = \{\mathcal{A}_\lambda\}$ is a (possibly non-uniform) family of circuits and $\mathcal{D} = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ is an ensemble, then we denote by $\mathcal{A}(\mathcal{D})$ the ensemble of the outputs of $\mathcal{A}_\lambda(x)$ when $x \leftarrow D_\lambda$. We say two ensembles $\mathcal{D} = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}' = \{D'_\lambda\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable (denoted by $\mathcal{D} \approx_c \mathcal{D}'$) if for every non-uniform polynomial time distinguisher \mathcal{A} we have $\mathcal{A}(\mathcal{D}) \approx_s \mathcal{A}(\mathcal{D}')$.

We denote by $[n]$ the set of integers $\{1, \dots, n\}$ and by $[:n]$ the set $\{0, 1, \dots, n-1\}$. By $(u_j)_{j \in [\ell]}$ we denote the tuple (u_1, \dots, u_ℓ) .

⁹ We stress that only our concrete instantiations require a trusted setup—our general composition framework does not.

2.1 Relations

Let $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of polynomial-time decidable relations R on pairs (x, w) where $x \in \mathcal{D}_x$ is called the *statement* or *input*, and $w \in \mathcal{D}_w$ the *witness*. We write $R(x, w) = 1$ to denote that R holds on (x, w) , else we write $R(x, w) = 0$. When discussing schemes that prove statements on committed values we assume that \mathcal{D}_w can be split in two subdomains $\mathcal{D}_u \times \mathcal{D}_\omega$. Finally we sometimes use an even finer grained specification of \mathcal{D}_u assuming we can split it over ℓ arbitrary domains $(\mathcal{D}_1 \times \dots \times \mathcal{D}_\ell)$ for some arity ℓ . In our security definitions we assume relations to be generated by a *relation generator* $\mathcal{RG}(1^\lambda)$ that, on input the security parameter 1^λ , outputs R together with some side information, an auxiliary input aux_R , that is given to the adversary. We define \mathcal{RG}_λ as the set of all relations that can be returned by $\mathcal{RG}(1^\lambda)$.

2.2 Commitment Schemes

We recall the notion of non-interactive commitment schemes.

Definition 2.1. *A commitment scheme is a tuple of algorithms $\text{Com} = (\text{Setup}, \text{Commit}, \text{VerCommit})$ that work as follows and satisfy the notions of correctness, binding and hiding defined below.*

- $\text{Setup}(1^\lambda) \rightarrow \text{ck}$ takes the security parameter and outputs a commitment key ck . This key includes descriptions of the input space \mathcal{D} , commitment space \mathcal{C} and opening space \mathcal{O} .
- $\text{Commit}(\text{ck}, u) \rightarrow (c, o)$ takes the commitment key ck and a value $u \in \mathcal{D}$, and outputs a commitment c and an opening o .
- $\text{VerCommit}(\text{ck}, c, u, o) \rightarrow b$ takes as input a commitment c , a value u and an opening o , and accepts ($b = 1$) or rejects ($b = 0$).

Correctness. *For all $\lambda \in \mathbb{N}$ and any input $u \in \mathcal{D}$ we have:*

$$\Pr [\text{ck} \leftarrow \text{Setup}(1^\lambda), (c, o) \leftarrow \text{Commit}(\text{ck}, u) : \text{VerCommit}(\text{ck}, c, u, o) = 1] = 1.$$

Binding. *For every polynomial-time adversary \mathcal{A} :*

$$\Pr \left[\begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda) \\ (c, u, o, u', o') \leftarrow \mathcal{A}(\text{ck}) \end{array} : \begin{array}{l} u \neq u' \wedge \text{VerCommit}(\text{ck}, c, u, o) = 1 \\ \wedge \text{VerCommit}(\text{ck}, c, u', o') = 1 \end{array} \right] = \text{negl}$$

Hiding. *For $\text{ck} \leftarrow \text{Setup}(1^\lambda)$ and $\forall u, u' \in \mathcal{D}$, the following two distributions are statistically close:*

$$\text{Commit}(\text{ck}, u) \approx \text{Commit}(\text{ck}, u')$$

2.3 Zero-Knowledge SNARKs

We recall the definition of (pre-processing) zero-knowledge succinct non-interactive arguments of knowledge (zkSNARKs, for short) [BCCT12, BCC⁺17].

Definition 2.2 (SNARK). *A SNARK for $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ is a triple of algorithms $\Pi = (\text{KeyGen}, \text{Prove}, \text{VerProof})$ that work as follows and satisfy the notions of completeness, succinctness and knowledge soundness defined below. If Π also satisfies zero-knowledge we call it a zkSNARK.*

- $\text{KeyGen}(R) \rightarrow (\text{ek}, \text{vk})$ takes the security parameter λ and a relation $R \in \mathcal{R}_\lambda$, and outputs a common reference string consisting of an evaluation and a verification key.
- $\text{Prove}(\text{ek}, x, w) \rightarrow \pi$ takes an evaluation key for a relation R , a statement x , and a witness w such that $R(x, w)$ holds, and returns a proof π .
- $\text{VerProof}(\text{vk}, x, \pi) \rightarrow b$ takes a verification key, a statement x , and either accepts ($b = 1$) or rejects ($b = 0$) the proof π .

Completeness. For any pair (x, w) satisfying the relation, the verifier always accepts the corresponding proof. Formally, $\forall \lambda \in \mathbb{N}$, $R \in \mathcal{R}_\lambda$ and (x, w) such that $R(x, w)$, it holds:

$$\Pr[(\text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R), \pi \leftarrow \text{Prove}(\text{ek}, x, w) : \text{VerProof}(\text{vk}, x, \pi) = 1] = 1$$

Succinctness. Π is said succinct if the running time of VerProof is $\text{poly}(\lambda)(\lambda + |x| + \log |w|)$ and the proof size is $\text{poly}(\lambda)(\lambda + \log |w|)$.

Knowledge Soundness. Let \mathcal{RG} be a relation generator such that $\mathcal{RG}_\lambda \subseteq \mathcal{R}_\lambda$. Π has knowledge soundness for \mathcal{RG} and auxiliary input distribution \mathcal{Z} , denoted $\text{KSND}(\mathcal{RG}, \mathcal{Z})$ for brevity, if for every (non-uniform) efficient adversary \mathcal{A} there exists a (non-uniform) efficient extractor \mathcal{E} such that $\Pr[\text{Game}_{\mathcal{RG}, \mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{KSND}} = 1] = \text{negl}$. We say that Π is knowledge sound if there exists benign \mathcal{RG} and \mathcal{Z} such that Π is $\text{KSND}(\mathcal{RG}, \mathcal{Z})$.

$$\frac{\text{Game}_{\mathcal{RG}, \mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{KSND}} \rightarrow b}{(R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda) \ ; \ \text{crs} := (\text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R) \ ; \ \text{aux}_Z \leftarrow \mathcal{Z}(R, \text{aux}_R, \text{crs}) \\ (x, \pi) \leftarrow \mathcal{A}(R, \text{crs}, \text{aux}_R, \text{aux}_Z) \ ; \ w \leftarrow \mathcal{E}(R, \text{crs}, \text{aux}_R, \text{aux}_Z) \ ; \ b = \text{VerProof}(\text{vk}, x, \pi) \wedge \neg R(x, w)}$$

Composable Zero-Knowledge. A scheme Π satisfies composable zero-knowledge for a relation generator \mathcal{RG} if there exists a simulator $\mathcal{S} = (\mathcal{S}_{\text{kg}}, \mathcal{S}_{\text{prv}})$ such that both following conditions hold for all adversaries \mathcal{A} :

KEYS INDISTINGUISHABILITY.

$$\begin{aligned} & \Pr [(R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), \text{crs} \leftarrow \text{KeyGen}(R) : \mathcal{A}(\text{crs}, \text{aux}_R) = 1] \\ & \approx \Pr [(R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R) : \mathcal{A}(\text{crs}, \text{aux}_R) = 1] \end{aligned}$$

PROOF INDISTINGUISHABILITY. For all (x, w) such that $R(x, w) = 1$,

$$\begin{aligned} & \Pr [(R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R) : \pi \leftarrow \text{Prove}(\text{ek}, x, w), \mathcal{A}(\text{crs}, \text{aux}_R, \pi) = 1] \\ & \approx \Pr [(R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R) : \pi \leftarrow \mathcal{S}_{\text{prv}}(\text{crs}, \text{td}_k, x), \mathcal{A}(\text{crs}, \text{aux}_R, \pi) = 1] \end{aligned}$$

Remark 2.1. In the notion of knowledge soundness defined above we consider two kinds of auxiliary inputs, aux_R generated together with the relation by \mathcal{RG} , and aux_Z that is generated from some distribution \mathcal{Z} that may depend on the common reference string that in turn depends on R . An example of this appears in our proof of Theorem B.1. Notice that although our notion is implied by a notion where auxiliary inputs can be arbitrary, our aim is a precise formalization of auxiliary inputs; this is useful to justify why certain auxiliary inputs should be considered benign, as required to avoid known impossibility results [BCPR14, BP15]. Finally, we also note that our notion is also implied by SNARKs that admit black-box extractors (as may be the case for those relying on random oracles [Mic00]).

zkSNARKs with Specializable Universal CRS

In the SNARK notion presented above, the common reference string generated by `KeyGen` is tied to a specific relation $R \in \mathcal{R}_\lambda$. A variant of this notion is that of SNARKs for universal relations in which the output of `KeyGen` depends only on the family \mathcal{R}_λ and can be used to prove and verify statements about any $R \in \mathcal{R}_\lambda$. Due to the practical concerns on the execution of `KeyGen`, SNARKs for universal relations are more convenient as one can reuse and amortize the cost of one setup. In a recent work, Groth et al. [GKM⁺18] introduced the notion of *zkSNARK with specializable universal common reference string*. In a nutshell, this notion formalizes the idea that key generation for R can be seen as the sequential combination of two steps: a first probabilistic algorithm that generates a CRS for the universal relation, and a second *deterministic* algorithm that specializes this universal CRS into one for a specific R . We remark that our UC SNARKs follow this model.

More formally, let \mathcal{R}_λ be a family of relations. The universal relation R^* for \mathcal{R}_λ defines a language with instances (R, x) such that $R^*(R, x, w)$ holds iff $R \in \mathcal{R}_\lambda$ and $R(x, w)$ holds.

A $\Pi = (\text{KeyGen}, \text{Prove}, \text{VerProof})$ is said a *zkSNARK with specializable universal common reference string* [GKM⁺18] if there exist algorithms `Derive`, `Prove*`, `VerProof*` such that:

- `Derive(crs, R) → crsR` is a *deterministic* algorithm that takes as input a $\text{crs} := (\text{ek}, \text{vk})$ produced by `KeyGen(R*)` and a relation $R \in \mathcal{R}_\lambda$, and outputs a specialized common reference string $\text{crs}_R := (\text{ek}_R, \text{vk}_R)$.
- `Prove(ek, (R, x), w) → π` runs $(\text{ek}_R, \text{vk}_R) \leftarrow \text{Derive}(\text{crs}, R)$ and returns $\pi \leftarrow \text{Prove}^*(\text{ek}_R, x, w)$.
- `VerProof(vk, (R, x), π) → b` runs $(\text{ek}_R, \text{vk}_R) \leftarrow \text{Derive}(\text{crs}, R)$ and returns $b \leftarrow \text{VerProof}^*(\text{vk}_R, x, \pi)$.

3 Building the LegoSNARK Framework

3.1 Commit and Prove SNARKs

In a nutshell, a *commit-and-prove SNARK* (CP-SNARK) is a SNARK that can prove knowledge of (x, w) such that $R(x, w)$ holds w.r.t. a witness $w = (u, \omega)$ and u opens a commitment c_u .¹⁰ Our formal definitions below add some syntactic sugar to this idea to explicitly handle relations where the input domain \mathcal{D}_u is more fine grained and splits over ℓ subdomains. For reasons that will shortly become clear, we call these subdomains *commitment slots*. This splitting is often natural (e.g., if u is a binary string, one can think of $u := (u_1, \dots, u_\ell)$ for suitable substrings), and it is crucial to exploit the compositional power of CP-SNARKs, as we show in Section 3.2. We assume the description of the splitting is part of R 's description.

Definition 3.1 (CP-SNARKs). Let $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of relations R over $\mathcal{D}_x \times \mathcal{D}_u \times \mathcal{D}_\omega$ such that \mathcal{D}_u splits over ℓ arbitrary domains $(\mathcal{D}_1 \times \dots \times \mathcal{D}_\ell)$ for some arity parameter $\ell \geq 1$. Let $\text{Com} = (\text{Setup}, \text{Commit}, \text{VerCommit})$ be a commitment scheme (as per Definition 2.1) whose input space \mathcal{D} is such that $\mathcal{D}_i \subset \mathcal{D}$ for all $i \in [\ell]$. A commit and prove zkSNARK for Com and $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ is a zkSNARK for a family of relations $\{\mathcal{R}_\lambda^{\text{Com}}\}_{\lambda \in \mathbb{N}}$ such that:

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

$$\bigwedge_{j \in [\ell]} \text{VerCommit}(\text{ck}, c_j, u_j, o_j) = 1 \wedge R(x, (u_j)_{j \in [\ell]}, \omega) = 1$$

¹⁰ Our notion assumes that only a portion of the witness is explicitly committed in c_u .

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

We denote a CP-SNARK as a triple of algorithms $\text{CP} = (\text{KeyGen}, \text{Prove}, \text{VerProof})$. For ease of exposition, in our constructions we adopt the syntax for CP's algorithms defined below.

- $\text{KeyGen}(\text{ck}, R) \rightarrow \text{crs} := (\text{ek}, \text{vk})$ generates the common reference string.
- $\text{Prove}(\text{ek}, x, (c_j)_{j \in [\ell]}, (u_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, \omega) \rightarrow \pi$ outputs the proof of correct commitment.
- $\text{VerProof}(\text{vk}, x, (c_j)_{j \in [\ell]}, \pi) \rightarrow b \in \{0, 1\}$ rejects or accepts the proof.

Remark 3.1 (Comparing with existing definitions). To define the Geppetto scheme [CFH⁺15] the authors define a notion of commit-and-prove SNARKs. Here we highlight the main differences between their definition and ours. First, our commitment key can be generated without fixing a priori a relation (or a set of relations, e.g., a multi-QAP). Second, in their model one needs to commit to data using a commitment key corresponding to a specific portion of the input (in their lingo a “bank”), whereas in our model one can just commit to a vector of data, and only at proving time one assigns that data to a specific input slot. Third, we do not require commitments to have a trapdoor. Our notion is closer to the one given by Lipmaa [Lip16] (although [Lip16] uses trapdoor commitments) and is in fact a specialized SNARK notion when considering relation families including verifying openings of commitments.

3.2 Composition Properties of CP-SNARKs

In this section, we formally show how the commit-and-prove capability can be used to combine different CP-SNARKs securely.

Conjunction of relations with shared inputs. Let $\{\mathcal{R}_\lambda^{(0)}\}_{\lambda \in \mathbb{N}}$ and $\{\mathcal{R}_\lambda^{(1)}\}_{\lambda \in \mathbb{N}}$ be two families of relations such that, for every $\lambda \in \mathbb{N}$ the input domains $\mathcal{D}_u^{(0)}$ and $\mathcal{D}_u^{(1)}$ of relations $R_0 \in \mathcal{R}_\lambda^{(0)}$ and $R_1 \in \mathcal{R}_\lambda^{(1)}$ respectively can split as follows: $\mathcal{D}_u^{(0)} := \mathcal{D}_0 \times \mathcal{D}_2$ and $\mathcal{D}_u^{(1)} := \mathcal{D}_1 \times \mathcal{D}'_2$ with $\mathcal{D}_2 = \mathcal{D}'_2$.¹¹ In other words we require these relations to share a commitment slot that we call the *shared slot*.

Given the above relation families, we define $\{\mathcal{R}_\lambda^\wedge\}_{\lambda \in \mathbb{N}}$ as the family of relations where for every $\lambda \in \mathbb{N}$, $\mathcal{R}_\lambda^\wedge = \{R_{R_0, R_1}^\wedge : R_0 \in \mathcal{R}_\lambda^{(0)}, R_1 \in \mathcal{R}_\lambda^{(1)}\}$ and $R_{R_0, R_1}^\wedge(x_0, x_1, u_0, u_1, u_2, w^*)$ is defined as follows:

$$R_{R_0, R_1}^\wedge(x_0, x_1, u_0, u_1, u_2, (w_0, w_1)) := R_0(x_0, u_0, u_2, w_0) \wedge R_1(x_1, u_1, u_2, w_1)$$

Let Com be a commitment scheme, for $b \in \{0, 1\}$ let CP_b be a CP-SNARK for Com and $\{\mathcal{R}_\lambda^{(b)}\}_{\lambda \in \mathbb{N}}$. In Figure 1 we show a construction of a CP-SNARK CP^\wedge for Com and $\{\mathcal{R}_\lambda^\wedge\}_{\lambda \in \mathbb{N}}$. It is also easy to see that if both CP_0 and CP_1 are CP-SNARKs with specializable universal CRS, then so is the resulting CP^\wedge .

Theorem 3.1. *If Com is a computationally binding commitment and, for $b \in \{0, 1\}$, CP_b is a zero-knowledge CP-SNARK for Com and relation family $\{\mathcal{R}_\lambda^{(b)}\}_{\lambda \in \mathbb{N}}$, then there is a zero-knowledge CP-SNARK CP^\wedge for Com and $\{\mathcal{R}_\lambda^\wedge\}_{\lambda \in \mathbb{N}}$.*

¹¹ Note such a splitting is rather general, as \mathcal{D}_2 and \mathcal{D}'_2 , or \mathcal{D}_0 , or \mathcal{D}_1 may be empty.

$\text{CP}^\wedge.\text{KeyGen}(\text{ck}, R_{R_0, R_1}^\wedge) :$	$\text{CP}^\wedge.\text{Prove}(\text{ek}^*, x_0, x_1, (c_j)_{j \in [3]}, (u_j)_{j \in [3]}, (o_j)_{j \in [3]}, \omega_0, \omega_1) :$	$\text{CP}^\wedge.\text{VerProof}(\text{vk}^*, x_0, x_1, (c_j)_{j \in [3]}, \pi^*) :$
$(\text{ek}_0, \text{vk}_0) \leftarrow \text{CP}_0.\text{KeyGen}(\text{ck}, R_0)$	$\pi_0 \leftarrow \text{CP}_0.\text{Prove}(\text{ek}_0, x_0, (c_0, c_2), (u_0, u_2), (o_0, o_2), \omega_0)$	$b_{\text{ok}}^{(0)} \leftarrow \text{CP}_0.\text{VerProof}(\text{vk}_0, x_0, (c_0, c_2), \pi_0)$
$(\text{ek}_1, \text{vk}_1) \leftarrow \text{CP}_1.\text{KeyGen}(\text{ck}, R_1)$	$\pi_1 \leftarrow \text{CP}_1.\text{Prove}(\text{ek}_1, x_1, (c_1, c_2), (u_1, u_2), (o_1, o_2), \omega_1)$	$b_{\text{ok}}^{(1)} \leftarrow \text{CP}_1.\text{VerProof}(\text{vk}_1, x_1, (c_1, c_2), \pi_1)$
$\text{ek}^* := (\text{ek}_b)_{b \in \{0,1\}}$	$\text{return } \pi^* := (\pi_b)_{b \in \{0,1\}}$	return $b_{\text{ok}}^{(0)} \wedge b_{\text{ok}}^{(1)}$
$\text{vk}^* := (\text{vk}_b)_{b \in \{0,1\}}$		
return $(\text{ek}^*, \text{vk}^*)$		

Figure 1: CP-SNARK construction for AND composition

Correctness and succinctness follow by inspection. Knowledge-soundness and zero-knowledge follow rather easily from the respective properties of the underlying schemes. In particular, for knowledge-soundness the basic idea is that in order for an adversary to break CP^\wedge it must break either one of the two underlying schemes, CP_0 , CP_1 , or the binding of the commitment scheme. We give a full proof of knowledge-soundness and zero-knowledge in Appendix A.

Functions composition. A CP-SNARK for conjunction of relations can be easily used for proving correctness of *composed functions*, e.g., proving that $\exists(y, w) : z = f(x, y, w)$, where $f(x, y, w) := h(g(x, w), y)$. Indeed, let $R_h(x', y, z) = 1$ iff $\exists(x', y) : h(x', y) = z$, and $R_g(x, x') = 1$ iff $\exists(x', w) : g(x, w) = x'$, then $\exists(y, w) : z = f(x, y, w)$ can be expressed as $R_h(x', y, z) \wedge R_g(x, x')$.

Disjunction of relations with shared inputs. We can reduce the case of OR composition to the conjunction construction above. For this we assume relations are defined over elements of a ring. For a relation $R(u)$ denote by $\hat{R}(u, t)$ the relation such that $\hat{R}(u, 0) = 1$ iff $R(u) = 1$ and $\hat{R}(u, t) = 1$ iff $R(u) = 0$ whenever $t \neq 0$. We can now express the disjunction of $R_0(u_0)$, $R_1(u_1)$ as $R_{R_0, R_1}^\vee(u_0, u_1, t_0, t_1) := \hat{R}_0(u_0, t_0) \wedge \hat{R}_1(u_1, t_1) \wedge t_0 t_1 = 0$. For this approach to work we need the proof systems for the two relations R_0, R_1 to support their modified version \hat{R}_0, \hat{R}_1 , which is the case for proof systems supporting general arithmetic or boolean circuits. Finally, we need a simple efficient proof system for the relation $R_{\text{mul}}(t_0, t_1) = 1$ iff $t_0 \cdot t_1 = 0$, where both t_0 and t_1 are committed in two different slots.

Composing more than two relations. By iterating the application of our Theorem 3.1 we can build CP-SNARKs that handle conjunctions and/or disjunctions of more than two relations. In order to maintain the succinctness property, one should apply composition only a small (e.g., constant, logarithmic) number of times. However, this is arguably the case when we deal with real-world heterogeneous computations. The following example scenarios consider heterogeneous computations that can be split naturally into two “homogeneous” components: square-and-multiply algorithms (splitting the relation into the conjunction of all the iterated squarings and the final inner product), aggregation queries to a database (that can be split in a “filter” and an “aggregate” component), proving a property P for a datum in a Merkle tree, as done in Zcash [BCG⁺14] (that can be split in a membership verification component and the property P , which could in turn be decomposed further).

3.3 Commit-Carrying SNARKs

In this section we define a variant of SNARKs that lies in between standard SNARKs and CP-SNARKs. We call these schemes *SNARKs with commit-carrying proofs* (or commit-carrying SNARKs,

cc-SNARKs for short). In a nutshell, a cc-SNARK is like a SNARK in which the proof contains a commitment to the portion u of the witness. Essentially the difference is that in cc-SNARKs we assume the extractor outputs the opening of the commitment returned along with the proof. Formalizing this idea requires to make explicit the commitment scheme associated to the SNARK, as well as the commitment key that is part of the common reference string. In the next section we discuss how many of the existing SNARK constructions satisfy this property. Later, in Section 3.5 we show that cc-SNARKs can be lifted to become full fledged, composable, CP-SNARKs. These two results together allow us to compose several existing SNARKs. We define commit-carrying SNARKs as follows:

Definition 3.2 (cc-SNARK). A commit-carrying zkSNARKs for $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ is a tuple of algorithms $\text{ccII} = (\text{KeyGen}, \text{Prove}, \text{VerProof}, \text{VerCommit})$ that work as follows and satisfy the notions of completeness, succinctness, knowledge soundness, zero knowledge and binding as defined below.

- $\text{KeyGen}(R) \rightarrow (\text{ck}, \text{ek}, \text{vk})$: the key generation takes as input the security parameter λ and a relation $R \in \mathcal{R}_\lambda$, and outputs a common reference string that includes a commitment key, an evaluation key and verification key.
- $\text{Prove}(\text{ek}, x, w) \rightarrow (c, \pi; o)$: the proving algorithm takes as input an evaluation key, a statement x and a witness $w := (u, \omega)$ such that the relation $R(x, u, \omega)$ holds, and it outputs a proof π , a commitment c and opening o such that $\text{VerCommit}(\text{ck}, c, u, o) = 1$.
- $\text{VerProof}(\text{vk}, x, c, \pi) \rightarrow b$: the verification algorithm takes a verification key, a statement x , a commitment c , and either accepts ($b = 1$) or rejects ($b = 0$) the proof π .
- $\text{VerCommit}(\text{ck}, c, u, o) \rightarrow b$: the commitment verification algorithm takes as input a commitment key, a commitment c , a message u and an opening o and accepts ($b = 1$) or rejects ($b = 0$).

cc-SNARKs can be seen as a less versatile version of CP-SNARKs (clearly, a CP-SNARK implies a cc-SNARK). In a cc-SNARK the commitment key depends on the relation taken by KeyGen , and a commitment is freshly created by the Prove algorithm and is tied to a single proof; in a CP-SNARK the commitment key is independent of relations and commitments can also be created independently and shared across different proofs. Furthermore, in the literature, there are examples of schemes that lie in between our notions of CP-SNARK and cc-SNARK; this is the case for commit and prove SNARKs in which the commitment key is relation-dependent, e.g., [CFH⁺15, Vee17].

Completeness. For any $\lambda \in \mathbb{N}$, $R \in \mathcal{R}_\lambda$ and (x, w) such that $R(x, w) = 1$, it holds

$$\Pr((\text{ck}, \text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R), (c, \pi; o) \leftarrow \text{Prove}(\text{ek}, x, w) : \text{VerProof}(\text{vk}, x, c, \pi)) = 1$$

Succinctness. ccII is said *succinct* if the running time of VerProof is $\text{poly}(\lambda) (\lambda + |x| + \log |w|)$ and the size of the proof is $\text{poly}(\lambda) \cdot (\lambda + \log |w|)$.

Knowledge Soundness. Let \mathcal{RG} be a relation generator such that $\mathcal{RG}_\lambda \subseteq \mathcal{R}_\lambda$. ccII satisfies knowledge soundness for \mathcal{RG} and auxiliary input distribution \mathcal{Z} , or $\text{ccKSND}(\mathcal{RG}, \mathcal{Z})$, if for every (non-uniform) efficient adversary \mathcal{A} there exists a (non-uniform) efficient extractor \mathcal{E} such that $\Pr[\text{Game}_{\mathcal{RG}, \mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{ccKSND}} = 1] = \text{negl}$. We say that ccII is knowledge sound if there exist benign \mathcal{RG} and \mathcal{Z} such that ccII is $\text{ccKSND}(\mathcal{RG}, \mathcal{Z})$.

$$\text{Game}_{\mathcal{RG}, \mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{ccKSNd}} \rightarrow b \in \{0, 1\}$$

$$(R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda)$$

$$\text{crs} := (\text{ck}, \text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R)$$

$$\text{aux}_Z \leftarrow \mathcal{Z}(R, \text{aux}_R, \text{crs}) \quad (x, c, \pi) \leftarrow \mathcal{A}(R, \text{crs}, \text{aux}_R, \text{aux}_Z) \quad (u, o, \omega) \leftarrow \mathcal{E}(R, \text{crs}, \text{aux}_R, \text{aux}_Z)$$

$$b \leftarrow \text{VerProof}(\text{vk}, x, c, \pi) = 1 \wedge (\text{VerCommit}(\text{ck}, c, u, o) = 0 \vee R(x, u, \omega) = 0)$$

Composable Zero-Knowledge. A scheme ccII has composable zero-knowledge for a relation generator \mathcal{RG} if for every adversary \mathcal{A} there exists a simulator $\mathcal{S} = (\mathcal{S}_{\text{kg}}, \mathcal{S}_{\text{prv}})$ such that both following conditions hold for all adversaries \mathcal{A} :

KEYS INDISTINGUISHABILITY.

$$\begin{aligned} & \Pr \left((R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), \text{crs} \leftarrow \text{KeyGen}(R) : \mathcal{A}(\text{crs}, \text{aux}_R) = 1 \right) \\ & \approx \Pr \left((R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R) : \mathcal{A}(\text{crs}, \text{aux}_R) = 1 \right) \end{aligned}$$

PROOF INDISTINGUISHABILITY.

$$\begin{aligned} \forall (x, w) : \Pr & \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R), (c, \pi; o) \leftarrow \text{Prove}(\text{ek}, x, w) \\ \mathcal{A}(\text{crs}, \text{aux}_R, c, \pi) = 1 \wedge R(x, w) = 1 \end{array} : \right] \\ & \approx \Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R), (c, \pi) \leftarrow \mathcal{S}_{\text{prv}}(\text{crs}, \text{td}_k, x) \\ \mathcal{A}(\text{crs}, \text{aux}_R, c, \pi) = 1 \wedge R(x, w) = 1 \end{array} : \right] \end{aligned}$$

Binding. For every polynomial-time adversary \mathcal{A} the following probability is $\text{negl}(\lambda)$:

$$\Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), \text{crs} := (\text{ck}, \text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R) \\ (c, u, o, u', o') \leftarrow \mathcal{A}(R, \text{crs}, \text{aux}_R) \end{array} : \begin{array}{l} \text{VerCommit}(\text{ck}, c, u', o') \wedge \\ \text{VerCommit}(\text{ck}, c, u, o) \wedge u \neq u' \end{array} \right]$$

Remark 3.2. While our definitions consider the case where the proof contains a commitment to a portion u of the witness $w = (u, \omega)$, notice that this partition of the witness is arbitrary and thus this notion also captures those constructions where the commitment is to the entire witness if one thinks of a void ω .

cc-SNARKs with Weak Binding

Let us now define a weaker variant of cc-SNARKs that differs from the one given in Definition 3.2 in that the underlying commitment scheme is not binding in the usual sense. Slightly more in detail, we consider the case where the commitment refers to the whole witness (i.e., ω is an empty string) and it is actually possible to find collisions for a given commitment as long as these collisions are among valid witnesses, or more precisely we require to be computationally infeasible to find two different witnesses that validly open the commitment such that one falsifies the relation and the other one satisfies it. Worth noting that our generic compiler can also turn weak cc-SNARKs into CP-SNARKs.

Definition 3.3 (cc-SNARKs with Weak Binding). We define cc-SNARKs with Weak Binding as in Definition 3.2 with two exceptions: we assume that the scheme is defined only for relations such that $\mathcal{D}_\omega = \emptyset$; we replace the binding property with the one below.

Weak Binding. \forall polynomial-time adversary \mathcal{A} and $u \neq u'$ the following probability is $\text{negl}(\lambda)$:

$$\Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda), \text{crs} := (\text{ck}, \text{ek}, \text{vk}) \leftarrow \text{KG}(R), (x, c, u, o, u', o', \pi) \leftarrow \mathcal{A}(R, \text{crs}, \text{aux}_R) \\ \text{VerCommit}(\text{ck}, c, u, o) \wedge \text{VerCommit}(\text{ck}, c, u', o') \wedge \text{VerProof}(\text{vk}, x, c, \pi) \wedge \neg R(x, u) \wedge R(x, u') \end{array} : \right]$$

3.4 Existing CP-SNARKs and cc-SNARKs

In this section, we provide a summary of existing schemes that can be explained, with no or little modification, under our CP-SNARK and cc-SNARK notions. In fact, existing QAP-based schemes [PHGR13, BCTV14, Gro16] are not fully binding but can satisfy our weak binding. In Appendix H.5 we prove that [Gro16] is a *weak* cc-SNARK.

Existing CP-SNARKs. The following list is a summary. Details supporting the following claims appear in Appendix H.

- Adaptive Pinocchio [Vee17] is a CP-SNARK for relations $R_Q(x, (u_j)_{j \in [\ell]}, \omega)$ where R_Q is a quadratic arithmetic program (QAP), and the commitment scheme is the extended Pedersen commitment of Groth [Gro10] in which the i -th basis of the commitment key is g^{x^i} for a random x .
- The scheme in [Lip16][Section 4] is a CP-SNARK for Hadamard product relations $R^{\text{had}}(\mathbf{a}, \mathbf{b}, \mathbf{c})$ over \mathbb{Z}_q^{3m} , i.e. R^{had} holds iff $\forall i \in [m] : a_i \cdot b_i = c_i$. In this case the commitment scheme is a variant of the extended Pedersen scheme where the i th basis of the commitment key is $g^{\ell_i(x)}$ for a random x and ℓ_i being the i -th Lagrange basis polynomial.
- zk-vSQL [ZGK⁺17b] is a CP-SNARK for relations $R((u_j)_{j \in [\ell]})$ where R is an arithmetic circuit, and the commitment is a polynomial commitment that, we observe (cf. Appendix H), can also be explained as a variant of extended Pedersen.

Existing cc-SNARKs. Geppetto [CFH⁺15] is a commit-and-prove SNARK for QAP relations $R_Q(x, u, \omega)$, with a *relation-dependent* commitment key. This scheme immediately yields a cc-SNARK where VerCommit is also a variant of extended Pedersen.

Existing Weak cc-SNARKs. There exist other schemes in the literature that fit the cc-SNARK syntax, but fail to satisfy the binding property. This is the case for some QAP-based schemes, such as Pinocchio [PHGR13, BCTV14] or the efficient SNARK of Groth [Gro16]. For the latter [Gro16] we prove in Appendix H.5 that it is a *weak* cc-SNARK for QAP relations $R_Q(u)$ QAP.¹² Worth noting that our generic compiler in the next section allows to turn also weak cc-SNARKs into CP-SNARKs.

A New Efficient cc-SNARK for QAPs. We show that the SNARK of [Gro16] can be modified to obtain a cc-SNARK for QAP relations $R_Q(u, \omega)$, where the witness portion committed in a fully binding way can be chosen (see Appendix H.5). Compared to the other cc-SNARKs for QAPs mentioned above, this scheme offers nearly optimal efficiency (essentially due to the fact that we start from [Gro16] whereas [CFH⁺15, Vee17] build on [PHGR13]).

¹² Using similar ideas we believe that such a result also holds for the Pinocchio variant in [BCTV14].

3.5 Bootstrapping our Framework

A key requirement to apply the composition results of the LegoSNARK framework is to start from CP-SNARKs that share the same commitment scheme. In practice this is not always the case (see for example the discussion in the previous section). In this section we propose a solution to this issue by giving a generic compiler for turning a cc-SNARK $ccII$ for a family of relations $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ into a CP-SNARK CP that supports the same relations and works for a given, global, commitment scheme Com. Incidentally, since a CP-SNARK CP for commitment Com' is also a cc-SNARK, our compiler can also turn CP into a CP-SNARK for another commitment Com.

As noted in the introduction one could solve the interoperability problem if the cc-SNARK (or even any SNARK) is sufficiently expressive so as to encode the commitment verification algorithm VerCommit in its relations (e.g., as a circuit). This approach of letting the SNARK take care of the commitment verification however has two main drawbacks. First, recall that in our vision, the cc-SNARK $ccII$ may be a proof system for a specialized task, and thus may not be able to express VerCommit in its language. Second, even if $ccII$ is expressive enough, such encodings of VerCommit (for various choices of schemes) are notoriously very expensive. Our approach to deal with this issue is to propose a slightly different methodology that shifts the problem of expressing a relation about VerCommit from $ccII$ to a CP-SNARK that is tailored to this problem. Our idea in brief: linking a proof-dependent commitment c' from $ccII$ to a general-purpose commitment c from a CP-SNARK. Specifically we rely on a CP-SNARK (from now on CP_{link}) able to prove that the two commitments, c' and c (actually a collection of c_j), open to the same value. In other words CP_{link} is a *minimal* tool able to turn a $ccII$ into a full fledged CP-SNARK CP that supports some general purpose commitment. The fact we require CP-SNARK to create a CP-SNARK is a curious aspect of this approach. What we require however is less than what we get: we only need to start from a *simple* scheme CP_{link} that handles a *specific* relation to create CP-SNARKs for *disparate* families of relations. Since CP_{link} is a simple object we can obtain from it efficient instantiations (as confirmed by our concrete construction proposed in Section 4.1).

Our cc-SNARK-lifting compiler. Let $ccII$ be a cc-SNARK for a family of relations $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ where, for every λ , $R \in \mathcal{R}_\lambda$ is over tuples in $\mathcal{D}_x \times \mathcal{D}_u \times \mathcal{D}_\omega$, and \mathcal{D}_u splits over ℓ subdomains ($\mathcal{D}_1 \times \dots \times \mathcal{D}_\ell$) for some arity parameter ℓ . Consider the commitment verification algorithm $ccII.VerCommit$. For any $\lambda \in \mathbb{N}$ and any $ck' \in \{ccII.KeyGen(R)\}_{R \in \mathcal{R}_\lambda}$, we define the relation R^{link} that has input space $\mathcal{D}_x^{\text{link}} = \mathcal{C}'$, and witness space $\mathcal{D}_\omega^{\text{link}} = \mathcal{D}_u^{\text{link}} \times \mathcal{D}_\omega^{\text{link}}$ such that $\mathcal{D}_u^{\text{link}} = \mathcal{D}_1 \times \dots \times \mathcal{D}_\ell$ and $\mathcal{D}_\omega^{\text{link}} := \mathcal{O}'$, where \mathcal{C}' and \mathcal{O}' are the commitment and opening space of the commitment of $ccII$. For compactness we represent R^{link} with $(ck', \mathcal{D}_x^{\text{link}}, \mathcal{D}_u^{\text{link}}, \mathcal{D}_\omega^{\text{link}})$. Then, R^{link} is defined as follows:

$$R^{\text{link}}(x^{\text{link}}, (u_j^{\text{link}})_{j \in [\ell]}, \omega^{\text{link}}) := ccII.VerCommit(ck', x^{\text{link}}, (u_j^{\text{link}})_{j \in [\ell]}, \omega^{\text{link}})$$

We remark that, above, $x^{\text{link}} \in \mathcal{C}'$ is a commitment for $ccII.VerCommit$ and $\omega^{\text{link}} \in \mathcal{O}'$ is its opening.

Let CP_{link} be a CP-SNARK for Com and a family of relations $\{\mathcal{R}_\lambda^{\text{link}}\}_{\lambda \in \mathbb{N}}$ such that for every $\lambda \in \mathbb{N}$ the relation R^{link} defined above is in $\mathcal{R}_\lambda^{\text{link}}$. In Table 2 we describe a CP-SNARK CP for $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ that works by using $ccII$ and CP_{link} .

The correctness of CP follows by that of the two schemes CP_{link} and $ccII$. The same holds for succinctness. In the following theorem we state how knowledge soundness and zero-knowledge of CP follow from the corresponding properties of CP_{link} and $ccII$. The formal statement appears in Appendix B, and proofs appear in Appendix B.1 and B.2 respectively.

$\text{CP.KeyGen}(\text{ck}, R) \rightarrow (\text{ek}, \text{vk})$	$\text{CP.Prove}(\text{ek}, x, (c_j, u_j, o_j)_{j \in [\ell]}, \omega) \rightarrow \pi := (c', \pi^{\text{link}}, \pi')$
$(\text{ck}', \text{ek}', \text{vk}') \leftarrow \text{ccII.KeyGen}(R)$	$(c', \pi', o') \leftarrow \text{ccII.Prove}(\text{ek}', x, (u_j)_{j \in [\ell]}; \omega); (x^{\text{link}}, \omega^{\text{link}}) := (c', o')$
$\text{Build } R^{\text{link}} \text{ from } (\text{ck}', \mathcal{D}_x^{\text{link}}, \mathcal{D}_u^{\text{link}}, \mathcal{D}_\omega^{\text{link}})$	$\pi^{\text{link}} \leftarrow \text{CP}_{\text{link}}.\text{Prove}(\text{ek}^{\text{link}}, x^{\text{link}}, (c_j)_{j \in [\ell]}, (u_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, \omega^{\text{link}})$
$(\text{ek}^{\text{link}}, \text{vk}^{\text{link}}) \leftarrow \text{CP}_{\text{link}}.\text{KeyGen}(\text{ck}, R^{\text{link}})$	$\text{CP.VerProof}(\text{vk}, x, (c_j)_{j \in [\ell]}, \pi) \rightarrow \{0, 1\}$
return $((\text{ck}', \text{ek}', \text{ek}^{\text{link}}), (\text{vk}', \text{vk}^{\text{link}}))$	$\text{CP}_{\text{link}}.\text{VerProof}(\text{vk}^{\text{link}}, c', (c_j)_{j \in [\ell]}, \pi^{\text{link}}) \wedge \text{ccII.VerProof}(\text{vk}', x, c', \pi')$

Figure 2: Generic Construction of CP from CP_{link} and ccII .

Theorem 3.2. *If ccII is a zk-cc-SNARK (or a weak cc-SNARK) for $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ and CP_{link} is a zk-CP-SNARK for $\{\mathcal{R}_\lambda^{\text{link}}\}_{\lambda \in \mathbb{N}}$, then the scheme CP in Figure 2 is a zk-CP-SNARK for $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$.*

4 CP-SNARKs for Pedersen-like Commitments

In this section we propose two CP-SNARKs that work for any commitment scheme whose verification algorithm is the same as the extended Pedersen commitment (essentially a multi-exponentiation). This class of commitments includes those underlying several existing SNARKs, such as *all* the ones we mentioned in section 3.4. Notable, this also includes the “classical” extension of Pedersen whose key is a set of random group elements, which can be sampled in a transparent way; in other words no trusted setup is needed for this commitment key.¹³

For vectors committed in this way, we show two schemes. Our first scheme (given in Section 4.1) allows to prove that another commitment, with the same verification algorithm but different key, opens to the same vector. This is essentially an efficient realization of the CP_{link} CP-SNARK needed in our compiler of Section 3.5, and that works for cc-SNARKs whose underlying commitment verification has the same structure as Pedersen. Our second scheme (given in Section 4.2) instead allows one to prove correctness of a linear function of the committed vector (i.e., that $\mathbf{x} = \mathbf{F} \cdot \mathbf{u}$).

In what follows we start by recalling facts and notation about bilinear groups and the Pedersen commitment.

Bilinear Groups. A bilinear group generator $\mathcal{BG}(1^\lambda)$ outputs $(q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e)$, where $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ are additive groups of prime order q , and $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is an efficiently computable, non-degenerate, bilinear map. In this paper, we consider Type-3 groups where it is assumed there is no efficiently computable isomorphism between \mathbb{G}_1 and \mathbb{G}_2 . We use bracket notation of [EHK⁺13], i.e., for $s \in \{1, 2, T\}$ and $a \in \mathbb{Z}_q$, we write $[a]_s$ to denote $a \cdot g_s \in \mathbb{G}_s$, where g_s is a fixed generator of \mathbb{G}_s . From an element $[a]_s \in \mathbb{G}_s$ and a scalar b it is possible to efficiently compute $[ab] \in \mathbb{G}_s$. Also, given elements $[a]_1 \in \mathbb{G}_1$ and $[b]_2 \in \mathbb{G}_2$, one can efficiently compute $[a \cdot b]_T$ by using the pairing $e([a]_1, [b]_2)$, that we compactly denote with $[a]_1 \cdot [b]_2$. Vectors and matrices are denoted in boldface. We use the bracket notation also for matrix operations, i.e., $[\mathbf{A}]_1 \cdot [\mathbf{B}]_2 = [\mathbf{A} \cdot \mathbf{B}]_T$. For a vector \mathbf{a} and for $i < j$ we denote by $\mathbf{a}_{[i,j]}$ its portion (a_i, \dots, a_j) .

¹³ The sampling of random group elements can be heuristically instantiated in the random oracle model by letting these elements be the output of a suitable hash function. The main advantage of this hash-based instantiation is that the commitment key has *constant-size* and no bound on the size of the vectors must be fixed a priori.

Pedersen Vector Commitment. Let us recall the extended Pedersen commitment scheme for vectors of size n . Here we consider an instantiation on a group \mathbb{G}_1 .

$\text{Ped.Setup}(1^\lambda)$: sample $[\mathbf{h}]_1 \leftarrow \mathbb{G}_1^{n+1}$ from a distribution \mathcal{D} , and output $\text{ck} := [\mathbf{h}]_1$;
 $\text{Ped.Commit}([\mathbf{h}]_1, \mathbf{w})$: sample $o \leftarrow \mathbb{Z}_q$ and return $(c, o) := ((o, \mathbf{w}^\top) \cdot [\mathbf{h}]_1, o)$;
 $\text{Ped.VerCommit}([\mathbf{h}]_1, c, \mathbf{w}, o)$: output 1 iff $c = (o, \mathbf{w}^\top) \cdot [\mathbf{h}]_1$.

Above \mathcal{D} is a probability distribution over the group elements that allows to argue that the scheme is perfectly hiding and computationally binding. For example, \mathcal{D} may be the uniform distribution, in which case we obtain the classical scheme that is binding under the discrete logarithm assumption, or \mathcal{D} may output powers of random values, e.g., $h_i = s^i$ for an $s \leftarrow \mathbb{Z}_q$, that has also been proved computationally binding under a suitable assumption.

In our constructions we only require the commitment scheme to have the same verification algorithm as Ped.VerCommit .

Tool: SNARK for Linear Subspaces. In our CP-SNARK constructions we make use of a SNARK for the linear subspace relation $R_{\mathbf{M}}([\mathbf{x}]_1, \mathbf{w})$ such that:

$$R_{\mathbf{M}}([\mathbf{x}]_1, \mathbf{w}) = 1 \iff [\mathbf{x}]_1 = [\mathbf{M}]_1 \cdot \mathbf{w} \in \mathbb{G}_1^l, \text{ where } [\mathbf{M}] \in \mathbb{G}_1^{l \times t}, \mathbf{w} \in \mathbb{Z}_q^t$$

Namely, given a fixed public matrix $[\mathbf{M}]_1$ and a public vector $[\mathbf{x}]_1$, one can prove knowledge of a vector \mathbf{w} such that $[\mathbf{x}]_1 = [\mathbf{M}]_1 \cdot \mathbf{w}$. We denote a SNARK for this family of relations with $\text{ss}\Pi$. A candidate scheme for $\text{ss}\Pi$ is the Kiltz-Wee QA-NIZK scheme Π'_{as} [KW15] that works in bilinear groups. As described in [KW15], the security of this scheme requires that $l > t$, which is not satisfied in our setting where matrices have always more columns than rows. This means that, when \mathbf{M} has full rank, $R_{\mathbf{M}}$ is satisfied for any $[\mathbf{x}]_1$. In fact, what we need is an argument of knowledge for this language. For this, by extending a recent result [FLSZ17], we show the knowledge soundness of Π'_{as} [KW15], without the $l > t$ restriction, under the discrete logarithm assumption, in the algebraic group model [FKL18]. We recall the scheme and its security statement in Appendix D. For knowledge soundness, the matrix $[\mathbf{M}]_1$ must be generated using a witness sampleable distribution \mathcal{D}_{mtx} , i.e., there must exist a polynomial time algorithm that samples \mathbf{M} in \mathbb{Z}_q such that $[\mathbf{M}]_1$ has the same distribution as the one sampled with \mathcal{D}_{mtx} . We note that this is satisfied by our use cases where \mathbf{M} includes bases of Pedersen-like commitment schemes.

4.1 CP-SNARK for Pedersen Verification

Our scheme CP_{link} is designed to work with, as global commitment scheme, any Com such that $\text{Com.VerCommit} = \text{Ped.VerCommit}$. Furthermore, it handles any cc-SNARK scheme $\text{cc}\Pi$ whose underlying commitment algorithm also follows Pedersen verification, i.e., $\text{cc}\Pi.VerCommit = \text{Ped.VerCommit}$. Let us stress that although the verification algorithm is the same the commitment keys are not. In particular, the key of Com is completely independent of the relations to be proven (e.g., are random group elements) whereas the key of $\text{cc}\Pi$ is relation-dependent.

More formally, let Com be a commitment scheme such that $\text{Com.VerCommit} = \text{Ped.VerCommit}$. We build a CP-SNARK CP_{link} for Com and for the following class of relations R^{link} . Fixed a security parameter λ (and the group setting for λ as well), R^{link} is over $(\mathcal{D}_x \times \mathcal{D}_1 \times \dots \times \mathcal{D}_\ell \times \mathcal{D}_\omega)$, where $\mathcal{D}_x = \mathbb{G}_1, \mathcal{D}_\omega = \mathbb{Z}_q$ and $\mathcal{D}_j = \mathbb{Z}_q^{n_j}$ for some n_j such that $\sum_j n_j = m$. R^{link} is parametrized by a commitment key $[\mathbf{f}]_1 \in \mathbb{G}_1^{m+1}$, and is defined as:

$$R^{\text{link}}(c', (\mathbf{u}_j)_{j \in [\ell]}, o') = 1 \iff c' \stackrel{?}{=} (o', \mathbf{u}_1^\top, \dots, \mathbf{u}_\ell^\top) \cdot [\mathbf{f}]_1$$

Before describing the construction in full detail, let us present the main ideas.

Let $\text{ck} = [\mathbf{h}]_1 \in \mathbb{G}_1^{n+1}$ be the key of the global commitment Com . In our CP_{link} the public inputs of the prover are ℓ commitments $(c_j)_{j \in [\ell]}$ and another commitment c' ; the witness is a set of openings $((\mathbf{u}_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]})$ for commitments $(c_j)_{j \in [\ell]}$, and an opening o' for c' . In particular, the prover must prove that

$$R_{\text{Ped}}^{\text{link}}(c', (c_j)_{j \in [\ell]}, (\mathbf{u}_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, o') = 1 \iff \bigwedge_{j \in [\ell]} c_j = (o_j, \mathbf{u}_j^\top) \cdot [\mathbf{h}_{[0, n_j]}]_1 \wedge c' = (o', \mathbf{u}_1^\top, \dots, \mathbf{u}_\ell^\top) \cdot [\mathbf{f}]_1$$

The description of our scheme CP_{link} follows:

$\text{CP}_{\text{link}}.\text{KeyGen}(\text{ck}, R^{\text{link}})$: parse $\text{ck} = [\mathbf{h}]_1 \in \mathbb{G}_1^{n+1}$, and let $R^{\text{link}} : \mathbb{G}_1 \times \mathcal{D}_1 \times \dots \times \mathcal{D}_\ell \times \mathbb{Z}_q$ be the relation defined above with $\text{ck}' = [\mathbf{f}]_1 \in \mathbb{G}_1^{m+1}$. Use $[\mathbf{h}]_1, [\mathbf{f}]_1$ and R^{link} to build a matrix \mathbf{M} as in equation (1). Compute $(\text{ek}, \text{vk}) \leftarrow \text{ss}\Pi.\text{KeyGen}([\mathbf{M}]_1)$ and return (ek, vk) .

$\text{CP}_{\text{link}}.\text{Prove}(\text{ek}, c', (c_j)_{j \in [\ell]}, (\mathbf{u}_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, o')$: define $[\mathbf{x}]_1$ and \mathbf{w} as in as in equation (1), compute $\pi \leftarrow \text{ss}\Pi.\text{Prove}(\text{ek}, [\mathbf{x}]_1, \mathbf{w})$ and return π .

$\text{CP}_{\text{link}}.\text{VerProof}(\text{vk}, c', (c_j)_{j \in [\ell]}, \pi)$: set $[\mathbf{x}]_1$ as in (1) and return $\text{ss}\Pi.\text{VerProof}(\text{vk}, [\mathbf{x}]_1, \pi)$.

The key idea of the construction is that this relation can be expressed as a linear subspace relation $R_{\mathbf{M}}([\mathbf{x}]_1, \mathbf{w})$ where $\mathbf{M}, \mathbf{x}, \mathbf{w}$ can be defined as follows from the inputs of $R_{\text{Ped}}^{\text{link}}$, with $l = \ell + 1$ and $t = m + \ell + 1$:

$$\underbrace{\begin{bmatrix} c_1 \\ \vdots \\ c_\ell \\ c' \end{bmatrix}}_{[\mathbf{x}]_1} = \underbrace{\begin{bmatrix} h_0 & 0 & \dots & 0 & 0 & \mathbf{h}_{[1, n_1]} & 0 & \dots & 0 \\ 0 & h_0 & \dots & 0 & 0 & 0 & \mathbf{h}_{[1, n_2]} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & h_0 & 0 & 0 & 0 & \dots & \mathbf{h}_{[1, n_\ell]} \\ 0 & 0 & \dots & 0 & f_0 & \mathbf{f}_{[1, n_1]} & \mathbf{f}_{[n_1+1, n_2]} & \dots & \mathbf{f}_{[n_{\ell-1}+1, n_\ell]} \end{bmatrix}}_{[\mathbf{M}]_1} \underbrace{\begin{pmatrix} o_1 \\ \vdots \\ o_\ell \\ o' \\ \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_\ell \end{pmatrix}}_{\mathbf{w}} \quad (1)$$

In the theorem below we show that CP_{link} is knowledge-sound and zero-knowledge assuming so is $\text{ss}\Pi$. We show the formal statement in Appendix C.1, where we also prove the security of CP_{link} based on that of $\text{ss}\Pi$. Appendix C.2 shows how to extend CP_{link} to handle a more general class of relations that essentially checks that a set of vectors $(\mathbf{u}_j)_{j \in [\ell]}$ is a *prefix*, of known length, of a vector \mathbf{u}' committed in c' .

EFFICIENCY. When using $\text{ss}\Pi$ from [KW15], the key generation algorithm outputs an evaluation key of $m + \ell + 1$ \mathbb{G}_1 elements and a verification key with $l + 1$ \mathbb{G}_2 elements. The prover cost is one multi-exponentiation of length $m + \ell + 1$ while the verifier needs $l + 1$ pairings to check the one group element composing the proof.

Theorem 4.1. *If $\text{ss}\Pi$ is $\text{KSND}(\text{ss}\Pi.\mathcal{RG}, \mathcal{Z})$ where \mathcal{Z} is an auxiliary input distribution, then the CP-SNARK construction CP_{link} given above is $\text{KSND}(\text{CP}_{\text{link}}.\mathcal{RG}, \mathcal{Z})$. Furthermore, if $\text{ss}\Pi$ is composable ZK for $\text{ss}\Pi.\mathcal{RG}$, then CP_{link} is composable ZK for $\text{CP}_{\text{link}}.\mathcal{RG}$.*

4.2 CP-SNARK for Linear Properties

In this section we show a CP-SNARK for the relation R^{lin} that checks linear properties of (committed) vectors: for a fixed public matrix $\mathbf{F} \in \mathbb{Z}_q^{n \times m}$, relation $R_{\mathbf{F}}^{\text{lin}}$ over public input $\mathbf{x} \in \mathbb{Z}_q^n$ and witness $\mathbf{u} \in \mathbb{Z}_q^m$, with $\mathbf{u} := (\mathbf{u}_j)_{j \in [\ell]}$ and $\mathbf{u}_j \in \mathbb{Z}_q^{n_j}$, holds iff $\mathbf{x} = \mathbf{F} \cdot \mathbf{u}$.

Our scheme, called $\text{CP}_{\text{lin}}^{\text{Ped}}$, is quite similar to CP_{link} and essentially consists into invoking $\text{ss}\Pi$ to prove that the above subspace relation holds. The full description of our scheme $\text{CP}_{\text{lin}}^{\text{Ped}}$ follows:

$\text{CP}_{\text{lin}}^{\text{Ped}}.\text{KeyGen}(\text{ck}, R_{\mathbf{F}}^{\text{lin}})$: parse $\text{ck} = [\mathbf{h}]_1 \in \mathbb{G}_1^{m+1}$. Use $[\mathbf{h}]_1$ and $R_{\mathbf{F}}^{\text{lin}}$ to build a matrix $[\mathbf{M}]$ as in equation (2). Compute $(\text{ek}, \text{vk}) \leftarrow \text{ss}\Pi.\text{KeyGen}([\mathbf{M}]_1)$ and return (ek, vk) .

$\text{CP}_{\text{lin}}^{\text{Ped}}.\text{Prove}(\text{ek}, \mathbf{x}, (c_j)_{j \in [\ell]}, (\mathbf{u}_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]})$: define $[\mathbf{x}']_1$ and \mathbf{w}' as in equation (2), and return $\pi \leftarrow \text{ss}\Pi.\text{Prove}(\text{ek}, [\mathbf{x}']_1, \mathbf{w}')$.

$\text{CP}_{\text{lin}}^{\text{Ped}}.\text{VerProof}(\text{vk}, \mathbf{x}, (c_j)_{j \in [\ell]}, \pi)$: set $[\mathbf{x}']_1$ as in (2) and return $\text{ss}\Pi.\text{VerProof}(\text{vk}, [\mathbf{x}']_1, \pi)$.

The scheme $\text{CP}_{\text{lin}}^{\text{Ped}}$ considers each \mathbf{u}_j to be committed using a commitment scheme Com such that $\text{Com.VerCommit} = \text{Ped.VerCommit}$, and whose key is $\text{ck} = [\mathbf{h}]_1 \in \mathbb{G}_1^{m^*+1}$, with $m^* \geq m$.¹⁴ The idea is to express such a commit-and-prove relation with the linear subspace relation $R_{\mathbf{M}}([\mathbf{x}']_1, \mathbf{w}')$ that holds iff $[\mathbf{x}']_1 = [\mathbf{M}]_1 \cdot \mathbf{w}'$, where $[\mathbf{x}']_1 \in \mathbb{G}_1^l$, $[\mathbf{M}]_1 \in \mathbb{G}_1^{l \times t}$ and $\mathbf{w}' \in \mathbb{Z}_q^t$ can be built from the inputs of $R_{\mathbf{F}}^{\text{lin}}$ as follows (for $l = \ell + n$ and $t = m + \ell$):

$$\begin{array}{c} \overbrace{\begin{bmatrix} c_1 \\ \vdots \\ c_\ell \\ \mathbf{x} \end{bmatrix}}^{[\mathbf{x}']_1} = \overbrace{\begin{bmatrix} h_0 & 0 & \dots & 0 & \mathbf{h}_{[1, n_1]} & 0 & \dots & 0 \\ 0 & h_0 & \dots & 0 & 0 & \mathbf{h}_{[1, n_2]} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & h_0 & 0 & 0 & \dots & \mathbf{h}_{[1, n_\ell]} \end{bmatrix}}^{[\mathbf{M}]_1} \cdot \overbrace{\begin{pmatrix} o_1 \\ \vdots \\ o_\ell \\ \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_\ell \end{pmatrix}}^{\mathbf{w}'} \end{array} \quad (2)$$

EFFICIENCY. When using $\text{ss}\Pi$ from [KW15], the prover cost is one multi-exponentiation of length $m + \ell$ while the verifier needs $\ell + |\mathbf{x}| + 1$ pairings. If \mathbf{x} is some fixed value, as in our applications, $|\mathbf{x}|$ of these pairings either disappear (if $\mathbf{x} = \mathbf{0}$) or can be precomputed. Furthermore, it is possible to see that the cost of KeyGen is $O(\ell \cdot t + n_F)$ where n_F is the number of nonzero entries of \mathbf{F} . Essentially this cost depends on the sparsity of the matrix; this is crucial in our applications where for example \mathbf{F} includes the \mathbf{W} matrices representing the linear constraints of a circuit [BCC⁺16].

We state the following theorem. We omit the proof, which is essentially the same as that of Theorem C.1.

Theorem 4.2. *Let $\mathbf{F} \in \mathbb{Z}_q^{n \times m}$ be a matrix from a distribution \mathcal{D}_{mtx} , and \mathcal{Z} be an auxiliary input distribution. If $\text{ss}\Pi$ is $\text{KSND}(\text{ss}\Pi.\mathcal{RG}, \mathcal{Z})$ where $\text{ss}\Pi.\mathcal{RG}$ is a relation generator that samples ck and $\mathbf{F} \leftarrow \mathcal{D}_{\text{mtx}}$, then the CP-SNARK construction $\text{CP}_{\text{lin}}^{\text{Ped}}$ given above is $\text{KSND}(\mathcal{D}_{\text{mtx}}, \mathcal{Z})$. Furthermore, if $\text{ss}\Pi$ is composable ZK for $\text{ss}\Pi.\mathcal{RG}$, then $\text{CP}_{\text{lin}}^{\text{Ped}}$ is composable ZK for \mathcal{D}_{mtx} .*

¹⁴ While in our description we use the same commitment key for every \mathbf{u}_j , our scheme easily extends to the case where different commitment keys are used.

5 Efficient CP-SNARKs for Polynomial Commitments

In this section we show a collection of zero-knowledge CP-SNARKs for a variety of relations over vectors committed using a specific commitment scheme from [ZGK⁺17b]. This scheme is for committing to multivariate polynomials and it can be used for vectors by converting them into their multilinear extension polynomials. Although this commitment scheme has a specially structured commitment key, its verification algorithm can be casted as a form of Pedersen verification; this means we can apply our results of Section 3.5 to turn all the CP-SNARKs in this section into ones for a standard Pedersen commitment, or to simply make them work under some common Pedersen-like scheme. Among the CP-SNARKs in this section, worth mentioning are one for Hadamard product and one for the self permutation relation. Notably these schemes have a CRS that is universal (and in some cases deterministically specializable).

5.1 Preliminaries and Building Blocks

We review the main building blocks of our constructions.

Polynomial Commitments

The specific commitment scheme we consider here is the polynomial commitment underlying the verifiable polynomial delegation (VPD) scheme of Zhang et al. [ZGK⁺17b]. In a nutshell, a VPD allows one to commit to multivariate polynomials and later prove their evaluations (also committed) at a public point. Here we show that their VPD scheme can be seen as a CP-SNARK for such polynomial commitment, for relations encoding polynomial evaluations. Namely, whereas in [ZGK⁺17b] VPD is presented as a single primitive, here we separate the commitment scheme from the argument system. With this simple change (together with a slightly stronger zero-knowledge notion) we can use our composition results to argue security when commitments are reused across different proofs.

Formally, we consider a commitment scheme whose message space \mathcal{D} includes both values in a finite field \mathbb{F} and a class \mathcal{F} of polynomials with coefficients in \mathbb{F} , with μ variables and maximal degree δ in each variable. We denote these partitions of $\mathcal{D} = \mathbb{F} \cup \mathcal{F}$ as $\mathcal{D}_{\text{pol}} = \mathcal{F}$ and $\mathcal{D}_{\text{val}} = \mathbb{F}$ and we use a flag `type` to differentiate between them so that $f \in \mathcal{F}$ when `type` = `pol`, and $f \in \mathbb{F}$ when `type` = `val`.¹⁵ In addition to satisfying the notion of Definition 2.1, we assume the scheme to be knowledge extractable and to have a trapdoor generation. For convenience, we summarize its definition below.

Definition 5.1 (Extractable Trapdoor Polynomial Commitments). *An extractable trapdoor polynomial commitment scheme for a class of polynomials \mathcal{F} is a tuple of algorithms $\text{PolyCom} = (\text{Setup}, \text{Commit}, \text{CheckCom}, \text{VerCommit})$ that work as follows.*

$\text{Setup}(1^\lambda) \rightarrow \text{ck}$: takes the security parameter and outputs a commitment key ck .

$\text{Commit}(\text{ck}, f, \text{type}) \rightarrow (c_f, o_f)$: takes the commitment key ck , a flag `type` $\in \{\text{pol}, \text{val}\}$ and an element $f \in \mathcal{D}_{\text{type}}$, and outputs a commitment c_f and an opening o_f . We use $\text{ComPoly}(\text{ck}, \cdot)$ and $\text{ComVal}(\text{ck}, \cdot)$ as shorthands for $\text{Commit}(\text{ck}, \cdot, \text{pol})$ and $\text{Commit}(\text{ck}, \cdot, \text{val})$ respectively. We also assume that `type` is part of c_f , namely it is not hidden.

$\text{CheckCom}(\text{ck}, c) \rightarrow b$: takes as input a commitment c and accepts it as valid ($b = 1$) or not ($b = 0$).

¹⁵ Note that the only ambiguity can occur when differentiating a degree-0 polynomial from a point.

$\text{VerCommit}(\text{ck}, c_f, f, o_f) \rightarrow b$: takes as input commitment c , element $f \in \mathcal{D}$ and opening o_f , and accepts ($b = 1$) or rejects ($b = 0$). If f is a degree-0 polynomial the same algorithm applies to commitments created using ComVal

PolyCom must satisfy correctness, binding and (perfect) hiding as in Definition 2.1 (with the additional requirements that correctness also implies that CheckCom accepts, and binding holds for adversarial commitments that are accepted by CheckCom). In addition PolyCom must satisfy the trapdoor and extractability properties defined below.

Trapdoor. There exists three algorithms $(\text{ck}, \text{td}) \leftarrow \mathcal{S}_{\text{ck}}(1^\lambda), (c, st) \leftarrow \text{TdCom}(\text{td}, \text{type})$ and $o \leftarrow \text{TdOpen}(\text{td}, st, c, f)$ such that: the distribution of the commitment key returned by \mathcal{S}_{ck} is perfectly/statistically close to the one of the key returned by Setup ; for any $\text{type} \in (\text{pol}, \text{val})$, any $f \in \mathcal{D}_{\text{type}}$, $(c, o) \approx (c', o')$ where $(c, o) \leftarrow \text{Commit}(\text{ck}, f, \text{type})$, $(c', st) \leftarrow \text{TdCom}(\text{td}, \text{type})$ and $o' \leftarrow \text{TdOpen}(\text{td}, st, c', f)$.

Extractability. PolyCom is knowledge extractable for auxiliary input distribution \mathcal{Z} if for every (non-uniform) efficient adversary \mathcal{A} there exists a (non-uniform) efficient extractor \mathcal{E} such that $\Pr[\text{Game}_{\mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{extr}} = 1] = \text{negl}$.

$$\begin{array}{l} \text{Game}_{\mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{extr}} \\ \hline \text{ck} \leftarrow \text{Setup}(1^\lambda) \ ; \ \text{aux}_{\mathcal{Z}} \leftarrow \mathcal{Z}(1^\lambda) \ ; \ c \leftarrow \mathcal{A}(\text{ck}, \text{aux}_{\mathcal{Z}}) \ ; \ (f, o) \leftarrow \mathcal{E}(\text{ck}, \text{aux}_{\mathcal{Z}}) \\ \text{return } \text{CheckCom}(\text{ck}, c) \stackrel{?}{=} 1 \ \wedge \ \text{VerCommit}(\text{ck}, c, f, o) \stackrel{?}{=} 0 \end{array}$$

Linearly Homomorphic Commitments. For the constructions presented in this section we assume that the commitments are linearly homomorphic. That is we assume existence of a deterministic algorithm $(c', o') \leftarrow \text{HomEval}(\text{ck}, g, (c_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]})$ such that, for a linear function $g : \mathbb{F}^\ell \rightarrow \mathbb{F}$, if $\text{VerCommit}(\text{ck}, c_j, a_j, o_j) = 1$ then $\text{VerCommit}(\text{ck}, c', g((a_j)_{j \in [\ell]}), o') = 1$. In the paper we assume HomEval takes in the vector of ℓ coefficients of g .

Zero-knowledge CP-SNARKs for PolyCom

Constructions in this section use the following existing CP-SNARKs for the scheme PolyCom :

- CP_{eq} : a CP-SNARK for relation $R^{\text{eq}}(u_1, u_2) := u_1 \stackrel{?}{=} u_2$, where $u_1, u_2 \in \mathbb{F}$.
- CP_{prd} : a CP-SNARK for relation $R^{\text{prd}}(u_1, u_2, u_3) := u_3 \stackrel{?}{=} u_1 \cdot u_2$, where $u_1, u_2, u_3 \in \mathbb{F}$.
- CP_{poly} : a CP-SNARK for the relation R^{poly} over $\mathcal{D}_x \times \mathcal{D}_1 \times \mathcal{D}_2$ where $\mathcal{D}_x = \mathbb{F}^\mu$, $\mathcal{D}_1 = \mathcal{F}$, $\mathcal{D}_2 = \mathbb{F}$ and $R^{\text{poly}}(\mathbf{x}, f, y) := y \stackrel{?}{=} f(\mathbf{x})$. For zero-knowledge, we assume that CP_{poly} satisfies a notion where the commitment key is generated in trapdoor mode and the CP_{poly} simulators $(\mathcal{S}_{\text{kg}}, \mathcal{S}_{\text{prv}})$ get access to the commitment trapdoor produced by \mathcal{S}_{ck} . Note that such notion is weaker than the one of Definition 3.1 but sufficient to argue that a scheme satisfying this notion is a cc-SNARK.

In Appendix E we show pairing-based constructions of PolyCom and CP_{poly} extracted from the verifiable polynomial delegation scheme of Zhang et al. [ZGK⁺17b]. As observed by Zhang et al. constructions for CP_{eq} and CP_{prd} can be obtained using standard techniques from classical sigma-protocols. Finally, we observe that all these schemes share the same (deterministic) KeyGen algorithm that, on input the commitment key ck , simply partitions the elements of ck into $\text{ek} = \text{ck}$ and $\text{vk} = \text{cvk}$, where cvk is a subset of the elements in ck that is sufficient to run algorithms CheckCom , ComVal and HomEval .

EFFICIENCY. Both proof-of-equality and proof-of-product (Appendix A in [WTas⁺17]) are built as Sigma protocols, where both prover and verifier run in constant time. They can be made non-interactive using the Fiat-Shamir heuristic [FS87]. The proof in CP_{eq} consists of one group element and one field element. In CP_{prd} , the prover sends $3\mathbb{G}_1 + 5\mathbb{F}$. The proof in CP_{poly} for polynomial evaluation CP_{poly} needs $2(\mu + 1)$ group elements, its verifier runs in $O(\mu)$ and the prover in time $O(m)$. Here, μ is the number of variables of the polynomial and m is its number of monomials. The KeyGen algorithm of CP_{poly} outputs an evaluation key of $(2(\delta + 1)^\mu + 3)\mathbb{G}_1 + (\mu + 3)\mathbb{G}_2$ elements and a subset verification key of size $(\mu + 3)\mathbb{G}_2$, where δ is the maximum degree in each variable of the committed polynomial. For clarification, note that by construction the public parameters \mathbb{P} within the commitment key are formed by two group elements per element in the set of all multisets of $\{1, \dots, \mu\}$ where each element appears at most δ times. Asymptotically, the crs of CP_{poly} contains $O\left(\binom{\mu\delta + \mu}{\mu\delta}\right)$ group elements. As will be explained, in our setting $\delta = 1$, which keeps the crs size small.

Multilinear Extensions

Given a function $f : \{0, 1\}^\mu \rightarrow \mathbb{F}$, its unique multilinear extension (MLE) is the (unique) multilinear polynomial $\tilde{f} : \mathbb{F}^\mu \rightarrow \mathbb{F}$ such that $f(\mathbf{b}) = \tilde{f}(\mathbf{b})$ for all $\mathbf{b} \in \{0, 1\}^\mu$. Such multilinear extension is defined as the following polynomial

$$\tilde{f}(X_1, \dots, X_\mu) = \sum_{\mathbf{b} \in \{0, 1\}^\mu} \chi_{\mathbf{b}}(X_1, \dots, X_\mu) \cdot f(\mathbf{b})$$

where $\chi_{\mathbf{b}}(X_1, \dots, X_\mu) = \prod_{j=1}^\mu \chi_{b_j}(X_j)$, $\chi_1(X) = X$ and $\chi_0(X) = 1 - X$. For a vector $\mathbf{u} \in \mathbb{F}^m$ (for some $m = 2^\mu$), its unique MLE is the MLE \tilde{u} of the function $u : \{0, 1\}^\mu \rightarrow \mathbb{F}$ such that, for every $0 \leq i \leq m - 1$ with $i = \sum_{j=0}^{\mu-1} i_j 2^j$, $u(i_0, \dots, i_{\mu-1}) = u_{i+1}$. Note that by using MLEs one can commit to a vector \mathbf{u} using PolyCom by committing to its MLE \tilde{u} , with maximum variable degree $\delta = 1$.

Let $eq : \{0, 1\}^\mu \times \{0, 1\}^\mu \rightarrow \{0, 1\}$ be the equality predicate ($eq(a, b) = 1$ iff $a = b$) and let \tilde{eq} be its MLE (which has a closed-form representation that allows evaluation in time $O(\mu)$ [Tha13]). We recall the following lemma from [Rot09] (as restated in [Tha13]):

Lemma 5.1 ([Rot09, Lemma 3.2.1]). *For any polynomial $h : \mathbb{F}^\mu \rightarrow \mathbb{F}$ extending $p : \{0, 1\}^\mu \rightarrow \mathbb{F}$ (i.e., such that $\forall \mathbf{b} \in \{0, 1\}^\mu : h(\mathbf{b}) = p(\mathbf{b})$), it holds*

$$\tilde{p}(\mathbf{X}) = \sum_{\mathbf{b} \in \{0, 1\}^\mu} \tilde{eq}(\mathbf{X}, \mathbf{b}) \cdot h(\mathbf{b}).$$

5.2 A CP-SNARK for Sum-Check

The sum-check protocol [LFKN92] is an interactive proof that allows a prover to convince a verifier of the validity of a statement of the form $t = \sum_{\mathbf{b} \in \{0, 1\}^\mu} g(\mathbf{b})$ where $g : \mathbb{F}^\mu \rightarrow \mathbb{F}$. The protocol consists of μ rounds, it is public coin, and the running time of the verifier in it is $O(\sum_{i=1}^\mu \text{deg}_i(g))$ plus the cost of evaluating g once (on a random point).

Here we propose a zero-knowledge variant of the sum-check protocol where both the polynomial g and the target value t are committed and one proves knowledge of these values such that $t = \sum_{\mathbf{b} \in \{0, 1\}^\mu} g(\mathbf{b})$. Precisely, we work with polynomials g defined as the product of $p + 1$ polynomials of the form $g(\mathbf{S}) = \prod_{i=0}^p g_i(\mathbf{S})$, such that all the g_i 's, except g_0 , are committed. Namely, we

show a CP-SNARK CP_{sc} for commitment scheme PolyCom and the relation $R^{\text{sc}}(\mathbf{x}, \mathbf{u})$, with $\mathbf{x} \in \mathcal{F}$ and $\mathbf{u} \in \mathbb{F} \times \mathcal{F}^p$, that is formally defined as:

$$R^{\text{sc}}(g_0, (t, (g_j)_{j \in [p]})) = 1 \iff g(\mathbf{S}) = \prod_{i=0}^p g_i(\mathbf{S}) \wedge t = \sum_{\mathbf{b} \in \{0,1\}^\mu} g(\mathbf{b})$$

Our scheme, dubbed CP_{sc} , is built as a generalization of the protocol recently proposed in [ZGK⁺17b, WTs⁺18] that works for a relation that is the same as the above one except that only t is committed while g is public to the verifier. For the reader familiar with the zero-knowledge sum-check protocol in [ZGK⁺17b, Construction 2], what we do here is to modify their protocol using the following ideas: whereas in [ZGK⁺17b] the verifier has access to g and computes a commitment to $g(\mathbf{s})$ for a random point \mathbf{s} on its own, in our case the verifier has access to a commitment c_g of g and we let the prover create a commitment to $g(\mathbf{s})$ and use CP_{poly} to prove its correctness with respect to c_g . More precisely, the verifier does not have a commitment to g but rather commitments to the factors of g . Hence our prover proceeds by additionally creating commitments to each $g_i(\mathbf{s})$, it proves their correct evaluations and then uses CP_{prd} to prove that $g(\mathbf{s}) = \prod_{i=0}^p g_i(\mathbf{s})$ with respect to these commitments. Making these changes results in a protocol that is the same as that in [ZGK⁺17b] except for the last round from the prover to the verifier. Indeed we can prove the security of our protocol by making a reduction to the one of [ZGK⁺17b]. In Figure 3 we give a detailed description of this protocol for the case $p = 2$; this is sufficient for our applications.

<p>Protocol Π_{sc}:</p> <hr/> <p>Common input: c_t, g_0, c_1, c_2; \mathcal{P}'s input: $(t, o_t, g_1, o_1, g_2, o_2)$</p> <p>$\mathcal{P} : g(\mathbf{S}) := \prod_{i=0}^2 g_i(\mathbf{S}), c_0 := c_t, t_0 := t, \rho_0 := o_t$, let $f(A_0, \dots, A_k) := A_0 + \sum_{j=0}^d A_j := (2, 1, \dots, 1)$</p> <p>for $i = 1 \dots \mu$:</p> <p style="padding-left: 20px;">$\mathcal{P} : h_i(X) := \sum_{b_{i+1}, \dots, b_\mu \in \{0,1\}} g(s_1, \dots, s_{i-1}, X, b_{i+1}, \dots, b_\mu) := \sum_{j=0}^d a_j X^j$</p> <p style="padding-left: 20px;">$\mathcal{P} : \text{compute } \{(\text{com}_{a_j}, \rho_{a_j}) \leftarrow \text{ComVal}(\text{ck}, a_j)\}_{j=0}^d, (\text{com}_{i-1}^*, \rho_{i-1}^*) \leftarrow \text{HomEval}(\text{ck}, f, \{\text{com}_{a_j}\}_{j=0}^d), \{\rho_{a_j}\}_{j=0}^d\}$</p> <p style="padding-left: 40px;">$\pi_{\text{eq}} \leftarrow \text{CP}_{\text{eq}}.\text{Prove}(\text{ck}, \text{com}_{i-1}, \text{com}_{i-1}^*, t_{i-1}, h_i(0) + h_i(1), \rho_{i-1}, \rho_{i-1}^*)$</p> <p style="padding-left: 20px;">$\mathcal{P} \rightarrow \mathcal{V} : \{\text{com}_{a_j}\}_{j=0}^d, \pi_{\text{eq}}$</p> <p style="padding-left: 20px;">$\mathcal{V} : \{\text{CheckCom}(\text{cvk}, \text{com}_{a_j})\}_{j=0}^d$, compute $(\text{com}_{i-1}^*, \cdot) \leftarrow \text{HomEval}(\text{ck}, f, \{\text{com}_{a_j}\}_{j=0}^d, \cdot)$</p> <p style="padding-left: 20px;">$\mathcal{V} : \text{CP}_{\text{eq}}.\text{VerProof}(\text{cvk}, \text{com}_{i-1}, \text{com}_{i-1}^*, \pi_{\text{eq}}), s_i \leftarrow \mathbb{F}, (\text{com}_i, \cdot) \leftarrow \text{HomEval}(\text{ck}, (1, s_i, \dots, s_i^d), \{\text{com}_{a_j}\}_{j=0}^d, \cdot)$</p> <p style="padding-left: 20px;">$\mathcal{V} \rightarrow \mathcal{P} : s_i \in \mathbb{F}$</p> <p style="padding-left: 20px;">$\mathcal{P} : t_i \leftarrow h_i(s_i), (\text{com}_i, \rho_i) \leftarrow \text{HomEval}(\text{ck}, (1, s_i, \dots, s_i^d), \{\text{com}_{a_j}\}_{j=0}^d, \{\rho_{a_j}\}_{j=0}^d)$</p> <p>endfor</p> <p style="padding-left: 20px;">$\mathcal{P} : \{(c'_j, o'_j) \leftarrow \text{ComVal}(\text{ck}, g'_j := g_j(\mathbf{s})), \pi_j \leftarrow \text{CP}_{\text{poly}}.\text{Prove}(\text{ek}, \mathbf{s}, (c_j, c'_j), (g_j, g'_j), (o_j, o'_j))\}_{j=1,2}$</p> <p style="padding-left: 20px;">$\mathcal{P} : (c_1^*, o_1^*) \leftarrow \text{HomEval}(\text{ck}, g_0(\mathbf{s}), c'_1, o'_1), \pi^* \leftarrow \text{CP}_{\text{prd}}.\text{Prove}(\text{ck}, (c_1^*, c_2^*, \text{com}_\mu), (g_0(\mathbf{s}) \cdot g'_1, g'_2, g(\mathbf{s})), (o_1^*, o_2^*, \rho_\mu))$</p> <p style="padding-left: 20px;">$\mathcal{P} \rightarrow \mathcal{V} : c'_1, c'_2, \pi_1, \pi_2, \pi^*$</p> <p style="padding-left: 20px;">$\mathcal{V} : \bigwedge_{j=1,2} \text{CheckCom}(\text{cvk}, c'_j) \wedge \text{CP}_{\text{poly}}.\text{VerProof}(\text{vk}, \mathbf{s}, c_j, c'_j, \pi_j)$</p> <p style="padding-left: 20px;">$\mathcal{V} : (c_1^*, \cdot) \leftarrow \text{HomEval}(\text{ck}, g_0(\mathbf{s}), c'_1, \cdot), \text{CP}_{\text{prd}}.\text{VerProof}(\text{vk}, (c_1^*, c_2^*, \text{com}_\mu), \pi^*)$</p>

Figure 3: Our sum-check protocol over committed result and polynomial with 3 factors ($p = 2$); in **black** are the steps identical to [ZGK⁺17b].

EFFICIENCY. In CP_{sc} , the verifier needs time $O(\mu)$ plus the time to compute $g_0(\mathbf{s})$. The prover's costs include the running time in the sum-check protocol and the creation of the CP_{poly} proofs. If the g_i are multilinear, $\text{CP}_{\text{poly}}.\text{Prove}$ time is $O(2^\mu)$. Also, from [Tha13], if the polynomials g_i allow for evaluation in $O(\mu)$ time or are MLE of vectors, the prover's cost in sum-check can be reduced to $O(2^\mu)$. More detailed, our verifier runs linearly on the number of variables of the polynomial and the prover time is linear on the number of monomials of the factors of the target polynomial. The degree- d polynomial $g(\mathbf{S})$ can have up to m monomials (in particular, $m \leq 2^\mu$). Note that in the i -th iteration, the prover evaluates the target polynomial $2^{\mu-i}$ times. This means 2^μ times in one whole execution of the scheme. For each variable of the polynomial, the prover sends one CP_{eq} proof and commitments to each nonzero coefficient of $h_i(X)$, at most $(d+1)$ of them. Finally, he sends commitments to the evaluations of the two factor polynomials, two CP_{poly} proofs and one CP_{prd} proof. The crs in this case is as long as the one for CP_{poly} , with $\delta = 1$ and then $(\delta + 1)^\mu \geq m$.

Theorem 5.1. *Assume PolyCom is an extractable linearly homomorphic commitment, CP_{poly} and CP_{prd} are zkSNARKs for relations R^{poly} and R^{prd} respectively, and Construction 2 in [ZGK⁺17b] is a ZK interactive argument for sum-check. Then there is a ZK interactive argument for relation R^{sc} . Furthermore, by applying the Fiat-Shamir heuristic we get a zkSNARK in the random oracle model, that we call CP_{sc} .*

5.3 A CP-SNARK for Hadamard Products

In this section we propose a CP-SNARK for PolyCom for the relation R^{had} over $(\mathbb{F}^m)^3$ such that:

$$R^{\text{had}}(\mathbf{u}_0, \mathbf{u}_1, \mathbf{u}_2) = 1 \iff \forall i \in [m] : u_{0,i} = u_{1,i} \cdot u_{2,i}$$

Let $m = 2^\mu$ and let $\tilde{u}_j : \mathbb{F}^\mu \rightarrow \mathbb{F}$ be the MLE of \mathbf{u}_j . Clearly, the relation holds iff for all $\mathbf{b} \in \{0, 1\}^\mu$ we have $\tilde{u}_0(\mathbf{b}) = \tilde{u}_1(\mathbf{b}) \cdot \tilde{u}_2(\mathbf{b})$. If the relation holds, observe that the polynomial $\tilde{u}_1(\mathbf{X}) \cdot \tilde{u}_2(\mathbf{X})$ is an extension of the vector \mathbf{u}_0 , but not a multilinear one. From Lemma 5.1 this equality holds:

$$\tilde{u}_0(\mathbf{X}) = \sum_{\mathbf{b} \in \{0, 1\}^\mu} \tilde{e}q(\mathbf{X}, \mathbf{b}) \cdot \tilde{u}_1(\mathbf{b}) \cdot \tilde{u}_2(\mathbf{b})$$

Without considering zero-knowledge, the main idea of our protocol is that, to check the above equality, the verifier starts by picking a random point $\mathbf{r} \leftarrow_{\$} \mathbb{F}^\mu$, and then the prover uses CP_{sc} to show that $t = \tilde{u}_0(\mathbf{r}) = \sum_{\mathbf{b} \in \{0, 1\}^\mu} g(\mathbf{b})$, where $g(\mathbf{S}) = \tilde{e}q(\mathbf{r}, \mathbf{S}) \cdot \tilde{u}_1(\mathbf{S}) \cdot \tilde{u}_2(\mathbf{S})$. Notice indeed that g can be written as the product of three polynomials $g(\mathbf{S}) := \prod_0^2 g_i(\mathbf{S})$, of which the first one is public: $g_1(\mathbf{S}) = \tilde{u}_1(\mathbf{S})$, $g_2(\mathbf{S}) = \tilde{u}_2(\mathbf{S})$ and $g_0(\mathbf{S}) := \tilde{e}q(\mathbf{r}, \mathbf{S})$. Finally, the prover also needs to convince the verifier that $t = \tilde{u}_0(\mathbf{r})$, which is done using a CP-SNARK CP_{poly} for proving correctness of polynomial evaluations. Therefore we build a CP-SNARK CP_{had} for R^{had} and PolyCom by using CP-SNARKs CP_{poly} , CP_{sc} for PolyCom as building blocks. Furthermore, we describe the scheme as a non-interactive one by letting $\mathbf{r} \leftarrow H((c_j)_{j \in [3]})$ using the random oracle model for H . The full scheme is given below.

EFFICIENCY. Computing π_0 takes time $O(m)$, and the same holds for π_{sc} . The latter follows by observing that the factors of $g(\mathbf{S})$ satisfy the good efficiency conditions for CP_{sc} , i.e., $\tilde{e}q(\mathbf{r}, \mathbf{s})$ can be computed in $O(\mu)$ time and \tilde{u}_1, \tilde{u}_2 are MLE of vectors of length $m = 2^\mu$. For similar reasons, the verifier's time is $O(\mu)$. More detailed, our CP_{had} that proves the result of Hadamard products $\mathbf{u}_0 = \mathbf{u}_1 \circ \mathbf{u}_2$ with $m = 2^\mu$ elements each. The prover runs linear in the number of monomials of

$\text{CP}_{\text{had}}.\text{KeyGen}(\text{ck}) \rightarrow (\text{ek}, \text{vk})$	$\text{CP}_{\text{had}}.\text{Prove}(\text{ek}, (c_j)_{j \in [3]}, (\mathbf{u}_j)_{j \in [3]}, (o_j)_{j \in [3]}) \rightarrow \pi := (c_t, \pi_0, \pi_{\text{sc}})$
$(\text{ek}_s, \text{vk}_s) \leftarrow \text{CP}_{\text{sc}}.\text{KeyGen}(\text{ck})$	$\mathbf{r} \leftarrow H((c_j)_{j \in [3]}) ; t \leftarrow \tilde{u}_0(\mathbf{r}) ; (c_t, o_t) \leftarrow \text{ComVal}(\text{ck}, t)$
$(\text{ek}_p, \text{vk}_p) \leftarrow \text{CP}_{\text{poly}}.\text{KeyGen}(\text{ck})$	$\pi_0 \leftarrow \text{CP}_{\text{poly}}.\text{Prove}(\text{ek}_p, \mathbf{r}, (c_0, c_t), (\tilde{u}_0, t), (o_0, o_t))$
$\text{ek} := (\text{ck}, \text{ek}_s, \text{ek}_p, H)$	$\pi_{\text{sc}} \leftarrow \text{CP}_{\text{sc}}.\text{Prove}(\text{ek}_s, \tilde{q}(\mathbf{r}, \mathbf{S}), (c_t, c_1, c_2), (t, o_t, \tilde{u}_1, o_1, \tilde{u}_2, o_2))$
$\text{vk} := (\text{cvk}, \text{vk}_s, \text{vk}_p, H)$	$\text{CP}_{\text{had}}.\text{VerProof}(\text{vk}, (c_j)_{j \in [3]}, \pi) \rightarrow b \in \{0, 1\}$
	$\mathbf{r} \leftarrow H((c_j)_{j \in [3]}) ; b \leftarrow \text{CP}_{\text{poly}}.\text{VerProof}(\text{vk}_p, \mathbf{r}, c_0, c_t, \pi_0)$
	$b \leftarrow b \wedge \text{CP}_{\text{sc}}.\text{VerProof}(\text{vk}_s, \tilde{q}(\mathbf{r}, \mathbf{S}), (c_t, c_1, c_2), \pi_{\text{sc}})$

Figure 4: CP-SNARK CP_{had} for relation R^{had}

$\tilde{u}_0(\mathbf{X})$, which is at most m , and the verifier time is linear in its number of variables μ . The prover sends one value commitment, one CP_{poly} proof and one CP_{sc} . Note the polynomial used inside sum-check is at most degree $d = 3$ in each variable. The crs includes the group elements output by $\text{CP}_{\text{poly}}.\text{KeyGen}$ with $\delta = 1$, and the description of the hash function $H : (\mathbb{G}_1, \mathbb{G}_1)^3 \rightarrow \mathbb{F}^\mu$ used as a random oracle to achieve noninteractivity.

We state the following result; its proof is in Appendix F.2.

Theorem 5.2. *In the random oracle model, assuming that PolyCom is an extractable trapdoor commitment, $\text{CP}_{\text{poly}}, \text{CP}_{\text{sc}}$ are zero-knowledge CP-SNARKs for PolyCom and relations R^{poly} and R^{sc} respectively, then the scheme CP_{had} described above is a zero-knowledge CP-SNARK for PolyCom and relation R^{had} .*

5.4 A CP-SNARK for Self Permutation

In this section we propose a CP-SNARK for PolyCom for the relation R_ϕ^{sfrm} defined below.

Definition 5.2 (Self permutation of a vector). *Let \mathcal{D} be some domain (e.g., a finite field \mathbb{F}), let n_0, \dots, n_ℓ be positive integers such that $\mathcal{D}_j := \mathcal{D}^{n_j}$ and let $m = \sum_{j=0}^\ell n_j$. Given a permutation $\phi : [m] \rightarrow [m]$, we define a relation R_ϕ^{sfrm} over $\mathcal{D}_0 \times \dots \times \mathcal{D}_\ell = \mathcal{D}^m$ such that:*

$$R_\phi^{\text{sfrm}}(\mathbf{x}, (\mathbf{u}_j)_{j \in [\ell]}) = 1 \iff \forall i \in [m] : y_i = y_{\phi(i)}, \text{ where } \mathbf{y} := (\mathbf{x}, (\mathbf{u}_j)_{j \in [\ell]})$$

Our scheme uses a probabilistic test to prove a permutation of vectors due to [Gro09, BCG⁺17]. For this we need of a CP-SNARK for proving that $t = \prod_{i=1}^m y_i$ with respect to a commitment to point t and vector \mathbf{y} . We call such a relation *internal product* R^{ipd} . A formal definition follows:

Definition 5.3 (Internal product). *Let n_1, \dots, n_ℓ be positive integers and let $m = \sum_{j=1}^\ell n_j$. We define the relation R^{ipd} over $\mathbb{F} \times \mathbb{F}^{n_1} \dots \times \mathbb{F}^{n_\ell}$ such that:*

$$R^{\text{ipd}}(u_0, \mathbf{y} := (\mathbf{u}_j)_{j \in [\ell]}) = 1 \iff u_0 \stackrel{?}{=} \prod_{j=1}^\ell \prod_{i=1}^{n_j} y'_{j,i}$$

We give a formal description of CP_{ipd} in Figure 19 and its computation complexity in Appendix G.

In what follows we present the main ideas to build a CP-SNARK for R^{sfprm} from one for R^{ipd} . Next, we discuss how a CP-SNARK for internal products can be instantiated.

Recall that the goal is to prove that, for a permutation $\phi : [m] \rightarrow [m]$ a committed vector \mathbf{y} satisfies $y_i = y_{\phi(i)}, \forall i \in [m]$. Consider the following vectors in \mathbb{F}^m , $\mathbf{1}, \mathbf{v} = (1, \dots, m)$, and $\phi = (\phi(1), \dots, \phi(m))$, and assume that the prover committed to \mathbf{y} . Let the verifier choose two random values $r, s \leftarrow \mathbb{F}$ and define the vectors $\mathbf{y}' := \mathbf{y} + r \cdot \mathbf{v} - s \cdot \mathbf{1}$ and $\mathbf{y}'' := \mathbf{y} + r \cdot \phi - s \cdot \mathbf{1}$.

If \mathbf{y} is a permutation of itself according to ϕ , then $(\mathbf{y} + r \cdot \phi)$ is a permutation of $(\mathbf{y} + r \cdot \mathbf{v})$ according to ϕ ; however, if \mathbf{y} is *not* a self-permutation according to ϕ then with overwhelming probability over the choice of r some of the entries of $\mathbf{y} + r \cdot \phi$ will not be in the vector $\mathbf{y} + r \cdot \mathbf{v}$. In our scheme the idea is to let the prover show that $\prod_i y'_i = z = \prod_i y''_i$ using CP_{ipd} on (z, \mathbf{y}') and (z, \mathbf{y}'') . However, if some entries of $\mathbf{y} + r \cdot \phi$ are not in $\mathbf{y} + r \cdot \mathbf{v}$, $\prod_i (y_i + r \cdot i - s) = \prod_i (y_i + r \cdot \phi(i) - s)$ holds with negligible probability over the choice of s by the Schwartz-Zippel lemma, thus a prover can be successful only by cheating with CP_{ipd} .

We consider an instantiation of CP_{ipd} based on Thaler's protocol for trees of multiplications [Tha13]. R^{ipd} can be expressed with an arithmetic circuit that is a tree of multiplications over $m = 2^\mu$ inputs. Thaler showed that for this specially regular circuit the CMT protocol can be adapted so that the prover and verifier run in time $O(m)$ and $O(\mu^2)$, respectively. To build a CP-SNARK for R^{ipd} , we thus modify the zk-vSQL protocol [ZGK⁺17b] so as to work over Thaler's protocol instead of CMT. The changes are quite minimal and mainly regard the equation that links the adjacent layers of the tree. We show this protocol in Appendix G.

One detail to be noted here is that such CP_{ipd} works with binary tree circuits, meaning that their input should be a power of two length, so we tweak our definition of the self-permutation relation accordingly. We must work on $\ell + 1$ vectors such that, each has length $n_j = 2^{\mu_j}$ (this is immediate since we commit to MLEs of vectors) but their concatenation has length $\sum_{j=0}^{\ell} n_j = m$ which may not be a power-of-two.

To solve this issue, we execute CP_{ipd} on each block and then aggregate the $\ell + 1$ committed results using a simple zero-knowledge argument for proving a product relation over three commitments, i.e., CP_{prd} . This results in about $\ell + 1$ calls to CP_{ipd} and CP_{prd} . Although this makes proofs grow with ℓ , we observe that in all our applications ℓ is some small constant, e.g., 8 – 10 in our arithmetic circuits encoding.

EFFICIENCY. From the efficiency observations about CP_{ipd} given above, we get that $\text{CP}_{\text{sfprm}}.\text{Prove}$ and $\text{CP}_{\text{sfprm}}.\text{VerProof}$ run in time $O(m)$ and $O(\log^2 m)$ respectively. More in depth, our CP_{sfprm} , built from CP_{ipd} , is used for proving that a vector $\mathbf{y} \in \mathbb{F}^m$ is a self-permutation. This scheme works for vectors whose components $\mathbf{y} := (\mathbf{x}, (\mathbf{u}_j)_{j \in [\ell]})$ have power-of-two length $n_j = 2^{\mu_j}$ such that $m = \sum_{j=0}^{\ell} n_j$, where ℓ is typically some small constant. For each of its $(\ell + 1)$ components, the prover runs 2 calls to CP_{ipd} and iteratively builds 2 proofs-of-product and a final call to CP_{eq} . This means the prover runs in time $O(m)$, the verifier runs in polylogarithmic time in m . The prover sends $(4\ell + 5)$ commitments, 1 opening, $2(\ell + 1)$ short CP_{ipd} proofs, 2ℓ CP_{prd} proofs and 1 proof-of-equality. Note that in each iteration, the length of the current vector is some $n_i < m$. This means that performing $(\ell + 1)$ calls to CP_{ipd} costs the prover $O(\sum_{i=0}^{\ell} n_i) = O(m)$. Conversely, the length of such $(\ell + 1)$ CP_{ipd} proofs is notably larger than the hypothetical length of one single CP_{ipd} proof over $\mathbf{y} \in \mathbb{F}^m$, except that our m is not necessarily a power of two as the tree of multiplications requires. Here, the length of these $(2\ell + 2)$ CP_{ipd} proofs is $\sum_{j=0}^{\ell} (11\mu_j^2 + 25\mu_j) + 2\ell + 2$ group elements and $\sum_{j=0}^{\ell} (\mu_j^2 + 15\mu_j)$ field elements. For brevity, we can express this kind of calculations asymptotically as $O(\log^2 m)$. Note

$\text{CP}_{\text{sfprm}}.\text{KeyGen}(\text{ck}) \rightarrow (\text{ek} := (\text{ck}, \text{ek}_p), \text{vk} := (\text{cvk}, \text{vk}_p)) :$
$(\text{ek}_p, \text{vk}_p) \leftarrow \text{CP}_{\text{ipd}}.\text{KeyGen}(\text{ck})$
$\text{CP}_{\text{sfprm}}.\text{Derive}((\text{ek}, \text{vk}), R_{\phi}^{\text{sfprm}}) \rightarrow (\text{ek}_{\phi}, \text{vk}_{\phi}) :$
$\text{for } j = 0 \dots \ell : \{ (c_{1,j}, o_{1,j}) \leftarrow \text{ComPoly}^*(\text{ck}, \tilde{1}_j); (c_{v,j}, o_{v,j}) \leftarrow \text{ComPoly}^*(\text{ck}, \tilde{v}_j); (c_{\phi,j}, o_{\phi,j}) \leftarrow \text{ComPoly}^*(\text{ck}, \tilde{\phi}_j) \}$
$\text{ek}_{\phi} := (\text{ek}, \{c_{1,j}, o_{1,j}, c_{v,j}, o_{v,j}, c_{\phi,j}, o_{\phi,j}\}_{j=0}^{\ell}, \phi) ; \text{vk}_{\phi} := (\text{vk}, \{c_{1,j}, c_{v,j}, c_{\phi,j}\}_{j=0}^{\ell})$
$\text{CP}_{\text{sfprm}}.\text{Prove}^*(\text{ek}_{\phi}, \mathbf{x}, (c_j)_{j \in [\ell]}, (\mathbf{u}_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}) \rightarrow \pi :$
$(r, s) \leftarrow H((c_{\phi,j})_{j \in [0, \ell]}, \mathbf{x}, (c_j)_{j \in [\ell]}) \text{ and let } \boldsymbol{\rho} = (1, r, -s) ; (c_0, o_0) \leftarrow \text{ComPoly}(\text{ck}, \tilde{x})$
$\text{for } j = 0 \dots \ell :$
$(c'_j, o'_j) \leftarrow \text{HomEval}(\text{ck}, \boldsymbol{\rho}, (c_j, c_{v,j}, c_{1,j}), (o_j, o_{v,j}, o_{1,j})) ; \mathbf{y}'_j := \mathbf{y}_j + r \cdot \mathbf{v}_j - s \cdot \mathbf{1}_j ; z'_j := \prod_{i=1}^{n_j} y'_{j,i}$
$(c''_j, o''_j) \leftarrow \text{HomEval}(\text{ck}, \boldsymbol{\rho}, (c_j, c_{\phi,j}, c_{1,j}), (o_j, o_{\phi,j}, o_{1,j})) ; \mathbf{y}''_j := \mathbf{y}_j + r \cdot \boldsymbol{\phi}_j - s \cdot \mathbf{1}_j ; z''_j := \prod_{i=1}^{n_j} y''_{j,i}$
$(c_{z'_j}, o_{z'_j}) \leftarrow \text{ComVal}(\text{ck}, z'_j) ; \pi'_j \leftarrow \text{CP}_{\text{ipd}}.\text{Prove}(\text{ek}_p, c_{z'_j}, (c'_{j,i})_{i \in [n_j]}, z'_j, (y'_{j,i})_{i \in [n_j]}, o_{z'_j}, (o'_{j,i})_{i \in [n_j]})$
$(c_{z''_j}, o_{z''_j}) \leftarrow \text{ComVal}(\text{ck}, z''_j) ; \pi''_j \leftarrow \text{CP}_{\text{ipd}}.\text{Prove}(\text{ek}_p, c_{z''_j}, (c''_{j,i})_{i \in [n_j]}, z''_j, (y''_{j,i})_{i \in [n_j]}, o_{z''_j}, (o''_{j,i})_{i \in [n_j]})$
$\text{if } j \stackrel{?}{=} 0 : \{ w'_0 := z'_0 ; w''_0 := z''_0 ; c_{w'_0} := c_{z'_0} ; c_{w''_0} := c_{z''_0} \}$
else :
$w'_j \leftarrow w'_{j-1} \cdot z'_j ; (c_{w'_j}, o_{w'_j}) \leftarrow \text{ComVal}(\text{ck}, w'_j) ; \pi_{w'_j} \leftarrow \text{CP}_{\text{prd}}.\text{Prove}(\text{ck}, c_{w'_{j-1}}, c_{z'_j}, c_{w'_j}, w'_{j-1}, z'_j, w'_j, o_{w'_{j-1}}, o_{z'_j}, o_{w'_j})$
$w''_j \leftarrow w''_{j-1} \cdot z''_j ; (c_{w''_j}, o_{w''_j}) \leftarrow \text{ComVal}(\text{ck}, w''_j) ; \pi_{w''_j} \leftarrow \text{CP}_{\text{prd}}.\text{Prove}(\text{ck}, c_{w''_{j-1}}, c_{z''_j}, c_{w''_j}, w''_{j-1}, z''_j, w''_j, o_{w''_{j-1}}, o_{z''_j}, o_{w''_j})$
endif
endfor
$\pi_z \leftarrow \text{CP}_{\text{eq}}.\text{Prove}(\text{ck}, c_{w'_\ell}, c_{w''_\ell}, w'_\ell, w''_\ell, o_{w'_\ell}, o_{w''_\ell})$
$\text{return } \pi := (c_0, o_0, \{c_{z'_j}, c_{z''_j}, c_{w'_j}, c_{w''_j}, \pi'_j, \pi''_j\}_{j=0}^{\ell}, \{\pi_{w'_j}, \pi_{w''_j}\}_{j=1}^{\ell}, \pi_z)$
$\text{CP}_{\text{sfprm}}.\text{VerProof}^*(\text{vk}_{\phi}, \mathbf{x}, (c_j)_{j \in [\ell]}, \pi) \rightarrow b :$
$(r, s) \leftarrow H((c_{\phi,j})_{j \in [0, \ell]}, \mathbf{x}, (c_j)_{j \in [\ell]}) \text{ and let } \boldsymbol{\rho} = (1, r, -s) ; b \leftarrow \text{VerCommit}(\text{cvk}, c_0, \tilde{x}, o_0)$
$\text{for } j = 0 \dots \ell :$
$(c'_j, \cdot) \leftarrow \text{HomEval}(\text{cvk}, \boldsymbol{\rho}, (c_j, c_{v,j}, c_{1,j}), \cdot)$
$(c''_j, \cdot) \leftarrow \text{HomEval}(\text{cvk}, \boldsymbol{\rho}, (c_j, c_{\phi,j}, c_{1,j}), \cdot)$
$b \leftarrow b \wedge \text{CheckCom}(\text{cvk}, c_{z'_j}) \wedge \text{CP}_{\text{ipd}}.\text{VerProof}(\text{vk}_p, c_{z'_j}, (c'_{j,i})_{i \in [n_j]}) \wedge \text{CheckCom}(\text{cvk}, c_{w'_j})$
$\wedge \text{CheckCom}(\text{cvk}, c_{z''_j}) \wedge \text{CP}_{\text{ipd}}.\text{VerProof}(\text{vk}_p, c_{z''_j}, (c''_{j,i})_{i \in [n_j]}) \wedge \text{CheckCom}(\text{cvk}, c_{w''_j})$
$\text{if } j \neq 0 : \{ b \leftarrow b \wedge \text{CP}_{\text{prd}}.\text{VerProof}(\text{cvk}, c_{w'_{j-1}}, c_{z'_j}, c_{w'_j}, \pi_{w'_j}) \wedge \text{CP}_{\text{prd}}.\text{VerProof}(\text{cvk}, c_{w''_{j-1}}, c_{z''_j}, c_{w''_j}, \pi_{w''_j}) \}$
endfor
$b \leftarrow b \wedge \text{CP}_{\text{eq}}.\text{VerProof}(\text{cvk}, c_{w'_\ell}, c_{w''_\ell}, \pi_z)$

Figure 5: CP-SNARK for specializable universal relation R^{sfprm}

its upper bound for any constant ℓ using the fact that $\sum_{j=0}^{\ell} \log^2 n_j < \ell \log^2 \max\{n_j\}_0^\ell < \ell \log^2 m$. Here the crs size is the same as that in CP_{ipd} : $(2m+3)\mathbb{G}_1 + (\mu+3)\mathbb{G}_2$ elements.

Theorem 5.3. *In the random oracle model, assuming that PolyCom is an extractable and linearly-homomorphic trapdoor commitment, CP_{ipd} , CP_{prd} are zero-knowledge CP-SNARKs for PolyCom and relations R^{ipd} and R^{prd} respectively, then CP_{sfprm} in Figure 5 is a zero-knowledge CP-SNARK for PolyCom and relation R^{sfprm} .*

5.5 A CP-SNARK for Linear Properties of Committed Vector

In this section we show a CP-SNARK for PolyCom that has a specializable universal CRS for relations $R_{\mathbf{F}}^{\text{lin}}(\mathbf{x}, \mathbf{u}) := \mathbf{x} \stackrel{?}{=} \mathbf{F} \cdot \mathbf{u}$ where $\mathbf{F} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{x} \in \mathbb{Z}_q^n$ and $\mathbf{u} \in \mathbb{Z}_q^m$. More precisely, our CP_{lin} works for a family of relations \mathcal{R} that includes all $R_{\mathbf{F}}^{\text{lin}}$ for all matrices $\mathbf{F} \in \mathbb{F}^{n \times m}$.

The scheme is based on the interactive proof for Matrix multiplication of Thaler [Tha13]. In a nutshell, we specialize this protocol to the case of a matrix-vector multiplication and we turn it into a ZK argument using ideas similar to those in [ZGK⁺17b].

Our scheme makes use of the building blocks defined in Section 5.1: a polynomial commitment scheme PolyCom, and CP-SNARKs CP_{poly} and CP_{sc} for the relations R^{poly} and R^{sc} respectively.

Review of Thaler’s Matrix Multiplication protocol. We begin by reviewing the idea of Thaler’s matrix multiplication protocol in our specific case of proving $\mathbf{x} = \mathbf{F} \cdot \mathbf{u}$. Let $\nu := \log n$, $\mu := \log m$. We let $\tilde{F} : \{0, 1\}^\nu \times \{0, 1\}^\mu \rightarrow \mathbb{Z}_q$ be the multilinear extension (MLE) of \mathbf{F} , i.e., the unique multilinear polynomial such that $\tilde{F}(i_1, \dots, i_\nu, j_1, \dots, j_\mu) = F_{i,j}$. Similarly, let \tilde{u} and \tilde{x} be the MLE of \mathbf{u} and \mathbf{x} respectively. The protocol exploits that the MLE \tilde{x} can also be expressed as $\tilde{x}(\mathbf{R}) = \sum_{\mathbf{b} \in \{0,1\}^\mu} \tilde{F}(\mathbf{R}, \mathbf{b}) \cdot \tilde{u}(\mathbf{b})$. In particular, since this MLE is unique, if \tilde{F} and \tilde{u} are MLE of \mathbf{F} and \mathbf{u} respectively, then \tilde{x} is a MLE of $\mathbf{x} = \mathbf{F} \cdot \mathbf{u}$. Next, starting from this observation, the verifier picks a random \mathbf{r} , and then starts a sum-check protocol where the prover convinces the verifier that $t = \tilde{x}(\mathbf{r}) = \sum_{\mathbf{b} \in \{0,1\}^\mu} g(\mathbf{b})$ for the polynomial $g(\mathbf{S}) := \tilde{F}(\mathbf{r}, \mathbf{S}) \cdot \tilde{u}(\mathbf{S})$. At the end of the sum-check the verifier instead of computing $g(\mathbf{s})$ directly, it gets it by evaluating $\tilde{F}(\mathbf{r}, \mathbf{s})$ and $\tilde{u}(\mathbf{s})$ and by computing their product.

The idea to turn the above protocol into a commit and prove argument is rather simple and consists into using a CP-SNARK for the sumcheck relation with a committed polynomial g , or more precisely for the case when a commitment to g is implicitly given through commitments to its factors (see the CP_{sc} scheme). To see this, let us write $g(\mathbf{S}) := \prod_0^p g_i(\mathbf{S})$, where $g_1(\mathbf{S}) := \tilde{F}(\mathbf{r}, \mathbf{S})$, $g_2(\mathbf{S}) = \tilde{u}(\mathbf{S})$, and $g_0(\mathbf{S}) := 1$ is the constant polynomial. A commitment to $\tilde{u}(\mathbf{S})$ is part of the statement, a commitment to $\tilde{F}(\mathbf{R}, \mathbf{S})$ can be generated when specializing the relation to \mathbf{F} in the Derive algorithm. However, note that CP_{sc} expects a commitment to a μ -variate polynomial, whereas \tilde{F} is in $\nu + \mu$ variables. For this, we let the prover commit to the partial evaluation of \tilde{F} on \mathbf{r} , i.e., to the polynomial $g_1(\mathbf{S})$ and uses this commitment and polynomial in CP_{sc} . Then, what is left to show is that such committed $g_1(\mathbf{S})$ is actually the partial evaluation of the other committed polynomial \tilde{F} . To prove this, the idea is that the verifier chooses a random $\sigma \leftarrow_{\$} \mathbb{F}^\mu$, and the prover uses CP_{poly} to prove that $g_1(\sigma) = \tilde{F}(\mathbf{r}, \sigma)$.

We show the full protocol CP_{lin} in Figure 6.

EFFICIENCY. Our CP_{lin} proves the result of a matrix-vector multiplication $\mathbf{F} \cdot \mathbf{u} = \mathbf{x}$ for $\mathbf{F} \in \mathbb{F}^{n \times m}$. The prover sends one CP_{sc} proof, two CP_{poly} proofs, three commitments and two field elements.

$\text{CP}_{\text{lin}}.\text{KeyGen}(\text{ck}) \rightarrow (\text{ek}, \text{vk}) :$	$\text{CP}_{\text{lin}}.\text{Derive}((\text{ek}, \text{vk}), \mathbf{F}) \rightarrow (\text{ek}_F, \text{vk}_F)$
$(\text{ek}_s, \text{vk}_s) \leftarrow \text{CP}_{\text{sc}}.\text{KeyGen}(\text{ck}) ; (\text{ek}_p, \text{vk}_p) \leftarrow \text{CP}_{\text{poly}}.\text{KeyGen}(\text{ck})$ $\text{ek} := (\text{ck}, \text{ek}_s, \text{ek}_p) ; \text{vk} := (\text{cvk}, \text{vk}_s, \text{ek}_p)$	$(c_F, o_F) \leftarrow \text{ComPoly}^*(\text{ck}, \tilde{F})$ $\text{ek}_F := (\text{ek}, c_F, \mathbf{F}, o_F) ; \text{vk}_F := (\text{vk}, c_F)$
$\text{CP}_{\text{lin}}.\text{Prove}^*(\text{ek}_F, \mathbf{x}, c_u, \mathbf{u}, o_u) \rightarrow \pi :$	
Let $g(\mathbf{S}) := \tilde{F}(\mathbf{r}, \mathbf{S}) \cdot \tilde{u}(\mathbf{S}) \equiv g_1(\mathbf{S}) \cdot \tilde{u}(\mathbf{S}) ; g_0(\mathbf{S}) := 1 ; (c_1, o_1) \leftarrow \text{ComPoly}(\text{ck}, g_1)$ $\mathbf{r} \leftarrow H_1(c_F, c_u, \mathbf{x}) ; \sigma \leftarrow H_2(c_F, c_1, \mathbf{r}) ; \mathbf{y}^* \leftarrow g_1(\sigma) ; (c^*, o^*) \leftarrow \text{ComVal}^*(\text{ck}, \mathbf{y}^*)$ $\pi_1 \leftarrow \text{CP}_{\text{poly}}.\text{Prove}(\text{ek}_p, \sigma, (c_1, c^*), (g_1, \mathbf{y}^*), (o_1, o^*))$ $\pi_F \leftarrow \text{CP}_{\text{poly}}.\text{Prove}(\text{ek}_p, (\mathbf{r}, \sigma), (c_F, c^*), (\tilde{F}, \tilde{F}(\mathbf{r}, \sigma)), (o_F, o^*))$ $t \leftarrow \tilde{x}(\mathbf{r}) ; (c_t, o_t) \leftarrow \text{ComVal}(\text{ck}, t) ; \pi_{\text{sc}} \leftarrow \text{CP}_{\text{sc}}.\text{Prove}(\text{ek}_s, g_0(\mathbf{S}), (c_t, c_1, c_u), (t, o_t, g_1, o_1, \tilde{u}, o_u))$ $\pi := (c_t, o_t, c_1, c^*, \mathbf{y}^*, \pi_1, \pi_F, \pi_{\text{sc}})$	
$\text{CP}_{\text{lin}}.\text{VerProof}^*(\text{vk}_F, \mathbf{x}, c_u, \pi) \rightarrow b \in \{0, 1\} :$	
$\mathbf{r} \leftarrow H_1(c_F, c_u, \mathbf{x}) ; t \leftarrow \tilde{x}(\mathbf{r}) ; \sigma \leftarrow H_2(c_F, c_1, \mathbf{r}) ; \text{Let } g_0(\mathbf{S}) := 1$ $b \leftarrow \text{VerCommit}(\text{cvk}, c_t, t, o_t) \wedge \text{VerCommit}^*(\text{cvk}, c^*, \mathbf{y}^*) \wedge \text{CP}_{\text{sc}}.\text{VerProof}(\text{vk}_p, g_0(\mathbf{S}), (c_t, c_1, c_u), \pi_{\text{sc}})$ $\wedge \text{CP}_{\text{poly}}.\text{VerProof}(\text{vk}_s, \sigma, (c_1, c^*), \pi_1) \wedge \text{CP}_{\text{poly}}.\text{VerProof}(\text{vk}_p, (\mathbf{r}, \sigma), (c_F, c^*), \pi_F)$	

Figure 6: CP-SNARK for specializable universal R^{lin}

Here the polynomial used inside sum-check is at most degree 2 in each of its $\mu = \log m$ variables. Here, $p = 2 = d$ because g_0 is the constant polynomial and does not increase the total maximum degree inside sum-check. Also, the number of variables for the first CP_{poly} proof over $g_1(\mathbf{S})$ is $\log m$ and the number of monomials is m . For the second one over $\tilde{F}(\mathbf{r}, \mathbf{S})$ with N monomials, the number of variables is $\log n + \log m = \log N$. The crs output by the derivation function includes matrix \mathbf{F} , whereas the output of $\text{CP}_{\text{lin}}.\text{KeyGen}$ has $(2 \cdot 2^{\nu+\mu} + 3)\mathbb{G}_1 + (\nu + \mu + 3)\mathbb{G}_2$ elements; that is, we are not considering the derived version including the description of matrix \mathbf{F} (as this is part of the statement).

Theorem 5.4. *In the random oracle model, assuming PolyCom is an extractable trapdoor commitment and CP_{poly} and CP_{sc} are zero-knowledge CP-SNARKs for PolyCom, then CP_{lin} in Figure 6 is a zero-knowledge CP-SNARK for PolyCom and relations R^{lin} .*

5.6 A CP-SNARK for Matrix Multiplication

In this section we propose a CP-SNARK for PolyCom for the relation R^{mm} over $\mathcal{D}_A \times \mathcal{D}_B \times \mathcal{D}_C$ where $\mathcal{D}_A = \mathbb{F}^{n \times n'}$, $\mathcal{D}_B = \mathbb{F}^{n' \times m}$, $\mathcal{D}_C = \mathbb{F}^{n \times m}$ and $R^{\text{mm}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) = 1 \iff \mathbf{C} = \mathbf{A} \cdot \mathbf{B}$. Namely, for two committed matrices \mathbf{A} and \mathbf{B} , one can prove that another committed matrix \mathbf{C} equals to their product.

The scheme is inspired by the interactive proof of matrix multiplication of Thaler [Tha13], making it a ZK argument with similar ideas to those in [ZGK⁺17b]. We build our scheme for the polynomial commitment scheme PolyCom and CP-SNARKs CP_{poly} and CP_{sc} for the relations R^{poly} and R^{sc} for factored polynomials.

We present this scheme for square matrices for readability, but the protocol can be simply adapted to the general form. Let matrices $\mathbf{A}, \mathbf{B}, \mathbf{C} \in \mathbb{F}^{n \times n}$, we build their multilinear extension

$\tilde{C} \in \mathbb{F}^\mu \times \mathbb{F}^\mu \rightarrow \mathbb{F}$ with $\mu = \log n$, which is the unique polynomial such that $\tilde{C}(i_1, \dots, i_\mu, j_1, \dots, j_\mu) = C_{i,j}$ if $\{i_k\}_{k=1}^\mu$ and $\{j_k\}_{k=1}^\mu$ are the binary representation of indexes i and j (resp. \tilde{A} and \tilde{B}).

Then, we can represent matrix multiplication as

$$\tilde{C}(\mathbf{I}, \mathbf{J}) = \sum_{\mathbf{b} \in \{0,1\}^\mu} \tilde{A}(\mathbf{I}, \mathbf{b}) \cdot \tilde{B}(\mathbf{b}, \mathbf{J})$$

where (\mathbf{I}, \mathbf{J}) could be seen as a single vector of length 2μ . However, we will stick to this notation instead as it makes clearer that multilinearity is conserved after the product of \tilde{A} and \tilde{B} .

The protocol works as follows. The prover evaluates \tilde{C} on a randomly chosen value $(\boldsymbol{\rho}|\boldsymbol{\sigma})$ and obtains a proof that the output t is indeed the result of the polynomial evaluation $t = \tilde{C}(\boldsymbol{\rho}, \boldsymbol{\sigma})$ using CP_{poly} . Then, the prover convinces the verifier that \tilde{C} is well-formed using CP_{sc} , i.e. $t = \sum_{\mathbf{b} \in \{0,1\}^\mu} g(\mathbf{b})$. Here $g(\mathbf{S}) = \prod_0^2 g_i(\mathbf{S})$, where $g_1(\mathbf{S}) := \tilde{A}(\boldsymbol{\rho}, \mathbf{S})$, $g_2(\mathbf{S}) := \tilde{B}(\mathbf{S}, \boldsymbol{\sigma})$ and $g_0(\mathbf{S}) := 1$ is the all-ones constant polynomial.

EFFICIENCY. The cost of this scheme is given by the complexity of CP_{poly} and CP_{sc} . The proving time of the former is linear in the number of monomials of the polynomial \tilde{C} , which is $2^{2\mu}$ by construction. Similarly, the latter's is linear in the monomials of $g(\mathbf{S})$, which is again $2^{2\mu}$. This makes a linear prover in the number of elements ($N = n^2$). The verifier runtime is linear in the number of variables of the polynomials (i.e. 2μ). The crs size is given by that of CP_{poly} for committed polynomials of length 2μ and $\delta = 1$ (because \tilde{A} , \tilde{B} and \tilde{C} are multilinear polynomials of $\log n + \log n$ variables). That is, it has linear length in the matrix size with $2n^2 + 3 \mathbb{G}_1$ and $2\mu + 3 \mathbb{G}_2$ elements. The proof involves one CP_{poly} proof ($4\mu + 2 \mathbb{G}_1$, with 2μ variables), one CP_{sc} proof ($11\mu + 11 \mathbb{G}_1$ and $\mu + 5 \mathbb{F}$, with μ variables) and one commitment ($2 \mathbb{G}_1$).

Theorem 5.5. *In the random oracle model, assuming that PolyCom is an extractable trapdoor commitment, CP_{poly} and CP_{sc} are zkSNARKs for PolyCom and relations R^{poly} and R^{sc} respectively, then the scheme CP_{mm} ¹⁶ described above is a zkSNARK for PolyCom and relation R^{mm} .*

CP-SNARK for Matrix Multiplication with Known Output

In this section we propose a variation of our CP_{mm} for PolyCom for the relation R^{mmp} over $\mathcal{D}_X \times \mathcal{D}_A \times \mathcal{D}_B$ where $\mathcal{D}_x = \mathbb{F}^{n \times m}$, $\mathcal{D}_A = \mathbb{F}^{n \times n'}$, $\mathcal{D}_B = \mathbb{F}^{n' \times m}$ and $R^{\text{mmp}}(\mathbf{X}, \mathbf{A}, \mathbf{B}) = 1 \iff \mathbf{X} = \mathbf{A} \cdot \mathbf{B}$. Namely, for two committed matrices \mathbf{A} and \mathbf{B} , one can prove that a public matrix \mathbf{X} equals to their product. This version is more efficient than the obvious solution of opening the commitment to \mathbf{C} in the CP_{mm} scheme.

These two versions only vary only subtly on the way the matrix \mathbf{X} is treated. Here, the verifier can check the correct evaluation on a random value $t \stackrel{?}{=} \tilde{X}(\boldsymbol{\rho}, \boldsymbol{\sigma})$, with no need of relying on CP_{poly} . We give the complete protocol in Figure 7 for completeness but we do not provide a formal proof, as its security is trivially implied by that of CP_{mm} .

The asymptotic complexity of the prover in this scheme is the same as that in CP_{mm} . In practice however, the prover is twice faster, as it will not run CP_{poly} . Conversely, the verifier will be slower because evaluating $\tilde{X}(\boldsymbol{\rho}, \boldsymbol{\sigma})$ is more costly than verifying a CP_{poly} proof (about $O(n^2)$ field operations vs. $O(\log n)$ group operations). Note here that evaluating the MLE of \mathbf{X} as

$$\tilde{X}(x_1, \dots, x_{2\mu}) := \sum_{\mathbf{b} \in \{0,1\}^{2\mu}} \chi_{\mathbf{b}}(x_1, \dots, x_{2\mu}) \cdot X(\mathbf{b})$$

¹⁶ This scheme was occasionally referred to as **LegoMM** in the proceedings version of this paper.

$\text{CP}_{\text{mm}}.\text{KeyGen}(\text{ck}) \rightarrow (\text{ek}, \text{vk}) :$ <hr/> $(\text{ek}_s, \text{vk}_s) \leftarrow \text{CP}_{\text{sc}}.\text{KeyGen}(ck)$ $(\text{ek}_p, \text{vk}_p) \leftarrow \text{CP}_{\text{poly}}.\text{KeyGen}(ck)$ return $(\text{ck}, \text{ek}_p, \text{ek}_s, H), (\text{cvk}, \text{vk}_p, \text{vk}_s, H)$ <hr/> $\text{CP}_{\text{mm}}.\text{VerProof}(\text{vk}, c_A, c_B, c_C, \pi) \rightarrow b \in \{0, 1\} :$ Define constant function $g_0(\mathbf{S}) := 1$ $b \leftarrow \text{CP}_{\text{poly}}.\text{VerProof}(\text{vk}_p, (\boldsymbol{\rho}, \boldsymbol{\sigma}), c_C, c_t, \pi_t)$ $\wedge \text{CP}_{\text{sc}}.\text{VerProof}(\text{vk}_s, g_0(\mathbf{S}), (c_t, c_A, c_B), \pi_{\text{sc}})$	$\text{CP}_{\text{mm}}.\text{Prove}(\text{ek}, c_A, c_B, c_C, \mathbf{A}, \mathbf{B}, \mathbf{C}, o_A, o_B, o_C,) \rightarrow \pi :$ <hr/> $(\boldsymbol{\rho} \boldsymbol{\sigma}) \leftarrow H(c_A, c_B, c_C) ; t \leftarrow \tilde{C}(\boldsymbol{\rho}, \boldsymbol{\sigma})$ $(c_t, o_t) \leftarrow \text{ComVal}(\text{ck}, t)$ Define constant function $g_0(\mathbf{S}) := 1$ Let $g(\mathbf{S}) := \tilde{A}(\boldsymbol{\rho}, \mathbf{S}) \cdot \tilde{B}(\mathbf{S}, \boldsymbol{\sigma}) \equiv g_1(\mathbf{S}) \cdot g_2(\mathbf{S})$ $\pi_t \leftarrow \text{CP}_{\text{poly}}.\text{Prove}(\text{ek}_p, (\boldsymbol{\rho}, \boldsymbol{\sigma}), (c_C, c_t), (\tilde{C}, t), (o_C, o_t))$ $\pi_{\text{sc}} \leftarrow \text{CP}_{\text{sc}}.\text{Prove}(\text{ek}_s, g_0(\mathbf{S}), (c_t, c_A, c_B), (t, o_t, \tilde{A}, o_A, \tilde{B}, o_B))$ return $\pi \leftarrow (c_t, \pi_t, \pi_{\text{sc}})$
$\text{CP}_{\text{mmp}}.\text{KeyGen}(\text{ck}) \rightarrow (\text{ek}, \text{vk}) :$ <hr/> $(\text{ek}_s, \text{vk}_s) \leftarrow \text{CP}_{\text{sc}}.\text{KeyGen}(ck)$ return $((\text{ck}, \text{ek}_s, H), (\text{cvk}, \text{vk}_s, H))$ <hr/> $\text{CP}_{\text{mmp}}.\text{VerProof}(\text{vk}, \mathbf{X}, c_A, c_B, \pi) \rightarrow b \in \{0, 1\} :$ <hr/> $(\boldsymbol{\rho}, \boldsymbol{\sigma}) \leftarrow H(X, c_A, c_B) ; t \leftarrow \tilde{X}(\boldsymbol{\rho}, \boldsymbol{\sigma}) ; (c_t, o_t) \leftarrow \text{ComVal}(\text{ck}, t) ; g_0(\mathbf{S}) := 1$ $b \leftarrow \text{CP}_{\text{sc}}.\text{VerProof}(\text{vk}_s, g_0(\mathbf{S}), (c_t, c_A, c_B), \pi)$	$\text{CP}_{\text{mmp}}.\text{Prove}(\text{ek}, c_A, c_B, \mathbf{X}, \mathbf{A}, \mathbf{B}, o_A, o_B) \rightarrow \pi :$ <hr/> $(\boldsymbol{\rho} \boldsymbol{\sigma}) \leftarrow H(X, c_A, c_B) ; t \leftarrow \tilde{X}(\boldsymbol{\rho}, \boldsymbol{\sigma}) ; (c_t, o_t) \leftarrow \text{ComVal}(\text{ck}, t)$ Let $g(\mathbf{S}) := \tilde{A}(\boldsymbol{\rho}, \mathbf{S}) \cdot \tilde{B}(\mathbf{S}, \boldsymbol{\sigma}) \equiv g_1(\mathbf{S}) \cdot g_2(\mathbf{S}) ; g_0(\mathbf{S}) := 1$ $\pi \leftarrow \text{CP}_{\text{sc}}.\text{Prove}(\text{ek}_s, g_0(\mathbf{S}), (c_t, c_A, c_B), (t, o_t, \tilde{A}, o_A, \tilde{B}, o_B))$

Figure 7: CP-SNARK for matrix multiplication with committed output (top) and CP-SNARK for matrix multiplication with known output (bottom)

where $\chi_{\mathbf{b}}(x_1, \dots, x_{2\mu}) := \prod_{j=1}^{2\mu} (\mathbf{b}_j \cdot x_j + (1 - \mathbf{b}_j)(1 - x_j))$, takes both the prover and the verifier $2^{2\mu} \cdot 2\mu$ operations naively. Following the strategy of [Tha13], the terms $\chi_{\mathbf{b}}$ can be precomputed offline so that computing each $\chi_{\mathbf{b}}(\mathbf{x}) \cdot X(\mathbf{b})$ takes a constant time and evaluating $\tilde{X}(\boldsymbol{\rho}, \boldsymbol{\sigma})$ becomes a quadratic-time task in n (as $2^{2\mu} = n^2$).

EFFICIENCY. Our CP_{mmp} proves matrix multiplication $\mathbf{X} = \mathbf{A} \cdot \mathbf{B}$ where the output \mathbf{X} is given in clear. We consider square matrices of $N = n \times n$ elements with $n = 2^\mu$. Both prover and verifier evaluate the \tilde{X} on a 2μ -length random point. By construction of the MLE $\tilde{X}(r_1, \dots, r_{2\mu}) = \sum_{\mathbf{b} \in \{0,1\}^{2\mu}} X(\mathbf{b}) \prod_{j=1}^{2\mu} \chi_{\mathbf{b}_j}(r_j)$, this carries a cost of $2^{2\mu} \cdot 2\mu = 2n^2 \log n = O(N \log n)$ field operations, which can be reduced to $O(N)$ through dynamic programming techniques [Tha13]. As the degree-2 polynomial $g(\mathbf{S})$ inside sum-check has μ variables, the proving algorithm in CP_{sc} involves $O(\mu)$ group operations. Checking π_{sc} requires $O(\mu)$ more group operations from the verifier. The proof contains one CP_{sc} proof consisting of $(11\mu + 11)\mathbb{G}_1$ and $(\mu + 5)\mathbb{F}$. The crs in this scheme is the same as the one in CP_{sc} for $p = 2 = d$ and $\delta = 1$, which also coincides with that of CP_{mm} . That is, it has linear length in the matrix size with $(2n^2 + 3)\mathbb{G}_1$ and $(2\mu + 3)\mathbb{G}_2$ elements. We do not include the description of the public matrix as this is part of the statement.

Theorem 5.6. *In the random oracle model, assuming that PolyCom is an extractable trapdoor commitment and CP_{sc} is a zkSNARK for PolyCom and relation R^{sc} , then the scheme CP_{mmp} is a zk-SNARK for PolyCom and relation R^{mmp} .*

6 LegoSNARK Applications and Evaluation

In this section we show how to use the modular commit-and-prove approach to obtain new CP-SNARKs for computations expressible by arithmetic circuits (ACs) and we discuss the resulting instantiations. Precisely, we show new CP-SNARKs for (1) arithmetic circuit satisfiability, and (2) parallel computation on joint inputs.

In both constructions the idea is to break the target problem into the conjunction of simpler relations with shared input. Once having done this, and assuming the existence of CP-SNARKs for these simpler relations and that share the same commitment scheme, we immediately obtain a CP-SNARK for the target problem by applying our composition Theorem 3.1. Furthermore, thanks to our lifting transformation of Section 3.5 sharing the same commitment scheme is not a restricting requirement.

6.1 Preliminaries and Building Blocks

We begin by formalizing some basic relations useful to express our target problems.

Equalities Among Vector Entries. A common building block in both schemes of this section is a system for proving that the entries of a vector satisfy a set of equalities between them. Namely, given a set S of pairs of indices (i, k) , we define a relation R_S^{veq} that holds for a vector \mathbf{u} iff $u_i = u_k$ for all $(i, k) \in S$.

Definition 6.1 (Relation for equalities among vector entries). *Let \mathcal{D} be some domain (e.g., a finite field \mathbb{F}), let n_0, \dots, n_ℓ be positive integers such that $\mathcal{D}_j := \mathcal{D}^{n_j}$ and let $m = \sum_{j=0}^{\ell} n_j$. Given a set $S = \{(i_1, k_1), \dots, (i_\ell, k_\ell)\} \subset [m] \times [m]$, we define a relation R_S^{veq} over $\mathcal{D}_0 \times \dots \times \mathcal{D}_\ell = \mathcal{D}^m$ such that: $R_S^{\text{veq}}(\mathbf{x}, (\mathbf{u}_j)_{j \in [\ell]}) = 1 \iff \forall (i, k) \in S : y_i = y_k$, where $\mathbf{y} := (\mathbf{x}, (\mathbf{u}_j)_{j \in [\ell]})$.*

In what follows, we discuss different ways to encode this relation.

The relation R_S^{veq} can be expressed using R_ϕ^{sfprm} in Definition 5.2 for an appropriate permutation ϕ that encodes S . The idea is that a set $S \subset [m] \times [m]$ can be seen as the description of an undirected graph with $2m$ vertices. From S it is possible to extract another set $S' \subset [m] \times [m]$ that contains a cycle $((i_1, k_1), \dots, (i_t, k_t))$ for every connected component of the graph represented by S . Taking the product of all the cycles in S' defines a permutation $\phi : [m] \rightarrow [m]$ such that $\forall (i, k) \in S : y_i = y_k$ iff $\forall j \in [m] : y_j = y_{\phi(j)}$. Then for such ϕ computed from S we have $R_S^{\text{veq}}(\mathbf{x}, (\mathbf{u}_j)_{j \in [\ell]}) \iff R_\phi^{\text{sfprm}}(\mathbf{x}, (\mathbf{u}_j)_{j \in [\ell]})$. We refer to [Gro09, BCG⁺17] for more details on the idea of this permutation encoding. Here is a small example. Consider an arbitrary m and assume $S = \{(1, 2), (1, 3), (3, 4), (6, 7)\}$. One can define a permutation ϕ_S over $[m]$ as: $\phi_S(1) = 2$, $\phi_S(2) = 3$, $\phi_S(3) = 4$, $\phi_S(4) = 1$, $\phi_S(6) = 7$, $\phi_S(7) = 6$, and $\phi_S(j) = j$ for all $8 \leq j \leq m$.

At this point one can either assume to have a proof system for R_ϕ^{sfprm} (as in Section 5.4) or to use an encoding of R_ϕ^{sfprm} based on linear constraints that can be obtained as follows. The idea is to define a relation on a vector $\mathbf{y} \in \mathbb{F}^m$ that is true iff $\mathbf{Z} \cdot \mathbf{y} = \mathbf{0}$, where $\mathbf{Z} \in \mathbb{F}^{m' \times m}$, with $m' \leq m$, is the matrix obtained by removing the zero rows from $(\mathbf{I} - \Sigma_\phi) \in \mathbb{F}^{m \times m}$, where Σ_ϕ is the permutation matrix representing ϕ . Then clearly $R_\phi^{\text{sfprm}}(\mathbf{x}, (\mathbf{u}_j)_{j \in [\ell]})$ holds iff $R_{\mathbf{Z}}^{\text{lin}}(\mathbf{0}, \mathbf{x}, (\mathbf{u}_j)_{j \in [\ell]})$ holds, where the relation R^{lin} , modelling the linear property over (committed) vectors, is formally defined as follows.

Definition 6.2 (Linear property relation). Let $n_1, n_2, m_1, \dots, m_\ell$ be integers such that $\{\mathcal{D}_{x,j} := \mathcal{D}^{n_j}\}_{j \in [1,2]}$, $\{\mathcal{D}_{u,j} := \mathcal{D}^{m_j}\}_{j \in [\ell]}$, and $m = n_2 + \sum_{j=1}^{\ell} m_j$. Given a matrix $\mathbf{F} \in \mathcal{D}^{n_1 \times m}$, we define a relation $R_{\mathbf{F}}^{\text{lin}}$ over $\mathcal{D}_{x,1} \times \mathcal{D}_{x,2} \times \mathcal{D}_{u,1} \times \dots \times \mathcal{D}_{u,\ell}$ such that:

$$R_{\mathbf{F}}^{\text{lin}}(\mathbf{x}_1, \mathbf{x}_2, (\mathbf{u}_j)_{j \in [\ell]}) = 1 \iff \mathbf{F} \cdot \mathbf{y} = \mathbf{x}_1, \text{ where } \mathbf{y} := (\mathbf{x}_2, (\mathbf{u}_j)_{j \in [\ell]})$$

Note that the above relation R^{lin} is slightly different from the one supported by $\text{CP}_{\text{lin}}^{\text{Ped}}$ of Section 4.2. The only difference is that in $\text{CP}_{\text{lin}}^{\text{Ped}}$ the linear function is not applied over public inputs. However, this small discrepancy can be easily solved by adding a commitment to the additional public input \mathbf{x}_2 and opening this commitment.

Summary of the Building Blocks. In Table 2 we recall the various commit-and-prove SNARKs presented Sections 4 and 5 along with a summary of their efficiency measures. First-level dependencies between the different building blocks can be found in the first column of the table. We wanted to show the minimal requirements to build such constructions, regardless of the inner instantiation of each modular component. That is, for each row we are pointing out the CP-SNARKs that appear only in the description of their respective protocol.

6.2 Arithmetic Circuit Satisfiability

Let us consider the problem of arithmetic circuit satisfiability.

Definition 6.3. Let $C : \mathbb{F}^{n_x} \times \mathbb{F}^{n_w} \rightarrow \mathbb{F}^l$ be an arithmetic circuit, where $n_x, n_w, l \in \mathbb{N}$ denote input, witness and output length. We define the arithmetic circuit satisfiability relation $R_C^{\text{ac}}(\mathbf{x}, \mathbf{w})$ as the set of pairs such that $C(\mathbf{x}, \mathbf{w}) = \mathbf{0}^l$.

CP-SNARKs		\mathcal{P}	\mathcal{V}	crs		\pi	
Dependencies	Scheme			\mathbb{G}_1	\mathbb{G}_2	\mathbb{G}_1	\mathbb{F}
	CP_{link}	$O(m + \ell)$	$O(\ell)$	$m + \ell + 1$	$\ell + 2$	1	0
	$\text{CP}_{\text{lin}}^{\text{Ped}}$	$O(m + \ell)$	$O(\ell + n)$	$m + \ell$	$n + \ell + 1$	1	0
$\text{CP}_{\text{eq}} \wedge \text{CP}_{\text{poly}} \wedge \text{CP}_{\text{prd}} \rightarrow$	CP_{sc}	$O(m)$	$O(\mu)$	$2(\delta + 1)^\mu + 3$	$\mu + 3$	$\mu(2d + 2p + 3) + 4p + 3$	$\mu + 5$
$\text{CP}_{\text{poly}} \wedge \text{CP}_{\text{sc}} \rightarrow$	CP_{lin}	$O(N)$	$O(\log N)$	$2N + 3$	$\log N + 3$	$2 \log N + 13\mu + 21$	$\mu + 7$
$\text{CP}_{\text{sc}} \rightarrow$	CP_{mmp}	$O(n^2)$	$O(n^2)$	$2n^2 + 3$	$2\mu + 3$	$11\mu + 11$	$\mu + 5$
$\text{CP}_{\text{poly}} \wedge \text{CP}_{\text{sc}} \rightarrow$	CP_{mm}	$O(n^2)$	$O(\mu)$	$2n^2 + 3$	$2\mu + 3$	$15\mu + 15$	$2\mu + 5$
$\text{CP}_{\text{poly}} \wedge \text{CP}_{\text{sc}} \rightarrow$	CP_{had}	$O(n)$	$O(\mu)$	$2n + 3$	$\mu + 3$	$15\mu + 15$	$\mu + 5$
$\text{CP}_{\text{eq}} \wedge \text{CP}_{\text{poly}} \wedge \text{CP}_{\text{prd}} \rightarrow$	CP_{ipd}	$O(n)$	$O(\mu^2)$	$2n + 3$	$\mu + 3$	$11\mu^2/2 + 25\mu/2 + 2$	$\mu^2/2 + 15\mu/2$
$\text{CP}_{\text{eq}} \wedge \text{CP}_{\text{ipd}} \wedge \text{CP}_{\text{prd}} \rightarrow$	CP_{sfprn}	$O(n)$	$O(\mu^2)$	$2n + 3$	$\mu + 3$	$\sum_{j=0}^{\ell} (11\mu_j^2 + 25\mu_j) + 16\ell + 13$	$\sum_{j=0}^{\ell} (\mu_j^2 + 15\mu_j) + 10\ell + 2$

Table 2: Direct dependencies and concrete theoretical costs of our CP-SNARKs. In CP_{sc} we denote: by m the number of monomials in the sumcheck polynomial; by μ the number of variables in the sumcheck polynomial (note $m \leq 2^\mu$); by δ the degree of the committed polynomial (if $\delta = 1$, it holds that $(\delta + 1)^\mu \geq m$); by d the maximum degree of each variable in the sumcheck polynomial; by p the number of polynomial factors committed. For all remaining schemes, we denote: by n the length of vectors, the order of square matrices or the largest dimension in a matrix; by μ the size of the multilinear extensions involved, which above it is always such that $\mu = \log n$; by N the size of the CP_{lin} and $\text{CP}_{\text{lin}}^{\text{Ped}}$ matrix, of dimension $2^\nu \times 2^\mu = n \times m$; by μ_j the logarithm of the length of the j -th set of variables; by ℓ the total number of commitments.

We show two solutions to model the above relation using a commit-and-prove paradigm. The first one relies on the encoding put forward by Bootle et al. [BCC⁺16] that reduces the relation R^{ac} to an Hadamard product and a set of linear constraints. The second one is similar to that of Groth [Gro09] (recently used in [BCG⁺17]) and encodes arithmetic circuit satisfiability using Hadamard products, additions and permutations of (committed) vectors.

Arithmetic Circuit Satisfiability through Hadamard and Linear Constraints

Following [BCC⁺16, BBB⁺17], an arithmetic circuit C can be described by a tuple $(n_x, n_u, N, \mathbf{W}_L, \mathbf{W}_R, \mathbf{W}_O, \mathbf{W}_x, \mathbf{W}_U, \mathbf{c})$ where n_x and n_u are the input and (committed) witness lengths respectively, N is the number of multiplication gates, and matrices $\mathbf{W}_L, \mathbf{W}_R, \mathbf{W}_O \in \mathbb{F}^{Q \times N}$, $\mathbf{W}_x \in \mathbb{F}^{Q \times n_x}$, $\mathbf{W}_U \in \mathbb{F}^{Q \times n_u}$ and vector $\mathbf{c} \in \mathbb{F}^Q$ describe a system of linear equations over the wires of C . Using such a definition, C is satisfied by (\mathbf{x}, \mathbf{u}) if there exist three vectors $\mathbf{u}_L^M, \mathbf{u}_R^M, \mathbf{u}_O^M \in \mathbb{F}^N$ such that

$$\mathbf{u}_L^M \circ \mathbf{u}_R^M = \mathbf{u}_O^M \wedge \mathbf{W}_L \cdot \mathbf{u}_L^M + \mathbf{W}_R \cdot \mathbf{u}_R^M + \mathbf{W}_O \cdot \mathbf{u}_O^M + \mathbf{W}_x \cdot \mathbf{x} + \mathbf{W}_U \cdot \mathbf{u} = \mathbf{c}$$

$$R_C^{\text{ac}}(\mathbf{x}, \mathbf{u}, \mathbf{u}_w) := R^{\text{had}}(\mathbf{u}_L^M, \mathbf{u}_R^M, \mathbf{u}_O^M) \wedge R_{\mathbf{F}}^{\text{lin}}(\mathbf{c}, (\mathbf{x}, \mathbf{u}, \mathbf{u}_L^M, \mathbf{u}_R^M, \mathbf{u}_O^M))$$

where $\mathbf{F} = (\mathbf{W}_x, \mathbf{W}_U, \mathbf{W}_L, \mathbf{W}_R, \mathbf{W}_O) \in \mathbb{F}^{Q \times (n_x + n_u + 3N)}$.

By the above definition of R_C^{ac} and our Theorem 3.1 we obtain the following corollary.

Corollary 6.1. *If there exist CP-SNARKs CP_{had} and CP_{lin} for a commitment scheme Com and for relations R^{had} and R^{lin} respectively, then there is a CP-SNARK LegoAC for Com and relation R_C^{ac} .*

INSTANTIATIONS. We evaluate two instantiations of LegoAC :

- LegoAC1 : from our $\text{CP}_{\text{lin}}^{\text{Ped}}$ (Section 4.2) and Lipmaa’s CP-SNARK for Hadamard products [Lip16]. LegoAC1 is a CP-SNARK for the commitment scheme of [Lip16], and its security holds in the generic group model (due to GGM security of $\text{CP}_{\text{lin}}^{\text{Ped}}$).
- LegoAC2 : from our $\text{CP}_{\text{lin}}^{\text{Ped}}$ (Section 4.2) and our CP_{had} (Section 5.3). This is a CP-SNARK for PolyCom , and its security holds in the GGM and random oracle model (the latter due to CP_{had}).

If needed, both schemes can be lifted to work with a standard Pedersen commitment using CP_{link} . Their complexity, summarized in Table 1, results from the combined efficiency of the building blocks plus the observation that the matrices $\mathbf{W}_L, \mathbf{W}_R, \mathbf{W}_O$ are sparse and with a number of nonzero entries linear in the number of circuit wires.

Arithmetic Circuit Satisfiability through Hadamard, Addition and Equalities

Any arithmetic circuit C consists of N_A addition gates, N_M multiplication gates, both of fan-in 2, and N_C multiplication-by-constant gates, of fan-in 1. Each gate has a left input, a right input and an output wire;¹⁷ also each output wire can be input to another gate. This means that C can be described by integers N_A, N_M, N_C , a vector $\mathbf{c} \in \mathbb{F}^{N_C}$ of constants, and the wiring information saying that the output wire of addition/multiplication i is the left/right input of addition/multiplication gate j . With such a representation $\exists \mathbf{w} : C(\mathbf{x}, \mathbf{w}) = \mathbf{0}^l$ can be encoded by showing the existence of an assignment to the inputs and outputs of C ’s gates that satisfies every gate, that is consistent with the wiring of C as well as with the public input \mathbf{x} and the output $\mathbf{0}$.

¹⁷ We model gates of fan-in 1 as having one single left input.

More formally, consider an arithmetic circuit $C : \mathbb{F}^{n_x} \times \mathbb{F}^{n_w} \rightarrow \mathbb{F}^l$ with N_A addition gates, N_M multiplication gates, and N_C multiplication by constant gates, and where we split the witness \mathbf{w} between committed witness $\mathbf{u} \in \mathbb{F}^{n_u}$ and free witness $\boldsymbol{\omega} \in \mathbb{F}^{n_\omega}$. Assume we arrange the wires of C so as to have, orderly: the n_x input wires, the n_u committed witness wires, the l output wires, the $3N_A$ left, right and output wires of the addition gates, the $3N_M$ left, right and output wires of the multiplication gates, and the $2N_C$ input and output wires of the multiplication-by-constant gates. All these wires can be indexed by integers from 1 to $m = n_x + n_u + l + 3(N_A + N_M) + 2N_C$, and the wiring information of C can be described by a set S of pairs $(i, k) \in [m] \times [m]$ indicating that the wire at position i is connected to the wire at position k .

Therefore we model an arithmetic circuit C with a tuple $(n_x, n_u, l, N_A, N_M, N_C, \mathbf{c}, S)$. Then proving $\exists(\mathbf{u}, \boldsymbol{\omega}) R_C^{\text{ac}}(\mathbf{x}, \mathbf{u}, \boldsymbol{\omega})$ can be done by proving the existence of a vector \mathbf{u}_w , that is the concatenation of vectors $\mathbf{u}_w := (\mathbf{u}_L^A, \mathbf{u}_R^A, \mathbf{u}_O^A, \mathbf{u}_L^M, \mathbf{u}_R^M, \mathbf{u}_O^M, \mathbf{u}_I^C, \mathbf{u}_O^C)$, such that

$$\begin{aligned} R_C^{\text{ac}}(\mathbf{x}, \mathbf{u}, \mathbf{u}_w) &:= R_{\text{add}}(\mathbf{u}_L^A, \mathbf{u}_R^A, \mathbf{u}_O^A) \wedge R^{\text{had}}(\mathbf{u}_I^C, \mathbf{c}, \mathbf{u}_O^C) \\ &\wedge R^{\text{had}}(\mathbf{u}_L^M, \mathbf{u}_R^M, \mathbf{u}_O^M) \wedge R_S^{\text{veq}}((\mathbf{x}, \mathbf{0}), \mathbf{u}, \mathbf{u}_w) \end{aligned}$$

where $R_{\text{add}}(\mathbf{u}_L^A, \mathbf{u}_R^A, \mathbf{u}_O^A)$ is the relation expressing the predicate $\mathbf{u}_L^A + \mathbf{u}_R^A \stackrel{?}{=} \mathbf{u}_O^A$, and $R^{\text{had}}(\mathbf{u}_L^M, \mathbf{u}_R^M, \mathbf{u}_O^M)$ is the Hadamard product relation $\mathbf{u}_L^M \circ \mathbf{u}_R^M \stackrel{?}{=} \mathbf{u}_O^M$ (i.e., $u_{L,j}^M \cdot u_{R,j}^M = u_{O,j}^M$ for all $j \in [3N_M]$).

If Com is a linearly homomorphic and extractable commitment scheme, a proof system for R_{add} comes for free. Therefore, by definition of R_C^{ac} and our Theorem 3.1 we obtain the following corollary.

Corollary 6.2. *If there exist CP-SNARKs CP_{had} and CP_{veq} for a linearly-homomorphic extractable commitment scheme Com and for relations R^{had} and R^{veq} respectively, then there is a CP-SNARK LegoUAC for Com and relation R_C^{ac} .*

Instantiating LegoUAC with Universal CRS. Both the schemes LegoAC1 and LegoAC2 considered earlier have a circuit-specific CRS due to the circuit-specific CRS of $\text{CP}_{\text{lin}}^{\text{ped}}$.¹⁸ The LegoUAC scheme obtained in the corollary above can be instantiated in such a way to have universal CRS of linear-size. To this end, we recall that R^{veq} can be expressing using R^{sfprm} , and therefore we evaluate an instantiation of LegoUAC with our CP_{had} and CP_{sfprm} schemes. Both schemes admit a universal CRS that can be deterministically specialized (due to specializing CP_{sfprm} 's CRS to the circuit-dependent permutation ϕ). The complexity of LegoUAC is depicted in Table 1 and stems from that of our CP_{had} and CP_{sfprm} .

6.3 Parallel Computation on Joint Inputs

We consider the problem of proving (in zero-knowledge) the correctness of a computation that consists in the parallel execution of the same subcomputation on (partially) shared inputs. Slightly more in detail, we consider relations $R^{\text{parjnt}}(u) := \bigwedge_{j=1}^N R'(u'_j)$ where each u'_j is a subset of the entries of u . This relation has several use cases. One example is proving knowledge of all the leaves of a Merkle tree of height N with respect to a public root; the corresponding relation can be seen as the parallel check of $2^N - 1$ hash verification relations (i.e., $R_H(x_1, x_2, y) := H(x_1, x_2) \stackrel{?}{=} y$) that share some of the inputs. Another example is proving correctness of the output of a sequential computation whose internal step is always the same (e.g., the square-and-multiply algorithm).

¹⁸ Using our CP_{lin} for PolyCom would give us an instantiation with a universal CRS, but unfortunately one of size $Q \cdot N$, that is quadratic in circuit size.




Figure 8: Inputs structures for parallel relations.

One way to deal with R^{parjnt} is by defining the arithmetic circuit that computes it (cf. Fig. 8b). The Hyrax system is particularly designed for parallel circuits [WTs⁺18]; they deal with non-parallel input by introducing a (non-parallel) redistribution layer (RDL) layer that *redistributes* the input and feeds it to the identical sub-circuits at the next level. Unfortunately an effect of using an RDL is that the verifier must pay an additional cost linear in the total width of the circuit. This makes verification time pretty high in applications like the Merkle tree example above.

Here we propose another natural modelling of relations with joint inputs, that is the simple conjunction of two relations: R^{par} that models fully parallel checks of some R' on *disjoint* inputs, and another relation that models the consistency of the shared inputs across the (fully) parallel executions. The advantage of this encoding is that R^{par} is now fully parallel and one could use for it a system for parallel computation without any caveat, whereas to check the consistency of shared input one can use a system for the R^{veq} relation from Definition 6.1. More formally, we define a parallel relation on disjoint inputs as follows.

Definition 6.4 (Parallel relation on disjoint inputs). For a relation R' over \mathcal{D}' and an integer $N \geq 1$, a parallel relation $R_{R'}^{\text{par}}$ on disjoint inputs is defined as $R_{R'}^{\text{par}}(\mathbf{u}) := \bigwedge_{j=1}^N R'(u_j)$, where $\mathbf{u} := (u_j)_{j \in [N]} \in (\mathcal{D}')^N$.

From R^{par} and R^{veq} we define a relation for parallel checks on joint inputs.

Definition 6.5 (Parallel relation on joint inputs). Let $n_0, n_1, n', N \in \mathbb{N}$ be integers such that $n', N \geq 1$ and $n_0, n_1 \geq 0$, and let $m = n_0 + n_1 + N \cdot n'$. Let \mathcal{D} be some domain, R' be a relation over $\mathcal{D}' := \mathcal{D}^{n'}$, and S a set of the form $S = \{(i_1, k_1), \dots, (i_l, k_l)\} \subset [m] \times [m]$. $R_{R',S}^{\text{parjnt}}$ is a relation over $\mathcal{D}_x \times \mathcal{D}_1 \times \mathcal{D}_2$, with $\mathcal{D}_x := \mathcal{D}^{n_0}$, $\mathcal{D}_1 := \mathcal{D}^{n_1}$ and $\mathcal{D}_2 := \mathcal{D}^{Nn'}$, such that:

$$R_{R',S}^{\text{parjnt}}(\mathbf{x}, \mathbf{u}_1, \mathbf{u}_2) := R_{R'}^{\text{par}}(\mathbf{u}_2) \wedge R_S^{\text{veq}}(\mathbf{x}, \mathbf{u}_1, \mathbf{u}_2)$$

Basically, $R_{R',S}^{\text{parjnt}}$ models the parallel checking of R' on N different subsets of the entries of $(\mathbf{x}, \mathbf{u}_1)$ (consisting of a public \mathbf{x} and committed \mathbf{u}_1) where such subsets are defined by the set S , and their concatenation is the vector \mathbf{u}_2 . Alternatively, if \mathbf{x}, \mathbf{u}_1 are empty, $R_{R',S}^{\text{parjnt}}$ models the parallel checking of R' on N different sets of inputs with some shared values (as specified by S).

From the definition of R^{parjnt} and our Theorem 3.1 we obtain the following corollary.

Corollary 6.3. If there exist CP-SNARKs CP_{par} and CP_{veq} for a commitment scheme Com relations R^{par} and R^{veq} respectively, then there is a CP-SNARK $\text{CP}_{\text{parjnt}}$ for Com and relations R^{parjnt} .

INSTANTIATIONS. Following the corollary above, we consider an instantiation of $\text{CP}_{\text{parjnt}}$ (that we call **LegoPar**) obtained as follows. As CP_{veq} we use our scheme $\text{CP}_{\text{lin}}^{\text{Ped}}$ using the encoding of R^{veq} with linear constraints. As CP_{par} we use an adaptation of Hyrax using the polynomial commitment PolyCom of zk-vSQL. We call **HyrPoly-Par** this scheme invoked on circuits without RDL (i.e., it supports R^{par}), and **HyrPoly-RDL** the same scheme for circuits with an RDL (i.e., it supports R^{parjnt}).

We compare the efficiency of **LegoPar** and **HyrPoly-RDL** on R^{parjnt} relations. Let d and G be depth and width of the arithmetic circuit evaluating R' . Proving time and proof size have the same complexity in both solutions; verifier time is $O(d(G+\log(NG)))$ in **LegoPar** and $O(d(G+\log(NG))+|u|+NG)$ in **HyrPoly-RDL**. We note that due to the use of $\text{CP}_{\text{lin}}^{\text{Ped}}$, the CRS of **LegoPar** becomes specific to the input wiring of R^{parjnt} , whereas in **HyrPoly-RDL** the CRS is just the commitment key. On the other hand, this one-time preprocessing allows the verifier to later check any number of proofs in shorter time.¹⁹ In Section 7 we discuss an experimental comparison based on an implementation.

7 Experimental Evaluation

We provide an implementation for **LegoSNARK**²⁰ that includes the following gadgets: our CP_{link} and $\text{CP}_{\text{lin}}^{\text{Ped}}$, the Hadamard product CP-SNARK of [Lip16], and the CP_{poly} from [ZGK⁺17b]²¹. **LegoSNARK** is written in C++; for polynomial operations and bilinear pairings we use the libraries underlying **libsark** [librk]. We executed our experiments on a virtual machine running Debian GNU/Linux with 8 Xeon Gold 6154 cores and 30 GB of RAM. We ran all tests single threaded. In our experiments, we tested the performance of some of our instantiations and compared to different baseline systems.

7.1 Commit-and-Prove SNARKs

We consider a generic application of proving commit-and-prove relations where commitments are created using the Pedersen scheme for vectors, i.e., proving $\exists(u, o, \omega) R(u, \omega) \wedge \text{VerCommit}(\text{ck}, c, u, o)$.

As baseline system, we use the Groth16 zkSNARK in **libsark** on the **libsark** gadget circuit for multi-scalar additions over a SNARK-friendly elliptic curve (to model the Pedersen computation). We call this **CPGro16**. We compare **CPGro16** to a CP-SNARK, **LegoGro16**, obtained by applying our cc-SNARK-lifting compiler with our CP_{link} scheme to the cc-variant of [Gro16], **ccGro16**, that we present in Appendix H.5. We measured the overhead of dealing with the commitment in both schemes (the R -dependent costs would be the same in both cases) at the increase of the committed vector’s size (from 8 to 2048).²² On the largest instance ($n = 2048$) **LegoGro16**’s proving time is $5\text{K}\times$ (0.08 vs. 428 s) faster than **CPGro16**, at the price of a verification that is $1.2\times$ slower (4.1 vs 3.4 ms), and a slightly larger proof (191 vs. 127 Bytes). **LegoGro16**’s CRS is also $7\text{K}\times$ shorter (130KB vs. 950MB).

In the case of **LegoGro16**, such overhead in proving time is essentially that of creating the additional element D of the proof that contains a commitment to the data and to create a CP_{link} proof to link D to the external commitment. The **LegoGro16** proof is longer because of these two additional elements of \mathbb{G}_1 . And for verification, the CP_{link} verification must be executed. With respect to the CRS, in **LegoGro16** we have the additional elements of the CRS needed to create D and the CP_{link} CRS, that is essentially one vector of n elements of \mathbb{G}_1 . In **CPGro16**, all the overhead in proving time and CRS is related to the size and degree of the QAP that models the computation

¹⁹ We do not see a way to run a similar preprocessing in **HyrPoly**. We evaluated the possibility to commit, in preprocessing, to the MLE of the RDL wiring so that the prover would compute this on behalf of the verifier and prove its correct evaluation using CP_{poly} . This idea however would require a commitment key quadratic in the circuit width, which is prohibitively large.

²⁰ The GitHub repository for the code is <https://github.com/imdea-software/legosnark>.

²¹ For this we adapted to our library the code provided by the authors of [ZGK⁺17b].

²² At $n = 4096$ **CPGro16** ran out of memory.

of the Pedersen commitment. This was done by selecting an appropriate gadget in `libsark`, which optimizes the task by selecting a suitable elliptic curve.

Table 3 shows our experimental results that compare the schemes `LegoGro16` and `CPGro16` with respect to the overhead for dealing with data committed using a Pedersen vector commitment. The experiments considered vectors of different length n .

7.2 Matrix Multiplication

We evaluate our CP_{mmp} scheme for matrix multiplication against a solution based on Groth16 [Gro16]. We remind the reader that in this version of matrix product relation the output matrix is given in the clear as part of the statement to be proven (rather than being committed as in CP_{mm}). Our scheme has a faster prover and smaller `crs` using an asymptotically more efficient verifier with a longer, but still succinct, proof. Our experiments confirm the theoretical costs of these schemes.

We evaluate both proving and verification time when delegating a square matrix multiplication with size $N = n \times n$ field elements, ranging from $n = 16$ to $n = 256$. We observe our scheme noticeably improves on proving time as our prover runs in linear time in the number of elements in the matrix (n^2), whereas Groth’s runs in quasicubic time in n . Even if our verifier is slower for smaller matrices, the $O(n^2)$ work in our scheme involves only *field operations* whereas in Groth16 one needs to do a $O(n^2)$ -wide multiexponentiation. On the largest instance we measured (square matrices with $n = 128$ rows and columns), our proving time is roughly $1300\times$ faster (109 seconds vs. 84 milliseconds) and verification time is $1.8\times$ faster (51 vs. 28 milliseconds). This is a tradeoff between the running time and the proof length: only 3 group elements in Groth16 vs $O(\log n)$ in our scheme (127 bytes vs. 32 kilobytes).

Table 4 shows concrete performance measurements of both schemes, showing a clear proving time improvement in our scheme.

7.3 LegoAC1 for Arithmetic Circuits

We tested our `LegoAC1` scheme (see Section 6.2) for arithmetic circuits and compared it to Groth16 as a baseline system. We considered two benchmark applications:

- (a) proving knowledge of a *SHA256 pre-image* on 512-bit inputs; for this we used the existing circuit gadgets implemented in `libsark` (for Groth16), and in `Bulletproofs [bulk1]` (for `LegoAC1`).
- (b) *matrix factoring*, i.e., proving knowledge of two $n \times n$ matrices A, B whose product is a public matrix C ; for this we designed suitable constraints systems, considering 32-bit integers entries and a varying $n = 16, 32, 64, 128$.

Overall, our experiments show that `LegoAC1` performs slightly worse than Groth16. For example, for SHA256 proving time is $1.2\times$ slower (0.7 vs. 0.9 s); verification is up to $2\times$ slower (0.9 vs. 1.8 ms) and improves with larger inputs; our key generation is about $5 - 6\times$ slower. Proof size is constant: 350B in `LegoAC1` and 127B in Groth16. Noteworthy that most of `LegoAC1` key generation time (about 70%) is taken by the corresponding algorithm for $\text{CP}_{\text{lin}}^{\text{Ped}}$; this is mainly due to an unoptimized technique for dealing with sparse matrices like the ones that encode the linear constraints $\mathbf{W}_L, \mathbf{W}_R, \mathbf{W}_O$, and we expect this to be improved in the future.

In a way this result is not surprising: Groth16 is an extremely optimized and well explored scheme, whereas for `LegoAC1` we believe that more optimizations could be explored (in a similar way as Groth16 optimized Pinocchio). More remarkably, `LegoAC1` *has a built-in commit-and-prove capability*, which means its proofs are done with respect to matrices that committed in a Pedersen

n	LegoGro16			CPGro16		
	\mathcal{KG} (ms)	\mathcal{P} (ms)	crs (KB)	\mathcal{KG} (s)	\mathcal{P} (s)	crs (MB)
8	3.677	3.044	0.51	3.928	1.185	3.653
16	5.949	4.202	1.02	7.307	2.252	7.305
32	10.90	5.201	2.04	13.78	4.461	14.61
64	19.37	8.979	4.08	26.04	8.685	29.22
128	32.49	15.58	8.16	50.69	16.50	58.44
256	57.76	19.50	16.32	102.8	33.02	116.9
512	117.8	30.84	32.64	292.0	65.42	233.7
1024	241.2	55.35	65.28	876.3	133.3	467.5
2048	466.6	84.09	130.6	1011	428.7	935.0

$ \pi $ (B)	191.13	127.38
\mathcal{V} (ms)	4.129	3.4

Table 3: Performance comparison of **LegoGro16** and **CPGro16**. Numbers for the two schemes are in different units. Those for **CPGro16** are three orders of magnitude larger.

	Time		Space	
	\mathcal{P}	\mathcal{V}	crs	π
CP _{mmp}	$O(n^2)$	$O(n^2)$	$O(n^2)$	$O(\log n)$
Groth16	$O(n^3 \log n)$	$O(n^2)$	$O(n^3)$	$O(1)$

n	CP _{mmp}			Groth16	
	\mathcal{P} (ms)	\mathcal{V} (ms)	$ \pi $ KB	\mathcal{P} (s)	\mathcal{V} (ms)
16	35.36	22.74	21	0.181	4.312
32	46.26	23.19	24	1.379	6.060
64	55.78	24.00	28	11.61	12.60
128	83.78	28.03	32	109.3	50.99
256	149.7	40.01	36		

Table 4: Comparing **CP_{mmp}** and **Groth16** for $n \times n$ matrices

n	LegoAC1			Groth16		
	\mathcal{KG} (s)	\mathcal{P} (s)	\mathcal{V} (ms)	\mathcal{KG} (s)	\mathcal{P} (s)	\mathcal{V} (ms)
16	1.105	0.278	3.097	0.210	0.150	1.662
32	7.569	1.680	4.697	1.227	0.957	3.696
64	52.86	11.90	10.73	8.848	7.177	9.686
128	419.8	89.70	35.71	69.21	58.60	34.83

$ \pi $ (B)	350.25	127.38
-------------	--------	--------

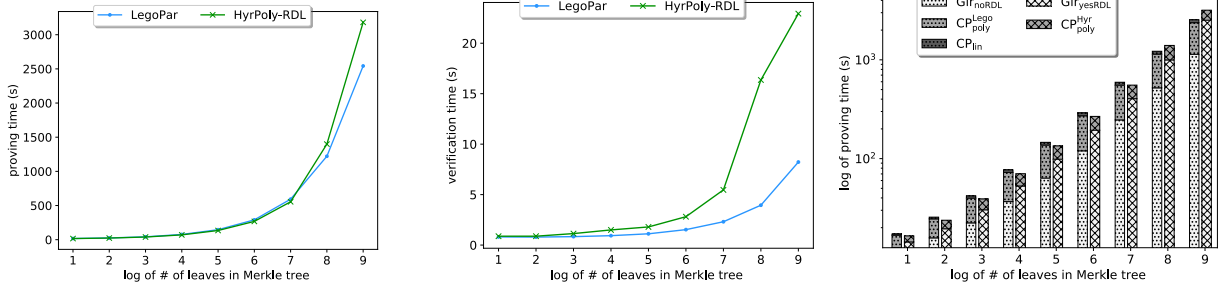
Table 5: Performance of **LegoAC1** comparing to **Groth16**

commitment (in a canonical vectorized form). This property is not present in **Groth16**, and can be useful in several applications.

For example, in the matrix factoring case, **LegoAC1** works with commitments to the three matrices that could be reused. This is a powerful feature as we could prove a statement like “ $B = A^{2^k}$ for a committed matrix A ” by doing k proofs, one for each squaring step (i.e., to show that $B_i = B_{i-1}^2$); this can be done by reusing the same CRS for one matrix factoring relation. In contrast, proving $B = A^{2^k}$ directly with **Groth16** would require a very large CRS and a memory intensive prover that would not scale for large k and n .

We give the experimental results that compare our **LegoAC1** commit-and-prove zkSNARK against the **Groth16** scheme, in the SHA256 and matrix factoring applications explained above. For SHA256, **Groth16** needs 1.9s for key generation of a CRS of 5.1MB, 0.7s for proving and 0.9ms for verification; **LegoAC1** needs 7.9s for key generation of a CRS of 6.2MB, 0.9s for proving and 1.8ms for verification. For matrix factoring, we used $n \times n$ matrices of 32-bit integers with $n \in \{16, 32, 64, 128\}$. Detailed timings are in Table 5.

Finally, we remark that **LegoAC1** is commit-and-prove, which means its proofs are done with respect to matrices that committed in a Pedersen commitment (in a canonical vectorized form).



(a) Proving time comparison for LegoPar and HyrPoly-RDL. (b) Verification time comparison for LegoPar and HyrPoly-RDL. (c) \mathcal{P} time (component-wise) for LegoPar (left) and HyrPoly-RDL (right).

Figure 9: Performance comparison of systems for parallel relations. Lower on the y axis is better (in (c), axis y is log-scale). We remind the reader that $\text{LegoPar} = \text{HyrPoly-Par} + \text{CP}_{\text{lin}} = (\text{Gir}_{\text{noRDL}} + \text{CP}_{\text{poly}}^{\text{Lego}}) + \text{CP}_{\text{lin}}$ and $\text{HyrPoly-RDL} = \text{Gir}_{\text{yesRDL}} + \text{CP}_{\text{poly}}^{\text{Hyr}}$.

7.4 Parallel Checks on Joint Inputs

We compare performances of our **LegoPar** system with a baseline system, i.e. **HyrPoly-RDL** (see Appendix F.5). Our choice of an Hyrax-based system for comparison is due to its optimization for parallel computations, and thus enhancing the latter implies refinements in all computations where Hyrax applies today.

Recall that **LegoPar** consists of our $\text{CP}_{\text{lin}}^{\text{Ped}}$ and **HyrPoly-Par**. To evaluate **HyrPoly-Par** and **HyrPoly-RDL** we executed separately the part concerning **PolyCom** and CP_{poly} , and the one that includes the **ZKGir⁺⁺** core. To benchmark the latter, we used the original Python code (appropriately modified for the commitment part) from the Hyrax project [gitax] (run using the JIT-compiling interpreter **PyPy** [pypy]).²³

We benchmarked **LegoPar** and **HyrPoly-RDL** on a highly parallel computation, that is proving knowledge of an assignment to the leaves of a Merkle tree [Mer88] (cf. Section ?? to see how it can be expressed using R^{parjnt}). We used SHA256 for the hash and a varying number of leaves (from 2 to 2^9). For this computation we generated two circuits using the Hyrax tool: one fully parallel to be fed to **HyrPoly-Par** and one with the RDL for **HyrPoly-RDL**. Recall that in **LegoPar** the RDL is checked using $\text{CP}_{\text{lin}}^{\text{Ped}}$. We finally note that the two largest inputs in our evaluation required extending the available RAM from 30 to 75GB for both schemes.

Results. Figure 9 compares the costs (proving and verification time) in the two schemes for repeated computation. Overall **LegoPar** is faster than **HyrPoly-RDL**, both in proving and verification time. On our largest input, proving in **LegoPar** is $1.25\times$ faster; verifying is more than $2.5\times$ faster. Verification is expected to become faster due to the asymptotic difference in the verification time.

- *Proving time:* On larger inputs **LegoPar** has a faster (up to $1.25\times$) proving time (Figs. 9a). In both schemes most of the computation is due to **ZKGir⁺⁺**: approximately 50% for **LegoPar** and 75% for **HyrPoly-RDL**. The higher time of **ZKGir⁺⁺** in **HyrPoly-RDL** is explained by the additional round for the RDL. On the other hand, **LegoPar** spends twice as much time for the proving step of CP_{poly} . This is because it evaluates a polynomial with twice as many terms, in turn requiring roughly twice the number of exponentiations. (This is due to the RDL output u_2 , on which

²³ Full integration of this component into our library is future work.

LegoPar operates, being twice as long as the RDL input u_1 (also the “bottom-layer” input), on which CP_{poly} runs in HyrPoly-RDL).

- *Verification time*: On larger inputs LegoPar has a shorter (up to $2.5\times$) verification time (Fig. 9b). This speedup is due to increase with larger inputs, as the verifier in HyrPoly-RDL has to perform an additional verification step for the RDL in ZKGir⁺⁺ (requiring a number of field operations roughly linear in the width of the circuit). On the other hand LegoPar performs the same step through a constant number of pairings (two) in CP_{veq} . In both schemes ZKGir⁺⁺ dominates running time (more than 99.5%)²⁴.

Discussion. Partly, the different performances we observed are due to specific features of the circuit chosen for benchmarks (we chose Merkle tree verification, due to its relevance in practice). In a circuit for parallel computation, at least two features, both related to the RDL, can have impact: (i) how “large” the output u_2 of the RDL is w.r.t. its input u_1 ; (ii) how “complex” the RDL is. A higher ratio $|u_2|/|u_1|$ will determine the difference in running time for the $\text{CP}_{\text{poly.Prove}}$ component. In our circuit of choice the ratio was 2.

8 Conclusions

We have described LegoSNARK, a framework for commit-and-prove zkSNARKs that comprises definitions, a general composition result, and a “lifting” construction. The LegoSNARK tools are useful as they enable designing zkSNARKs in a modular way (due to the framework of definitions and the composition theorem) and they allow to efficiently add commit-and-prove capabilities to a variety of existing schemes thus made interoperable. Furthermore we have proposed efficient proof gadgets for specialized relations and shown how to combine them into succinct proof systems for more complex relations. We have described instantiations of these new proof systems and evaluated them against prior work. The results show they have competitive performances. Specifically they show slightly worse (but still acceptable) performances in some applications (general arithmetic circuits) and significant improvements in others (commit-ahead-of-time systems, parallel computations).

A limitation of our current instantiations is their reliance on pairing-based systems with a trusted setup. Interestingly in some cases this is only needed to generate the commitment key of PolyCom. We believe this is doable by a large-scale MPC ceremony similar to the powers-of-tau round 1 of [BGM17] since the CRS includes only monomials in the exponent. It is future work to explore this direction. Nonetheless we note that this limitation is not inherent. The basic results of the framework (i.e., Section 3) are general enough to be instantiated in the future with schemes without trust assumptions. Finally, another future work direction is investigating new and more efficient proof gadgets CP-SNARKs for specialized relations and test them in specific applications.

Acknowledgements

Research leading to these results has been supported by the Spanish Government under projects Datamantium (ref. RTC-2016-4930-7), SCUM (ref. RTI2018-102043-B-I00), and ERC2018-092822, by the Madrid Regional Government under project BLOQUES (ref. S2018/TCS-4339) and by Protocol Labs. The project that gave rise to these results received the support of a fellowship from “la Caixa” Foundation (ID 100010434). The fellowship code is LCF/BQ/ES18/11670018.

²⁴ This is why we do not show a detailed bar plot for each component as for proving time.

References

- AGM18. Shashank Agrawal, Chaya Ganesh, and Payman Mohassel. Non-interactive zero-knowledge proofs for composite statements. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 643–673. Springer, Heidelberg, August 2018.
- AHIV17. Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 2087–2104. ACM Press, October / November 2017.
- AJ18. Kurt M. Alonso and Jordi Herrera Joancomartí. Monero - privacy in the blockchain. Cryptology ePrint Archive, Report 2018/535, 2018. <https://eprint.iacr.org/2018/535>.
- BBB⁺17. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Efficient range proofs for confidential transactions. Technical report, Cryptology ePrint Archive, Report 2017/1066, 2017. <https://eprint.iacr.org/2017/1066>, 2017.
- BBFR15. Michael Backes, Manuel Barbosa, Dario Fiore, and Raphael M. Reischuk. ADSNARK: Nearly practical and privacy-preserving proofs on authenticated data. In *2015 IEEE Symposium on Security and Privacy*, pages 271–286. IEEE Computer Society Press, May 2015.
- BCC⁺16. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 327–357. Springer, Heidelberg, May 2016.
- BCC⁺17. Nir Bitansky, Ran Canetti, Alessandro Chiesa, Shafi Goldwasser, Huijia Lin, Aviad Rubinfeld, and Eran Tromer. The hunting of the SNARK. *Journal of Cryptology*, 30(4):989–1066, October 2017.
- BCCT12. Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In Shafi Goldwasser, editor, *ITCS 2012*, pages 326–349. ACM, January 2012.
- BCG⁺13. Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108. Springer, Heidelberg, August 2013.
- BCG⁺14. Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014.
- BCG⁺17. Jonathan Bootle, Andrea Cerulli, Essam Ghadafi, Jens Groth, Mohammad Hajiabadi, and Sune K. Jakobsen. Linear-time zero-knowledge proofs for arithmetic circuit satisfiability. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part III*, volume 10626 of *LNCS*, pages 336–365. Springer, Heidelberg, December 2017.
- BCI⁺13. Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. Succinct non-interactive arguments via linear interactive proofs. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 315–333. Springer, Heidelberg, March 2013.
- BCPR14. Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In David B. Shmoys, editor, *46th ACM STOC*, pages 505–514. ACM Press, May / June 2014.
- BCS16. Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 31–60. Springer, Heidelberg, October / November 2016.
- BCTV14. Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 781–796. USENIX Association, August 2014.
- BFR⁺13. Benjamin Braun, Ariel J. Feldman, Zuoqiang Ren, Srinath Setty, Andrew J. Blumberg, and Michael Walfish. Verifying computations with state. In *Proc. of the ACM SOSP*, 2013.
- BGM17. Sean Bowe, Ariel Gabizon, and Ian Miers. Scalable multi-party computation for zk-snark parameters in the random beacon model. Cryptology ePrint Archive, Report 2017/1050, 2017. <https://eprint.iacr.org/2017/1050>.
- BP15. Elette Boyle and Rafael Pass. Limits of extractability assumptions with distributional auxiliary input. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 236–261. Springer, Heidelberg, November / December 2015.

- BSBHR18. Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. Cryptology ePrint Archive, Report 2018/046, 2018. <https://eprint.iacr.org/2018/046>.
- bulk1. <https://github.com/apoelstra/secp256k1-mw/tree/bulletproofs>, libsecp256k1.
- CDG⁺17. Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Reicherberger, Daniel Slamanig, and Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1825–1842. ACM Press, October / November 2017.
- CFH⁺15. Craig Costello, Cédric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. Geppetto: Versatile verifiable computation. In *2015 IEEE Symposium on Security and Privacy*, pages 253–270. IEEE Computer Society Press, May 2015.
- CGM16. Melissa Chase, Chaya Ganesh, and Payman Mohassel. Efficient zero-knowledge proof of algebraic and non-algebraic statements with applications to privacy preserving credentials. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 499–530. Springer, Heidelberg, August 2016.
- CLOS02. Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *34th ACM STOC*, pages 494–503. ACM Press, May 2002.
- CMT12. Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Practical verified computation with streaming interactive proofs. In Shafi Goldwasser, editor, *ITCS 2012*, pages 90–112. ACM, January 2012.
- EHK⁺13. Alex Escala, Gottfried Herold, Eike Kiltz, Carla Ràfols, and Jorge Villar. An algebraic framework for Diffie-Hellman assumptions. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147. Springer, Heidelberg, August 2013.
- FFG⁺16. Dario Fiore, Cédric Fournet, Esha Ghosh, Markulf Kohlweiss, Olga Ohrimenko, and Bryan Parno. Hash first, argue later: Adaptive verifiable computations on outsourced data. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1304–1316. ACM Press, October 2016.
- FKL18. Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Heidelberg, August 2018.
- FLSZ17. Prastudy Fauzi, Helger Lipmaa, Janno Siim, and Michal Zajac. An efficient pairing-based shuffle argument. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part II*, volume 10625 of *LNCS*, pages 97–127. Springer, Heidelberg, December 2017.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO '86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.
- GGPR13. Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.
- gitax. <https://github.com/hyraxZK>, Hyrax.
- GKM⁺18. Jens Groth, Markulf Kohlweiss, Mary Maller, Sarah Meiklejohn, and Ian Miers. Updatable and universal common reference strings with applications to zk-SNARKs. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 698–728. Springer, Heidelberg, August 2018.
- GKR08. Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 113–122. ACM Press, May 2008.
- GKS17. Daniel Günther, Ágnes Kiss, and Thomas Schneider. More efficient universal circuit constructions. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part II*, volume 10625 of *LNCS*, pages 443–470. Springer, Heidelberg, December 2017.
- GMO16. Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. ZKBoo: Faster zero-knowledge for Boolean circuits. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016*, pages 1069–1083. USENIX Association, August 2016.
- GMR89. Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- GMW87. Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987.

- Gro09. Jens Groth. Linear algebra with sub-linear zero-knowledge arguments. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 192–208. Springer, Heidelberg, August 2009.
- Gro10. Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 321–340. Springer, Heidelberg, December 2010.
- Gro16. Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.
- GW11. Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108. ACM Press, June 2011.
- IKO07. Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. Efficient arguments without short pcps. In *Proceedings of the Twenty-Second Annual IEEE Conference on Computational Complexity, CCC '07*, pages 278–291, Washington, DC, USA, 2007. IEEE Computer Society.
- IKOS07. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 21–30. ACM Press, June 2007.
- Kil89. J. Kilian. Uses of randomness in algorithms and protocols. PhD Thesis. Massachusetts Institute of Technology, 1989.
- Kil92. Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *24th ACM STOC*, pages 723–732. ACM Press, May 1992.
- KPP⁺14. Ahmed E. Kosba, Dimitrios Papadopoulos, Charalampos Papamanthou, Mahmoud F. Sayed, Elaine Shi, and Nikos Triandopoulos. TRUESET: Faster verifiable set computations. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 765–780. USENIX Association, August 2014.
- KW15. Eike Kiltz and Hoeteck Wee. Quasi-adaptive NIZK for linear subspaces revisited. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 101–128. Springer, Heidelberg, April 2015.
- LFKN92. Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, October 1992.
- librk. <https://github.com/scipr-lab/libsnark>, libsnark.
- Lip12. Helger Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 169–189. Springer, Heidelberg, March 2012.
- Lip16. Helger Lipmaa. Prover-efficient commit-and-prove zero-knowledge SNARKs. In David Pointcheval, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 16*, volume 9646 of *LNCS*, pages 185–206. Springer, Heidelberg, April 2016.
- Mer88. Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, pages 369–378. Springer, Heidelberg, August 1988.
- Mic94. Silvio Micali. CS proofs (extended abstracts). In *35th FOCS*, pages 436–453. IEEE Computer Society Press, November 1994.
- Mic00. Silvio Micali. Computationally sound proofs. *SIAM Journal on Computing*, 30(4):1253–1298, 2000.
- NY90. Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd ACM STOC*, pages 427–437. ACM Press, May 1990.
- Par15. Bryan Parno. A note on the unsoundness of vntinyRAM’s SNARK. Cryptology ePrint Archive, Report 2015/437, 2015. <http://eprint.iacr.org/2015/437>.
- Ped92. Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 129–140. Springer, Heidelberg, August 1992.
- PHGR13. Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013.
- pypPy. <https://pypy.org>, PyPy.
- Rot09. Guy Rothblum. Delegating computation reliably: paradigms and constructions, 2009. PhD thesis.
- RRR16. Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In Daniel Wichs and Yishay Mansour, editors, *48th ACM STOC*, pages 49–62. ACM Press, June 2016.
- Sch91. Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, January 1991.

- Tha13. Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 71–89. Springer, Heidelberg, August 2013.
- Val76. Leslie G. Valiant. Universal circuits (preliminary report). In *STOC*, pages 196–203. ACM, 1976.
- Vee17. Meelof Veeningen. Pinocchio-based adaptive zk-SNARKs and secure/correct adaptive function evaluation. In Marc Joye and Abderrahmane Nitaj, editors, *AFRICACRYPT 17*, volume 10239 of *LNCS*, pages 21–39. Springer, Heidelberg, May 2017.
- WJB⁺17. Riad S. Wahby, Ye Ji, Andrew J. Blumberg, abhi shelat, Justin Thaler, Michael Walfish, and Thomas Wies. Full accounting for verifiable outsourcing. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 2071–2086. ACM Press, October / November 2017.
- WSR⁺15. Riad S. Wahby, Srinath T. V. Setty, Zuocheng Ren, Andrew J. Blumberg, and Michael Walfish. Efficient RAM and control flow in verifiable outsourced computation. In *NDSS 2015*. The Internet Society, February 2015.
- WTas⁺17. Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zksnarks without trusted setup. Cryptology ePrint Archive, Report 2017/1132, 2017. <https://eprint.iacr.org/2017/1132>.
- WTs⁺18. Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zk-SNARKs without trusted setup. In *2018 IEEE Symposium on Security and Privacy*, pages 926–943. IEEE Computer Society Press, May 2018.
- ZGK⁺17a. Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. vSQL: Verifying arbitrary SQL queries over dynamic outsourced databases. In *2017 IEEE Symposium on Security and Privacy*, pages 863–880. IEEE Computer Society Press, May 2017.
- ZGK⁺17b. Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. A zero-knowledge version of vsql. Cryptology ePrint Archive, Report 2017/1146, 2017. <https://eprint.iacr.org/2017/1146>.

A Security proof of CP-SNARK composition

In this section we provide a proof of Theorem 3.1. We first define relation generators and auxiliary input generators for this construction.

$$\begin{array}{l}
\text{Aux}^{\mathcal{RG}}(1^\lambda) : \\
(R_0, \text{aux}_R^{(0)}) \leftarrow \mathcal{RG}_0(1^\lambda) \\
(R_1, \text{aux}_R^{(1)}) \leftarrow \mathcal{RG}_1(1^\lambda) \\
\mathbf{return} (R_b, \text{aux}_R^{(b)})_{b \in \{0,1\}}
\end{array}
\quad
\begin{array}{l}
\overline{\mathcal{RG}}_b(1^\lambda) : \\
(R_b, \text{aux}_R^{(b)})_{b \in \{0,1\}} \leftarrow \text{Aux}^{\mathcal{RG}}(1^\lambda) \\
\overline{\text{aux}}_R^{(b)} := (R_{1-b}, (\text{aux}_R^{(b)})_{b \in \{0,1\}}) \\
\mathbf{return} (R_b, \overline{\text{aux}}_R^{(b)})
\end{array}
\quad
\begin{array}{l}
\mathcal{Z}^*((\text{ck}, R_{R_0, R_1}^\wedge), (\text{ek}^*, \text{vk}^*), (\text{aux}_R, \text{aux}'_R)) : \\
(\text{aux}_Z^{(b)})_{b \in \{0,1\}} \leftarrow \text{Aux}^{\mathcal{Z}}(\text{ck}, (\text{crs}_b, R_b, \text{aux}_R^{(b)})_{b \in \{0,1\}}) \\
\mathbf{return} (\text{aux}_Z^{(b)})_{b \in \{0,1\}}
\end{array}$$

$$\begin{array}{l}
\text{Aux}^{\mathcal{Z}}(\text{ck}, (\text{crs}_b, R_b, \text{aux}_R^{(b)})_{b \in \{0,1\}}) : \\
\text{aux}_Z^{(0)} \leftarrow \mathcal{Z}_0(\text{ck}, R_0, \text{crs}_0, \text{aux}_R^{(0)}) \\
\text{aux}_Z^{(1)} \leftarrow \mathcal{Z}_1(\text{ck}, R_1, \text{crs}_1, \text{aux}_R^{(1)}) \\
\mathbf{return} (\text{aux}_Z^{(b)})_{b \in \{0,1\}}
\end{array}
\quad
\begin{array}{l}
\mathcal{RG}^*(1^\lambda) : \\
(R_b, \text{aux}_R^{(b)})_{b \in \{0,1\}} \leftarrow \text{Aux}^{\mathcal{RG}}(1^\lambda) \\
\mathbf{return} (R_{R_0, R_1}^\wedge, (\text{aux}_R^{(b)})_{b \in \{0,1\}})
\end{array}
\quad
\begin{array}{l}
\overline{\mathcal{Z}}_b(\text{ck}, R_b, \text{crs}_b, \overline{\text{aux}}_R^{(b)}) : \\
\text{Parse } \overline{\text{aux}}_R \text{ as } (R_{1-b}, (\text{aux}_R^{(b)})_{b \in \{0,1\}}) \\
\text{crs}_{1-b} \leftarrow \text{CP}_{1-b}.\text{KeyGen}(\text{ck}, R_{1-b}) \\
\{\text{aux}_Z^{(b)} \leftarrow \text{Aux}^{\mathcal{Z}}(\text{ck}, (\text{crs}_b, R_b, \text{aux}_R^{(b)})_{b \in \{0,1\}})\} \\
\mathbf{return} \overline{\text{aux}}_Z^{(b)} := (\text{crs}_{1-b}, (\text{aux}_Z^{(b)})_{b \in \{0,1\}})
\end{array}$$

Figure 10: Relation and Auxiliary Input Generators for AND Composition Construction

A.1 Proof of Knowledge Soundness

We state the following lemma.

Lemma A.1. *If Com is computationally binding, and if CP_b is $\text{KSND}(\overline{\mathcal{RG}}_b, \overline{\mathcal{Z}}_b)$ (where $\overline{\mathcal{RG}}_b, \overline{\mathcal{Z}}_b$ are defined in terms of $\mathcal{RG}_b, \mathcal{Z}_b$ in Figure 10) for $b \in \{0, 1\}$, then the scheme CP^\wedge in Figure 1 is $\text{KSND}(\mathcal{RG}^*, \mathcal{Z}^*)$ where $\mathcal{RG}^*, \mathcal{Z}^*$ are as defined in Figure 10.*

Proof Let \mathcal{A}^* be an adversary against the soundness of CP^\wedge with respect to \mathcal{RG}^* and \mathcal{Z}^* . Now for $b \in \{0, 1\}$ consider adversary \mathcal{A}_b (defined in Figure 11) against CP_b with respect to $\overline{\mathcal{RG}}_b$ and $\overline{\mathcal{Z}}_b$. By the fact that CP_b is $\text{KSND}(\overline{\mathcal{RG}}_b, \overline{\mathcal{Z}}_b)$ there exists an extractor \mathcal{E}_b such that $\Pr[\text{Game}_{\overline{\mathcal{RG}}_b, \overline{\mathcal{Z}}_b, \mathcal{A}_b, \mathcal{E}_b}^{\text{KSND}} = 1]$ is negligible.

We define an extractor \mathcal{E}^* for CP^\wedge in Figure 11, and we claim is such that $\Pr[\text{Game}_{\mathcal{RG}^*, \mathcal{Z}^*, \mathcal{A}^*, \mathcal{E}^*}^{\text{KSND}} = 1]$. First observe that with overwhelming probability the values u_2 and u'_2 in \mathcal{E}^* are equal, conditioned to the openings being all correct for their respective commitments (i.e., conditioned to VerCommit returning 1 on each of them). In fact, if it were otherwise, we could then break the binding of Com (as done in the proof of Theorem B.1).

We now define the following notations:

$$\begin{aligned} \{\text{GdCom}(c_b, u_b, o_b) := \text{Com.VerCommit}(\text{ck}, c_b, u_b, o_b) = 1\}_{b \in \{0, 1\}} \\ \text{GdCom}(c_2, u_2, o_2) := \text{Com.VerCommit}(\text{ck}, c_2, u_2, o_2) = 1 \\ \text{GdCom}(c_2, u'_2, o'_2) := \text{Com.VerCommit}(\text{ck}, c_2, u'_2, o'_2) = 1 \end{aligned}$$

For $b \in \{0, 1\}$, by the soundness properties of CP_b and the definition of $\mathcal{E}_b, \mathcal{E}^*$ we have that p_b , as defined below, is negligible.

$$p_b := \Pr[b_{\text{ok}}^{(b)} \wedge (\neg \text{GdCom}(c_b, u_b, o_b) \vee \neg \text{GdCom}(c_2, u_2, o_2) \vee R_b(x_b, u_b, u_2, \omega_b) = 0)]$$

where all the symbols above are as defined in the construction of \mathcal{E}^* . Now we can observe that

$$\begin{aligned} \Pr[\text{Game}_{\mathcal{RG}^*, \mathcal{Z}^*, \mathcal{A}^*, \mathcal{E}^*}^{\text{KSND}} = 1] &= \dots \\ &= \Pr[b_{\text{ok}}^{(0)} \wedge b_{\text{ok}}^{(1)} \wedge (\neg \text{GdCom}(c_0, u_0, o_0) \vee \neg \text{GdCom}(c_1, u_1, o_1) \vee \neg \text{GdCom}(c_2, u_2, o_2) \\ &\quad \vee R_0(x_0, u_0, u_2; \omega_0) = 0 \vee R_1(x_1, u_1, u_2; \omega_1) = 0)] \\ &\leq \Pr[b_{\text{ok}}^{(0)} \wedge (\neg \text{GdCom}(c_0, u_0, o_0) \vee \neg \text{GdCom}(c_2, u_2, o_2) \vee R_0(u_0, u_2, \omega_0) = 0)] + \\ &\quad \Pr[b_{\text{ok}}^{(1)} \wedge (\neg \text{GdCom}(c_1, u_1, o_1) \vee \neg \text{GdCom}(c_2, u'_2, o'_2) \vee R_1(u_1, u'_2, \omega_1) = 0)] + \text{negl}(\lambda) \leq p_0 + p_1 + \text{negl}(\lambda) \leq \text{negl}(\lambda) \end{aligned}$$

where in the last two inequalities we used our earlier observations on the openings of u_2 and u'_2 and p_0 and p_1 being negligible respectively. \square

$\mathcal{A}_b(\text{ck}, (\text{crs}_b, R_b), \overline{\text{aux}}_R^{(b)}, \overline{\text{aux}}_Z^{(b)}) :$	$\mathcal{E}^*(\text{ck}, ((\text{crs}_b)_{b \in \{0, 1\}}, R_{R_0, R_1}^\wedge), \text{aux}_R^{(b)}, \text{aux}_Z^{(b)}) :$
Parse $\overline{\text{aux}}_R^{(b)}$ as $(R_{1-b}, (\text{aux}_R^{(b)})_{b \in \{0, 1\}})$	$\overline{\text{aux}}_R^{(b)} := (R_{1-b}, (\text{aux}_R^{(b)})_{b \in \{0, 1\}})$ for $b \in \{0, 1\}$
Parse $\overline{\text{aux}}_Z^{(b)}$ as $(\text{crs}_{1-b}, (\text{aux}_Z^{(b)})_{b \in \{0, 1\}})$	$\overline{\text{aux}}_Z^{(b)} := (\text{crs}_{1-b}, (\text{aux}_Z^{(b)})_{b \in \{0, 1\}})$ for $b \in \{0, 1\}$
$(x_0, x_1, (c_j)_{j \in [3]}, \pi^* := (\pi_b)_{b \in \{0, 1\}})$	$((x_0, u_0, u_2), (o_0, o_2), \omega_0) \leftarrow \mathcal{E}_0(\text{ck}, (\text{crs}_0, R_0), \overline{\text{aux}}_R^{(0)}, \overline{\text{aux}}_Z^{(0)})$
$\leftarrow \mathcal{A}^*(\text{ck}, (\text{crs}_0, \text{crs}_1, R_{R_0, R_1}^\wedge),$	$((x_1, u_1, u'_2), (o_1, o'_2), \omega_1) \leftarrow \mathcal{E}_1(\text{ck}, (\text{crs}_1, R_1), \overline{\text{aux}}_R^{(1)}, \overline{\text{aux}}_Z^{(1)})$
$(\text{aux}_R^{(b)})_{b \in \{0, 1\}}, (\text{aux}_Z^{(b)})_{b \in \{0, 1\}})$	return $((x_b)_{b \in \{0, 1\}}, (u_j)_{j \in [3]}, (o_j)_{j \in [3]}, (\omega_b)_{b \in \{0, 1\}})$
return (x_b, c_b, c_2, π_b)	

Figure 11: Adversary and Extractor for Proof of Lemma A.1

A.2 Proof of Zero-Knowledge

We state the following lemma.

Lemma A.2. *If CP_b is zero-knowledge for Com and $\overline{\mathcal{RG}}_b$ for $b \in \{0, 1\}$, then the scheme CP^\wedge in Figure ?? is a zero-knowledge CP-SNARK for Com and \mathcal{RG}^* (where relation generators are defined in Figure 10).*

Proof We construct the following two simulators for \mathcal{RG}^* from simulators for CP_0, CP_1 . Then ZK follows through a standard hybrid argument.

$$\frac{\mathcal{S}_{\text{kg}}^*(\text{ck}, R_{R_0, R_1}^\wedge)}{\text{for } b \in \{0, 1\} : (\text{crs}_b, \text{td}_k^{(b)}) \leftarrow \mathcal{S}_{\text{kg}}^{(b)}(\text{ck}, R_b) \\ \text{return } (\text{crs}^* := (\text{crs}_b)_{b \in \{0, 1\}}, \text{td}_k^* := (\text{td}_k^{(b)})_{b \in \{0, 1\}})}{\quad} \frac{\mathcal{S}_{\text{prv}}^*((\text{crs}_b)_{b \in \{0, 1\}}, (\text{td}_k^{(b)})_{b \in \{0, 1\}}, (x_b)_{b \in \{0, 1\}}, (c_j)_{j \in [3]})}{\text{for } b \in \{0, 1\} : \pi_b \leftarrow \mathcal{S}_{\text{prv}}^{(b)}(\text{crs}_b, \text{td}_k^{(b)}, x_b, (c_b, c_2)) \\ \text{return } (\pi_b)_{b \in \{0, 1\}}}$$

□

B Proofs for the General Compiler

Theorem B.1. *Let $\text{CP}.\mathcal{RG}$ be a relation generator such that $\text{CP}.\mathcal{RG}_\lambda \subseteq \mathcal{R}_\lambda$, and let $\text{CP}.\mathcal{Z}$ be an auxiliary input distribution. Then the scheme CP in Table 2 is $\text{KSND}(\text{CP}.\mathcal{RG}, \text{CP}.\mathcal{Z})$ and composable zero-knowledge for $\text{CP}.\mathcal{RG}$ whenever: (i) ccII is $\text{ccKSND}(\text{ccII}.\mathcal{RG}, \text{ccII}.\mathcal{Z})$ and composable zero-knowledge for $\text{ccII}.\mathcal{RG}$, (ii) CP_{link} is $\text{KSND}(\text{CP}_{\text{link}}.\mathcal{RG}, \text{CP}_{\text{link}}.\mathcal{Z})$ and composable zero-knowledge for $\text{CP}_{\text{link}}.\mathcal{RG}$, where the relation generators and auxiliary input distributions $\text{ccII}.\mathcal{RG}$, $\text{ccII}.\mathcal{Z}$, $\text{CP}_{\text{link}}.\mathcal{RG}$, $\text{CP}_{\text{link}}.\mathcal{Z}$ are the ones in Figure 12. This result also holds when ccII is a cc-SNARK with weak binding (Definition 3.3).*

$$\begin{array}{l} \text{CP}_{\text{link}}.\mathcal{RG}(1^\lambda) : \\ \hline (R, \text{aux}_R) \leftarrow \text{CP}.\mathcal{RG}(1^\lambda) \\ (\text{ck}', \text{ek}', \text{vk}') \leftarrow \text{ccII}.\text{KeyGen}(R) \\ R^{\text{link}} := (\text{ck}', \mathcal{D}_x^{\text{link}}, \mathcal{D}_u^{\text{link}}, \mathcal{D}_\omega^{\text{link}}) \\ \text{aux}_R^{\text{link}} := (\text{ek}', \text{vk}', R, \text{aux}_R) \\ \text{return } (R^{\text{link}}, \text{aux}_R^{\text{link}}) \end{array} \quad \begin{array}{l} \text{CP}_{\text{link}}.\mathcal{Z}((\text{ck}, R^{\text{link}}), \text{aux}_R^{\text{link}}, \text{crs}^{\text{link}}) : \\ \hline \text{Parse } \text{aux}_R^{\text{link}} \text{ as } (\text{ek}', \text{vk}', R, \text{aux}_R) ; \text{ Parse } \text{crs}^{\text{link}} \text{ as } (\text{ek}^{\text{link}}, \text{vk}^{\text{link}}) \\ \text{Get } \text{ck}' \text{ from } R^{\text{link}} ; \text{ek} := (\text{ck}', \text{ek}', \text{ek}^{\text{link}}) ; \text{vk} := (\text{vk}', \text{vk}^{\text{link}}) \\ \text{return } \text{aux}_Z^{\text{link}} \leftarrow \text{CP}.\mathcal{Z}((\text{ck}, R), \text{aux}_R, (\text{ek}, \text{vk})) \end{array}$$

$$\begin{array}{l} \text{ccII}.\mathcal{RG}(1^\lambda) : \\ \hline (R, \text{aux}_R) \leftarrow \text{CP}.\mathcal{RG}(1^\lambda) \\ \text{ck} \leftarrow \text{CP}.\text{Setup}(1^\lambda) \\ \text{return } (R, \text{aux}_R' := (\text{ck}, \text{aux}_R)) \end{array} \quad \begin{array}{l} \text{ccII}.\mathcal{Z}(R, \text{aux}_R', \text{crs}') : \\ \hline \text{Parse } \text{crs}' \text{ as } (\text{ck}', \text{ek}', \text{vk}') \text{ and } \text{aux}_R' \text{ as } (\text{ck}, \text{aux}_R) ; R^{\text{link}} := (\text{ck}', \mathcal{D}_x^{\text{link}}, \mathcal{D}_u^{\text{link}}, \mathcal{D}_\omega^{\text{link}}) \\ (\text{ek}^{\text{link}}, \text{vk}^{\text{link}}) \leftarrow \text{CP}_{\text{link}}.\text{KeyGen}(\text{ck}, R^{\text{link}}) ; \text{ek} := (\text{ck}', \text{ek}', \text{ek}^{\text{link}}) ; \text{vk} := (\text{vk}', \text{vk}^{\text{link}}) \\ \text{aux}_Z \leftarrow \text{CP}.\mathcal{Z}((\text{ck}, R), \text{aux}_R, (\text{ek}, \text{vk})) \\ \text{return } \text{aux}_Z' := (\text{ek}^{\text{link}}, \text{vk}^{\text{link}}, \text{aux}_Z) \end{array}$$

Figure 12: Relation and Auxiliary Input Generators for Theorem B.1

B.1 Proof of Knowledge Soundness

Proof First, recall that proving the knowledge soundness of a CP-SNARK scheme CP for relation generator $\text{CP}.\mathcal{RG}$ means proving the knowledge soundness of CP as a SNARK for the corresponding

relation generator $\text{CP}.\mathcal{RG}_{\text{Com}}$ that, we recall, honestly generates the commitment key $\text{ck} \leftarrow \text{Setup}(1^\lambda)$ and generates (R, aux_R) using $\text{CP}.\mathcal{RG}$ and outputs $((\text{ck}, R), \text{aux}_R)$.

Our proof proceeds in the following steps.

First, assume there exists an adversary $\text{CP}.\mathcal{A}$ against scheme CP that runs in the experiment $\text{Game}_{\text{CP}.\mathcal{RG}_{\text{Com}}, \text{CP}.\mathcal{Z}}^{\text{KSND}}$ and outputs a tuple $(x, (c_j)_{j \in [\ell]}, \pi)$ such that $\text{CP}.\text{VerProof}(\text{vk}, x, (c_j)_{j \in [\ell]}, \pi) = 1$. Then, from such $\text{CP}.\mathcal{A}$ we can build:

1. an adversary $\text{ccII}.\mathcal{A}$ against ccII that runs in the experiment $\text{Game}_{\text{ccII}.\mathcal{RG}, \text{ccII}.\mathcal{Z}}^{\text{ccKSND}}$ (with the relation and auxiliary input generators $\text{ccII}.\mathcal{RG}, \text{ccII}.\mathcal{Z}$ defined in Fig. 12), and outputs (x, c', π') ;
2. an adversary $\text{CP}_{\text{link}}.\mathcal{A}$ against CP_{link} that runs in the experiment $\text{Game}_{\text{CP}_{\text{link}}.\mathcal{RG}_{\text{Com}}, \text{CP}_{\text{link}}.\mathcal{Z}}^{\text{KSND}}$ (with the relation and auxiliary input generators $\text{CP}_{\text{link}}.\mathcal{RG}_{\text{Com}}, \text{CP}_{\text{link}}.\mathcal{Z}$ defined in Fig. 12), and that outputs $(c', (c_j)_{j \in [\ell]}, \pi^{\text{link}})$;

The two adversaries $\text{ccII}.\mathcal{A}, \text{CP}_{\text{link}}.\mathcal{A}$ are defined below. By looking at the way their inputs are sampled in their respective games $\text{Game}_{\text{ccII}.\mathcal{RG}, \text{ccII}.\mathcal{Z}}^{\text{ccKSND}}$ and $\text{Game}_{\text{CP}_{\text{link}}.\mathcal{RG}_{\text{Com}}, \text{CP}_{\text{link}}.\mathcal{Z}}^{\text{KSND}}$, and how the relation and auxiliary input generators are defined, the input received by $\text{CP}.\mathcal{A}$ in both simulations (the one by $\text{ccII}.\mathcal{A}$ and the one by $\text{CP}_{\text{link}}.\mathcal{A}$) is distributed identically as the input $\text{CP}.\mathcal{A}$ would receive in $\text{Game}_{\text{CP}.\mathcal{RG}_{\text{Com}}, \text{CP}.\mathcal{Z}}^{\text{KSND}}$.

$\text{ccII}.\mathcal{A}(R, \text{crs}', \text{aux}'_R, \text{aux}'_Z)$:

Parse aux'_R as $(\text{ck}, \text{aux}_R)$; aux'_Z as $(\text{ek}^{\text{link}}, \text{vk}^{\text{link}}, \text{aux}_Z)$; crs' as $(\text{ck}', \text{ek}', \text{vk}')$
 $\text{ek} := (\text{ck}', \text{ek}', \text{ek}^{\text{link}})$; $\text{vk} := (\text{vk}', \text{vk}^{\text{link}})$; $(x, (c_j)_{j \in [\ell]}, \pi := (c', \pi^{\text{link}}, \pi')) \leftarrow \text{CP}.\mathcal{A}((\text{ck}, R), (\text{ek}, \text{vk}), \text{aux}_R, \text{aux}_Z)$
return (x, c', π')

$\text{CP}_{\text{link}}.\mathcal{A}((\text{ck}, R^{\text{link}}), \text{crs}^{\text{link}}, \text{aux}'_R, \text{aux}'_Z)$:

Parse aux'_R as $(\text{ek}', \text{vk}', R, \text{aux}_R)$; crs^{link} as $(\text{ek}^{\text{link}}, \text{vk}^{\text{link}})$; aux'_Z as aux_Z ; R^{link} as $(\text{ck}', \mathcal{D}_x^{\text{link}}, \mathcal{D}_u^{\text{link}}, \mathcal{D}_\omega^{\text{link}})$
 $\text{ek} := (\text{ck}', \text{ek}', \text{ek}^{\text{link}})$; $\text{vk} := (\text{vk}', \text{vk}^{\text{link}})$; $(x, (c_j)_{j \in [\ell]}, \pi := (c', \pi^{\text{link}}, \pi')) \leftarrow \text{CP}.\mathcal{A}((\text{ck}, R), (\text{ek}, \text{vk}), \text{aux}_R, \text{aux}_Z)$
return $(c', (c_j)_{j \in [\ell]}, \pi^{\text{link}})$

Second, observe that:

- If ccII is $\text{ccKSND}(\text{ccII}.\mathcal{RG}, \text{ccII}.\mathcal{Z})$ then for every $\text{ccII}.\mathcal{A}$ there exists an extractor $\text{ccII}.\mathcal{E}$ that returns $((u'_j)_{j \in [\ell]}, o', w')$ such that $\Pr[\text{Game}_{\text{ccII}.\mathcal{RG}, \text{ccII}.\mathcal{Z}, \text{ccII}.\mathcal{A}, \text{ccII}.\mathcal{E}}^{\text{ccKSND}} = 1]$ is negligible.
- If CP_{link} is $\text{KSND}(\text{CP}_{\text{link}}.\mathcal{RG}_{\text{Com}}, \text{CP}_{\text{link}}.\mathcal{Z})$ then for every $\text{CP}_{\text{link}}.\mathcal{A}$ there exists extractor $\text{CP}_{\text{link}}.\mathcal{E}$ that returns $((u_j^{\text{link}})_{j \in [\ell]}, (o_j^{\text{link}})_{j \in [\ell]}, \omega^{\text{link}})$ such that the following probability is negligible $\Pr[\text{Game}_{\text{CP}_{\text{link}}.\mathcal{RG}_{\text{Com}}, \text{CP}_{\text{link}}.\mathcal{Z}, \text{CP}_{\text{link}}.\mathcal{A}, \text{CP}_{\text{link}}.\mathcal{E}}^{\text{KSND}} = 1]$.

Hence, let $\text{ccII}.\mathcal{E}$ and $\text{CP}_{\text{link}}.\mathcal{E}$ be the extractors corresponding to our adversaries $\text{ccII}.\mathcal{A}$ and $\text{CP}_{\text{link}}.\mathcal{A}$ respectively. From the existence of the two extractors $\text{ccII}.\mathcal{E}$ and $\text{CP}_{\text{link}}.\mathcal{E}$ we construct extractor $\text{CP}.\mathcal{E}$ as below.

$\text{CP.}\mathcal{E}((\text{ck}, R), (\text{ek}, \text{vk}), \text{aux}_R, \text{aux}_Z) :$

Parse ek as $(\text{ck}', \text{ek}', \text{ek}'^{\text{link}})$; vk as $(\text{vk}', \text{vk}'^{\text{link}})$; $\text{crs}' := (\text{ck}', \text{ek}' \text{vk}')$
 $\text{aux}'_R := (\text{ck}, \text{aux}_R)$; $\text{aux}'_Z := (\text{ek}^{\text{link}}, \text{vk}^{\text{link}}, \text{aux}_Z)$; $((u'_j)_{j \in [\ell]}, o', \omega') \leftarrow \text{ccII.}\mathcal{E}(R, \text{crs}', \text{aux}'_R, \text{aux}'_Z)$
 $R^{\text{link}} := (\text{ck}', \mathcal{D}_x^{\text{link}}, \mathcal{D}_u^{\text{link}}, \mathcal{D}_\omega^{\text{link}})$; $\text{aux}_R^{\text{link}} := (\text{ek}', \text{vk}', R^*, \text{aux}_R)$; $\text{aux}_Z^{\text{link}} := \text{aux}_Z$
 $((u_j^{\text{link}})_{j \in [\ell]}, (o_j^{\text{link}})_{j \in [\ell]}, \omega^{\text{link}}) \leftarrow \text{CP}_{\text{link}}.\mathcal{E}((\text{ck}, R^{\text{link}}), \text{crs}^{\text{link}}, \text{aux}_R^{\text{link}}, \text{aux}_Z^{\text{link}})$
return $((u_j^{\text{link}})_{j \in [\ell]}, (o_j^{\text{link}})_{j \in [\ell]}, \omega')$

Combining the steps above, we have shown that for any CP adversary $\text{CP.}\mathcal{A}$ there exists a corresponding extractor $\text{CP.}\mathcal{E}$. We are left to prove that $\Pr[\text{Game}_{\text{CP.}\mathcal{R}\mathcal{G}_{\text{Com}}, \text{CP.}\mathcal{Z}, \text{CP.}\mathcal{A}, \text{CP.}\mathcal{E}}^{\text{KSND}} = 1] = \text{negl}$. Recall that the output of $\text{CP.}\mathcal{A}$ is of the form $(x, (c_j)_{j \in [\ell]}, \pi)$ with $\pi = (c', \pi^{\text{link}}, \pi')$, and for $\text{CP.}\mathcal{E}$ is of the form $((u_j^{\text{link}})_{j \in [\ell]}, (o_j^{\text{link}})_{j \in [\ell]}, \omega')$.

For convenience we use the following shorter notations about “good proofs” and “good commitments”:

$$\begin{aligned} \text{GdPf}(\pi') &:= \text{ccII.VerProof}(\text{vk}', x, c', \pi') = 1 \\ \text{GdPf}(\pi^{\text{link}}) &:= \text{CP}_{\text{link}}.\text{VerProof}(\text{vk}^{\text{link}}, c', (c_j)_{j \in [\ell]}, \pi^{\text{link}}) = 1 \\ \text{GdCom}(c_j, u_j^{\text{link}}) &:= \text{VerCommit}(\text{ck}, c_j, u_j^{\text{link}}, o_j^{\text{link}}) = 1 \\ \text{GdCom}'(c', u^{\text{link}}) &:= \text{ccII.VerCommit}(\text{ck}', c', (u_j^{\text{link}})_{j \in [\ell]}, \omega^{\text{link}}) = 1 \\ \text{GdCom}'(c', u') &:= \text{ccII.VerCommit}(\text{ck}', c', (u'_j)_{j \in [\ell]}, o') = 1 \\ R^{\text{link}}(x^{\text{link}}, u^{\text{link}}, \omega^{\text{link}}) &:= \text{ccII.VerCommit}(\text{ck}', x^{\text{link}}, (u_j^{\text{link}})_{j \in [\ell]}, \omega^{\text{link}}) \end{aligned}$$

Let us define the following events:

$$\begin{aligned} \text{bad} &:= \left(\bigvee_{j \in [\ell]} \neg \text{GdCom}(c_j, u_j^{\text{link}}) \vee \neg R(x, u^{\text{link}}, \omega') \right) \\ \text{bad}' &:= (\neg \text{GdCom}'(c', u') \vee \neg R(x, u', \omega')) ; \\ \text{bad}^\circ &:= \left(\bigvee_{j \in [\ell]} \neg \text{GdCom}(c_j, u_j^{\text{link}}) \vee \neg \text{GdCom}'(c', x^{\text{link}}) \vee \neg R^{\text{link}}(x^{\text{link}}, u^{\text{link}}, \omega^{\text{link}}) \right) \end{aligned}$$

By the knowledge soundness of CP_{link} and ccII we have that $\Pr[\text{GdPf}(\pi^{\text{link}}) \wedge \text{bad}^\circ] = \text{negl}(\lambda)$ and $\Pr[\text{GdPf}(\pi') \wedge \text{bad}'] = \text{negl}(\lambda)$, and we abbreviate $n_\lambda := \text{negl}(\lambda)$ for convenience. Let us now first consider the case when cc-SNARK is binding and observe that:

$$\begin{aligned} &\Pr[\text{Game}_{\text{CP.}\mathcal{R}\mathcal{G}_{\text{Com}}, \text{CP.}\mathcal{Z}, \text{CP.}\mathcal{A}, \text{CP.}\mathcal{E}}^{\text{KSND}} = 1] \\ &= \Pr[\text{GdPf}(\pi') \wedge \text{GdPf}(\pi^{\text{link}}) \wedge \text{bad}] \end{aligned} \tag{1}$$

$$\leq \Pr[\text{GdPf}(\pi^{\text{link}}) \wedge \text{bad}^\circ] + \Pr[\text{GdPf}(\pi') \wedge \text{bad} \wedge R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}}) \bigwedge_{j \in [\ell]} \text{GdCom}(c_j, u_j^{\circ})] \tag{2}$$

$$\leq \Pr[\text{GdPf}(\pi') \wedge \neg R(x, u^{\text{link}}, \omega') \wedge R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}})] + n_\lambda \tag{3}$$

$$\leq \Pr[\text{GdPf}(\pi') \wedge \neg R(x, u^{\text{link}}, \omega') \wedge R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}}) \wedge (\neg R^{\text{link}}(c', u', o') \vee u' = u^{\text{link}})] + \tag{4}$$

$$\begin{aligned} &\Pr[R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}}) \wedge R^{\text{link}}(c', u', o') \wedge u' \neq u^{\text{link}}] + n_\lambda \\ &\leq \Pr[\text{GdPf}(\pi') \wedge \neg R(x, u^{\text{link}}, \omega') \wedge R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}}) \wedge (\neg R^{\text{link}}(c', u', o') \vee u' = u^{\text{link}})] + n_\lambda \end{aligned} \tag{5}$$

$$\begin{aligned} &\leq \Pr[\text{GdPf}(\pi') \wedge (\neg R(x, u', \omega') \wedge R^{\text{link}}(c', u', \omega^{\text{link}})) \vee \\ &\quad (\neg R^{\text{link}}(c', u', o') \wedge \neg R(x, u^{\text{link}}, \omega') \wedge R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}}))] + n_\lambda \end{aligned} \tag{6}$$

$$\leq \Pr[\text{GdPf}(\pi') \wedge (\neg R(x, u', \omega') \vee \neg R^{\text{link}}(c', u', o'))] + n_\lambda \tag{7}$$

$$\leq \text{negl}(\lambda) \tag{8}$$

Above, (1) follows by spelling out the winning condition of the experiment considering our construction of CP.VerCommit ; (2) follows first partitioning over bad° and then by observing that $\neg\text{bad}^\circ := R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}}) \wedge_{j \in [\ell]} \text{GdCom}(c_j, u_j^{\text{link}})$; (3) follows by knowledge soundness of CP_{link} ; (4) follows after partitioning on the event $R^{\text{link}}(c', u', o') \wedge u' \neq u^{\text{link}}$; (5) is by the binding property of the commitment of ccII ; ²⁵ (7) holds by using that $\Pr[((E_1 \wedge E'_1) \vee (E_2 \wedge E'_2))] \leq \Pr(E_1 \vee E_2)$; finally, (8) follows by knowledge soundness of ccII .

The case of weak binding. Let us now consider the case in which ccII has only weak binding. In this case the commitment returned by ccII.Prove refers to the whole witness $w = u$, which in the previous proof means that the value ω' returned by ccII.E is empty.

To show that with this change the adversary and extractor still have negligible probability of making the knowledge soundness experiment output 1, we closely follow the analysis we already carried out by equations 1 through 8 above. We slightly deviate after (3) and obtain

$$\begin{aligned}
& \Pr[\text{Game}_{\text{CP}, \mathcal{RG}_{\text{Com}}, \text{CP.Z}, \text{CP.A}, \text{CP.E}}^{\text{KSND}} = 1] \\
& \quad \vdots \\
& \leq \Pr[\text{GdPf}(\pi') \wedge \neg R(x, u^{\text{link}}) \wedge R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}})] + \text{negl}(\lambda) \\
& \leq \Pr[\text{GdPf}(\pi') \wedge \neg R(x, u^{\text{link}}) \wedge R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}}) \wedge (\neg R^{\text{link}}(c', u', o') \vee u' = u^{\text{link}})] + \\
& \quad \Pr[\text{GdPf}(\pi') \wedge \neg R(x, u^{\text{link}}) \wedge R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}}) \wedge R^{\text{link}}(c', u', o') \wedge u' \neq u^{\text{link}}] + n_\lambda
\end{aligned} \tag{3}$$

For the case $u' = u^{\text{link}}$ we proceed exactly as before. For the case $u' \neq u^{\text{link}}$, defining $\text{comsOpen} := R^{\text{link}}(c', u^{\text{link}}, \omega^{\text{link}}) \wedge R^{\text{link}}(c', u', o')$, we have

$$\begin{aligned}
& \Pr[\text{GdPf}(\pi') \wedge \neg R(x, u^{\text{link}}) \wedge \text{comsOpen} \wedge u' \neq u^{\text{link}}] \\
& \leq \Pr[\text{GdPf}(\pi') \wedge \neg R(x, u^{\text{link}}) \wedge \text{comsOpen} \wedge u' \neq u^{\text{link}} \wedge R(u', w')] + n_\lambda \\
& \leq \text{negl}(\lambda)
\end{aligned}$$

where the two inequalities follow respectively from knowledge soundness and weak binding of ccII . \square

B.2 Proof of Zero-Knowledge

$\mathcal{S}_{\text{kg}}(\text{ck}, R)$	$\mathcal{S}_{\text{prv}}(\text{crs}, \text{td}_k, x, (c_j)_{j \in [\ell]})$
$(\text{crs}', \text{td}'_k) \leftarrow \mathcal{S}'_{\text{kg}}(R)$; Parse crs' as $(\text{ck}', \text{ek}', \text{vk}')$ $(\text{crs}^{\text{link}}, \text{td}^{\text{link}}_k) \leftarrow \mathcal{S}^\circ_{\text{kg}}((\text{ck}', \mathcal{D}_u^{\text{link}}, \mathcal{D}_w^{\text{link}}))$ return $(\text{crs} := (\text{crs}^{\text{link}}, \text{crs}'), \text{td}_k := (\text{td}^{\text{link}}_k, \text{td}'_k))$	Parse crs as $(\text{crs}^{\text{link}}, \text{crs}')$; crs' as $(\text{ck}', \text{ek}', \text{vk}')$; td_k as $(\text{td}^{\text{link}}_k, \text{td}'_k)$ $(c', \pi') \leftarrow \mathcal{S}'_{\text{prv}}(\text{crs}', \text{td}'_k, x)$; $\pi^{\text{link}} \leftarrow \mathcal{S}^\circ_{\text{prv}}(\text{crs}^{\text{link}}, \text{td}^{\text{link}}_k, c', (c_j)_{j \in [\ell]})$ return $\pi := (c', \pi^{\text{link}}, \pi')$

Figure 13: Zero-knowledge simulators for our generic CP.

²⁵ We can do this through an adversary that would first run \mathcal{A} and \mathcal{E} and then return $(c', (u^{\text{link}}, \omega^{\text{link}}), (u', o'))$.

$\mathcal{HS}_{\text{kg}}(\text{ck}, R)$	$\mathcal{HS}_{\text{prv}}(\text{crs}, \text{td}_k, \mathbf{x}, \mathbf{w})$
$\text{crs}' \leftarrow \text{ccII.KeyGen}(R)$	Parse \mathbf{x} as $(x, (c_j)_{j \in [\ell]})$; crs as $(\text{crs}^{\text{link}}, \text{crs}')$; crs' as $(\text{ck}', \text{ek}', \text{vk}')$
Parse crs' as $(\text{ck}', \text{ek}', \text{vk}')$	Parse \mathbf{w} as $((u_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, \omega)$; td_k as $(\text{td}_k^{\text{link}}, \text{td}_k')$
$(\text{crs}^{\text{link}}, \text{td}_k^{\text{link}}) \leftarrow \mathcal{S}_{\text{kg}}^\circ(\text{ck}', \mathcal{D}_x^{\text{link}}, \mathcal{D}_u^{\text{link}}, \mathcal{D}_\omega^{\text{link}})$	$(c', \pi', o') \leftarrow \text{ccII.Prove}(\text{ek}', x, (u_j)_{j \in [\ell]}, \omega)$
$\text{crs} := (\text{crs}^{\text{link}}, \text{crs}')$; $\text{td}_k := \text{td}_k^{\text{link}}$	$\pi^{\text{link}} \leftarrow \mathcal{S}_{\text{prv}}^\circ(\text{crs}^{\text{link}}, \text{td}_k^{\text{link}}, c', (c_j)_{j \in [\ell]})$
return $(\text{crs}, \text{td}_k)$	return $(c', \pi^{\text{link}}, \pi')$

Figure 14: Hybrids for proof of ZK of Theorem B.1 (differences with original simulators in blue).

Proof Let \mathcal{A} be an adversary. Since the scheme CP_{link} is zero-knowledge there exists a simulator $\mathcal{S}^\circ = (\mathcal{S}_{\text{kg}}^\circ, \mathcal{S}_{\text{prv}}^\circ)$ such that keys and proof indistinguishability hold for \mathcal{A} as in Definition 2.2. Similarly, since the scheme ccII is zero-knowledge²⁶ there exists a simulator $\mathcal{S}' = (\mathcal{S}'_{\text{kg}}, \mathcal{S}'_{\text{prv}})$ such that keys and proof indistinguishability hold for \mathcal{A} as in Definition 3.3. In Figure 13 we show simulators $\mathcal{S} = (\mathcal{S}_{\text{kg}}, \mathcal{S}_{\text{prv}})$ for the CP scheme of Figure 2, and below we argue that keys and proof indistinguishability hold for such simulators.

PROOF INDISTINGUISHABILITY. fixed arbitrary \mathcal{A} , $x, (c_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, (u_j)_{j \in [\ell]}, \omega$, we define three hybrids (Figure ??): $\mathcal{H}_0, \mathcal{H}_1$ and \mathcal{H}_{sim} , and claim that $\mathcal{H}_0 \approx \mathcal{H}_1 \approx \mathcal{H}_{\text{sim}}$, which, by definition of the hybrids, implies proof indistinguishability. We skip the proof of the claim as it follows from a standard hybrid argument.

Below we use the same notation as in Definition 3.1: define $\mathbf{x} := (x, (c_j)_{j \in [\ell]})$, $\mathbf{w} := ((u_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, \omega)$; the relation \mathbf{R} over pairs (\mathbf{x}, \mathbf{w}) both tests commitment openings and the underlying relation R . \mathcal{H}_0 is defined as the probability that an adversary outputs 1 when a proof is computed through CP.Prove . This is the same as in Definition 2.2 for the case in which \mathcal{A} takes in input an actual proof:

$$\mathcal{H}_0 := \Pr \left[\begin{array}{l} (\mathbf{R}, \text{aux}_R) \leftarrow \mathcal{RG}_{\text{Com}}(1^\lambda) ; (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(\mathbf{R}) ; \pi \leftarrow \text{CP.Prove}(\text{crs}, \mathbf{x}, \mathbf{w}) \\ \mathbf{R}(\mathbf{x}, \mathbf{w}) = 1 \wedge \mathcal{A}(\text{crs}, \text{aux}_R, \pi) = 1 \end{array} \right]$$

In \mathcal{H}_1 we replace the sub-proof π^{link} for CP_{link} with its respective simulated version (see Figure 14 for a definition of $\mathcal{HS}_{\text{prv}}$):

$$\mathcal{H}_1 := \Pr \left[\begin{array}{l} (\mathbf{R}, \text{aux}_R) \leftarrow \mathcal{RG}_{\text{Com}}(1^\lambda) ; (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(\mathbf{R}) ; \pi \leftarrow \mathcal{HS}_{\text{prv}}(\text{crs}, \text{td}_k, \mathbf{x}, \mathbf{w}) \\ \mathbf{R}(\mathbf{x}, \mathbf{w}) = 1 \wedge \mathcal{A}(\text{crs}, \text{aux}_R, \pi) = 1 \end{array} \right]$$

We define \mathcal{H}_{sim} as the simulated proof output as in the standard zero-knowledge experiment (Definition 2.2). We point out that the only change from \mathcal{H}_ℓ consists in replacing the actual proof for ccII with its simulated version:

$$\mathcal{H}_{\text{sim}} := \Pr \left[\begin{array}{l} (\mathbf{R}, \text{aux}_R) \leftarrow \mathcal{RG}_{\text{Com}}(1^\lambda) ; (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(\mathbf{R}) ; \pi \leftarrow \mathcal{S}_{\text{prv}}(\text{crs}, \text{td}_k, \mathbf{x}) \\ \mathbf{R}(\mathbf{x}, \mathbf{w}) = 1 \wedge \mathcal{A}(\text{crs}, \text{aux}_R, \pi) = 1 \end{array} \right]$$

Figure 15: Hybrids for proof indistinguishability of CP.

²⁶ We notice that for this proof we only need the zero-knowledge of ccII , and it does not matter if ccII has binding or weak binding.

Keys indistinguishability: we proceed by a standard hybrid argument. Consider the hybrid simulator \mathcal{HS}_{kg} in Figure 14. By construction of \mathcal{HS}_{kg} and the keys indistinguishability for $\mathcal{S}'_{\text{kg}}, \mathcal{S}_{\text{kg}}^\circ$ we have that:

$$\begin{aligned} & \Pr [(\text{ck}, R, \text{aux}_R) \leftarrow \mathcal{RG}_{\text{Com}}(1^\lambda), \text{crs} \leftarrow \text{CP.KeyGen}(\text{ck}, R) = 1 : \mathcal{A}(\text{ck}, \text{crs}, \text{aux}_R) = 1] \\ & \approx \Pr [(\text{ck}, R, \text{aux}_R) \leftarrow \mathcal{RG}_{\text{Com}}(1^\lambda), (\text{crs}, \text{td}_k) \leftarrow \mathcal{HS}_{\text{kg}}(\text{ck}, R) : \mathcal{A}(\text{ck}, \text{crs}, \text{aux}_R) = 1] \\ & \approx \Pr [(\text{ck}, R, \text{aux}_R) \leftarrow \mathcal{RG}_{\text{Com}}(1^\lambda), (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(\text{ck}, R) : \mathcal{A}(\text{ck}, \text{crs}, \text{aux}_R) = 1] \end{aligned}$$

□

C Supplementary Results on CP_{link}

This section contains the security proof and an extension of the CP_{link} scheme.

C.1 Proof of CP_{link} Security

The following theorem shows that CP_{link} is knowledge-sound and zero-knowledge assuming so is $\text{ss}\Pi$.

Theorem C.1. *Let $\text{CP}_{\text{link}}.\mathcal{RG}$ be a relation generator and $\text{CP}_{\text{link}}.\mathcal{Z}$ be an auxiliary input distribution. If $\text{ss}\Pi$ is $\text{KSND}(\text{ss}\Pi.\mathcal{RG}, \text{ss}\Pi.\mathcal{Z})$ where $\text{ss}\Pi.\mathcal{RG}$ is a relation generator as in Figure 16 and $\text{ss}\Pi.\mathcal{Z} = \text{CP}_{\text{link}}.\mathcal{Z}$, then the CP-SNARK construction CP_{link} given above is $\text{KSND}(\text{CP}_{\text{link}}.\mathcal{RG}, \text{CP}_{\text{link}}.\mathcal{Z})$. Furthermore, if $\text{ss}\Pi$ is composable ZK for $\text{ss}\Pi.\mathcal{RG}$, then CP_{link} is composable ZK for $\text{CP}_{\text{link}}.\mathcal{RG}$.*

$$\begin{array}{l} \text{ss}\Pi.\mathcal{RG}(1^\lambda) \rightarrow ([\mathbf{M}]_1, \text{aux}_R^{\text{link}}) \\ \hline [\mathbf{h}]_1 \leftarrow \text{Ped.Setup}(1^\lambda) \text{ using distribution } \mathcal{D} \\ (R^{\text{link}}, \text{aux}_R^{\text{link}}) \leftarrow \text{CP}_{\text{link}}.\mathcal{RG}(1^\lambda) \\ \text{Define } [\mathbf{M}]_1 \text{ from } [\mathbf{h}]_1, R^{\text{link}} \end{array}$$

Figure 16: Relation generator on which we base $\text{ss}\Pi$ security.

Knowledge Soundness. Consider an arbitrary adversary \mathcal{A} against CP_{link} . From \mathcal{A} we can construct an adversary \mathcal{A}' against $\text{ss}\Pi$ as follows.

$$\begin{array}{l} \mathcal{A}'([\mathbf{M}]_1, \text{crs}, \text{aux}_R, \text{aux}_Z) : \\ \hline \text{Extract } [\mathbf{f}]_1, [\mathbf{h}]_1 \text{ from } [\mathbf{M}]_1 ; ([\mathbf{x}]_1, \pi) \leftarrow \mathcal{A}([\mathbf{h}]_1, R^{\text{link}}, \text{crs}, \text{aux}_R, \text{aux}_Z) ; \text{Parse } [\mathbf{x}]_1 \text{ as } ((c_j)_{j \in [\ell]}, c') \\ \text{return } (c', (c_j)_{j \in [\ell]}, \pi) \\ \hline \mathcal{E}([\mathbf{h}]_1, R^{\text{link}}, \text{crs}, \text{aux}_R, \text{aux}_Z) : \\ \text{Compute matrix } [\mathbf{M}]_1 ; \mathbf{w} \leftarrow \text{ss}\Pi.\mathcal{E}([\mathbf{M}]_1, \text{crs}, \text{aux}_R, \text{aux}_Z) ; \text{Parse } \mathbf{w} \text{ as } ((o_j)_{j \in [\ell]}, o', (\mathbf{u}_j)_{j \in [\ell]}) \\ \text{return } ((\mathbf{u}_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, o') \end{array}$$

By knowledge soundness of $\text{ss}\Pi$, for every such \mathcal{A}' there is an extractor $\text{ss}\Pi.\mathcal{E}$, that we can use to build the above extractor \mathcal{E} for \mathcal{A} . In particular, the knowledge soundness of $\text{ss}\Pi$ and the definition of \mathbf{M} give us that \mathcal{E} 's output is such that the following probability is negligible:

$$\Pr(\text{ss}\Pi.\text{VerProof}(\text{vk}, (c_j)_{j \in [\ell]}, c') = 1 \wedge (\bigvee_{j \in [\ell]} (c_j \neq (o_j, \mathbf{u}_j^\top) \cdot [\mathbf{h}_{[0, n_j]}]_1) \vee c' \neq (o', \mathbf{u}_1^\top, \dots, \mathbf{u}_\ell^\top) \cdot [\mathbf{f}]_1))$$

Hence we can conclude that $\Pr[\text{Game}_{\text{CP}_{\text{link}}.\mathcal{R}\mathcal{G}, \text{CP}_{\text{link}}.\mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{KSND}} = 1] = \Pr[\text{GdPf} \wedge (\text{BadComm} \vee \text{BadRel})] \leq \text{negl}(\lambda)$ using: $\text{GdPf} := \text{CP}_{\text{link}}.\text{VerProof}(\text{vk}, c', (c_j)_{j \in [\ell]}, \pi) = 1$, $\text{BadComm} := \bigvee_{j \in [\ell]} c_j \neq (o_j, \mathbf{u}_j^\top) \cdot [\mathbf{h}_{[0, n_j]}]_1$, $\text{BadRel} := c' \neq (o', \mathbf{u}_1^\top, \dots, \mathbf{u}_\ell^\top) \cdot [\mathbf{f}]_1$.

Zero-Knowledge. From the zero-knowledge property of $\text{ss}\Pi$ we know there exists a simulator $\text{ss}\Pi.\mathcal{S} = (\text{ss}\Pi.\mathcal{S}_{\text{kg}}, \text{ss}\Pi.\mathcal{S}_{\text{prv}})$ such that keys and proof indistinguishability hold for an arbitrary \mathcal{A} as in Definition 2.2. We now define the following key simulator $\text{CP}_{\text{link}}.\mathcal{S}_{\text{kg}}$ such that $\text{CP}_{\text{link}}.\mathcal{S}_{\text{kg}}([\mathbf{h}]_1, R^{\text{link}}) := \text{ss}\Pi.\mathcal{S}_{\text{kg}}([\mathbf{M}]_1)$. Keys indistinguishability follows directly from the assumption on $\text{ss}\Pi.\mathcal{S}_{\text{kg}}$. Analogously, we obtain proof indistinguishability by defining a proof simulator $\text{CP}_{\text{link}}.\mathcal{S}_{\text{prv}}$ such that $\text{CP}_{\text{link}}.\mathcal{S}_{\text{prv}}(\text{crs}, \text{td}_k, c', (c_j)_{j \in [\ell]}) := \text{ss}\Pi.\mathcal{S}_{\text{prv}}(\text{crs}, \text{td}_k, [\mathbf{x}]_1)$, with $[\mathbf{x}]_1 = ((c_j)_{j \in [\ell]}, c')$.

C.2 An extension of CP_{link} for Prefixes of a Committed Vector

Fixed a security parameter λ (and the bilinear group setting for λ as well), $R_{\text{pre}}^{\text{link}}$ is a relation over $(\mathcal{D}_x \times \mathcal{D}_1 \times \dots \times \mathcal{D}_\ell \times \mathcal{D}_\omega)$, where $\mathcal{D}_x = \mathbb{G}_1$, $\mathcal{D}_\omega = \mathbb{Z}_q^{n_\omega+1}$ and $\mathcal{D}_j = \mathbb{Z}_q^{n_j}$ for some n_j such that $n_\omega + \sum_j n_j = m$. $R_{\text{pre}}^{\text{link}}$ is parametrized by a commitment key $[\mathbf{f}]_1 \in \mathbb{G}_1^{m+1}$, and is defined as:

$$R_{\text{pre}}^{\text{link}}(c', (\mathbf{u}_j)_{j \in [\ell]}, (\mathbf{u}_{\ell+1}, o')) = 1 \iff c' \stackrel{?}{=} (o', \mathbf{u}_1^\top, \dots, \mathbf{u}_{\ell+1}^\top) \cdot [\mathbf{f}]_1$$

Similarly to the case of R^{link} , this relation can be expressed as a linear subspace relation, $R_{\mathbf{M}}([\mathbf{x}]_1, \mathbf{w})$, where \mathbf{M} , \mathbf{x} , \mathbf{w} are as follows:

$$\begin{array}{c} \overbrace{\begin{bmatrix} c_1 \\ \vdots \\ c_\ell \\ c' \end{bmatrix}}^{\mathbf{x}} = \overbrace{\begin{bmatrix} h_0 & 0 & \dots & 0 & 0 & \mathbf{h}_{[1, n_1]} & 0 & \dots & 0 & 0 \\ 0 & h_0 & \dots & 0 & 0 & 0 & \mathbf{h}_{[1, n_2]} & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & h_0 & 0 & 0 & 0 & \dots & \mathbf{h}_{[1, n_\ell]} & 0 \\ 0 & 0 & \dots & 0 & f_0 & \mathbf{f}_{[1, n_1]} & \mathbf{f}_{[n_1+1, n_2]} & \dots & \mathbf{f}_{[n_{\ell-1}+1, n_\ell]} & \mathbf{f}_{[n_\ell+1, n_{\ell+1}]} \end{bmatrix}}^{\mathbf{M}} \overbrace{\begin{pmatrix} o_1 \\ \vdots \\ o_\ell \\ o' \\ \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_{\ell+1} \end{pmatrix}}^{\mathbf{w}} \end{array}$$

Given the above encoding, it is straightforward to extend our scheme CP_{link} to support the relation $R_{\text{pre}}^{\text{link}}$ instead of R^{link} .

D A zkSNARK for Linear Subspaces

Here we recall the QA-NIZK scheme for linear subspaces Π'_{as} of Kiltz and Wee [KW15], in the MDDH setting where $k = 1$.

$$\text{ss}\Pi.\text{KeyGen}([\mathbf{M}]_1 \in \mathbb{G}_1^{l \times t}): \mathbf{k} \leftarrow_{\mathcal{S}} \mathbb{Z}_q^l, a \leftarrow_{\mathcal{S}} \mathbb{Z}_q; \mathbf{P} := \mathbf{M}^\top \mathbf{k}; \mathbf{C} := a \cdot \mathbf{k}$$

return (ek := $[\mathbf{P}]_1 \in \mathbb{G}_1^t$, vk := $([\mathbf{C}]_2, [a]_2) \in \mathbb{G}_2^l \times \mathbb{G}_2$)

ss Π .Prove(ek, $[\mathbf{x}]_1, \mathbf{w}$): **return** $[\pi]_1 \leftarrow \mathbf{w}^\top [\mathbf{P}]_1 \in \mathbb{G}_1$

ss Π .VerProof(vk, $[\mathbf{x}]_1, [\pi]_1$): check that $[\mathbf{x}]_1^\top \cdot [\mathbf{C}]_2 = [\pi]_1 \cdot [a]_2$

ss Π .S_{kg}(1^λ): run as ss Π .KeyGen and output $\text{td}_k = \mathbf{k}$ and (ek, vk)

ss Π .S_{prv}($\text{td}_k, [\mathbf{x}]_1$): **return** $[\pi]_1 \leftarrow \mathbf{k}^\top [\mathbf{x}]_1$

In the following theorem we prove the knowledge soundness of the scheme given above. The proof holds under the discrete logarithm assumption in the algebraic group model of [FKL18]; this can also be interpreted as a proof in the (bilinear) generic group model. We also note that a similar proof about the use of this scheme in a non-falsifiable setting [KW15] also appeared in [FLSZ17].

Theorem D.1. *Assume that \mathcal{D}_{mtx} is a witness sampleable matrix distribution. Then, under the discrete logarithm assumption, in the algebraic group model, the QA-NIZK Π'_{as} in [KW15] (in the MDDH setting $k = 1$) is a knowledge-sound SNARK for linear subspace relations with matrices from \mathcal{D}_{mtx} .*

Proof Consider an algebraic adversary \mathcal{A} against the knowledge soundness of ss Π . Its input consists of the matrix $[\mathbf{M}]_1$ and the associated auxiliary input aux , along with the common reference string $[\mathbf{P}]_1, [\mathbf{C}]_2, [a]_2$. Let $[\mathbf{z}]_1$ be a vector that contains \mathbf{M} and the portion of aux that has elements from the group \mathbb{G}_1 , and also assume $[\mathbf{z}]$ includes $[1]_1$. \mathcal{A} returns a pair $([\mathbf{x}]_1, [\pi]_1)$ along with coefficients that “explain” these elements as linear combinations of its input in the group \mathbb{G}_1 . Let these coefficients be:

$$\begin{aligned} [\mathbf{x}]_1 &:= \mathbf{X}_0 [\mathbf{P}]_1 + \mathbf{X}_1 [\mathbf{z}]_1 = \mathbf{X}_0 [\mathbf{M}^\top \mathbf{k}]_1 + \mathbf{X}_1 [\mathbf{z}]_1 \\ [\pi]_1 &:= \boldsymbol{\pi}_0^\top [\mathbf{P}]_1 + \boldsymbol{\pi}_1^\top [\mathbf{z}]_1 = \boldsymbol{\pi}_0^\top [\mathbf{M}^\top \mathbf{k}]_1 + \boldsymbol{\pi}_1^\top [\mathbf{z}]_1 \end{aligned}$$

We define the extractor \mathcal{E} to be the algorithm that runs the algebraic \mathcal{A} and returns $\mathbf{w} := \boldsymbol{\pi}_0$, i.e., the coefficients of $[\pi]_1$ corresponding to \mathbf{P} . Next, we have to show that the probability that the output of $(\mathcal{A}, \mathcal{E})$ satisfies verification while $\mathbf{x} \neq \mathbf{M}\mathbf{w}$ is negligible. In other words, assume that the output of \mathcal{A} is such that:

$$[\mathbf{x}]_1^\top \cdot [a \cdot \mathbf{k}]_2 = [\pi]_1 \cdot [a]_2 \quad \text{and} \quad [\mathbf{x}]_1 \neq [\mathbf{M}]_1 \boldsymbol{\pi}_0$$

If \mathcal{A} returns such a tuple with non-negligible probability, we show how to build an algorithm \mathcal{B} that on input $([\mathbf{k}]_1, [\mathbf{k}]_2)$ outputs nonzero elements $\mathbf{A} \in \mathbb{Z}_q^{l \times l}$, $\mathbf{b} \in \mathbb{Z}_q^l, c \in \mathbb{Z}_q$ such that

$$\mathbf{k}^\top \mathbf{A} \mathbf{k} + \mathbf{k}^\top \mathbf{b} + c = 0$$

Such a \mathcal{B} can in turn be reduced to an algorithm \mathcal{B}' that solves discrete log, i.e., on input $([\alpha]_1, [\alpha]_2)$ return α .

Algorithm $\mathcal{B}([\mathbf{k}]_1, [\mathbf{k}]_2)$ proceeds as follows. First, it uses \mathcal{D}_{mtx} to sample $([\mathbf{M}]_1, \text{aux})$ along with its \mathbb{G}_1 witness (i.e., a vector \mathbf{z} of entries in \mathbb{Z}_q). Second, it samples $a \leftarrow_{\$} \mathbb{Z}_q$ and runs $\mathcal{A}([\mathbf{z}, \mathbf{P}]_1, [a, a \cdot \mathbf{k}]_2)$ (notice that \mathcal{A} 's input can be efficiently simulated). Third, once received the output of \mathcal{A} , \mathcal{B} sets $\mathbf{A} := \mathbf{X}_0 \mathbf{M}^\top$, $\mathbf{b} := \mathbf{X}_1 \mathbf{z} - \mathbf{M} \boldsymbol{\pi}_0$ and $c = -\boldsymbol{\pi}_1^\top \mathbf{z}$. Notice that

$$\begin{aligned} \mathbf{k}^\top \mathbf{A} \mathbf{k} + \mathbf{k}^\top \mathbf{b} + c &= \mathbf{k}^\top \mathbf{X}_0 \mathbf{M}^\top \mathbf{k} + \mathbf{k}^\top \mathbf{X}_1 \mathbf{z} - \mathbf{k}^\top \mathbf{M} \boldsymbol{\pi}_0 - \boldsymbol{\pi}_1^\top \mathbf{z} \\ &= \mathbf{k}^\top \mathbf{X}_0 \mathbf{M}^\top \mathbf{k} + \mathbf{k}^\top \mathbf{X}_1 \mathbf{z} - \pi \\ &= \mathbf{k}^\top \mathbf{x} - \pi = 0 \end{aligned}$$

Also, one among \mathbf{A}, \mathbf{b} and c must be nonzero. Indeed, if they are all zero then $\mathbf{X}_1 \mathbf{z} - \mathbf{M} \boldsymbol{\pi}_0 = 0$, that is $\mathbf{x} = \mathbf{M} \boldsymbol{\pi}_0$, which contradicts our assumption on \mathcal{A} 's output.

To finish the proof, we show how the above problem can be reduced to discrete log in asymmetric groups, i.e., \mathcal{B}' on input $([\alpha]_1, [\alpha]_2)$ returns α . \mathcal{B}' samples $\mathbf{r}, \mathbf{s} \in \mathbb{Z}_q^l$ and implicitly sets $\mathbf{k} := \alpha \cdot \mathbf{r} + \mathbf{s}$. It is easy to see that $([\mathbf{k}]_1, [\mathbf{k}]_2)$ can be efficiently simulated with a distribution identical to the one expected by \mathcal{B} . Next, given a solution $(\mathbf{A}, \mathbf{b}, c)$ such that $\mathbf{k}^\top \mathbf{A} \mathbf{k} + \mathbf{k}^\top \mathbf{b} + c = 0$ one can find $a', b', c' \in \mathbb{Z}_q$ such that:

$$\begin{aligned} 0 &= (\alpha \mathbf{r} + \mathbf{s})^\top \mathbf{A} (\alpha \mathbf{r} + \mathbf{s}) + (\alpha \mathbf{r} + \mathbf{s})^\top \mathbf{b} + c \\ &= \alpha^2 (\mathbf{r}^\top \mathbf{A} \mathbf{r}) + \alpha \cdot (\mathbf{r}^\top \mathbf{A} \mathbf{s} + \mathbf{s}^\top \mathbf{A} \mathbf{r} + \mathbf{r}^\top \mathbf{b}) + (\mathbf{s}^\top \mathbf{A} \mathbf{s} + \mathbf{s}^\top \mathbf{b} + c) \\ &= a' \alpha^2 + b' \alpha + c' \end{aligned}$$

In particular, with overwhelming probability (over the choice of \mathbf{s} that is information theoretically hidden from \mathcal{B} 's view) $c' \neq 0$. From this solution \mathcal{B}' can solve the system and extract α . \square

E A Construction of PolyCom and CP_{poly} from zk-vSQL

We show a pairing-based construction of the commitment PolyCom and CP-SNARK CP_{poly} that are “extracted” from the verifiable polynomial delegation scheme of Zhang et al. [ZGK⁺17b]. Basically, we separate the algorithms related to committing from the ones related to proving and verifying evaluations of committed polynomials. Except for that, the only noticeable difference is that in our case we can prove that c_y opens to $y = f(\mathbf{x})$ (with respect to c_f which opens to f) for a given c_y instead of one that is freshly generated at proving time. As we show below, this difference would matter only for zero-knowledge, for which we give a proof a slightly different than the one in [ZGK⁺17b].

Setup(1^λ): let \mathcal{F} be μ -variate polynomials of degree d in each variable. Sample $\alpha, \beta, s_1, \dots, s_{\mu+1} \leftarrow_{\$} \mathbb{Z}_q$ uniformly at random, compute $\mathbb{P} = \{[\prod_{i \in W} s_i, \alpha \prod_{i \in W} s_i]_1\}_{W \in \mathcal{W}_{\mu,d}}$, and output $\text{ck} = (\mathbb{P}, [s_{\mu+1}, \alpha s_{\mu+1}, \beta s_{\mu+1}]_1, [\alpha, \beta, s_1, \dots, s_{\mu+1}]_2)$.

ComPoly(ck, f) $\rightarrow (c_f, o_f)$: sample $o_f \leftarrow_{\$} \mathbb{Z}_q$, compute $c_{f,1} = [f(s_1, \dots, s_\mu) + o_f s_{\mu+1}]_1$, $c_{f,2} = [\alpha (f(s_1, \dots, s_\mu) + o_f s_{\mu+1})]_1$ and output $c_f = (c_{f,1}, c_{f,2})$.

ComVal(ck, y) $\rightarrow (c_y, o_y)$: sample $o_y \leftarrow_{\$} \mathbb{Z}_q$, compute $c_{y,1} = [y + o_y s_{\mu+1}]_1$, $c_{y,2} = [\beta (y + o_y s_{\mu+1})]_1$ and output $c_y = (c_{y,1}, c_{y,2})$.

CheckCom(ck, c): we assume one knows the **type** for which c was created. If **type** = **pol**, output 1 iff $c_1 \cdot [\alpha]_2 = c_2 \cdot [1]_2$. If **type** = **val**, output 1 iff $c_1 \cdot [\beta]_2 = c_2 \cdot [1]_2$.

VerCommit(ck, c, f, o) $\rightarrow b$: output $c_1 \stackrel{?}{=} [f(s_1, \dots, s_\mu) + o s_{\mu+1}]_1$.

Theorem E.1 ([ZGK⁺17b]). *Under the $(\mu + 1)\delta$ -Strong Diffie-Hellman and the (δ, μ) -Extended Power Knowledge of Exponent assumptions (see [ZGK⁺17b]), PolyCom is an extractable trapdoor polynomial commitment.*

The proof of the theorem follows from Theorem 1 in [ZGK⁺17b]. The only property that is not proved there is the trapdoor property, which is however straightforward to see if one considers a simulator \mathcal{S}_{ck} that sets the values $\alpha, \beta, s_1, \dots, s_{\mu+1}$ as trapdoor.

Next, we show a CP-SNARK for polynomial evaluation relations R^{poly} :

$\text{CP}_{\text{poly}}.\text{KeyGen}(\text{ck})$: set $\text{ek} := \text{ck}$ and $\text{vk} := ([\alpha, \beta, s_1, \dots, s_{\mu+1}]_2)$

$\text{CP}_{\text{poly}}.\text{Prove}(\text{ek}, \mathbf{x}, f, y, o_f, o_y)$: sample $o_1, \dots, o_\mu \leftarrow \mathbb{Z}_q$; find polynomials q_i such that $f(Z_1, \dots, Z_\mu) + o_f Z_{\mu+1} - (y + o_y Z_{\mu+1}) = \sum_{i=1}^{\mu} (Z_i - x_i)(q_i(Z_1, \dots, Z_\mu) + o_i Z_{\mu+1}) + X_{\mu+1}(o_f - o_y - \sum_{i=1}^{\mu} o_i(Z_i - x_i))$.

For $i = 1$ to μ , compute $c_i := (c_{i,1}, c_{i,2}) = [q_i(s_1, \dots, s_\mu) + o_i s_{\mu+1}, \alpha(q_i(s_1, \dots, s_\mu) + o_i s_{\mu+1})]_1$,
 $c_{\mu+1} := (c_{\mu+1,1}, c_{\mu+1,2}) = [o_f - o_y - \sum_{i=1}^{\mu} o_i(s_i - x_i), \alpha(o_f - o_y - \sum_{i=1}^{\mu} o_i(s_i - x_i))]_1$. Output
 $\pi := (c_1, \dots, c_{\mu+1})$.

$\text{CP}_{\text{poly}}.\text{VerProof}(\text{vk}, \mathbf{x}, c_f, c_y, \pi)$: parse $\pi := (c_1, \dots, c_{\mu+1})$, output $(c_{f,1} - c_{y,1}) \cdot [1]_2 = c_{\mu+1,1} \cdot [s_{\mu+1}]_2 \sum_{i=1}^{\mu} c_{i,1} \cdot [(s_i - x_i)]_2$ and $\text{CheckCom}(\text{vk}, c_f) \wedge \text{CheckCom}(\text{vk}, c_y) \wedge_{i=1}^{\mu+1} \text{CheckCom}(\text{vk}, c_i)$.

Theorem E.2 ([ZGK⁺17b]). *Under the $(\mu + 1)\delta$ -Strong Diffie-Hellman and the (δ, μ) -Extended Power Knowledge of Exponent assumptions (see [ZGK⁺17b]), CP_{poly} is a zero-knowledge CP-SNARK for R^{poly} .*

Correctness and knowledge soundness are immediate from Theorem 1 in [ZGK⁺17b]. The only difference is in the zero-knowledge property. For this, consider the following proof simulator algorithm, $\mathcal{S}_{\text{prv}}(\text{td}, \mathbf{x}, c_f, c_y)$: for $i = 1$ to μ , sample $c_{i,1} \leftarrow \mathbb{G}_1$ and compute $c_{i,2} = \alpha \cdot c_{i,1}$. Next, compute $c_{\mu+1,1}$ such that $(c_{f,1} - c_{y,1}) \cdot [1]_2 = c_{\mu+1,1} \cdot [s_{\mu+1}]_2 + \sum_{i=1}^{\mu} c_{i,1} \cdot [(s_i - x_i)]_2$ holds and set $c_{\mu+1,2} \leftarrow \beta \cdot c_{\mu+1,1}$. It is straightforward to check that proofs created by \mathcal{S}_{prv} are identically distributed to the ones returned by $\text{CP}_{\text{poly}}.\text{Prove}$.

F Additional Material on CP-SNARKs for PolyCom

In this section we present more CP-SNARKs for PolyCom.

F.1 Proof of our CP_{sc}

We give a full description of the interactive protocol in Figure 3.

Proof We show the security of our protocol by reducing it to the one of [ZGK⁺17b, Construction 2]. For this let us recall the following theorem from [ZGK⁺17b]:

Theorem F.1 ([ZGK⁺17b, Theorem 2]). *For any μ -variate total-degree- d polynomial $g : \mathbb{F}^\mu \rightarrow \mathbb{F}$ with m non-zero coefficients, assuming Com is an extractable linearly homomorphic commitment scheme, and CP_{eq} is a zero-knowledge non-interactive argument for testing equality of commitments for Com , then there is an interactive argument for the relation*

$$\text{VerCommit}(\text{ck}, c_t, t, o_t) = 1 \wedge t = \sum_{\mathbf{b} \in \{0,1\}^\mu} g(\mathbf{b})$$

Moreover, we recall below the last two steps of Construction 2 in [ZGK⁺17b] (i.e., Construction 2 is the same as in our Figure 3 with the blue part replaced by the following steps):

- 1: Common input: c_t, g ; \mathcal{P} 's input: (t, o_t)
- 2: $\mathcal{P} : (c_\mu^*, o_\mu^*) \leftarrow \text{ComVal}(\text{ck}, g(\mathbf{s}))$; $\pi^* \leftarrow \text{CP}_{\text{eq}}.\text{Prove}(\text{ck}, (c_\mu^*, \text{com}_\mu), g(\mathbf{s}), (o_\mu^*, \rho_\mu))$
- 3: $\mathcal{P} \rightarrow \mathcal{V} : c_\mu^*, o_\mu^*, \pi^*$
- 4: $\mathcal{V} : \text{VerCommit}(\text{cvk}, c_\mu^*, g(\mathbf{s}), o_\mu^*) \wedge \text{CP}_{\text{eq}}.\text{VerProof}(\text{vk}, (c_\mu^*, \text{com}_\mu), \pi^*)$

For knowledge soundness, the idea of the proof is that for any adversary \mathcal{A} against CP_{sc} we can create an adversary \mathcal{B} against Construction 2 in [ZGK⁺17b].

Similarly to [ZGK⁺17b], we begin by observing that the commitments c_1, c_2 as well as all the commitments com_{a_j} 's sent during the μ rounds are extractable. By extractability, for any successful \mathcal{A} there exists an extractor $\mathcal{E}_{\mathcal{A}}$ that, on the same input of \mathcal{A} , outputs with all but negligible probability valid openings of all these commitments. Thus we define \mathcal{B} as the adversary that executes $(\mathcal{A}, \mathcal{E}_{\mathcal{A}})$, obtains g_1, g_2 , reconstructs the polynomial $g(\mathbf{S})$, and then keeps executing \mathcal{A} until the end of the protocol, forwarding its messages to its challenger. This is done until the last step where \mathcal{A} sends c'_1, c'_2, π^* . Notice that \mathcal{B} also has the commitments com_{a_j} sent by \mathcal{A} in step μ as well as their openings extracted through $\mathcal{E}_{\mathcal{A}}$. Thus, \mathcal{B} can compute homomorphically the commitment com_{μ} and its opening.

With this knowledge, \mathcal{B} executes the last two lines in Figure 3 (acting as the verifier): if all verifications pass and \mathcal{B} has an opening of com_{μ} to $g(\mathbf{s})$, then it executes the lines 2–4 above and sends $(c_{\mu}^*, o_{\mu}^*, \pi^*)$ to its challenger.

If all verifications pass but \mathcal{B} has an opening of com_{μ} to a value different from $g(\mathbf{s})$, then it must be the case that \mathcal{A} cheated in one of the proofs π_1, π_2, π^* . By the knowledge soundness of CP_{poly} and CP_{prd} this however occurs only with negligible probability.

To show zero-knowledge, we build a simulator that can simulate the verifier's view without knowing the prover's input. Our simulator is the same as the one in [ZGK⁺17b] up to their step (d). For step (e), we let our simulator additionally create commitments (c'_1, c'_2) to dummy values and then run the ZK simulators of CP_{poly} and CP_{prd} to simulate proofs (π_1, π_2, π^*) . By the proof in [ZGK⁺17b], the verifier's transcript except for the last message $(c'_1, c'_2, \pi_1, \pi_2, \pi^*)$ is indistinguishable from an honest one. The indistinguishability with respect to the last message follows immediately from the zero-knowledge CP_{poly} and CP_{prd} . \square

F.2 Proof of Security of CP_{had}

Proof Let \mathcal{A}_{had} be the adversary against CP_{had} that, on input $(\text{ck}, \text{ek}_s, \text{ek}_p)$ and interacting with the random oracle H , returns a statement $(c_j)_{j \in [3]}$ and a proof π that verifies correctly. For any such \mathcal{A}_{had} we can define a non-interactive adversary $\mathcal{A}_{\text{had}}^*$ that additionally takes as input a sequence of random values \mathbf{r}_i , for $i = 1$ to Q , such that \mathbf{r}_i is used to answer the i -th query of \mathcal{A}_{had} to the random oracle H . For any \mathcal{A}_{had} making Q queries to H there exists an index $i \in [0, Q]$ such that the commitments $(c_j)_{j \in [3]}$ returned at the end of its execution were queried to H in the i -th query (letting $i = 0$ being the case in which they were not asked at all). From the above adversary $\mathcal{A}_{\text{had}}^*$ we can define \mathcal{A}_{com} as the non-uniform adversary that on input $(\text{ck}, \text{ek}_s, \text{ek}_p, \mathbf{r}_1, \dots, \mathbf{r}_{i-1})$ runs \mathcal{A}_{had} (in the same way as $\mathcal{A}_{\text{had}}^*$ does) up to its i -th query $H((c_j)_{j \in [3]})$ and returns $(c_j)_{j \in [3]}$. By the extractability of the commitment, for \mathcal{A}_{com} there exists an extractor Ext_{com} that on the same input of \mathcal{A}_{com} outputs openings $(\tilde{u}_j)_{j \in [3]}, (o_j)_{j \in [3]}$. We define the extractor \mathcal{E}_{had} to be the one that runs Ext_{com} and returns its output. Notice that by the extractability of PolyCom it holds $\text{VerCommit}(\text{ck}, c_j, \tilde{u}_j, o_j)$ for $j = 0, 1, 2$ with all but negligible probability.

Next, we need to argue that this adversary-extractor pair $(\mathcal{A}_{\text{had}}, \mathcal{E}_{\text{had}})$ has negligible probability of winning in the knowledge soundness experiment. From $\mathcal{A}_{\text{had}}^*$ we can define two adversaries \mathcal{A}_p and \mathcal{A}_{sc} against CP_{poly} and CP_{sc} respectively, and by using the knowledge soundness of the two CP-SNARKs we have that for each of these adversaries there is a corresponding extractor that gives us a value t such that $\tilde{u}_0(\mathbf{r}) = t$ and $t = \sum_{\mathbf{b} \in \{0,1\}^{\mu}} \tilde{e}q(\mathbf{r}, \mathbf{b}) \cdot \tilde{u}_1(\mathbf{b}) \cdot \tilde{u}_2(\mathbf{b})$ hold respectively with all but negligible probability. Furthermore, the binding of PolyCom implies that the values and openings

for all the commitments $(c_j)_{j \in [3]}$, c_t obtained using these extractors are all the same with all but negligible probability (otherwise we could define a reduction against the binding of PolyCom).

Since $\text{VerCommit}(\text{ck}, c_j, \tilde{u}_j, o_j)$ for $j = 0, 1, 2$, the only way for the adversary to win is when the relation R^{had} is not satisfied. Since we have vectors in MLE form, the check of relation R^{had} can be equivalently written as $\forall \mathbf{b} \in \{0, 1\}^\mu : \tilde{u}_0(\mathbf{b}) \stackrel{?}{=} \tilde{u}_1(\mathbf{b}) \cdot \tilde{u}_2(\mathbf{b})$. Let us define the polynomial $\tilde{u}_0^*(\mathbf{X}) = \sum_{\mathbf{b} \in \{0, 1\}^\mu} \tilde{e}q(\mathbf{X}, \mathbf{b}) \cdot \tilde{u}_1(\mathbf{b}) \cdot \tilde{u}_2(\mathbf{b})$; essentially $\tilde{u}_0^*(\mathbf{X})$ is the MLE of the vector that should correctly verify the R^{had} relation. In particular, by lemma 5.1, $\tilde{u}_0^*(\mathbf{X})$ agrees with $\tilde{u}_1(\mathbf{X}) \cdot \tilde{u}_2(\mathbf{X})$ on all boolean points. Thus, if the relation does not hold we must have $\tilde{u}_0^*(\mathbf{X}) \neq \tilde{u}_0(\mathbf{X})$. However, from above we have that $\tilde{u}_0(\mathbf{r}) = \tilde{u}_0^*(\mathbf{r})$ holds. Notice that from the construction of \mathcal{E}_{had} , the polynomials $\tilde{u}_0(\mathbf{X}), \tilde{u}_1(\mathbf{X}), \tilde{u}_2(\mathbf{X})$ are independent from \mathbf{r} (this is because the extractor \mathcal{E}_{com} that returned this polynomial did not have $\mathbf{r} = \mathbf{r}_i$ among its inputs), and $\tilde{u}_0^*(\mathbf{X})$ is fully determined from $\tilde{u}_1(\mathbf{X}), \tilde{u}_2(\mathbf{X})$. Therefore, by the Schwartz-Zippel lemma, the event $\tilde{u}_0^*(\mathbf{X}) \neq \tilde{u}_0(\mathbf{X}) \wedge \tilde{u}_0(\mathbf{r}) = \tilde{u}_0^*(\mathbf{r})$ occurs with negligible probability over the random choice of \mathbf{r} .

The zero-knowledge of CP_{had} relies on the hiding of PolyCom and the zero-knowledge of CP_{poly} and CP_{sc} . Building simulators \mathcal{S}_{kg} and \mathcal{S}_{prv} for CP_{had} from the corresponding simulators for CP_{poly} and CP_{sc} is fairly straightforward and is omitted here. \square

F.3 Proof of CP_{sfprm}

Proof Let $\mathcal{A}_{\text{sfprm}}$ be the adversary against CP_{sfprm} that, on input (ck, ek_p) and interacting with the random oracle H , returns a statement $(\phi, x, (c_j)_{j \in [\ell]})$ and a proof π that verifies correctly. For any such $\mathcal{A}_{\text{sfprm}}$ we can define a non-interactive adversary $\mathcal{A}_{\text{sfprm}}^*$ that additionally takes as input a sequence of random values (r_i, s_i) , for $i = 1$ to Q , such that (r_i, s_i) are used to answer the i -th query of $\mathcal{A}_{\text{sfprm}}$ to the random oracle H . For any $\mathcal{A}_{\text{sfprm}}$ making Q queries to H there exists an index $i \in [0, Q]$ such that for the relation statement $(\phi, x, (c_j)_{j \in [\ell]})$ returned at the end of its execution, $((c_{\phi, j})_{j \in [0, \ell]}, \mathbf{x}, (c_j)_{j \in [\ell]})$ was queried to H in the i -th query (letting $i = 0$ being the case in which they were not asked at all, and $c_{\phi, j}$ be deterministically derived from ϕ). From the above adversary $\mathcal{A}_{\text{sfprm}}^*$ we can define \mathcal{A}_{com} as the non-uniform adversary that on input $(\text{ck}, \text{ek}_p, r_1, s_1, \dots, r_{i-1}, s_{i-1})$ runs $\mathcal{A}_{\text{sfprm}}$ (in the same way as $\mathcal{A}_{\text{sfprm}}^*$ does) up to its i -th query $H((c_{\phi, j})_{j \in [0, \ell]}, \mathbf{x}, (c_j)_{j \in [\ell]})$ and returns $(c_j)_{j \in [\ell]}$. By the extractability of the commitment, for \mathcal{A}_{com} there exists an extractor \mathcal{E}_{com} that on the same input of \mathcal{A}_{com} outputs openings $(\tilde{u}_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}$. We define the extractor $\mathcal{E}_{\text{sfprm}}$ to be the one that runs \mathcal{E}_{com} and returns its output. Notice that by the extractability of PolyCom it holds $\text{VerCommit}(\text{ck}, c_j, \tilde{u}_j, o_j)$ for $j = 0, 1, 2$ with all but negligible probability.

Next, we need to argue that this adversary-extractor pair $(\mathcal{A}_{\text{sfprm}}, \mathcal{E}_{\text{sfprm}})$ has negligible probability of winning in the knowledge soundness experiment. Recall that we have $\text{VerCommit}(\text{ck}, c_j, \tilde{u}_j, o_j)$ for $j \in [\ell]$ and, by the linear homomorphic property of PolyCom, for all $j \in [0, \ell]$, c'_j and c''_j are commitments to the MLE of $\mathbf{y}'_j := \mathbf{y}_j + r \cdot \mathbf{v}_j - s \cdot \mathbf{1}_j$ and $\mathbf{y}''_j := \mathbf{y}_j + r \cdot \phi_j - s \cdot \mathbf{1}_j$ respectively. Also, in order for the adversary to be successful it must be the case that the relation does not hold, i.e., \mathbf{y} is *not* a self-permutation according to ϕ . Notice that the vector \mathbf{y} is independent of (r, s) since it was returned by \mathcal{E}_{com} without having these values in its view. This allows us to argue that with overwhelming probability over the choice of r it is the case that at least one of the entries of $\mathbf{y} + r \cdot \phi$ is not in $\mathbf{y} + r \cdot \mathbf{v}$. Moreover, when these vectors have different entries the equation $\prod_i (y_i + r \cdot i - s) = \prod_i (y_i + r \cdot \phi(i) - s)$ holds with negligible probability over the choice of s by the Schwartz-Zippel lemma.

Hence we have that with all but negligible probability $\prod_i (y_i + r \cdot i - s) \neq \prod_i (y_i + r \cdot \phi(i) - s)$, which means that at one of the statements in the CP_{ipd} , CP_{prd} or CP_{eq} proofs is not correct. We can reduce these cases to the knowledge soundness of CP_{ipd} , CP_{prd} or CP_{eq} using a fairly standard reduction, in which from an adversary $\mathcal{A}_{\text{sfrpm}}^*$ that falls into the above conditions (i.e., an (r, s) that cause the above inequality) we build either an adversary \mathcal{A}_{ipd} against CP_{ipd} , or an adversary \mathcal{A}_{prd} against CP_{prd} or an \mathcal{A}_{eq} against CP_{eq} .

The zero-knowledge of CP_{sfrpm} follows from the hiding of PolyCom (for creating dummy commitments $(c_{z'_j}, c_{z''_j})_{j \in [0 \dots \ell]}$) and the zero-knowledge of all the underlying CP-SNARKs. \square

F.4 Proof of CP_{lin}

Proof Let \mathcal{A}_{lin} be the adversary against CP_{lin} that, on input $(\text{ck}, \text{ek}_s, \text{ek}_p)$ and interacting with the random oracles H_1, H_2 , returns a statement $(\mathbf{F}, \mathbf{x}, c_u)$ and a proof π that verifies correctly. For any such \mathcal{A}_{lin} we can define a non-interactive adversary $\mathcal{A}_{\text{lin}}^*$ that additionally takes as input a sequence of random values $\{\mathbf{r}_i\}_i, \{\boldsymbol{\sigma}_j\}_j$, for $i = 1$ to Q_1 and $j = 1$ to Q_2 , such that \mathbf{r}_i (resp. $\boldsymbol{\sigma}_j$) is used to answer the i -th (resp. j -th) query of \mathcal{A}_{lin} to the random oracle H_1 (resp. H_2). For any \mathcal{A}_{lin} making Q_1 queries to H_1 there exists an index $i \in [0, Q_1]$ such that for the statement $(\mathbf{F}, \mathbf{x}, c_u)$ returned at the end of its execution the i -th query to H_1 (letting $i = 0$ being the case in which they were not asked at all) is (c_F, \mathbf{x}, c_u) . From the above adversary $\mathcal{A}_{\text{lin}}^*$ we can define \mathcal{A}_{com} as the non-uniform adversary that on input $(\text{ck}, \text{ek}_s, \mathbf{r}_1, \dots, \mathbf{r}_{i-1})$ runs \mathcal{A}_{lin} (in the same way as $\mathcal{A}_{\text{lin}}^*$ does) up to its i -th query $H(c_F, \mathbf{x}, c_u)$ and returns c_u . By the extractability of the commitment, for \mathcal{A}_{com} there exists an extractor \mathcal{E}_{com} that on the same input of \mathcal{A}_{com} outputs an opening \tilde{u}, o_u . We define the extractor \mathcal{E}_{lin} to be the one that runs Ext_{com} and returns its output. Notice that by the extractability of PolyCom it holds $\text{VerCommit}(\text{ck}, c_u, \tilde{u}, o_u)$ with all but negligible probability.

Next, we need to argue that this adversary-extractor pair $(\mathcal{A}_{\text{lin}}, \mathcal{E}_{\text{lin}})$ has negligible probability of winning in the knowledge soundness experiment. In a similar way as we argued extractability of c_u , we can show that it is possible to extract the polynomial g_1 that correctly opens c_1 .

Recall that the adversary is successful if the verifications pass and the relation does not hold, i.e., $\mathbf{F} \cdot \mathbf{u} \neq \mathbf{x}$. Considering MLEs, this means there is some $\mathbf{a} \in \{0, 1\}^\nu$ such that

$$\tilde{x}(\mathbf{a}) \neq \sum_{\mathbf{b} \in \{0, 1\}^\mu} \tilde{F}(\mathbf{a}, \mathbf{b}) \tilde{u}(\mathbf{b}).$$

This means that the following polynomial inequality holds:

$$\tilde{x}(\mathbf{R}) \neq \sum_{\mathbf{b} \in \{0, 1\}^\mu} \tilde{F}(\mathbf{R}, \mathbf{b}) \cdot \tilde{u}(\mathbf{b})$$

First, we argue that with all but negligible probability over the choice of \mathbf{r} we have $t = \tilde{x}(\mathbf{r}) \neq \sum_{\mathbf{b} \in \{0, 1\}^\mu} \tilde{F}(\mathbf{r}, \mathbf{b}) \tilde{u}(\mathbf{b})$. Indeed, \mathbf{r} is random and independent from $\mathbf{x}, \tilde{F}, \tilde{u}$ and the two polynomials would be equal when evaluated on \mathbf{r} with probability at most $\nu/|\mathbb{F}|$ by Schwartz-Zippel. Thus we can continue the proof assuming that $t \neq \sum_{\mathbf{b} \in \{0, 1\}^\mu} \tilde{F}(\mathbf{r}, \mathbf{b}) \cdot \tilde{u}(\mathbf{b})$.

Next, consider that for the extracted g_1 there are two possible cases: (i) $g_1(\mathbf{S}) = \tilde{F}(\mathbf{r}, \mathbf{S})$, and (ii) $g_1(\mathbf{S}) \neq \tilde{F}(\mathbf{r}, \mathbf{S})$.

If (i) occurs, then we can immediately build an adversary against the soundness of CP_{sc} .

If (ii) occurs, consider two subcases: (ii.a) $g_1(\boldsymbol{\sigma}) = \tilde{F}(\mathbf{r}, \boldsymbol{\sigma})$, and (ii.b) $g_1(\boldsymbol{\sigma}) \neq \tilde{F}(\mathbf{r}, \boldsymbol{\sigma})$. However, by Schwartz-Zippel (ii.a) occurs with negligible probability $\mu/|\mathbb{F}|$ over the choice of $\boldsymbol{\sigma}$. And if (ii.b)

occurs then it is possible to do a reduction to the soundness of CP_{poly} (since at least one of the claims $y^* = g_1(\sigma)$ or $y^* = \tilde{F}(r, \sigma)$ is false).

The zero-knowledge of CP_{lin} follows immediate from the zero-knowledge of CP_{sc} . \square

F.5 A CP-SNARK for Data-Parallel Computations

In this section we discuss how a CP-SNARK for relations R^{par} and R^{parjnt} , and for the commitment scheme PolyCom of [ZGK⁺17b] can be obtained by merging ideas from [ZGK⁺17b] and [WTs⁺18]. Such a merge of techniques was hinted possible in [WTs⁺18]. Here we give more details on how such a scheme looks like. The main motivation of studying such a scheme is that the commitment part of the proof (and similarly a factor of the verification time) is $O(\log |w|)$, instead of $O(\sqrt{|w|})$.

An Abstract Version of Hyrax. Hyrax [WTs⁺18] is a zero-knowledge proof, based on discrete log in the random oracle model that is based on the CMT protocol [CMT12]. Hyrax extends CMT, which is particularly suited for circuits composed of parallel identical basic blocks, by supporting non-determinism in zero-knowledge, as well as including other optimizations. Its basic structure as an interactive protocol: (i) the prover creates a commitment c_w to the witness w (a vector of field elements); (ii) the parties run a ZK variant of CMT (including optimizations from Giraffe++ [WJB⁺17]); (iii) the prover “links” together the outputs of steps (i) and (ii). For this, it must prove that the MLE \tilde{w} of the witness in c_w evaluated on a random point q_d is equal to another value y committed in ζ .

In Figure 17 we formalize this structure via a generic use of a commitment scheme for polynomials and a proof system for proving the correct evaluations of committed polynomials. For these two tools we use the syntax formalized in Appendix F. We call this scheme **Hyrax-Abstract**. It is clear from the security proof of [WTs⁺18] that one could rephrase their security statement so that **Hyrax-Abstract** has witness extended emulation if PolyCom is an extractable commitment and CP_{poly} is a NIZK argument of knowledge for polynomial evaluations.

Instantiating Hyrax-Abstract with PolyCom. We call **Hyrax – PolyCom** the instantiation of **Hyrax-Abstract** with the PolyCom commitment and CP_{poly} argument from [ZGK⁺17b] as described in Appendix E. This is essentially the only difference with the original Hyrax scheme that uses (an extension of) a matrix commitment of size $O(|w|^{1/l})$ and Bulletproof for proving polynomial evaluations with $O(|w|^{(l-1)/l})$ verification time. In **HyrPoly** there is instead a succinct commitment (of constant size) and a verification time, in step (iii), of $O(\log(|w|))$.

Using HyrPoly for Data-Parallel Computations. Hyrax, and in particular its Gir^{++} core protocol, is designed to work on arithmetic circuits of fan-in two, consisting of N identical sub-computations, each having d layers and width at most G . For this class of circuits, considering Hyrax’s cost analysis combined with the costs of PolyCom commitment and CP_{poly} , we have that in **HyrPoly**: the verifier runs in time $O(|x| + |y| + dG + \lambda d \log(NG))$ and proofs have length $O(\lambda d \log(NG))$.

It is easy to see that the relation $R^{\text{par}}((u_j)_{j \in [N]}) := \bigwedge_{j=1}^N R'(u_j)$ can be modeled with an arithmetic circuit C consisting of N copies of a circuit C' that outputs 0 on u_j iff $R'(u_j)$ holds.

If we instead consider a parallel relation with joint inputs, i.e., $R^{\text{parjnt}}(u) := \bigwedge_{j=1}^N R'(u'_j)$ where each u'_j is a subset of the entries of u , a corresponding circuit can be built by taking the parallel C as for R^{par} , and by adding one layer – called *redistribution layer* (RDL) in [WTs⁺18] – that appropriately duplicates and redistributes wires from the input layer to the input wires of each C'

<p><u>Hyrax-Abstract.Setup(1^λ) \rightarrow ck :</u></p> <p style="padding-left: 20px;">$ck \leftarrow \text{PolyCom.Setup}(1^\lambda)$</p> <p><u>Hyrax-Abstract.Prove(ek, u) \rightarrow π :</u></p> <p style="padding-left: 20px;">$(c_{\tilde{u}}, o_{\tilde{u}}) \leftarrow \text{ComPoly}(ck, \tilde{u})$</p> <p style="padding-left: 20px;">$(\pi_{\text{core}}, q_d, \zeta) \leftarrow \text{ZK-Gir}^{++}\text{Core}_{\mathcal{P}}(ek, u)$</p> <p style="padding-left: 20px;">$y \leftarrow \tilde{u}(q_d); (c_y, o_y) \leftarrow \text{ComVal}(ck, y)$</p> <p style="padding-left: 20px;">$\pi_{\text{eval}} \leftarrow \text{CP}_{\text{poly}}.\text{Prove}(ek, q_d, (c_{\tilde{u}}, c_y), (\tilde{u}, y), (o_{\tilde{u}}, o_y))$</p> <p style="padding-left: 20px;">$\pi_{\text{eq}} \leftarrow \text{NIPoK-Eq}_{\mathcal{P}}(c_y, \zeta)$</p> <p style="padding-left: 20px;">$\pi \leftarrow (c_{\tilde{u}}, \pi_{\text{core}}, c_y, \pi_{\text{eval}}, \pi_{\text{eq}})$</p>	<p><u>Hyrax-Abstract.KeyGen(ck) \rightarrow (ek, vk) :</u></p> <p style="padding-left: 20px;">$(ek, vk) \leftarrow \text{CP}_{\text{poly}}.\text{KeyGen}(ck)$</p> <p><u>Hyrax-Abstract.VerProof(vk, $c_{\tilde{u}}$, π_{core}, c_y, π_{eval}, π_{eq}) :</u></p> <p style="padding-left: 20px;">$(q_d, \zeta) \leftarrow \text{ZK-Gir}^{++}\text{Core}_{\mathcal{V}}(vk, \pi_{\text{core}})$</p> <p style="padding-left: 20px;">Run and test $\text{CheckCom}(vk, c_{\tilde{u}})$ and $\text{CheckCom}(vk, c_y)$</p> <p style="padding-left: 20px;">Run and test $\text{CP}_{\text{poly}}.\text{VerProof}(vk, q_d, c_{\tilde{u}}, c_y, \pi_{\text{eval}})$</p> <p style="padding-left: 20px;">Run and test $\text{NIPoK-Eq}_{\mathcal{V}}(\pi_{\text{eq}}, c_y, \zeta)$</p> <p style="padding-left: 20px;">Accept if all tests above pass</p>
---	---

Figure 17: Pseudocode for Hyrax-Abstract.

copy. In the case of using an RDL, the verifier of Hyrax, and also in our HyrPoly scheme, incurs an additional overhead in running time of the verifier $O(|x| + |u| + NG)$. Essentially, for this break of parallelism the verifier must pay a cost in the total width of the circuit.

For the sake of our experiments, we call HyrPoly-Par the HyrPoly scheme executed on fully parallel circuits (no RDL), and we call HyrPoly-RDL the version of Hyrax – PolyCom executed with circuits with an RDL.

G A CP-SNARK for Internal Products from Thaler’s Protocol

In this section we show how to modify the zk-vSQL protocol of [ZGK⁺17b] with a special class of circuits that simply consist of a tree of multiplications. The basic idea is to replace the CMT protocol over homomorphic commitment schemes proposed in [ZGK⁺17b] with an analogous version of the protocol proposed by Thaler [Tha13] for the specific case of trees of multiplications. The advantage of this encoding is to bring the prover runtime linear in the number of gates in the circuit.

We first explain some preliminaries and then present this construction.

G.1 CMT Protocol

The CMT protocol [CMT12] is a variant of the GKR protocol [GKR08] where the prover runs in time $\mathcal{O}(S \log S)$, where S is the size of the circuit. This protocol provides a proof that an element is the output of a circuit evaluated over a certain input. That is $y = C(\mathbf{x})$, where C is a circuit of depth d , \mathbf{x} are the wires of layer d and y is claimed to be the output wire of the first layer 0. In short, the prover reduces recursively a claim on layer i to another claim on layer $(i + 1)$, until he obtains a publicly verifiable claim on the input. In order to do that, both prover and verifier engage in a sum-check protocol for each layer, using one polynomial representing the values of the wires in layer i . Its multilinear extension links layer i (of size s_i) to layer $(i + 1)$ by a summation of wiring predicates as follows

$$\tilde{V}_i(\mathbf{q}) = \sum_{\substack{\mathbf{b} \in \{0,1\}^{s_i} \\ \mathbf{l}, \mathbf{r} \in \{0,1\}^{s_{i+1}}} } g_{\mathbf{q}}^{(i)}(\mathbf{b}, \mathbf{r}, \mathbf{l}) := \sum_{\substack{\mathbf{b} \in \{0,1\}^{s_i} \\ \mathbf{l}, \mathbf{r} \in \{0,1\}^{s_{i+1}}} } \tilde{\beta}_i(\mathbf{q}, \mathbf{b}) \cdot (\widetilde{\text{add}}_{i+1}(\mathbf{l}, \mathbf{r}, \mathbf{b}) \cdot (\tilde{V}_{i+1}(\mathbf{l}) + \tilde{V}_{i+1}(\mathbf{r})) + \widetilde{\text{mul}}_{i+1}(\mathbf{l}, \mathbf{r}, \mathbf{b}) \cdot \tilde{V}_{i+1}(\mathbf{l}) \cdot \tilde{V}_{i+1}(\mathbf{r}))$$

where \tilde{V}_i returns the value of one gate, $\tilde{\beta}_i(\mathbf{q}, \mathbf{b}) = \mathbf{q} \stackrel{?}{=} \mathbf{b}$ is a selector function, and $\widetilde{\text{opn}}_i(\mathbf{l}, \mathbf{r}, \mathbf{b})$ checks whether the value of gate \mathbf{b} at layer i is the result of an $\text{opn} \in \{\text{add}, \text{mul}\}$ addition or multiplication gate with \mathbf{l} and \mathbf{r} being its left and right inputs in the $(i + 1)$ -th layer.

The standard version of the protocol suggests that for each layer of the circuit the verifier has to check two claims. This results in $O(2^d)$ calls to the sum-check protocol. However, an ingenious technique shows how to use a single claim per layer using a line through both values. Then the verifier chooses one random point on which they perform a single sum-check invocation per layer, resulting in $O(d)$ calls.

G.2 Thaler’s Protocol for Trees of Multiplications

In [Tha13], Thaler proposes another variation of the CMT/GKR protocol [GKR08, CMT12] for some specific classes of circuits, allowing for a logarithmic factor reduction in the prover’s runtime. One of his protocols takes advantage of circuits where all gates perform the same operation, and whose wires are settled in a binary tree structure. He denotes these regular circuits by *trees of multiplications* or *additions*. This section only shows the notation of the former one due to its suitability for our construction of CP_{sfprm} . Nonetheless, moving to the addition case is straightforward.

The main difference that will be discussed here is a different polynomial for sum-check, as well as the notation of the wiring predicates. Thaler’s protocol assumes highly structured wiring in order to reduce the number of arguments of the predicates. Namely, given a gate at layer i with label $\mathbf{b} \in \{0, 1\}^{s_i}$, we assume its value is the result of a multiplication of gates of layer $(i + 1)$ with labels $(\mathbf{b}|0) \in \{0, 1\}^{s_{i+1}}$ and $(\mathbf{b}|1) \in \{0, 1\}^{s_{i+1}}$. This means, the number of inputs to the circuit is a power of two and each layer has half the size of its preceding one. On this basis, the resulting polynomial of each layer is much simpler as shown below:

$$\tilde{V}_i(\mathbf{q}) = \sum_{\mathbf{b} \in \{0, 1\}^{s_i}} g_{\mathbf{q}}^{(i)}(\mathbf{b}) = \sum_{\mathbf{b} \in \{0, 1\}^{s_i}} \tilde{\beta}_i(\mathbf{q}, \mathbf{b}) \cdot \tilde{V}_{i+1}(\mathbf{b}|0) \cdot \tilde{V}_{i+1}(\mathbf{b}|1)$$

This tweak, together with a series of precomputations of $\tilde{\beta}_i(\mathbf{q}, \mathbf{b})$ and $\tilde{V}_{i+1}(\mathbf{b})$ values allows to obtain a linear-time prover.

G.3 Adapting zk-vSQL to Thaler’s Protocol

Here we show how to change the CMT protocol over homomorphic commitments in [ZGK⁺17b, Construction 3] in order to work with circuits that are a tree of multiplication gates using Thaler’s representation [Tha13] to achieve faster prover runtime. From the point of view of security, this modification of [ZGK⁺17b, Construction 3] does not require any significant change; essentially a proof would be a rewrite of the one in [ZGK⁺17b]. The precise description of the protocol is however interesting and therefore we give it for completeness in Figure 18.

Let $C : \mathbb{F}^m \rightarrow \mathbb{F}$ be a depth- d binary tree of multiplications such that $C(\mathbf{y}) = z$ represents the operation $z = \prod_{i=1}^m y_i$ where m is a power of two, and let $\text{ck} \leftarrow \text{Setup}(1^\lambda)$ be a commitment key of a linearly homomorphic commitment scheme. The protocol in Figure 18 allows a prover \mathcal{P} to convince a verifier \mathcal{V} that $C(\mathbf{y}) = z$ with respect to \mathbf{y} and z committed in $\{c_{y_j}\}_{j \in \{1 \dots m\}}$ and c_z .

As in [ZGK⁺17b], let CP_{eq} (resp. CP_{prd}) be a zero-knowledge argument of knowledge for testing equality of two committed values (resp. the product relation between three commitments).

A Succinct Zero Knowledge Argument for R^{prd} . In Figure 19 we give the succinct version of the protocol TTM^{Com} presented in Figure 18. The protocol is almost identical to Construction 4 in

TTM^{Com} :

```

1 : Common input:  $\text{cvk}$  ;  $r_0 = 0$  ;  $c_0 := c_z$  ;  $(c_{y_j})_{j \in \{1 \dots m\}}$ 
2 :  $\mathcal{P}$  input:  $\text{ck}$  ;  $t_0 := z$  ;  $o_0 := o_z$  ;  $\mathbf{y}$  ;  $(o_{y_j})_{j \in \{1 \dots m\}}$ 
3 : for  $i = 0 \dots d - 1$  :
4 :   Run Step 1 of Construction 2 [ZGK+17b] (sum-check over homomorphic commitments)
5 :   on the claim  $t_i = V_i(\mathbf{r}_i) = \sum_{\mathbf{b} \in \{0,1\}^{s_i}} g_{\mathbf{r}_i}^{(i)}(\mathbf{b})$ 
6 :   At the end of Step 1,  $\mathcal{P}$  and  $\mathcal{V}$  hold  $\mathbf{r}'_i \in \mathbb{F}^{s_i}$  and commitment  $c'_i$  to  $t'_i = g_{\mathbf{r}'_i}^{(i)}(\mathbf{r}'_i)$ 
7 :    $\mathcal{P}$  : Claims that  $\text{VerCommit}(\text{cvk}, c'_i, t'_i, o'_i) = 1$ 
8 :    $\mathcal{P}$  :  $(c_R, o_R) \leftarrow \text{ComVal}(\text{ck}, v_R := \tilde{V}_{i+1}(\mathbf{r}'_i|0))$  ;  $(c_L, o_L) \leftarrow \text{ComVal}(\text{ck}, v_L := \tilde{V}_{i+1}(\mathbf{r}'_i|1))$ 
9 :    $\mathcal{P}$  :  $(c^*, o^*) \leftarrow \text{ComVal}(\text{ck}, v^* := v_L \cdot v_R)$ 
10 :    $\mathcal{P} \rightarrow \mathcal{V}$  :  $c_R, c_L, c^*$ 
11 :    $\mathcal{P}$  and  $\mathcal{V}$  run  $\text{CP}_{\text{prd}}(\text{ck}, (c_L, c_R, c^*); ((v_L, v_R, v^*), (o_L, o_R, o^*)))$ 
12 :    $\mathcal{P}$  :  $(c_i^*, o_i^*) \leftarrow \text{HomEval}(\text{cvk}, \tilde{\beta}_i(\mathbf{r}_i, \mathbf{r}'_i), c^*, o^*)$ 
13 :    $\mathcal{V}$  :  $(c_i^*, \cdot) \leftarrow \text{HomEval}(\text{cvk}, \tilde{\beta}_i(\mathbf{r}_i, \mathbf{r}'_i), c^*, \cdot)$ 
14 :    $\mathcal{P}$  and  $\mathcal{V}$  run  $\text{CP}_{\text{eq}}(\text{ck}, (c'_i, c_i^*); (t'_i, (o'_i, o_i^*)))$ 
15 :    $\mathcal{P}$  : Computes  $\{(c_{\ell_j}, o_{\ell_j}) \leftarrow \text{ComVal}(\text{ck}, \ell_j)\}_{j=0}^{s_{i+1}}$  where  $\ell(\rho) = \tilde{V}_{i+1}(\mathbf{r}'_i|\rho)$  for  $\rho \in \mathbb{F}$  and  $\{\ell_j\}_{j=0}^{s_{i+1}}$  its coefficients
16 :    $\mathcal{P} \rightarrow \mathcal{V}$  :  $\{c_{\ell_j}\}_{j \in \{0 \dots s_{i+1}\}}$ 
17 :    $\mathcal{P}$  :  $c_{\ell(0)} \leftarrow c_{\ell_0}$  ;  $(c_{\ell(1)}, o_{\ell(1)}) \leftarrow \text{HomEval}(\text{cvk}, (1, \dots, 1), \{c_{\ell_j}, o_{\ell_j}\}_{j \in [0, s_{i+1}]})$ 
18 :    $\mathcal{V}$  :  $c_{\ell(0)} \leftarrow c_{\ell_0}$  ;  $(c_{\ell(1)}, \cdot) \leftarrow \text{HomEval}(\text{cvk}, (1, \dots, 1), \{c_{\ell_j}\}_{j=0}^{s_{i+1}}, \cdot)$ 
19 :    $\mathcal{P}$  and  $\mathcal{V}$  run  $\text{CP}_{\text{eq}}(\text{ck}, (c_R, c_{\ell(0)}); (v_R, o_R, o_{\ell(0)}))$  ;  $\text{CP}_{\text{eq}}(\text{ck}, (c_L, c_{\ell(1)}); (v_L, o_L, o_{\ell(1)}))$ 
20 :    $\mathcal{V} \rightarrow \mathcal{P}$  :  $r''_i \leftarrow_{\$} \mathbb{F}$  and define  $\mathbf{r}_{i+1} \leftarrow (\mathbf{r}'_i | r''_i)$ 
21 :    $\mathcal{V}$  :  $(c_{i+1}, \cdot) \leftarrow \text{HomEval}(\text{ck}, (1, r''_i, \dots, r''_i^{s_{i+1}}), \{c_{\ell_j}\}_{j \in [0, s_{i+1}]}, \cdot)$ 
22 :    $\mathcal{P}$  :  $\mathbf{r}_{i+1} \leftarrow (\mathbf{r}'_i | r''_i)$  ;  $t_{i+1} \leftarrow \tilde{V}_{i+1}(\mathbf{r}_{i+1})$  ;  $(c_{i+1}, o_{i+1}) \leftarrow \text{HomEval}(\text{ck}, (1, r''_i, \dots, r''_i^{s_{i+1}}), \{c_{\ell_j}, o_{\ell_j}\}_{j \in [0, s_{i+1}]})$ 
23 : endfor
24 :  $\mathcal{P} \rightarrow \mathcal{V}$  :  $\mathbf{y}$  ;  $(o_{y_j})_{j \in \{1 \dots m\}}$  ;  $o_0$ 
25 :  $\mathcal{V}$  :  $\{\text{VerCommit}(\text{cvk}, c_{y_j}, y_j, o_{y_j})\}_{j=1}^m$  ;  $\text{VerCommit}(\text{cvk}, c_0, t_0, o_0)$ 
26 :  $\mathcal{V}$  :  $(c_y^*, o_y^*) \leftarrow \text{ComVal}(\text{ck}, \tilde{V}_y(\mathbf{r}_d))$  where  $\text{MLE}(V_y(j) = y_j) = \tilde{V}_y$ 
27 :  $\mathcal{V} \rightarrow \mathcal{P}$  :  $o_y^*$ 
28 :  $\mathcal{P}$  and  $\mathcal{V}$  run  $\text{CP}_{\text{eq}}(\text{ck}, (c_y^*, c_d); (\tilde{V}_y(\mathbf{r}_d), o_y^*, o_d))$ 

```

Figure 18: Thaler’s tree of multiplications over homomorphic commitment schemes. Main differences from [ZGK⁺17b, Construction 3] in blue

[ZGK⁺17b] except for a few simplifications due to the fact that in our case the input and output of the circuit are assumed to be already committed and these commitments are known to the verifier, and that all the input is committed (i.e., there is no public input). Basically, the idea is that prover and verifier run the TTM^{Com} protocol until they get to the end of the last round (line 23). Then the last lines of TTM^{Com}, in which the prover opens the commitments to input and output and the verifier gets convinced that c_d opens to $\tilde{y}(\mathbf{r}_d)$, are replaced with a step that does the same: the prover uses CP_{poly} to prove that c_d opens to $\tilde{y}(\mathbf{r}_d)$ with respect to the commitment c_y . For the

Preprocessing: generate the commitment key
 $(\text{ck}, \text{cvk}) \leftarrow \text{PolyCom.Setup}(1^\lambda)$ for m -variate multilinear polynomials.
 $(\text{ek}, \text{vk}) \leftarrow \text{CP}_{\text{poly}}.\text{Setup}(\text{ck})$

Evaluation: on common input (c_y, c_z) ; \mathcal{P} input (z, \tilde{y}, o_y)
 $\mathcal{V} : \text{CheckCom}(\text{vk}, c_y) \wedge \text{CheckCom}(\text{vk}, c_z)$
 $\mathcal{P}, \mathcal{V} : \text{Execute TTM}^{\text{Com}}$ until line **23** :
Both hold \mathbf{r}_d, c_d ; \mathcal{P} holds an opening o_d of $\tilde{V}_d(\mathbf{r}_d) = \tilde{y}(\mathbf{r}_d)$
 $\mathcal{P} \rightarrow \mathcal{V} : \pi_y \leftarrow \text{CP}_{\text{poly}}.\text{Prove}(\text{ek}, \mathbf{r}_d, (c_y, c_d), (\tilde{y}, \tilde{y}(\mathbf{r}_d)), (o_y, o_d))$
 $\mathcal{V} : \text{CP}_{\text{poly}}.\text{VerProof}(\text{vk}, \mathbf{r}_d, c_{\tilde{V}_d}, c_d, \pi_d)$

Figure 19: Succinct zero-knowledge argument for TTM^{Com}

polynomial commitments and the proof system for their evaluations we use our notation of Section F.

EFFICIENCY. Our CP_{ipd} is a succinct zero-knowledge argument for R^{prd} that uses a variant of Thaler’s protocol for trees of multiplications [Tha13] over homomorphic commitment schemes [ZGK⁺17b, Construction 4]. Here, we compute a proof of the product of the elements of a vector $\mathbf{y} \in \mathbb{F}^m$ where $m = 2^\mu$. This is encoded as a depth- μ circuit C of size $S = (m - 1)$ with m inputs and 1 output element. By the regularity of the circuit, here the number of gates of each layer is double the size of the previous one $S_i = 2S_{i-1}$, meaning that $\log S_{i+1} = s_{i+1} = s_i + 1$ where s_i is the number of variables of the target polynomial at layer i . Since the polynomial used inside CP_{sc} is a product of three polynomials, each of its s_i variables will be at most degree 3. Considering that the output layer has an only gate, then the sum for the whole circuit of the number of variables of all target polynomials can be computed as $\sum_{i=0}^{\mu-1} s_i = \sum_{i=0}^{\mu-1} 2^i = \frac{2^\mu - 1}{1} = \frac{2^\mu - 1}{1}$. The SuccinctZK – TTM construction shows that the proof consists of $(\frac{\mu^2 - \mu}{2} + 3\mu)$ CP_{eq} proofs, μ CP_{prd} proofs, 1 CP_{poly} proofs and $\frac{5\mu^2 - \mu}{2}$ commitments. The prover requires linear time in the circuit size and the verifier runs in quadratic time in the circuit depth. Its crs has length $(2 \cdot 2^\mu + 3)\mathbb{G}_1 + (\mu + 3)\mathbb{G}_2$. We refer the reader to Table 2 for a summary.

H Commit and Prove SNARKs from existing schemes

In this section we give details supporting our claims of Section 3.4.

Background on Quadratic Arithmetic Programs. Since several of the SNARKs considered in this section rely on quadratic arithmetic programs [GGPR13] here we recall this notion.

Definition H.1 (QAP [GGPR13]). A Quadratic Arithmetic Program (QAP) $\mathcal{Q} = (\mathcal{A}, \mathcal{B}, \mathcal{C}, t(Z))$ of size m and degree d over a finite field \mathbb{F} is defined by three sets of polynomials $\mathcal{A} := \{a_i(Z)\}_{i=0}^m$, $\mathcal{B} := \{b_i(Z)\}_{i=0}^m$, $\mathcal{C} := \{c_i(Z)\}_{i=0}^m$ of degree $\leq d - 1$, and a target degree- d polynomial $t(Z)$. Given \mathcal{Q} we

define a relation $R_{\mathcal{Q}}$ over pairs $(\mathbf{x}, \mathbf{w}) \in \mathbb{F}^n \times \mathbb{F}^{m-n}$ that holds iff there exists a polynomial $h(X)$ (of degree at most $d - 2$) such that:

$$\left(\sum_{k=0}^m y_k \cdot a_k(Z) \right) \cdot \left(\sum_{k=0}^m y_k \cdot b_k(Z) \right) = \left(\sum_{k=0}^m y_k \cdot c_k(Z) \right) + h(Z)t(Z) \quad (9)$$

where $y_0 = 1$, $y_k = x_k$ for all $k = 1$ to n , and $y_k = w_{k-n}$ for $k = n + 1$ to m .

H.1 “Adaptive Pinocchio” [Vee17]

The Adaptive Pinocchio scheme proposed in [Vee17] yields a CP-SNARK for QAP relations $R_{\mathcal{Q}}(\mathbf{x}, \mathbf{u}, \omega)$. First, note that [Vee17] already presents the scheme as a commit-and-prove SNARK for QAP relations $R_{\mathcal{Q}}(\mathbf{u}_1, \dots, \mathbf{u}_\ell, \omega)$, and for an extractable trapdoor commitment scheme, which is the one proposed by Groth in [Gro10]. Second, observe that the commitment key consists of two vectors $\mathbf{S} := [1, s, s^2, \dots, s^d]_1$, $\mathbf{S}' := [\alpha, \alpha s, \alpha s^2, \dots, \alpha s^d]_2$, for random $s, \alpha \leftarrow_{\$} \mathbb{Z}_q$, and the commitment to \mathbf{u}_j is a pair $(C, C') = (r, \mathbf{u}_j^\top) \cdot (\mathbf{S}, \mathbf{S}')$. To see how this implies a CP-SNARK for $R_{\mathcal{Q}}(\mathbf{x}, \mathbf{u}, \omega)$, consider $\ell = 2$ so that the first input \mathbf{u}_1 is used for the public input \mathbf{x} (the corresponding commitment can be a dummy one) and the second one for the actual committed value \mathbf{u} . Also, to fit our syntax let C be the actual commitment whereas C' is part of the proof.

H.2 Lipmaa’s Hadamard Product Argument [Lip16]

The product argument proposed by Lipmaa in [Lip16] is a commit-and-prove SNARK for the Hadamard product relation $R^{\text{had}}(\mathbf{a}, \mathbf{b}, \mathbf{c})$. In this case the commitment key ck are two vectors $\mathbf{S} := [Z(\chi), \ell_1(\chi), \dots, \ell_m(\chi)]_1^\top$ and $\mathbf{S}' := [\gamma Z(\chi), \gamma \ell_1(\chi), \dots, \gamma \ell_m(\chi)]_2^\top$, for random $\chi, \gamma \leftarrow_{\$} \mathbb{Z}_q$, where, for m a power of two and ω the m -th root of unity modulo q , $Z(X) = \prod_{i=1}^m (X - \omega^{i-1})$ and $\ell_i(X)$ is the i -th Lagrange basis polynomial (such distribution of ck guarantees binding under the m -PDL assumption [Lip12, Lip16]). A commitment to \mathbf{a} is a pair $(A_1, A_2) = (r_a, \mathbf{a}^\top) \cdot (\mathbf{S}, \mathbf{S}')$ (and similarly to \mathbf{b}, \mathbf{c}). As in the previous section, to fit our CP-SNARK syntax we can think of A_1, B_1, C_1 as the actual commitments and let their “knowledge components” as part of the proof.

H.3 zk-vSQL [ZGK⁺17b]

The zk-vSQL protocol [ZGK⁺17b] is a CP-SNARK for relations $R((u_j)_{j \in [\ell]})^{27}$ where R is an arithmetic circuit (that we assume to output some constant, e.g., 0, on acceptance), and for the commitment scheme PolyCom introduced in [ZGK⁺17b] and recalled in Appendix E.²⁸ The commit and prove capability is immediate by the construction and security of [ZGK⁺17b]. In what follows we observe that their commitments can also be seen as a variant of extended Pedersen commitment. This observation is crucial to see that we can apply our lifting transformation using our CP_{link} scheme to zk-vSQL. Let us recall that for an input $\mathbf{u} \in \mathbb{Z}_q^m$ (for some $m = 2^\mu$), its commitment is $\text{Compoly}(\text{ck}, \tilde{u})$ where \tilde{u} is the multilinear extension of \mathbf{u} (cf. Section 5.1 about multilinear extensions). In particular, such MLE is the following μ -variate multilinear polynomial:

²⁷ Precisely, although the scheme in [ZGK⁺17b] is described with a single u , the same technique used in its predecessor [ZGK⁺17a] trivially allows to let it work with multiple commitments.

²⁸ Here we are considering the non-interactive version in the random oracle model obtained after applying the Fiat-Shamir transform.

$$\tilde{u}(X_1, \dots, X_\mu) = \sum_{i=0}^{m-1} \chi_i(X_1, \dots, X_\mu) \cdot u_{i+1}$$

Since c returned by $\text{ComPoly}(\text{ck}, \tilde{u}, \rho)$ is defined as $[\tilde{u}(s_1, \dots, s_\mu) + \rho s_{\mu+1}]_1$ and the common reference string includes the monomials $[\prod_{j \in W} s_j]_1$ for all possible subsets of indices W needed to evaluate such a \tilde{u} , c can also be seen as a Pedersen commitment $c = (\rho, u^\top) \cdot [s_{\mu+1}, \chi_0(s_1, \dots, s_\mu), \dots, \chi_{m-1}(s_1, \dots, s_\mu)]_1^\top = (\rho, u^\top) \cdot \text{ck}$, where the elements $[\chi_i(s_1, \dots, s_\mu)]_1$ can be publicly computed from the existing key. Note that this commitment is binding. This can be seen via a simple reduction to the soundness of the polynomial delegation protocol in [ZGK⁺17b]. The idea is that from an adversary that opens the commitment to two different polynomials \tilde{u}_1, \tilde{u}_2 one can sample a random t such that with overwhelming probability $y_1 = \tilde{u}_1(t) \neq \tilde{u}_2(t) = y_2$, honestly compute a proof for the evaluation of $y_1 = \tilde{u}_1(t)$ and then claim this is an evaluation for $\tilde{u}_2(t)$.

H.4 Geppetto [CFH⁺15]

The Geppetto scheme [CFH⁺15] yields a cc-SNARK for QAP relations $R_{\mathcal{Q}}(\mathbf{x}, \boldsymbol{\omega})$ where $\mathbf{x} \in \mathbb{Z}_q^n$ and $\boldsymbol{\omega} = (\mathbf{u}, \boldsymbol{\omega})$ with $\mathbf{u} \in \mathbb{Z}_q^{n'}$, $\boldsymbol{\omega} \in \mathbb{Z}_q^{m-n-n'}$ for some integers n, n' . We recall that Geppetto is a SNARK for MultiQAP relations. A polynomial MultiQAP is a tuple $\mathcal{MQ} = (\ell, \mathcal{J}, \mathcal{A}, \mathcal{B}, \mathcal{C}, t(Z))$ such that $(\mathcal{A}, \mathcal{B}, \mathcal{C}, t(Z))$ is a QAP, and $\mathcal{J} = \{I_0, \dots, I_{\ell-1}\}$ is a partition of $[m]$. Let $R_{\mathcal{MQ}}$ denote the relation corresponding to \mathcal{MQ} . To model $R_{\mathcal{Q}}(\mathbf{x}, \mathbf{u}, \boldsymbol{\omega})$ we consider a MultiQAP where $\ell = 3$ and where the partition \mathcal{J} consists of $I_0 = [n]$, $I_1 = \{n+1, \dots, n+n'\}$ and $I_2 = \{n+n'+1, \dots, m\}$ such that I_0 and I_1 are in the binding subset S .

To see how Geppetto yields a cc-SNARK for such family of relations, we consider the following straightforward modification:

$\text{ccGep.KeyGen}(R_{\mathcal{Q}}) \rightarrow (\text{ck}, \text{ek}, \text{vk})$: run $(EK, VK) \leftarrow \text{Geppetto.KeyGen}(R_{\mathcal{MQ}})$; set $\text{ek} = EK$, $\text{vk} = VK$ and let ck be subset of EK consisting of $[r_y t(s), r_c c_{n+1}(s), \dots, r_c c_{n+n'}(s)]_1^\top \in \mathbb{G}_1^{n'+1}$.

$\text{ccGep.VerCommit}(\text{ck}, c, \mathbf{u}, o) \rightarrow b$: output 1 iff $(o, \mathbf{u}^\top) \cdot \text{ck} = c$.

$\text{ccGep.Prove}(\text{ek}, \mathbf{x}, \mathbf{u}, \boldsymbol{\omega}) \rightarrow (c, \pi; o)$:

Compute commitments:

$C_0 \leftarrow \text{Geppetto.Commit}(EK_0, \mathbf{x}, 0)$,²⁹ $C_1 \leftarrow \text{Geppetto.Commit}(EK_1, \mathbf{u}, o_1)$, $C_2 \leftarrow \text{Geppetto.Commit}(EK_2, \boldsymbol{\omega}, o_2)$

Compute the proof $\pi' \leftarrow \text{Geppetto.Prove}(EK, (\mathbf{x}, \mathbf{u}, \boldsymbol{\omega}), (o, o_1, o_2))$.

Parse C_1 as $(C_{1,1}, C_{1,\alpha}, C_{1,\beta}) \in \mathbb{G}_1^3$.

Output $c = C_{1,1}$, $\pi = (C_{1,\alpha}, C_{1,\beta}, C_2, \pi')$, and $o = o_1$.

$\text{ccII.VerProof}(\text{vk}, \mathbf{x}, c, \pi) \rightarrow b$: recompute $C_0 \leftarrow \text{Geppetto.Commit}(EK_0, \mathbf{x}, 0)$; reconstruct $C_1 \leftarrow (c, C_{1,\alpha}, C_{1,\beta})$; check $\text{Geppetto.Verify}(VK_j, C_j)_{j=1,2}$; check $\text{Geppetto.Verify}(VK, C_0, C_1, C_2, \pi')$.

We claim that assuming Geppetto is a commit-and-prove SNARK for MultiQAPs (according to the commit-and-prove definition in [CFH⁺15]), then the scheme ccGep described above is a cc-SNARK for QAP relations $R_{\mathcal{Q}}(\mathbf{x}, \mathbf{u}, \boldsymbol{\omega})$.

The correctness of ccGep immediately follows from the one of Geppetto, and the same holds for knowledge soundness. Indeed, notice that the knowledge soundness satisfied by Geppetto provides extractability of the commitment's openings. The perfect zero-knowledge of ccGep follow from the zero-knowledge of Geppetto and the perfect hiding of its commitments. Finally, we observe that by

²⁹ Setting randomness 0 here is essentially a trick to let this commitment correspond to the public input of the relation.

Def. 10 in [CFH⁺15] the polynomials $\{c_k(x)\}_{k \in I_1}$ are linearly independent; thus for a random s , the vector $[r_c t(s), r_c c_{n+1}(s), \dots, r_c c_{n+n'}(s)]_1$ defines a Pedersen commitment key whose distribution guarantees the binding property under the d -SDH assumption.

H.5 cc-SNARKs based on Groth’s SNARK

In this section we show that the SNARK of [Gro16] is a *weak* cc-SNARK, and then that it can be modified to obtain an efficient cc-SNARK (with binding commitments). Below we start by giving a background on non-interactive linear proofs, that are instrumental for presenting the scheme.

Split Non-Interactive Linear Proofs of Degree 2. This notion, dubbed NILP for brevity, was introduced by Groth [Gro16] as a refinement of the linear interactive proofs defined in [BCI⁺13]. A NILP is a triple of algorithms (**LinSetup**, **ProofMatrix**, **Test**) working as follows. **LinSetup** takes in a relation R (e.g., a QAP) and outputs two vectors $\sigma_1 \in \mathbb{F}^{\mu_1}, \sigma_2 \in \mathbb{F}^{\mu_2}$. **ProofMatrix** on input a relation R and a pair (x, w) outputs two matrices $(\Pi_1, \Pi_2) \in \mathbb{F}^{k_1 \times \mu_1} \times \mathbb{F}^{k_2 \times \mu_2}$ so that a proof (π_1, π_2) is computed as $(\Pi_1 \cdot \sigma_1, \Pi_2 \cdot \sigma_2)$. **Test** on input a relation R and a statement x outputs a collection of matrices $T_1, \dots, T_\eta \in \mathbb{F}^{(\mu_1+k_1) \times (\mu_2+k_2)}$ such that a proof (π_1, π_2) is accepted iff $(\sigma_1^\top, \pi_1^\top) \cdot T_i \cdot (\sigma_2^\top, \pi_2^\top) = 0$ for all $i = 1$ to η . A NILP is required to satisfy completeness, statistical knowledge soundness and zero-knowledge. Informally, completeness says that honestly computed proofs for true statements are accepted. Knowledge soundness says that there must exist an extractor algorithm that on input R, x and a prover strategy (Π_1, Π_2) outputs a witness w such that the probability that $(\Pi_1 \cdot \sigma_1, \Pi_2 \cdot \sigma_2)$ is accepted while $R(x, w) = 0$ is negligible (over the random choices of **LinSetup**). Finally, (perfect) zero-knowledge states requires to show a simulator that with knowledge of $(\sigma_1, \sigma_2, R, x)$ outputs proofs (π_1, π_2) that have the same distribution as honestly generated ones.

Groth’s zkSNARK [Gro16] is a *weak* cc-SNARK for QAP relations $R_Q(\mathbf{u})$. First, we recall the scheme from [Gro16]: this scheme is obtained by instantiating the generic pairing-based construction of Figure 20 with the Non-Interactive Linear Proof (NILP) in Figure 21.

Recall that for our claim we only consider the case of QAP relations where \mathbf{x} is void and the witness is $\mathbf{w} = \mathbf{u}$ (i.e., the commitment is to the entire witness). This is enough to instantiate our compiler of Section 3.5. To see why this scheme is a weak cc-SNARK for QAP relations $R_Q(\mathbf{u})$ we make the following observations.

First, let the commitment c to \mathbf{u} be the value $[A]_1 = r[\delta]_1 + \sum_{k=0}^m u_k \cdot [a_k(\tau)]_1 + [\alpha]_1$; this means that ck is $[\delta, \{a_k(\tau)\}, \alpha]_1$ where α, δ, τ are random. Second, for knowledge soundness we observe that from the existing security proof we can also extract the opening r of $[A]_1$. What is left to argue is the binding of such commitment. Since the $\{a_k(Z)\}_k$ polynomials are not necessarily linearly independent (see, e.g., [Par15]) the commitment key ck does not guarantee binding. However, we

$\text{KeyGen}(R_Q)$	$\text{Prove}(\sigma, R_Q, x, w)$	$\text{VerProof}(\sigma, x, \pi)$
$(\sigma_1, \sigma_2) \leftarrow \text{LinSetup}(R_Q)$ return $\sigma := ([\sigma_1]_1, [\sigma_2]_2)$	$(\Pi_1, \Pi_2) \leftarrow \text{ProofMatrix}(R_Q, x, w)$ $[\pi_1]_1 \leftarrow \Pi_1 \cdot [\sigma_1]_1$; $[\pi_2]_2 \leftarrow \Pi_2 \cdot [\sigma_2]_2$ return $\pi = ([\pi_1]_1, [\pi_2]_2)$	$T_1, \dots, T_\eta \leftarrow \text{Test}(R_Q, x)$ return 1 iff $\forall i \in [\eta]$: $[0]_T = \begin{pmatrix} [\sigma_1]_1 \\ [\pi_1]_1 \end{pmatrix} \cdot T_i \cdot \begin{pmatrix} [\sigma_2]_2 \\ [\pi_2]_2 \end{pmatrix}$

Figure 20: Groth’s generic SNARK in asymmetric groups from a split NILP.

$\text{LinSetup}(R_Q) \rightarrow (\sigma_1, \sigma_2)$
$\alpha, \beta, \gamma, \delta, \tau \leftarrow_{\$} \mathbb{F}^*$
$\sigma_1 := \left(1, \alpha, \delta, \{\tau^i\}_{i=0}^{d-1}, \left\{ \frac{1}{\gamma}(\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau)) \right\}_{i=0}^n, \left\{ \frac{1}{\delta}(\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau)) \right\}_{i=n+1}^m, \left\{ \frac{1}{\delta} \tau^i t(\tau) \right\}_{i=0}^{d-2} \right)$
$\sigma_2 := \left(1, \beta, \gamma, \delta, \{\tau^i\}_{i=0}^{d-1} \right)$
$\text{ProofMatrix}(R_Q, \mathbf{x}, \mathbf{w}) \rightarrow (\Pi_1, \Pi_2)$
Compute $h(Z)$ and define \mathbf{y} from (\mathbf{x}, \mathbf{w}) as in (9) ; $r, s \leftarrow_{\$} \mathbb{F}$
Let $\Pi_1 \in \mathbb{F}^{3 \times (m+2d+3)}, \Pi_2 \in \mathbb{F}^{1 \times (d+4)}$ s.t. $(A, C)^\top = \Pi_1 \cdot \sigma_1, B = \Pi_2 \cdot \sigma_2$ with
$A := \alpha + \sum_{k=0}^m y_k \cdot a_k(\tau) + r\delta; \quad B := \beta + \sum_{k=0}^m y_k \cdot b_k(\tau) + s\delta$
$C := \sum_{k=n+1}^m y_k \cdot \frac{\beta a_k(\tau) + \alpha b_k(\tau) + c_k(\tau)}{\delta} + \sum_{i=0}^{d-2} h_i \frac{\tau^i t(\tau)}{\delta} + As + Br - rs\delta$
$\text{Test}(R_Q, \mathbf{x}) \rightarrow T$
Let $T \in \mathbb{F}^{(m+2d+5) \times (d+5)}$ encode the quadratic test: $A \cdot B = \alpha \cdot \beta + C \cdot \delta + \gamma \left(\sum_{k=0}^n \frac{x_k}{\gamma} (\beta a_k(\tau) + \alpha b_k(\tau) + c_k(\tau)) \right)$

Figure 21: Groth’s NILP for a QAP relation $R_Q(\mathbf{x}, \mathbf{w})$.

can show as follows that the scheme satisfies *weak binding*. In a nutshell, this means that it is computationally infeasible to open $[A]_1$ to two different witnesses \mathbf{u} and \mathbf{u}' with $R_Q(\mathbf{u}) \neq R_Q(\mathbf{u}')$.

Notice that from the two different valid openings (\mathbf{u}, r) and (\mathbf{u}', r') of $[A]_1$ we can easily rule out two cases. The first case is the one where $r \neq r'$: this can be immediately reduced to finding the discrete log δ . The second case is the one when $r = r'$ and $\sum_k (u_k - u'_k) a_k(Z)$ is a nonzero polynomial: this can be reduced to finding the discrete log τ (which is known as PDL problem [Lip12]), as τ can be computed by factoring this polynomial. Therefore we are left with the case when $\sum_k (u_k - u'_k) a_k(Z)$ is the zero polynomial, yet $\mathbf{u} \neq \mathbf{u}'$. We argue that it cannot be that $R_Q(\mathbf{u}) \neq R_Q(\mathbf{u}')$. Indeed, the existing proof [Gro16][Theorem 1] shows that equalities $A = \alpha + r\delta + \sum_{k=0}^m C_k \cdot a_k(\tau)$ and $B = \beta + s\delta + \sum_{k=0}^m C_k \cdot b_k(\tau)$ hold, where $\{C_k\}_{k=0}^m$ are the same coefficients of the term $\sum_{k=0}^m C_k \cdot \frac{\alpha b_k(\tau) + \beta a_k(\tau) + c_k(\tau)}{\delta}$ in C . Therefore, if the commitment A opens to \mathbf{u}' then it must be the case that $C_k = u'_k$, but in this case the QAP would be satisfied (i.e., $R_Q(\mathbf{u}') = 1$) contradicting that \mathbf{u}' is an invalid witness for R_Q .

A new cc-SNARK for QAP relations $R_Q(\mathbf{u}, \omega)$. Here we show how we can modify the zkSNARK of [Gro16] in order to obtain a cc-SNARK for proving the satisfiability of QAP relations of the form $R_Q(\mathbf{u}, \omega)$, that is a scheme where there is a binding commitment to a portion, \mathbf{u} , of the witness and where the public input is void.³⁰

In our construction we consider an *augmented QAP* (in the sense of [BCTV14]), which is a QAP as in Definition H.1 with the additional property that the polynomials $a_k(X)$ for $k = 0$ to n are *linearly independent*.

³⁰ It is possible to extend this construction to support non-empty public inputs. For simplicity we keep public input void as our interest is to use this scheme in order to obtain a full fledged CP-SNARK though our compiler of Section 3.5 together with the CP_{link} scheme. In such a case, CP_{link} can take care of showing that a given prefix of \mathbf{u} is the public input.

$\text{LinSetup}(R_{\mathcal{Q}}) \rightarrow (\boldsymbol{\sigma}_1, \boldsymbol{\sigma}_2)$
$\alpha, \beta, \gamma, \delta, \eta, \tau \leftarrow_{\$} \mathbb{F}^*$
$\boldsymbol{\sigma}_1 := \left(1, \alpha, \delta, \{\tau^i\}_{i=0}^{d-1}, \left\{ \frac{\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau)}{\gamma} \right\}_{i=0}^n, \frac{\eta}{\gamma}, \left\{ \frac{\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau)}{\delta} \right\}_{i=n+1}^m, \left\{ \frac{1}{\delta} \tau^i t(\tau) \right\}_{i=0}^{d-2}, \frac{\eta}{\delta} \right)$
$\boldsymbol{\sigma}_2 := \left(1, \beta, \gamma, \delta, \{\tau^i\}_{i=0}^{d-1} \right)$
$\text{ProofMatrix}(R_{\mathcal{Q}}, \boldsymbol{w}) \rightarrow (\Pi_1, \Pi_2)$
Let $\boldsymbol{w} := (\boldsymbol{u}, \boldsymbol{\omega})$. Compute $h(Z)$ as in (9) ; $r, s, v \leftarrow_{\$} \mathbb{F}$
Let $\Pi_1 \in \mathbb{F}^{3 \times (m+2d+5)}$, $\Pi_2 \in \mathbb{F}^{1 \times (d+4)}$ s.t. $(A, C, D)^\top = \Pi_1 \cdot \boldsymbol{\sigma}_1, B = \Pi_2 \cdot \boldsymbol{\sigma}_2$ and
$A := \alpha + \sum_{k=0}^m w_k \cdot a_k(\tau) + r\delta \quad ; \quad B := \beta + \sum_{k=0}^m w_k \cdot b_k(\tau) + s\delta \quad ; \quad D := \sum_{k=0}^n w_k \cdot \frac{1}{\gamma} (\beta a_k(\tau) + \alpha b_k(\tau) + c_k(\tau)) + v \frac{\eta}{\gamma}$
$C := \sum_{k=n+1}^m w_k \cdot \frac{\beta a_k(\tau) + \alpha b_k(\tau) + c_k(\tau)}{\delta} - v \frac{\eta}{\delta} + \sum_{i=0}^{d-2} h_i \frac{\tau^i t(\tau)}{\delta} + As + Br - rs\delta$
$\text{Test}(R_{\mathcal{Q}}) \rightarrow T$
Define $T \in \mathbb{F}^{(m+2d+7) \times (d+5)}$ encoding the following quadratic test: $A \cdot B = \alpha \cdot \beta + C \cdot \delta + D \cdot \gamma$

Figure 22: Our NILP for an augmented QAP relation $R_{\mathcal{Q}}(\boldsymbol{u}, \boldsymbol{\omega})$, to be used to obtain ccGro16.

Our new cc-SNARK, called **ccGro16**, is the scheme obtained by instantiating the generic SNARK construction of [Gro16] recalled in Figure 20 with the NILP that we describe in Figure 22. To match the cc-SNARK syntax we let the commitment be the proof element $[D]_1$. Clearly, $[D]_1$ can be seen as a Pedersen commitment for the key $\text{ck} = [\frac{\eta}{\gamma}, \{\frac{1}{\gamma}(\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau))\}_{i=0}^n]_1$. By the linear independence of the $a_i(Z)$ polynomials the binding of this commitment can be reduced to the PDL assumption. Correctness and knowledge soundness follow from the proof of the generic construction in [Gro16], assuming that the construction in Figure 22 is a NILP. We show that this is the case in the following theorem.

Theorem H.1. *The construction in Figure 22 is a NILP with perfect completeness, perfect zero-knowledge and statistical knowledge soundness against affine provers.*

Proof Perfect completeness is easy to verify. For perfect zero-knowledge, we define the simulator that samples $A, B, D \leftarrow_{\$} \mathbb{F}$ at random and then finds C so that the verification test is satisfied. This shows that real and simulated proofs are identically distributed.

For knowledge soundness, let $(\Pi_1, \Pi_2) \in \mathbb{F}^{3 \times (m+2d+5)} \times \mathbb{F}^{1 \times (d+4)}$ be an affine prover strategy. First, we define the extractor as the one that returns as witness component $\boldsymbol{\omega}$ the entries of the second row of Π_1 corresponding to the terms $\{\frac{1}{\delta}(\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau))\}_{i=n+1}^m$, and as witness component \boldsymbol{u} the entries in the third row of Π_1 corresponding to the terms $\{\frac{1}{\gamma}(\beta a_i(\tau) + \alpha b_i(\tau) + c_i(\tau))\}_{i=1}^n$. For the cc-SNARK knowledge soundness, we need to additionally extract the commitment opening that can be taken from the entry in the third row of Π_1 corresponding to the term $\frac{\eta}{\gamma}$.

Once defined the extractor, we need to show that the probability that the proof verifies and the relation $R_{\mathcal{Q}}(\boldsymbol{u}, \boldsymbol{\omega})$ does not hold is negligible. This proof is essentially identical to the one used for the NILP of [Gro16]. \square