

Proof-Carrying Data without Succinct Arguments

Benedikt Bünz

benedikt@cs.stanford.edu

Stanford University

Alessandro Chiesa

alexch@berkeley.edu

UC Berkeley

William Lin

will.lin@berkeley.edu

UC Berkeley

Pratyush Mishra

pratyush@berkeley.edu

UC Berkeley

Nicholas Spooner

nspooner@bu.edu

Boston University

December 31, 2020

Abstract

Proof-carrying data (PCD) is a powerful cryptographic primitive that enables mutually distrustful parties to perform distributed computations that run indefinitely. Known approaches to construct PCD are based on succinct non-interactive arguments of knowledge (SNARKs) that have a succinct verifier or a succinct accumulation scheme for their proofs.

In this paper we show how to obtain PCD without relying on SNARKs. We construct a PCD scheme given any non-interactive argument of knowledge (e.g., with linear-size proofs) that has a *split accumulation scheme*, which is a weak form of accumulation that we introduce.

We additionally construct a transparent non-interactive argument of knowledge for R1CS whose accumulation is verifiable via a *constant number of group and field operations*. This leads, via the random oracle heuristic and our result above, to efficiency improvements for PCD. Along the way, we construct a split accumulation scheme for a simple polynomial commitment scheme based on Pedersen commitments.

Our results are supported by a modular and efficient implementation.

Keywords: proof-carrying data; accumulation schemes; recursive proof composition

Contents

1	Introduction	3
1.1	Contributions	3
2	Techniques	7
2.1	Accumulation: atomic vs split	7
2.2	PCD from split accumulation	9
2.3	NARK with split accumulation based on DL	10
2.4	Split accumulation for Pedersen polynomial commitments	13
2.5	Implementation and evaluation	16
3	Preliminaries	18
3.1	Non-interactive arguments in the ROM	18
3.2	Proof-carrying data	19
3.3	Instantiating the random oracle	20
3.4	Post-quantum security	20
3.5	Commitment schemes	20
3.6	Polynomial commitments	21
4	Split accumulation schemes for relations	22
4.1	Accumulation schemes for certain predicates	23
5	PCD from arguments of knowledge with split accumulation	25
5.1	Construction	26
5.2	Completeness	26
5.3	Knowledge soundness	27
5.4	Efficiency	29
6	Accumulating Pedersen polynomial commitments	30
6.1	Accumulation scheme for PC_{HC}	31
6.2	Proof of Theorem 6.1	32
6.3	Knowledge soundness	32
7	Implementation	38
8	Evaluation	39
8.1	Comparing polynomial commitments based on DLs	39
8.2	Comparing accumulation schemes based on DLs	39
A	Split accumulation scheme for R1CS	41
	Acknowledgements	43
	References	43

1 Introduction

Proof-carrying data (PCD) [CT10] is a powerful cryptographic primitive that enables mutually distrustful parties to perform distributed computations that run indefinitely, while ensuring that every intermediate state of the computation can be efficiently verified. A special case of PCD is *incrementally-verifiable computation* (IVC) [Val08]. PCD has found applications in enforcing language semantics [CTV13], verifiable MapReduce computations [CTV15], image authentication [NT16], blockchains [Mina; KB20; BMRS20; CCDW20], and others. Given the theoretical and practical relevance of PCD, it is an important research question to build efficient PCD schemes from minimal cryptographic assumptions.

PCD from succinct verification. The canonical construction of PCD is via *recursive composition* of succinct non-interactive arguments (SNARGs) [BCCT13; BCTV14; COS20]. Informally, a proof that the computation was executed correctly for t steps consists of a proof of the claim “the t -th step of the computation was executed correctly, and there exists a proof that the computation was executed correctly for $t - 1$ steps”. The latter part of the claim is expressed using the SNARG verifier itself. This construction yields secure PCD (with IVC as a special case) provided the SNARG satisfies an adaptive knowledge soundness property (i.e., is a SNARK). Efficiency requires the SNARK to have sublinear-time verification, achievable via SNARKs for machine computations [BCCT13] or preprocessing SNARKs for circuit computations [BCTV14; COS20].

Requiring sublinear-time verification, however, significantly restricts the choice of SNARK, which limits what we can achieve for PCD. These restrictions have practical implications: the concrete efficiency of recursion is limited by the use of expensive curves for pairing-based SNARKs [BCTV14] or heavy use of cryptographic hash functions for hash-based SNARKs [COS20].

PCD from accumulation. Recently, [BCMS20] gave an alternative construction of PCD using SNARKs that have succinct *accumulation schemes*; this developed and formalized a novel approach for recursion sketched in [BGH19]. Informally, rather than being required to have sublinear-time verification, the SNARK is required to be accompanied by a cryptographic primitive that enables “postponing” the verification of SNARK proofs by way of an accumulator that is updated at each recursion step. The main efficiency requirement on the accumulation scheme is that the accumulation procedure must be succinctly verifiable, and in particular the accumulator itself must be succinct.

Requiring a SNARK to have a succinct accumulation scheme is a weaker condition than requiring it to have sublinear-time verification. This has enabled constructing PCD from SNARKs that do *not* have sublinear-time verification [BCMS20]. This yields PCD constructions from assumptions and with efficiency properties that were not previously achieved. Practitioners have exploited this freedom to design implementations of recursive composition with improved practical efficiency [Halo20; Pickles20].

Our motivation. The motivation of this paper is twofold. First, can PCD be built from a weaker primitive than SNARKs with succinct accumulation schemes? If so, can we leverage this to obtain PCD constructions with improved *concrete* efficiency?

1.1 Contributions

We make theory and systems contributions that advance the state of the art for PCD: (1) We introduce *split accumulation schemes for relations*, a cryptographic primitive that relaxes prior notions of accumulation. (2) We obtain PCD from any non-interactive argument of knowledge that satisfies this weaker notion of accumulation; surprisingly, this allows for arguments with no succinctness whatsoever. (3) We construct a non-interactive argument of knowledge based on discrete logarithms (and random oracles) whose accumulation verifier uses a constant number of group operations (improving over logarithmically many of prior

accumulation schemes in this setting). (4) We contribute libraries with generic implementation of PCD via accumulation from [BCMS20] and this paper and corresponding accumulation schemes.

We elaborate on each of these contributions next.

(1) Split accumulation for relations. Recall from [BCMS20] that an accumulation scheme for a predicate $\Phi: X \rightarrow \{0, 1\}$ enables proving/verifying that each input in an infinite stream q_1, q_2, \dots satisfies the predicate Φ , by augmenting the stream with *accumulators*. Informally, for each i , the prover produces a new accumulator acc_i from the input q_i and the old accumulator acc_{i-1} ; the verifier can check that the triple (acc_{i-1}, q_i, acc_i) is a valid accumulation step, much more efficiently than running Φ on q_i . At any time, the decider can validate acc_i , which establishes that for all $j \leq i$ it was the case that $\Phi(q_j) = 1$. The accumulator size (and hence the running time of the three algorithms) cannot grow in the number of accumulation steps.

We extend this notion in two orthogonal ways. First we consider relations $\Phi: X \times W \rightarrow \{0, 1\}$ and now for a stream of instances $q_1.x, q_2.x, \dots$ the goal is to establish that there exist witnesses $q_1.w, q_2.w, \dots$ such that $\Phi(q_i.x, q_i.w) = 1$ for each i . Second we consider accumulators acc_i that are split into an instance part $acc_i.x$ and a witness part $acc_i.w$ with the restriction that the accumulation verifier only gets to see the instance part (and possibly an auxiliary accumulation proof). We refer to this notion as *split accumulation for relations*, and refer (for contrast) the notion from [BCMS20] as *atomic accumulation for languages*.

The purpose of these extensions is to enable us to consider accumulation schemes in which predicate witnesses and accumulator witnesses are large while still requiring the accumulation verifier to be succinct (it receives short predicate instances and accumulator instances but not large witnesses). We will see that such accumulation schemes are both simpler and cheaper, while still being useful for primitives such as PCD.

(2) PCD via split accumulation. A non-interactive argument has a split accumulation scheme if the relation corresponding to its verifier has a split accumulation scheme (we make this precise later). We show that any non-interactive argument of knowledge (NARK) having a split accumulation scheme whose *accumulation verifier* is sublinear can be used to build a proof-carrying data (PCD) scheme, *even if the NARK does not have sublinear argument size*. This significantly broadens the class of non-interactive arguments from which PCD can be built, and is the first result to obtain PCD from non-interactive arguments that need not be succinct.

Theorem 1 (informal). *There is an efficient transformation that compiles any NARK with a split accumulation scheme into a PCD scheme. If the NARK and its split accumulation scheme are zero knowledge, then the PCD scheme is also zero knowledge.*

Similarly to all PCD results known to date, the above theorem holds in a model where all parties have access to a common reference string, *but no oracles*. (The construction makes non-black-box use of the accumulation scheme verifier, and the theorem does not carry over to the random oracle model.)

A corollary of Theorem 1 is that any NARK with a split accumulation scheme can be “bootstrapped” into a SNARK for machine computations. (PCD implies IVC and, further assuming collision-resistant hashing, also efficient SNARKs for machine computations [BCCT13].) This is surprising: an argument with decidedly weak efficiency properties implies an argument with succinct proofs and succinct verification!

See Section 2.2 for a summary of the ideas behind Theorem 1, and Section 5 for technical details.

(3) NARK with split accumulation based on DL. Our Theorem 1 motivates the question of whether we can leverage the weaker condition on the argument DL system to improve the efficiency of PCD. Our focus is on minimizing the cost of the accumulation verifier for the argument system, because it is the one component that is used in a black-box way, and thus typically determines concrete efficiency. Towards this end, we present a NARK with split accumulation based on discrete logarithms, with a *constant-size* accumulation verifier; the NARK has a transparent (public-coin) setup.

Theorem 2 (informal). *In the random oracle model and assuming the hardness of the discrete logarithm problem, there exists a (transparent) NARK for RICS and a corresponding split accumulation scheme with the following efficiency:*

NARK			split accumulation scheme			
prover time	verifier time	argument size	prover time	verifier time	decider time	accumulator size
$O(m) \mathbb{G}$	$O(m) \mathbb{G}$	$O(1) \mathbb{G}$	$O(m) \mathbb{G}$	$O(1) \mathbb{G}$	$O(m) \mathbb{G}$	$ \text{acc.x} = O(1) \mathbb{G} + O(1) \mathbb{F}$
$O(m \log m) \mathbb{F}$	$O(m) \mathbb{F}$	$O(m) \mathbb{F}$	$O(m) \mathbb{F}$	$O(1) \mathbb{F}$	$O(m) \mathbb{F}$	$ \text{acc.w} = O(m) \mathbb{F}$

Above, m the number of constraints in the RICS instance, \mathbb{G} denotes group scalar multiplications or group elements, and \mathbb{F} denotes field operations or field elements.

The NARK construction from Theorem 2 is particularly simple: it is obtained by applying the Fiat–Shamir transformation to a sigma protocol for RICS based on Pedersen commitments (and linear argument size). The only “special” feature about the construction is that, as we prove, it has a very efficient split accumulation scheme for the relation corresponding to its verifier. By heuristically instantiating the random oracle, we can apply Theorem 1 (and [BCCT13]) to obtain a SNARK for machines from this modest starting point.

We find it informative to compare Theorem 2 and SNARKs with atomic accumulation based on discrete logarithms [BCMS20]:

- the SNARK’s argument size is $O(\log m)$ group elements, *much less* than the NARK’s $O(m)$ field elements;
- the SNARK’s accumulator verifier uses $O(\log m)$ group scalar multiplications and field operations, *much more* than the NARK’s $O(1)$ group scalar multiplications and field operations.

Therefore Theorem 2 offers a tradeoff that minimizes the cost of the accumulator at the expense of argument size. (As we shall see later, this tradeoff has concrete efficiency advantages.)

Our focus on argument systems based on discrete logarithms is motivated by the fact that they can be instantiated based on efficient curves suitable for recursion: the Tweedle [BGH19] or Pasta [Hop20] curve cycles, which follow the curve cycle technique for efficient recursion [BCTV14]. This focus on discrete logarithms is a choice made for this paper, and we believe that our ideas can lead to efficiency improvements to recursion in other settings (e.g., pairing-based and hash-based arguments) and leave these to future work.

(4) Implementation and evaluation. We contribute a set of Rust libraries that realize PCD via accumulation via modular combinations of interchangeable components: (a) generic interfaces for atomic and split accumulation; (b) generic construction of PCD from arguments with atomic and split accumulation; (c) accumulation schemes for polynomial commitments based on the inner product argument or a trivial use of the Pedersen commitment (this latter underlies our NARK), including constraints for the accumulation verifiers. Practitioners interested in PCD will find these libraries useful for prototyping and comparing different types of recursion (and, e.g., may help decide if current systems based on atomic recursion [Halo20; Pickles20] are better off via split recursion or not).

We additionally conduct experiments to evaluate our implementation, and establish that for constraint systems of interest, recursion based on our split accumulation scheme is much cheaper in practice than recursion based on prior atomic accumulation schemes.

Remark 1.1 (concurrent work). A concurrent work [BDFG20] studies similar questions as this paper. Below we summarize the similarities and the differences between the two papers.

Similarities. Both papers are motivated by the goal of reducing the cost of recursive arguments. The main object of study in [BDFG20] is additive polynomial commitment schemes (PC schemes), for which [BDFG20] considers different types of *aggregation schemes*: (1) *public* aggregation in [BDFG20] is closely related to atomic accumulation specialized to PC schemes from a prior work [BCMS20]; and (2) *private*

aggregation in [BDFG20] is closely related to split accumulation specialized to PC schemes from this paper. Moreover, the private aggregation scheme for additive PC schemes in [BDFG20] is similar to our split accumulation scheme for Pedersen PC schemes (overviewed in Section 2.4 and detailed in Section 6). The protocols differ in how efficiency depends on the n claims to aggregate/accumulate: the verifier in [BDFG20] uses $n + 1$ group scalar multiplications while ours uses $2n$. (Informally, [BDFG20] first randomly combines claims and then evaluates at a random point, while we first evaluate at a random point and then randomly combine claims.)

Differences. The two papers develop distinct, and complementary, directions.

The focus of [BDFG20] is to design protocols for any additive PC scheme (and, even more generally, any PC scheme with a linear combination scheme), including the aforementioned private aggregation protocol and a compiler that endows a given PC scheme with zero knowledge.

In contrast, our focus is to formulate a definition of split accumulation for general relation predicates that (a) we demonstrate suffices to construct PCD, and (b) in the random oracle model, we can also demonstrably achieve via a split accumulation scheme for Pedersen commitments. We emphasize that our definitions are materially different from the case of atomic accumulation in [BCMS20], and necessitate careful consideration of technicalities such as the flavor of adaptive knowledge soundness, which algorithms can be allowed to query oracles, and so on. Hence, we cannot simply rely on the existing foundations for atomic accumulation of [BCMS20] in order to infer the correct definitions and security reductions for split accumulation. Overall, our theoretical work enables us to achieve the first construction of PCD without succinct arguments, and also to obtain a novel NARK for RICS with constant-size accumulation verifier.

We stress that the treatment of accumulation at a higher level of abstraction than for PC schemes is essential to prove theorems about PCD. There are no known techniques that enable provably constructing PCD from an aggregation/accumulation scheme for a PC scheme (unlike what is stated in [BDFG20]), because using the PC scheme would involve using the random oracle in a way that cannot be recursively proved. Instead, one can prove a theorem in the standard model starting from an aggregation/accumulation scheme for a NARK, as we do. Separately, similar to the disconnect in [BCMS20], we do not have accumulation schemes for NARKs in the standard model so instead we prove an analogous result in the ROM.

Another major difference is that we additionally contribute a comprehensive and modular implementation of protocols from [BCMS20] and this paper, and conduct an evaluation for the discrete logarithm setting. This supports the asymptotic improvements with measured improvements in concrete efficiency.

2 Techniques

We summarize the main ideas behind our results. In Section 2.1 we discuss the new notion of split accumulation for relation predicates, and compared it with the notion of atomic accumulation for language predicates from [BCMS20]. In Section 2.2 we discuss the proof of Theorem 1. In Section 2.3 we discuss the proof of Theorem 2; for this we rely, among other things, on a new result about split accumulation for a Pedersen-based polynomial commitment (Theorem 3), which we discuss in Section 2.4. Finally in Section 2.5 we elaborate on our implementation and evaluation. Figure 1 illustrates the relation between our results. The rest of the paper contains technical details, and we provide pointers to relevant sections along the way.

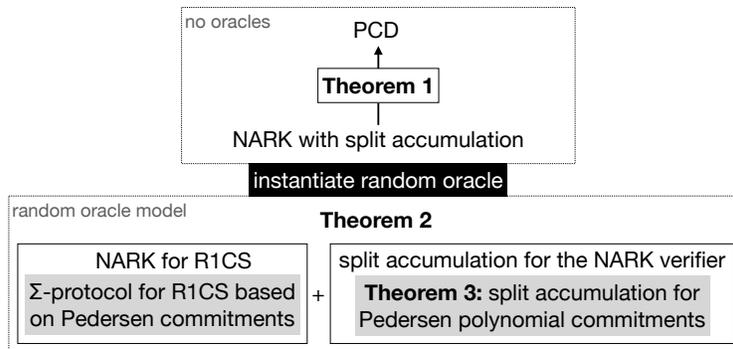


Figure 1: Diagram showing the relation between our results. Gray boxes within a result are notable subroutines.

2.1 Accumulation: atomic vs split

We review the notion of accumulation from [BCMS20], which we refer to as *atomic accumulation*, and then describe the weaker notion that we consider, which we call *split accumulation*.

Atomic accumulation for languages. An *accumulation scheme for a language predicate* $\Phi: X \rightarrow \{0, 1\}$ consists of a tuple of algorithms (P, V, D) , known as the prover, verifier, and decider, that enable proving/verifying statements of the form $\Phi(q_1) \wedge \Phi(q_2) \wedge \dots$ more efficiently than simply invoking the predicate Φ on each input.

This is done by starting from an initial (“empty”) accumulator acc_0 , the prover is used to accumulate the first input q_1 to produce a new accumulator $\text{acc}_1 \leftarrow P(\text{acc}_0, q_1)$; then the prover is used again to accumulate the second input q_2 to produce a new accumulator $\text{acc}_2 \leftarrow P(\text{acc}_1, q_2)$; and so on.

Each accumulator produced so far enables efficient verification of the predicate on all inputs that went into the accumulator. For example, to establish that $\Phi(q_1) \wedge \dots \wedge \Phi(q_t) = 1$ it suffices to check that:

- the verifier accepts each accumulation step: $V(\text{acc}_0, q_1, \text{acc}_1) = 1$, $V(\text{acc}_1, q_2, \text{acc}_2) = 1$, and so on; and
- the decider accepts the final accumulator: $D(\text{acc}_t) = 1$.

Qualitatively, this replaces the naive cost $t \cdot |\Phi|$ with the new cost $t \cdot |V| + |D|$. This is beneficial when the verifier is much cheaper than checking the predicate directly and the decider is not much costlier than checking the predicate directly. Crucially, the verifier and decider costs (and, in particular, the accumulator size) should not grow with the number t of accumulation steps (which need not be known in advance).

The properties of an accumulation scheme are summarized in the following informal definition.

Definition 2.1 (informal). An **accumulation scheme** for a predicate $\Phi: X \rightarrow \{0, 1\}$ consists of a triple of algorithms (P, V, D) , known as the prover, verifier, and decider, that satisfies the following properties.

- **Completeness:** For all accumulators acc and predicate inputs $q \in X$, if $D(\text{acc}) = 1$ and $\Phi(q) = 1$, then for $\text{acc}^* \leftarrow P(\text{acc}, q)$ it holds that $V(\text{acc}, q, \text{acc}^*) = 1$ and $D(\text{acc}^*) = 1$.
- **Soundness:** For every efficiently-generated accumulators acc, acc^* and predicate input $q \in X$, if $D(\text{acc}^*) = 1$ and $V(\text{acc}, q, \text{acc}^*) = 1$ then, with all but negligible probability, $\Phi(q) = 1$ and $D(\text{acc}) = 1$.

The above definition omits many details, such as the ability to accumulate multiple inputs and multiple accumulators in one step, or features required for zero knowledge such as a validity proof for the accumulator (produced by P and checked by V but not passed on like the accumulator). We refer the reader to [BCMS20] for more details.

The aspect that we wish to highlight here is the following: in order for the verifier to be much cheaper than the predicate ($|V| \ll |\Phi|$) it must be that the accumulator itself is much smaller than the predicate ($|\text{acc}| \ll |\Phi|$) because the verifier receives the accumulator as input. (And if the accumulator is accompanied by a validity proof π_V then this proof must also be small.)

We refer to this setting as *atomic accumulation* because the entirety of the accumulator is treated as one short monolithic string. In contrast, next we consider a relaxed variant where this is not the case, and will enable us to consider new instantiations that ultimately lead to interesting theoretical and practical results.

Split accumulation for relations. We propose a relaxed notion of accumulation scheme. A *split accumulation scheme* for a relation predicate $\Phi: X \times W \rightarrow \{0, 1\}$ is again a tuple of algorithms (P, V, D) as before. The difference is that: (a) an input q to Φ is split into a short instance part $q.x$ and a (possibly) long witness part $q.w$; (b) an accumulator acc is split into a short instance part acc.x and a (possibly) long witness part acc.w ; (c) the verifier only needs the short parts of inputs and accumulators to verify an accumulation step, along with a short validity proof instead of the long witness parts.

As before, the prover is used to accumulate a new input q_i into a prior accumulator acc_{i-1} to obtain a new accumulator and validity proof $(\text{acc}_i, \pi_{V,i}) \leftarrow P(\text{acc}_{i-1}, q_i)$. Different from before, however, we wish to establish that there exists (or, more precisely, a party knows) witnesses $q_1.w, \dots, q_t.w$ such that $\Phi(q_1.x, q_1.w) \wedge \dots \wedge \Phi(q_t.x, q_t.w) = 1$, and for this it suffices to check that:

- the verifier accepts each accumulation step given only the short instance parts: $V(\text{acc}_0.x, q_1.x, \text{acc}_1.x, \pi_{V,1}) = 1, V(\text{acc}_1.x, q_2.x, \text{acc}_2.x, \pi_{V,2}) = 1$, and so on; and
- the decider accepts the final accumulator (made of both the instance and witness part): $D(\text{acc}_t) = 1$.

Again the naive cost $t \cdot |\Phi|$ is replaced with the new cost $t \cdot |V| + |D|$, but now it could be that an accumulator is, e.g., as large as $|\Phi|$; we only need the *instance part* of the accumulator (and predicate inputs) to be short.

The security property of a split accumulation scheme involves an extractor that outputs a long witness part from a short instance part and proof, and is reminiscent to the knowledge soundness of a succinct non-interactive argument. Turning this high level description into a working definition requires some care, however, and we view this as a contribution of this paper. (By “working definition” we mean a definition that we can provably fulfill under concrete hardness assumptions, and provably suffices for recursive composition.) Informally the security definition could be summarized as follows.

Definition 2.2 (informal). A **split accumulation scheme** for a predicate $\Phi: X \times W \rightarrow \{0, 1\}$ consists of a triple of algorithms (P, V, D) that satisfies the following properties.

- **Completeness:** For all accumulators acc and predicate inputs $q \in X \times W$, if $D(\text{acc}) = 1$ and $\Phi(q) = 1$, then for $(\text{acc}^*, \pi_V^*) \leftarrow P(\text{acc}, q)$ it holds that $V(\text{acc.x}, q.x, \text{acc}^*.x, \pi_V^*) = 1$ and $D(\text{acc}^*) = 1$.
- **Knowledge:** For every efficiently-generated old accumulator instance acc.x , old input instance $q.x$, accumulation proof π_V^* , and new accumulator acc^* , if $D(\text{acc}^*) = 1$ and $V(\text{acc.x}, q.x, \text{acc}^*.x, \pi_V^*) = 1$ then, with all but negligible probability, an efficient extractor can find old witness parts acc.x and $q.x$ such that $\Phi((q.x, q.w)) = 1$ and $D((\text{acc.x}, \text{acc.w})) = 1$.

One can verify that split accumulation is indeed a relaxation of atomic accumulation: any atomic accumulation scheme is (trivially) a split accumulation scheme with empty witnesses. Crucially, however, a split accumulation scheme alleviates a major restriction of atomic accumulation, namely, that the accumulator itself (and a predicate input) has to be short.

See Section 4 for details about the definition of split accumulation.

Next, in Section 2.2 we show that split accumulation suffices for recursive composition (which has surprising theoretical consequences) and then in Section 2.3 we present a NARK with split accumulation scheme based on discrete logarithms.

2.2 PCD from split accumulation

We summarize the main ideas behind Theorem 1, which obtains proof-carrying data (PCD) from any NARK that has a split accumulation scheme. To ease exposition, in this summary we focus on IVC, which can be viewed as the special case where a circuit F is repeatedly applied. That is, we wish to incrementally prove a claim of the form “ $F^T(z_0) = z_T$ ” where F^T denotes F composed with itself T times.

Prior work: recursion via atomic accumulation. Our starting point is a theorem from [BCMS20] that obtains PCD from any SNARK that has an atomic accumulation scheme. The IVC construction implied by that theorem is roughly follows.

- The *IVC prover* receives as input a previous instance z_i , proof π_i , and accumulator acc_i ; then accumulates (z_i, π_i) with acc_i to obtain a new accumulator acc_{i+1} ; and finally generates a SNARK proof π_{i+1} of the claim: “ $z_{i+1} = F(z_i)$, and there exist a proof π_i and an accumulator acc_i such that the accumulation verifier accepts $((z_i, \pi_i), \text{acc}_i, \text{acc}_{i+1})$ ”, expressed as a circuit R (see Fig. 2, middle box). The IVC proof for z_{i+1} is $(\pi_{i+1}, \text{acc}_{i+1})$.
- The *IVC verifier* validates an IVC proof (π_i, acc_i) for z_i by running the SNARK verifier on $((z_i, \text{acc}_i), \pi_i)$ and running the accumulation scheme decider on acc_i .

In each iteration we maintain the invariant that if acc_i is a valid accumulator (according to the decider) and π_i is a valid proof, then the computation is correct up to the i -th step.

Note that while it would suffice to prove that “ $z_{i+1} = F(z_i)$, π_i is a valid proof, and acc_i is a valid accumulator”, we cannot afford to do so. Indeed: (i) proving that π_i is a valid proof requires proving a statement about the argument verifier, which may not be sublinear; and (ii) proving that acc_i is a valid accumulator requires proving a statement about the decider, which may not be sublinear. Instead of proving this claim directly, we “defer” it by having the prover accumulate (z_i, π_i) into acc_i to obtain a new accumulator acc_{i+1} . The soundness property of the accumulation scheme ensures that if acc_{i+1} is valid and the accumulation verifier accepts $((z_i, \pi_i), \text{acc}_i, \text{acc}_{i+1})$, then π_i is a valid proof and acc_i is a valid accumulator. Thus all that remains to maintain the invariant is for the prover to prove that the accumulation verifier accepts; this is possible provided that the *accumulation verifier* is sublinear.

Our construction: recursion via split accumulation. Our construction naturally extends the above idea in the setting of NARKs with split accumulation schemes. Indeed, the only difference to the above construction is that the proof π_{i+1} generated by the IVC prover is for the statement “ $z_{i+1} = F(z_i)$, and there exist a proof *instance* $\pi_{i.\mathbb{x}}$ and an accumulator *instance* $\text{acc}_{i.\mathbb{x}}$ such that the accumulation verifier accepts $((z_i, \pi_{i.\mathbb{x}}), \text{acc}_{i.\mathbb{x}}, \text{acc}_{i+1.\mathbb{x}})$ ”, and accordingly the IVC verifier runs the NARK verifier on $((z_i, \text{acc}_{i.\mathbb{x}}), \pi_i)$. This is illustrated in Fig. 2 (lower box). Note that the circuit R itself is unchanged from the atomic case; the difference is in whether we pass the *entire* proof and accumulators or just the \mathbb{x} part.

Proving that this relaxation yields a secure construction is more complex. Similar to prior work, the proof of security proceeds via a recursive extraction argument, as we explain next.

For an atomic accumulation scheme ([BCMS20]), one maintains the following extraction invariant: the i -th extractor outputs $(z_i, \pi_i, \text{acc}_i)$ such that π_i is valid according to the SNARK, acc_i is valid according to the decider, and $F^{T-i}(z_i) = z_T$. The T -th “extractor” is simply the malicious prover, and we can obtain the i -th extractor by applying the knowledge guarantee of the SNARK to the $(i + 1)$ -th extractor. That the invariant is maintained is implied by the soundness guarantee of the atomic accumulation scheme.

For a split accumulation scheme, we want to maintain the same extraction invariant; however, the extractor for the NARK will only yield $(z_i, \pi_i.\mathbb{x}, \text{acc}_i.\mathbb{x})$, and not the corresponding witnesses. This is where we make use of the extraction property of the split accumulation scheme itself. Specifically, we interleave the knowledge guarantees of the NARK and accumulation scheme as follows: the i -th NARK extractor is obtained from the $(i + 1)$ -th accumulation extractor using the knowledge guarantee of the NARK, and the i -th accumulation extractor is obtained from the i -th NARK extractor using the knowledge guarantee of the accumulation scheme. We take the malicious prover to be the T -th accumulation extractor.

From sketch to proof. In Section 5, we give the formal details of our construction and a proof of correctness. In particular, we show how to construct PCD, a more general primitive than IVC. In the PCD setting, rather than each computation step having a single input z_i , it receives m inputs from different nodes. Proving correctness hence requires proving that *all* of these inputs were computed correctly. For our construction, this entails checking m proofs and m accumulators. To do this, we extend the definition of an accumulation scheme to allow accumulating multiple instance-proof pairs and multiple “old” accumulators.

We also note that the application to PCD leads to other definitional considerations, which are similar to those that have appeared in previous works [COS20; BCMS20]. In particular, the knowledge soundness guarantee for both the NARK *and* the accumulation scheme should be of the stronger “witness-extended emulation with auxiliary input and output” type used in previous work. Additionally, the underlying construction of split accumulation achieves only expected polynomial-time extraction (in the ROM), and so the recursive extraction technique requires that we are able to extract from expected-time adversaries.

Remark 2.3 (flavors of PCD). The recent advances in PCD from accumulation achieve weaker forms than PCD from succinct verification, and formally these results are incomparable. (Starting from weaker assumptions they obtain weaker conclusions.) The essential feature that all these works achieve is that the efficiency of PCD algorithms is independent of the number of nodes in the PCD computation, which is how PCD is defined (see Section 3.2). That said, prior work on PCD from succinct verification [BCCT13; BCTV14; COS20] additionally guarantees that verifying a PCD proof is sublinear in a node’s computation; and prior work on PCD from atomic accumulation [BCMS20] merely ensures that a PCD proof has size (but not necessarily verification time) that is sublinear in a node’s computation. The PCD scheme obtained in this paper does not have these additional features: a PCD proof has size that is linear in a node’s computation.

2.3 NARK with split accumulation based on DL

We summarize the main ideas behind Theorem 2, which provides, in the discrete logarithm setting with random oracles, a NARK for RICS that has a split accumulation scheme whose accumulation verifier uses a constant number of group scalar multiplications. We describe the NARK and then explain at high level how we accumulate the verifier’s computation. This in turn motivates accumulation for polynomial commitments based on Pedersen commitments, which we address in the next subsection (Section 2.4).

recursion circuit via succinct verification	$R((\text{ivk}, z_i), (z_{i-1}, \pi_{i-1})) :$ <ul style="list-style-type: none"> • $z_i = F(z_{i-1})$ • $\text{SNARK.V}(\text{ivk}, z_{i-1}, \pi_{i-1}) = 1$
recursion circuit via atomic accumulation	$R((\text{avk}, z_i, \text{acc}_i), (z_{i-1}, \pi_{i-1}, \text{acc}_{i-1}, \pi_{V,i})) :$ <ul style="list-style-type: none"> • $z_i = F(z_{i-1})$ • $\text{ACC.V}(\text{avk}, ((\text{avk}, z_{i-1}, \text{acc}_{i-1}), \pi_{i-1}), \text{acc}_{i-1}, \text{acc}_i, \pi_{V,i}) = 1$
recursion circuit via split accumulation	$R((\text{avk}, z_i, \text{acc}_{i.\mathbb{X}}), (z_{i-1}, \pi_{i-1.\mathbb{X}}, \text{acc}_{i-1.\mathbb{X}}, \pi_{V,i})) :$ <ul style="list-style-type: none"> • $z_i = F(z_{i-1})$ • $\text{ACC.V}(\text{avk}, ((\text{avk}, z_{i-1}, \text{acc}_{i-1.\mathbb{X}}), \pi_{i-1.\mathbb{X}}), \text{acc}_{i-1.\mathbb{X}}, \text{acc}_{i.\mathbb{X}}, \pi_{V,i}) = 1$

Figure 2: Comparison of circuits used to realize recursion with different techniques.

We highlight here that both the NARK and the accumulation scheme are particularly simple compared to other protocols in the SNARK literature (especially with regard to recursion!), and view this as a significant advantage for potential practical deployments.

NARK for RICS. Recall that RICS is a standard generalization of arithmetic circuit satisfiability where the “circuit description” is given by coefficient matrices, as below; there “ \circ ” denotes the entry-wise product.

Definition 2.4 (RICS problem). *Given a finite field \mathbb{F} , coefficient matrices $A, B, C \in \mathbb{F}^{m \times n}$ and an instance vector $x \in \mathbb{F}^\ell$, is there a witness vector $w \in \mathbb{F}^{n-\ell}$ such that $z := (x, w) \in \mathbb{F}^n$ and $Az \circ Bz = Cz$?*

The NARK for RICS in Theorem 2 is obtained by applying the Fiat–Shamir transformation to a sigma protocol for RICS based on Pedersen commitments. We provide the sigma protocol in Figure 3. The public parameters pp are those for the Pedersen commitment to have message space \mathbb{F}^{m+1} (i.e., pp contains a group description and $m + 1$ random group elements) and Commit its commitment function; also, $V_H(X)$ denotes the vanishing polynomial of a subset $H \subseteq \mathbb{F}$ of size m .

The efficiency is evident from the construction: quasilinear prover time (due to polynomial arithmetic); linear argument size (the first message is succinct but the third message is not); and linear verifier time (more precisely, linear number of group scalar multiplications). In more detail, denoting by k the number of non-zero entries in the coefficient matrices $A, B, C \in \mathbb{F}^{m \times n}$, the asymptotic efficiency is as follows:

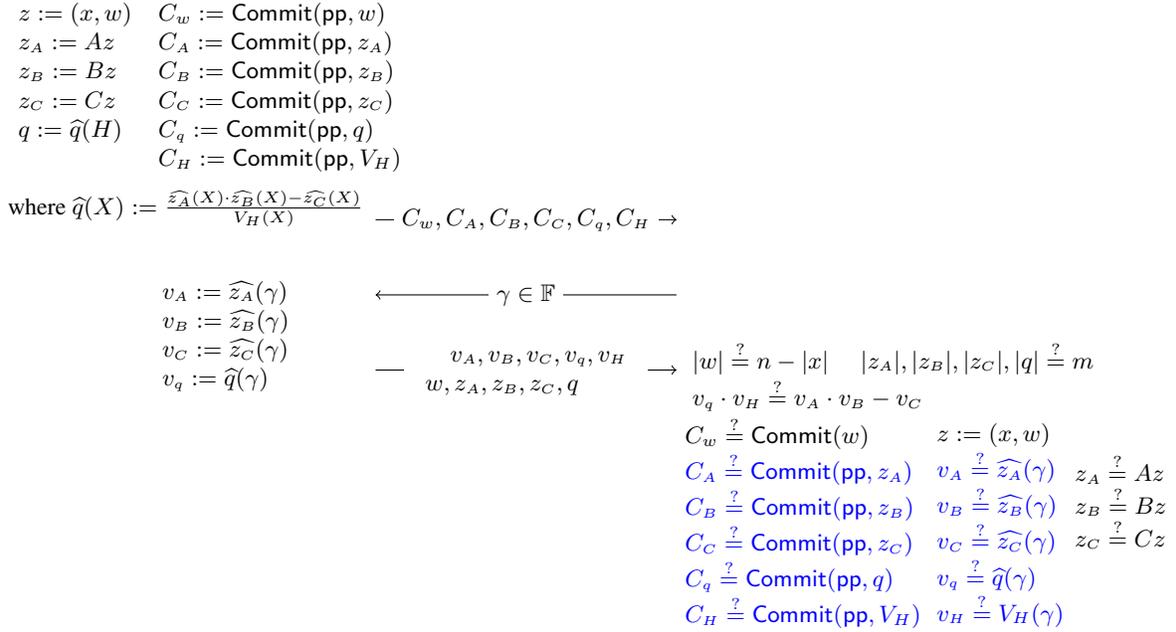
NARK prover time	NARK verifier time	NARK argument size
$O(m + n) \mathbb{G}$	$O(m + n) \mathbb{G}$	$O(1) \mathbb{G}$
$O(m \log m + n + k) \mathbb{F}$	$O(m + n + k) \mathbb{F}$	$O(m + n) \mathbb{F}$

The sigma protocol may superficially appear useless because it has linear argument size and is not zero knowledge as stated (though it could be modified to be); in fact, the prover commits to the witness (and more) in the first message, only to send “everything” to the verifier in the next round, so that the verifier is (almost) checking the RICS relation itself!

How is this “better” than the prover simply sending the witness as a non-interactive proof?

The key point is that we do not know how to obtain a split accumulation scheme for the computation that directly checks the witness but we do know how to obtain one for the NARK verifier in Figure 3.

Split accumulation for the NARK. We describe a split accumulation scheme for the NARK for RICS with the efficiency as reported in the table below (where we additionally report the number of oracle calls).

**Figure 3:** The sigma protocol for RICS that underlies the NARK for RICS.

accumulation prover time (per accumulated NARK)	accumulation verifier time (per accumulated NARK)	decider time	accumulator size
$O(m+n) \mathbb{G}$	$O(1) \mathbb{G}$	$O(m+n) \mathbb{G}$	$ \text{acc.x} = O(1) \mathbb{G} + O(1) \mathbb{F}$
$O(m+n+k) \mathbb{F}$	$O(1) \mathbb{F}$	$O(m+n+k) \mathbb{F}$	$ \text{acc.w} = O(m+n) \mathbb{F}$
$O(1) \text{RO}$	$O(1) \text{RO}$	-	-

We view the first message of the protocol as a succinct instance and the third message of the protocol as a non-succinct “opening” witness aimed at making the NARK verifier accept. We now sketch how to design a split accumulation scheme for the relation implied by this instance-witness split of the NARK verifier; technical details are in Appendix A.

The NARK verifier performs four types of checks:

1. a test for the polynomial equation $\hat{q}(X) \cdot V_H(X) = \widehat{z}_A(X) \cdot \widehat{z}_B(X) - \widehat{z}_C(X)$;
2. checks for the linear equations $z_A = Az, z_B = Bz, z_C = Cz$;
3. checking a commitment ($C_w \stackrel{?}{=} \text{Commit}(w)$);
4. several checks of the form “ $C = \text{Commit}(f)$ and $\hat{f}(\gamma) = v$ ” (highlighted in blue in Figure 3).

The first check is cheap so does not have to be accumulated (the accumulation verifier directly performs this check). The second and third checks are linear and so can be easily (split) accumulated via random linear combinations. The last type of check is concerned with *evaluation claims about committed polynomials*.¹

¹Note that it may appear odd that these include even a claim about the vanishing polynomial V_H because, for suitable choices of H , can be evaluated in only $\log |H| = \log m$ field operations directly. This is because, similarly to an optimization in [CCDW20], by having the prover commit to the vanishing polynomial, we can ensure later on that the accumulation verifier only has to perform $O(1)$ field operations rather than $O(\log m)$.

This requires some work: designing a highly-efficient split accumulation scheme for the “trivial” polynomial commitment scheme based on Pedersen commitments — we discuss this in the next section.

Remark 2.5. If one replaced the (succinct) polynomial commitment scheme that underlies the preprocessing SNARK in [CHMMVW20] with a (non-succinct) Pedersen polynomial commitment scheme then one would obtain a NARK for RICS with a split accumulation scheme whose accumulation verifier *is* of constant size but other asymptotics would be worse compared to Theorem 2. First, the cryptographic costs and the quasilinear costs of the NARK and accumulation scheme would also grow in the number k of non-zero entries in the coefficient matrices, which can be much larger than m and n (asymptotically and concretely). Second, the NARK verifier and accumulation decider would additionally use a quasilinear number of field operations due to FFTs (to run the the indexer from that paper). Finally, in addition to poorer asymptotics, this approach would lead to a concretely more expensive accumulation verifier and overall a more complex protocol.

2.4 Split accumulation for Pedersen polynomial commitments

Any commitment implies a “trivial” polynomial commitment scheme: commit to the polynomial (in coefficient or evaluation form), and then reveal the entire polynomial (the receiver can then check the evaluation claim directly). We fix the commitment to be a Pedersen commitment (what we used in Section 2.3), which implies that commitments are linear in the underlying polynomials. We prove that this simple polynomial commitment scheme has a very efficient split accumulation scheme in the random oracle model.

Theorem 3 (informal). *The Pedersen polynomial commitment scheme has a split accumulation scheme AS_{HC} that is secure in the random oracle model (and assuming the hardness of the discrete logarithm problem). Verifying accumulation requires 2 group scalar multiplications and $O(1)$ field additions/multiplications per claim, and results in an accumulator whose instance part is 1 group element and 2 field elements and whose witness part is d field elements. (See Table 1.)*

In Table 1 we compare our the efficiency of the split accumulation scheme AS_{HC} for the (non-succinct) Pedersen PC scheme PC_{HC} and the efficiency of the atomic accumulation scheme AS_{IPA} in [BCMS20] for the (succinct) PC scheme PC_{IPA} based on the inner-product argument on cyclic groups [BCCGP16; BBBPWM18; WTSTW18]. The takeaway is that accumulation verifier for AS_{HC} is significantly cheaper than the accumulation verifier for AS_{IPA} .

Technical details are in Section 6; in the rest of this section we sketch the ideas behind Theorem 3.

accumulation scheme	type	assumption	accumulation prover (per claim)	accumulation verifier (per claim)	accumulation decider	accumulator size instance	witness
AS_{IPA} [BCMS20]	atomic	DLOG + RO †	$O(\log d) \mathbb{G}$ $O(d) \mathbb{F}$ [+ $O(d) \mathbb{G}$ per accumulation]	$O(\log d) \mathbb{G}$ $O(\log d) \mathbb{F}$ $O(\log d) \text{RO}$	$O(d) \mathbb{G}$ $O(d) \mathbb{F}$	1 \mathbb{G} $O(\log d) \mathbb{F}$	0
AS_{HC} [this work]	split	DLOG + RO	$O(d) \mathbb{G}$ $O(d) \mathbb{F}$	2 \mathbb{G} $O(1) \mathbb{F}$ 2 RO	$O(d) \mathbb{G}$ $O(d) \mathbb{F}$	1 \mathbb{G} 2 \mathbb{F}	$d \mathbb{F}$

Table 1: Efficiency comparison between the atomic accumulation scheme AS_{IPA} for PC_{IPA} in [BCMS20] and the split accumulation scheme AS_{HC} for PC_{HC} in this work. Above \mathbb{G} denotes group scalar multiplications or group elements, and \mathbb{F} denotes field operations or field elements. (†: AS_{IPA} relies on knowledge soundness of PC_{IPA} , which results from applying the Fiat–Shamir transformation to a logarithmic-round protocol. The security of this protocol has only been proven via a superpolynomial-time extractor [BMMTV19] or in the algebraic group model [GT20].)

2.4.1 A simple linear polynomial commitment scheme

A polynomial commitment scheme is a cryptographic primitive that enables one to produce a commitment C to a polynomial p , and then to prove that this committed polynomial evaluates to a claimed value v at a desired point z . (See Section 3.6 for a definition.)

Any commitment scheme `Commit` implies a “trivial” commitment scheme: in the commit phase, the sender commits to the coefficients of the polynomial p ; in the reveal phase, the sender sends the entire polynomial p as opening, and the receiver can itself check that $p(z) = v$ as claimed.² (We are assuming that the message space of `Commit` is \mathbb{F}^n where \mathbb{F} is the field over which p is defined and that $n \geq \deg(p) + 1$.)

If additionally we know that `Commit` is linear (i.e., for every a and b in \mathbb{F}^n , $\text{Commit}(a) + \text{Commit}(b) = \text{Commit}(a+b)$) then the trivial construction yields a (trivial) linear polynomial commitment scheme. Typically the commitment space is a cryptographic group \mathbb{G} (so addition is over the group), and a natural example is the Pedersen commitment scheme, which is secure provided the discrete logarithm problem is hard in \mathbb{G} .

We refer to the foregoing trivial scheme as PC_{HC} , and next discuss accumulation for it.

2.4.2 Split accumulation for PC_{HC}

An (atomic or split) accumulation scheme for a PC scheme accumulates claims of the form “I know a polynomial p of degree at most d such that C is a commitment to p and $p(z) = v$ ”. We summarize such an *evaluation claim* via the tuple (C, z, v, d) . Below we summarize our split accumulation scheme for PC_{HC} .

First we describe a simple public-coin interactive reduction for combining two or more evaluation claims into a single evaluation claim, and then explain how this interactive reduction gives rise to the split accumulation scheme. We prove security in the random oracle model, using an expected-time extractor.

Batching evaluation claims. First consider two evaluation claims (C_1, z, v_1, d) and (C_2, z, v_2, d) for the *same* evaluation point z (and degree d). We can use a random challenge $\alpha \in \mathbb{F}$ to combine these claims into one claim (C', z, v', d) where $C' := C_1 + \alpha C_2$ and $v' := v_1 + \alpha v_2$. If either of the original claims does not hold then, with high probability over the choice of α , neither does the new claim. This idea extends to any number of claims for the same evaluation point, by taking $C' := \sum_i \alpha^i C_i$ and $v' := \sum_i \alpha^i v_i$.

Next consider two evaluation claims (C_1, z_1, v_1, d) and (C_2, z_2, v_2, d) at (possibly) different evaluation points z_1 and z_2 . We explain how these can be combined into four claims all at the *same* point. Below we use the fact that $p(z) = v$ if and only if there exists a polynomial $w(X)$ such that $p(X) = w(X) \cdot (X - z) + v$.

Let $p_1(X)$ and $p_2(X)$ be the polynomials “inside” C_1 and C_2 , respectively, that are known to the prover.

1. The prover computes the witness polynomials $w_1 := \frac{p_1(X) - v_1}{X - z_1}$ and $w_2 := \frac{p_2(X) - v_2}{X - z_2}$ and sends the commitments $W_1 := \text{Commit}(w_1)$ and $W_2 := \text{Commit}(w_2)$.
2. The verifier sends a random evaluation point $z^* \in \mathbb{F}$.
3. The prover computes and sends the evaluations $y_1 := p_1(z^*)$, $y_2 := p_2(z^*)$, $y'_1 := w_1(z^*)$, $y'_2 := w_2(z^*)$.
4. The verifier checks the relation between each witness polynomial and the original polynomial at the random evaluation point z^* :

$$y_1 = y'_1(z^* - z_1) + y_1 \quad \text{and} \quad y_2 = y'_2(z^* - z_2) + y_2 .$$

²Alternatively, and this is what we actually use in Section 2.3 (though we ignore the difference in this high level overview), the sender commits to the evaluation of the polynomial, and subsequently the receiver checks that extending the revealed evaluation table at z yields the claimed value v . Over suitable domains this operation is also linear time.

Next the verifier outputs four evaluation claims for $p_1(z^*) = y_1, p_2(z^*) = y_2, w_1(z^*) = y'_1, w_2(z^*) = y'_2$:

$$(C_1, z^*, y_1, d), (C_2, z^*, y_2, d), (W_1, z^*, y'_1, d), (W_2, z^*, y'_2, d).$$

More generally, we can reduce m evaluation claims at m points to $2m$ evaluation claims all at the same point.

By combining the two techniques, one obtains a public-coin interactive reduction from any number of evaluation claims (regardless of evaluation points) to a single evaluation claim.

Split accumulation. The batching protocol described above yields a split accumulation scheme for PC_{HC} (and, more generally, any linear polynomial commitment scheme) in the random oracle model. The accumulated predicate Φ takes inputs q whose instance part is an evaluation claim $q.\text{x} = (C, z, v, d)$ and whose witness part is a polynomial $q.\text{w} = p(X)$, and accepts if and only if $C = \text{Commit}(p)$ and $p(z) = v$. An accumulator acc has the same form as a predicate input: the instance part is an evaluation claim and the witness part is a polynomial. Next we describe the algorithms of the accumulation scheme.

- The accumulation prover P runs the interactive reduction by relying on the random oracle to generate the random verifier messages (i.e., it applies the Fiat–Shamir transformation to the reduction), in order to combine the instance parts of old accumulators and inputs to obtain the instance part of a new accumulator. Then P also combines the committed polynomials using the same linear combinations in order to derive the new committed polynomial, which is the witness part of the new accumulator.
- The accumulation verifier V checks that the challenges were correctly computed from the random oracle, and performs the checks of the reduction (the claims were correctly combined and that the proper relation between each y_i, y'_i, z_i, z^* holds).
- The accumulation decider D reads the accumulation in its entirety and checks that the polynomial (the witness part) satisfies the evaluation claim (the instance part). (Here the random oracle is not used.)

Efficiency. The efficiency claimed in Theorem 3 (and Table 1) is evident from the construction. The accumulation prover P computes $n + m$ commitments to polynomials when combining n old accumulators and m predicate inputs (all polynomials are for degree at most d). An accumulator consists of $O(d)$ group elements, with the (short) instance part consisting of 1 group element and 2 field elements. The accumulator decider D computes 1 commitment (and 1 polynomial evaluation at 1 point) in order to validate an accumulator. Finally, the cost of running the accumulator verifier V is dominated by $2(n + m)$ scalar multiplication of the linear commitments.

2.4.3 Security

While proving the completeness property of our split accumulation scheme is straightforward, proving its knowledge property is less so. Given an adversary that produces evaluation claims (C_i, z_i, v_i, d_i) , a single claim (C, z, v, d) which passes the checks in the reduction procedure, and a polynomial s with $s(z) = v$ to which C is a commitment, we obtain an extractor which outputs polynomials p_i with $p_i(z_i) = v_i$ to which C_i is a commitment. Our security proof (detailed in Section 6.3) works in the random oracle model, assuming only the hardness of the discrete logarithm problem.

The extractor obtains $2n$ polynomials $s^{(j)}$ for the same evaluation point z^* but distinct challenges α_j , where n is the number of evaluation claims. The checks in the reduction procedure imply that $s^{(j)} = \sum_i \alpha_j^i p_i + \sum_i \alpha_j^{n+i} w_i$, where w_i is the witness corresponding to p_i ; hence we can recover the p_i, w_i by solving a linear system. We then use a variant of the zero-finding game lemma from [BCMS20] (see

Section 6.3.3) to show that if a particular polynomial equation on p_i, w_i holds at the point z^* obtained from the random oracle, it must with overwhelming probability be an identity. Applying this to the equation induced by the reduction shows that the original claims hold with high probability.

The proof crucially relies on the ability to efficiently generate the $2n$ related proofs. For this we formulate a new expected-time forking lemma in the random oracle model, which is informally stated below (and detailed in Section 6.3.2). For this application, the set L will be valid accumulators and proofs.

Lemma 1 (informal). *Let L be an efficiently recognizable set. Suppose there exists an expected polynomial time algorithm with access to a random oracle ρ that outputs, with probability δ , a tuple $(q, \alpha, y) \in L$ such that $\rho(q) = \alpha$. Then there exists an expected polynomial time extractor E such that for $N = \text{poly}(\lambda)$, E outputs N related tuples $(q, \alpha_i, y_i) \in L$ for distinct α_i with probability at least $\delta - \text{negl}(\lambda)$.*

This forking lemma differs from prior forking lemmas in three significant ways. First, it is in the random oracle model rather than the interactive setting (unlike [BCCGP16]). Second, we can obtain any polynomial number of accepting transcripts in polynomial time with only negligible loss in success probability (unlike [BN06]). Finally, it holds even if the adversary itself runs in expected (as opposed to strict) polynomial time. This is important for our application to PCD where the extractor of one recursive step becomes the adversary of the next recursive step.

2.5 Implementation and evaluation

We elaborate on our implementation and evaluation of accumulation schemes and their application to PCD.

The case for a PCD framework. Different PCD constructions offer different trade-offs. The tradeoffs are both about asymptotics (see Remark 2.3) and about practical concerns, as we review below.

- *PCD from sublinear verification* [BCCT13; BCTV14; COS20] is typically instantiated via preprocessing SNARKs based on pairings.³ This route offers excellent verifier time (a few milliseconds regardless of the computation at a PCD node) but requires a circuit-specific trusted setup (which complicates deployment) and cycles of pairing-friendly elliptic curves (which are costly in terms of group arithmetic and size).
- *PCD from atomic accumulation* [BCMS20] can, e.g., be instantiated via SNARKs based on cyclic groups [BGH19]. This route offers a transparent setup (easy to deploy) and logarithmic-size arguments (a few kilobytes even for large computations), using cycles of standard elliptic curves (more efficient than their pairing-friendly counterparts). On the other hand, this route yields linear verification times (expensive for large computations) and logarithmic costs for accumulation (increasing the cost of recursion).
- *PCD from split accumulation* (this work) can, e.g., be instantiated via NARKs based on cyclic groups. This route still offers a transparent setup and allows using cycles of standard elliptic curves. Moreover, it offers constant costs for accumulation, but at the expense of argument size, which is now linear.

It would be desirable to have a *single framework that supports different PCD constructions via a modular composition of simpler building blocks*. It would allow for replacing older building blocks with new ones. It would enable prototyping different PCD constructions for different applications (which may have different needs) and thereby enable practitioners to make informed choices about which PCD construction is best for them. It would facilitate auditing of what would otherwise be a complex cryptographic system with many intermixed layers. (Realizing even a single PCD construction is a substantial implementation task.) Finally, it would be useful for applications to be separated from the underlying recursion via a common PCD interface.

³Instantiations based hashes are also possible [COS20] but are (post-quantum and) less efficient.

Implementation (Section 7). The foregoing considerations motivated our implementation efforts for PCD. Our code base has two main parts, one for realizing accumulation schemes and another for realizing PCD from accumulation (which we integrated with PCD from succinct verification).

- *Framework for accumulation.* We designed a modular framework for (atomic and split) accumulation schemes, and use it to provide a common interface for accumulation schemes for popular polynomial commitments (and thus for SNARKs that use these), including: (a) the atomic accumulation scheme AS_{AGM} in [BCMS20] for the PC scheme PC_{AGM} ; (b) the atomic accumulation scheme AS_{IPA} in [BCMS20] for the PC scheme PC_{IPA} ; (c) the split accumulation scheme AS_{HC} in this paper for PC_{HC} . Our framework also provides a generic method for defining RICS constraints for the verifiers of these accumulation schemes; we leverage this to implement RICS constraints for all these accumulation schemes.
- *PCD from accumulation.* We use the foregoing framework to implement a generic construction of PCD from accumulation. We support the PCD construction of [BCMS20] (which uses atomic accumulation) and the PCD construction in this paper (which uses split accumulation). Our code builds on and extends an existing PCD library.⁴ Our implementation is modular: it takes as ingredients an implementation of any NARK, an implementation of any accumulation scheme for that NARK, and constraints for the accumulation verifier, and produces a concrete PCD construction. We use our concrete instantiations of these ingredients to construct PCD that is based on atomic accumulation of PC_{AGM} and PC_{IPA} , and split accumulation of PC_{HC} .

Evaluation for DL setting (Section 8). When realizing PCD in practice the main goal is to “minimize the cost of recursion”, that is, to minimize the number of constraints that need to be recursively proved in each PCD step (excluding the constraints for the application) without hurting other parameters too much (prover time, argument size, and so on). We evaluate our implementation with respect to this goal, with a focus on understanding the trade-offs between atomic and split accumulation *in the discrete logarithm setting*.

The DL setting is of particular interest to practitioners, as it leads to systems with a transparent (public-coin) setup that can be based on efficient cycles of (standard) elliptic curves [BGH19; Hop20]; indeed, some projects are developing real-world systems that use PCD in the DL setting [Halo20; Pickles20]. The main drawback of the DL setting is that verification time (and sometimes argument size) is linear in a PCD node’s computation. This inefficiency is, however, tolerable if a PCD node’s computation is not too large, as is the case in the aforementioned projects. (Especially so when taking into account the disadvantages of PCD based on pairings, which involves relying on a circuit-specific trusted setup and more expensive curve cycles.)

We compare two routes for PCD in the DL setting:

- recursion based on the atomic accumulation scheme for the PC scheme PC_{IPA} [BCMS20];
- recursion based on the split accumulation scheme for PC_{HC} (Section 6).

Our evaluation focuses on these polynomial commitment schemes and their corresponding accumulation schemes, as they comprise the primary cost of recursion in popular accumulation-friendly (S)NARKs, and demonstrates that the constraint cost of the AS_{HC} accumulation verifier is 8 to 20 times cheaper than that of the AS_{IPA} accumulation verifier.

We emphasize that the key aspect to focus on is the relative costs of the two components. This is because the cost of both components is dominated by the cost of many common subcomponents, and so improvements in these subcomponents will preserve the relative cost. For example, applying existing techniques [Halo20; Pickles20] for optimizing scalar multiplications should benefit both schemes in a similar way.

⁴<https://github.com/arkworks-rs/pcd>

3 Preliminaries

Indexed relations. An *indexed relation* \mathcal{R} is a set of triples $(\mathfrak{i}, \mathfrak{x}, \mathfrak{w})$ where \mathfrak{i} is the index, \mathfrak{x} is the instance, and \mathfrak{w} is the witness; the corresponding *indexed language* $\mathcal{L}(\mathcal{R})$ is the set of pairs $(\mathfrak{i}, \mathfrak{x})$ for which there exists a witness \mathfrak{w} such that $(\mathfrak{i}, \mathfrak{x}, \mathfrak{w}) \in \mathcal{R}$. For example, the indexed relation of satisfiable boolean circuits consists of triples where \mathfrak{i} is the description of a boolean circuit, \mathfrak{x} is a partial assignment to its input wires, and \mathfrak{w} is an assignment to the remaining wires that makes the circuit to output 0.

Security parameters. For simplicity of notation, we assume that all public parameters have length at least λ , so that algorithms which receive such parameters can run in time $\text{poly}(\lambda)$.

Random oracles. We denote by $\mathcal{U}(\lambda)$ the set of all functions that map $\{0, 1\}^*$ to $\{0, 1\}^\lambda$. We denote by $\mathcal{U}(\ast)$ the set $\bigcup_{\lambda \in \mathbb{N}} \mathcal{U}(\lambda)$. A *random oracle* with security parameter λ is a function $\rho: \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ sampled uniformly at random from $\mathcal{U}(\lambda)$.

3.1 Non-interactive arguments in the ROM

A tuple of algorithms $\text{ARG} = (\mathcal{G}, \mathcal{I}, \mathcal{P}, \mathcal{V})$ is a (preprocessing) *non-interactive argument* in the random oracle model (ROM) for an indexed relation \mathcal{R} if the following properties hold.

- **Completeness.** For every adversary \mathcal{A} ,

$$\Pr \left[\begin{array}{c|c} (\mathfrak{i}, \mathfrak{x}, \mathfrak{w}) \notin \mathcal{R} & \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow \mathcal{G}^\rho(1^\lambda) \\ (\mathfrak{i}, \mathfrak{x}, \mathfrak{w}) \leftarrow \mathcal{A}^\rho(\text{pp}) \\ (\text{ipk}, \text{ivk}) \leftarrow \mathcal{I}^\rho(\text{pp}, \mathfrak{i}) \\ \pi \leftarrow \mathcal{P}^\rho(\text{ipk}, \mathfrak{x}, \mathfrak{w}) \end{array} \\ \vee \\ \mathcal{V}^\rho(\text{ivk}, \mathfrak{x}, \pi) = 1 & \end{array} \right] = 1 .$$

- **Soundness.** For every polynomial-size adversary $\tilde{\mathcal{P}}$,

$$\Pr \left[\begin{array}{c|c} (\mathfrak{i}, \mathfrak{x}) \notin \mathcal{L}(\mathcal{R}) & \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow \mathcal{G}^\rho(1^\lambda) \\ (\mathfrak{i}, \mathfrak{x}, \pi) \leftarrow \tilde{\mathcal{P}}^\rho(\text{pp}) \\ (\text{ipk}, \text{ivk}) \leftarrow \mathcal{I}^\rho(\text{pp}, \mathfrak{i}) \end{array} \\ \wedge \\ \mathcal{V}^\rho(\text{ivk}, \mathfrak{x}, \pi) = 1 & \end{array} \right] \leq \text{negl}(\lambda) .$$

The above formulation of completeness allows $(\mathfrak{i}, \mathfrak{x}, \mathfrak{w})$ to depend on the random oracle ρ and public parameters pp , and the above formulation of soundness allows $(\mathfrak{i}, \mathfrak{x})$ to depend on the random oracle ρ and public parameters pp .

Our PCD construction makes use of the stronger property of *knowledge soundness*, and optionally also the property of (statistical) *zero knowledge*. We define both of these properties below. Note that this definition is stronger the standard definition of knowledge soundness; this is required to prove post-quantum security in Theorem 5.2. This stronger definition is similar to the notion of *witness-extended emulation* [Lin03]. We refer to an argument with knowledge soundness as a NARK (non-interactive argument of knowledge) whereas an argument that just satisfies soundness is a NARG.

Knowledge soundness. We say that $\text{ARG} = (\mathcal{G}, \mathcal{I}, \mathcal{P}, \mathcal{V})$ has *knowledge soundness* (with respect to auxiliary input distribution \mathcal{D}) if for every expected polynomial time adversary $\tilde{\mathcal{P}}$ there exists an expected

polynomial time extractor \mathcal{E} such that for every set Z ,

$$\Pr \left[\begin{array}{l} (\text{pp}, \text{ai}, \vec{\mathbf{i}}, \vec{\mathbf{x}}, \text{ao}) \in Z \\ \wedge \forall j \in [\ell], (\mathbf{i}_j, \mathbf{x}_j, \mathbf{w}_j) \in \mathcal{R} \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow \mathcal{G}(1^\lambda) \\ \text{ai} \leftarrow \mathcal{D}(\text{pp}) \\ (\vec{\mathbf{i}}, \vec{\mathbf{x}}, \vec{\mathbf{w}}, \text{ao}) \leftarrow \mathcal{E}_{\tilde{\mathcal{P}}}(\text{pp}, \text{ai}) \end{array} \right] \\ \geq \Pr \left[\begin{array}{l} (\text{pp}, \text{ai}, \vec{\mathbf{i}}, \vec{\mathbf{x}}, \text{ao}) \in Z \\ \wedge \forall j \in [\ell], \mathcal{V}(\text{ivk}_j, \mathbf{x}_j, \pi_j) = 1 \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow \mathcal{G}(1^\lambda) \\ \text{ai} \leftarrow \mathcal{D}(\text{pp}) \\ (\vec{\mathbf{i}}, \vec{\mathbf{x}}, \vec{\pi}, \text{ao}) \leftarrow \tilde{\mathcal{P}}^\rho(\text{pp}, \text{ai}) \\ \forall j \in [\ell], (\text{ipk}_j, \text{ivk}_j) \leftarrow \mathcal{I}^\rho(\text{pp}, \mathbf{i}_j) \end{array} \right] - \text{negl}(\lambda) .$$

Zero knowledge. We say that $\text{ARG} = (\mathcal{G}, \mathcal{I}, \mathcal{P}, \mathcal{V})$ has (statistical) zero knowledge if there exists a probabilistic polynomial-time simulator \mathcal{S} such that for every polynomial-size honest adversary \mathcal{A} the distributions below are computationally indistinguishable:

$$\left\{ (\rho, \text{pp}, \mathbf{i}, \mathbf{x}, \pi) \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow \mathcal{G}^\rho(1^\lambda) \\ (\mathbf{i}, \mathbf{x}, \mathbf{w}) \leftarrow \mathcal{A}^\rho(\text{pp}) \\ (\text{ipk}, \text{ivk}) \leftarrow \mathcal{I}^\rho(\text{pp}, \mathbf{i}) \\ \pi \leftarrow \mathcal{P}^\rho(\text{ipk}, \mathbf{x}, \mathbf{w}) \end{array} \right\} \text{ and } \left\{ (\rho[\mu], \text{pp}, \mathbf{i}, \mathbf{x}, \pi) \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ (\text{pp}, \tau) \leftarrow \mathcal{S}^\rho(\mathbf{i}, \mathbf{x}) \\ (\mathbf{i}, \mathbf{x}, \mathbf{w}) \leftarrow \mathcal{A}^\rho(\text{pp}) \\ (\pi, \mu) \leftarrow \mathcal{S}^\rho(\text{pp}, \mathbf{i}, \mathbf{x}, \tau) \end{array} \right\} .$$

Above, $\rho[\mu]$ is the function that, on input x , equals $\mu(x)$ if μ is defined on x , or $\rho(x)$ otherwise. This definition uses explicitly-programmable random oracles [BR93]. (Non-interactive zero knowledge with non-programmable random oracles is impossible for non-trivial languages [Pas03; BCS16].)

3.2 Proof-carrying data

A triple of algorithms $\text{PCD} = (\mathbb{G}, \mathbb{I}, \mathbb{P}, \mathbb{V})$ is a (preprocessing) *proof-carrying data scheme* (PCD scheme) for a class of compliance predicates \mathbb{F} if the properties below hold.

Definition 3.1. A **transcript** \mathbb{T} is a directed acyclic graph where each vertex $u \in V(\mathbb{T})$ is labeled by local data $z_{\text{loc}}^{(u)}$ and each edge $e \in E(\mathbb{T})$ is labeled by a message $z^{(e)} \neq \perp$. The **output** of a transcript \mathbb{T} , denoted $\text{o}(\mathbb{T})$, is $z^{(e)}$ where $e = (u, v)$ is the lexicographically-first edge such that v is a sink.

Definition 3.2. A vertex $u \in V(\mathbb{T})$ is φ -**compliant** for $\varphi \in \mathbb{F}$ if for all outgoing edges $e = (u, v) \in E(\mathbb{T})$:

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

We say that \mathbb{T} is φ -**compliant** if all of its vertices are φ -compliant.

Completeness. For every adversary \mathcal{A} ,

$$\Pr \left[\begin{array}{l} \left(\varphi \in \mathbb{F} \wedge (\forall i, z_i = \perp \vee \forall i, \mathbb{V}(\text{ivk}, z_i, \pi_i) = 1) \wedge \right. \\ \quad \left. \varphi(z, z_{\text{loc}}, z_1, \dots, z_m) \text{ accepts} \right) \\ \quad \downarrow \\ \quad \mathbb{V}(\text{ivk}, z, \pi) = 1 \end{array} \middle| \begin{array}{l} \mathbb{PP} \leftarrow \mathbb{G}(1^\lambda) \\ (\varphi, z, z_{\text{loc}}, [z_i, \pi_i]_{i=1}^m) \leftarrow \mathcal{A}(\mathbb{PP}) \\ (\text{ipk}, \text{ivk}) \leftarrow \mathbb{I}(\mathbb{PP}, \varphi) \\ \pi \leftarrow \mathbb{P}(\text{ipk}, z, z_{\text{loc}}, [z_i, \pi_i]_{i=1}^m) \end{array} \right] = 1 .$$

Knowledge soundness. We say that $\text{PCD} = (\mathbb{G}, \mathbb{I}, \mathbb{P}, \mathbb{V})$ has knowledge soundness (with respect to auxiliary input distribution \mathcal{D}) if for every expected polynomial-time (non-uniform) adversary $\tilde{\mathbb{P}}$, there exists an expected polynomial-time extractor $\mathbb{E}_{\tilde{\mathbb{P}}}$ such that for every set Z ,

$$\Pr \left[\begin{array}{l} \varphi \in \mathbb{F} \wedge (\mathbb{P}\mathbb{P}, \text{ai}, \varphi, \text{o}(\mathbb{T}), \text{ao}) \in Z \\ \wedge \mathbb{T} \text{ is } \varphi\text{-compliant} \end{array} \middle| \begin{array}{l} \mathbb{P}\mathbb{P} \leftarrow \mathbb{G}(1^\lambda) \\ \text{ai} \leftarrow \mathcal{D}(\mathbb{P}\mathbb{P}) \\ (\varphi, \mathbb{T}, \text{ao}) \leftarrow \mathbb{E}_{\tilde{\mathbb{P}}}(\mathbb{P}\mathbb{P}, \text{ai}) \end{array} \right] \\ \geq \Pr \left[\begin{array}{l} \varphi \in \mathbb{F} \wedge (\mathbb{P}\mathbb{P}, \text{ai}, \varphi, \text{o}, \text{ao}) \in Z \\ \wedge \mathcal{V}(\text{ivk}, \text{o}, \pi) = 1 \end{array} \middle| \begin{array}{l} \mathbb{P}\mathbb{P} \leftarrow \mathbb{G}(1^\lambda) \\ \text{ai} \leftarrow \mathcal{D}(\mathbb{P}\mathbb{P}) \\ (\varphi, \text{o}, \pi, \text{ao}) \leftarrow \tilde{\mathbb{P}}(\mathbb{P}\mathbb{P}, \text{ai}) \\ (\text{ipk}, \text{ivk}) \leftarrow \mathbb{I}(\mathbb{P}\mathbb{P}, \varphi) \end{array} \right] - \text{negl}(\lambda) .$$

Zero knowledge. We say that $\text{PCD} = (\mathbb{G}, \mathbb{I}, \mathbb{P}, \mathbb{V})$ has (statistical) zero knowledge if there exists a probabilistic polynomial-time simulator \mathbb{S} such that for all honest adversaries \mathcal{A} the distributions below are statistically close:

$$\left\{ (\mathbb{P}\mathbb{P}, \varphi, z, \pi) \middle| \begin{array}{l} \mathbb{P}\mathbb{P} \leftarrow \mathbb{G}(1^\lambda) \\ (\varphi, z, z_{\text{loc}}, [z_i, \pi_i]_{i=1}^m) \leftarrow \mathcal{A}(\mathbb{P}\mathbb{P}) \\ (\text{ipk}, \text{ivk}) \leftarrow \mathbb{I}(\mathbb{P}\mathbb{P}, \varphi) \\ \pi \leftarrow \mathbb{P}(\text{ipk}, \varphi, z, z_{\text{loc}}, [z_i, \pi_i]_{i=1}^m) \end{array} \right\} \text{ and } \left\{ (\mathbb{P}\mathbb{P}, \varphi, z, \pi) \middle| \begin{array}{l} (\mathbb{P}\mathbb{P}, \tau) \leftarrow \mathbb{S} \\ (\varphi, z, z_{\text{loc}}, [z_i, \pi_i]_{i=1}^m) \leftarrow \mathcal{A}(\mathbb{P}\mathbb{P}) \\ \pi \leftarrow \mathbb{S}(\varphi, z, \tau) \end{array} \right\} .$$

Here, an adversary is honest if its output satisfies the implicant of the completeness condition with probability 1, namely: $\varphi \in \mathbb{F}$, $\varphi(z, z_{\text{loc}}, z_1, \dots, z_m) = 1$, and either for all i , $z_i = \perp$, or for all i , $\mathbb{V}(\text{ivk}, z_i, \pi_i) = 1$.

Efficiency. The generator \mathbb{G} , prover \mathbb{P} , indexer \mathbb{I} and verifier \mathbb{V} run in polynomial time. A proof π has size $\text{poly}(\lambda, |\varphi|)$; in particular, it is not permitted to grow with each application of \mathbb{P} .

3.3 Instantiating the random oracle

Almost all of the results in this paper are proved in the *random oracle model*, and so we give definitions which include random oracles. The single exception is our construction of proof-carrying data, in Section 5.1. We do not know how to build PCD schemes which are secure in the random oracle model from any standard assumption. Instead, we show that assuming the existence of a non-interactive argument with security in the standard (CRS) model, we obtain a PCD scheme which is also secure in the standard (CRS) model.

For this reason, the definition of PCD above is stated in the standard model (without oracles). We do not explicitly define non-interactive arguments in the standard model; the definition is easily obtained by removing the random oracle from the definition presented in Section 3.1.

3.4 Post-quantum security

The definitions of both non-interactive arguments (in the standard model) and proof-carrying data can be strengthened, in a straightforward way, to express post-quantum security. In particular, we replace “polynomial-size circuit” and “polynomial-time algorithm” with their quantum analogues. Since we do not prove post-quantum security of any construction in the random oracle model, we do not discuss the quantum random oracle model.

3.5 Commitment schemes

A commitment scheme $\text{CM} = (\text{Setup}, \text{Trim}, \text{Commit})$ enables one to create binding commitments to messages.

- CM.Setup , on input a *message format* L , outputs public parameters pp ; this specifies a message universe \mathcal{M}_{pp} and a commitment universe \mathcal{C}_{pp} .
- CM.Trim , on input public parameters pp and a *trim specification* ℓ , outputs a commitment key ck containing a description of a message space $\mathcal{M}_{\text{ck}} \subseteq \mathcal{M}_{\text{pp}}$ (corresponding to ℓ).
- CM.Commit , on input a commitment key ck , a message $m \in \mathcal{M}_{\text{ck}}$ and randomness ω , outputs a commitment $C \in \mathcal{C}_{\text{pp}}$.

CM is binding if, for every message format L such that $|L| = \text{poly}(\lambda)$, and for every efficient adversary \mathcal{A} , the following holds.

$$\Pr \left[\begin{array}{c} m_1 \in \mathcal{M}_{\text{ck}_1}, m_2 \in \mathcal{M}_{\text{ck}_2}, m_1 \neq m_2 \\ \wedge \\ \text{CM.Commit}(\text{ck}_1, m_1; \omega_1) = \text{CM.Commit}(\text{ck}_2, m_2; \omega_2) \end{array} \middle| \begin{array}{c} \text{pp} \leftarrow \text{CM.Setup}^\rho(1^\lambda, L) \\ ((\ell_1, m_1, \omega_1), (\ell_2, m_2, \omega_2)) \leftarrow \mathcal{A}^\rho(\text{pp}) \\ \text{ck}_1 \leftarrow \text{CM.Trim}^\rho(\text{pp}, \ell_1) \\ \text{ck}_2 \leftarrow \text{CM.Trim}^\rho(\text{pp}, \ell_2) \end{array} \right] \leq \text{negl}(\lambda) .$$

Note that $m_1 \neq m_2$ is well-defined since $\mathcal{M}_{\text{ck}_1}, \mathcal{M}_{\text{ck}_2} \subseteq \mathcal{M}_{\text{pp}}$.

3.6 Polynomial commitments

A polynomial commitment scheme is a cryptographic primitive that enables a sender to commit to a polynomial p over a field \mathbb{F} and then later prove the correct evaluation of the polynomial at a desired point. In more detail, a polynomial commitment scheme PC is a tuple of algorithms (Setup , Trim , Commit , Open , Check) with the following syntax and properties:

- $\text{PC.Setup}(1^\lambda, D) \rightarrow \text{pp}$. On input a security parameter λ (in unary), and a maximum degree bound $D \in \mathbb{N}$, PC.Setup samples public parameters pp_{PC} . The parameters contain the description of a finite field \mathbb{F} (which has size that is super-polynomial in λ).
- $\text{PC.Trim}(\text{pp}, d) \rightarrow (\text{ck}, \text{rk})$. On input public parameters pp_{PC} , and supported degree d , PC.Trim deterministically computes a key pair (ck, rk) that is specialized to d .
- $\text{PC.Commit}(\text{ck}, p; \omega) \rightarrow C$. On input ck , a univariate polynomial p over the field \mathbb{F} such that $\deg(p) \leq \text{ck}.d$, PC.Commit outputs a commitment C to the polynomial p . The randomness ω is used if the commitment C is hiding.
- $\text{PC.Open}(\text{ck}, p, C, z; \omega) \rightarrow \pi$. On input the commitment key ck , a univariate polynomial p over the field \mathbb{F} , a commitment C to p , an evaluation point z , and commitment randomness ω , PC.Open outputs an evaluation proof π .
- $\text{PC.Check}(\text{rk}, C, z, v, \pi) \rightarrow b$. On input the receiver key rk , a commitment C , an evaluation point z , a claimed evaluation v , and an evaluation proof π , PC.Check if π attests that the polynomial p committed in C has degree at most $\text{rk}.d$, and evaluates to v at z .

A polynomial commitment scheme PC must be such that $(\text{PC.Setup}, \text{PC.Trim}, \text{PC.Commit})$ is a (binding) commitment scheme for bounded-degree polynomials over a field. The message format L is equal to the maximum degree bound D ; the message universe is the set of polynomials over some field \mathbb{F} of degree at most D . The trim specification ℓ is equal to the supported degree d ; the corresponding message space is the set of polynomials over \mathbb{F} of degree at most d .

In this work we will not use any property of the PC.Open and PC.Check algorithms; their inclusion here is merely for consistency with previous definitions.

4 Split accumulation schemes for relations

Let $\Phi: (\{0, 1\}^*)^3 \rightarrow \{0, 1\}$ be a relation predicate. Let \mathcal{H} be a randomized oracle algorithm, which outputs predicate parameters pp_Φ .

An **accumulation scheme** for (Φ, \mathcal{H}) is a tuple of algorithms $\text{AS} = (G, I, P, V, D)$ of which I, P, V have access to the same random oracle ρ . The algorithms have the following syntax and properties.

Syntax. The algorithms comprising AS have the following syntax:

- *Generator:* On input a security parameter λ (in unary), G samples and outputs public parameters pp .
- *Indexer:* On input public parameters pp , predicate parameters pp_Φ (generated by \mathcal{H}), and a predicate index i_Φ , I deterministically computes and outputs a triple $(\text{apk}, \text{avk}, \text{dk})$ consisting of an accumulator proving key apk , an accumulator verification key avk , and a decision key dk .⁵
- *Accumulation prover:* On input the accumulator proving key apk , inputs $[q_i = (q.\text{x}_i, q.\text{w}_i)]_{i=1}^n$, and old accumulators $[\text{acc}_j = (\text{acc}.\text{x}_j, \text{acc}.\text{w}_j)]_{j=1}^m$, P outputs a new accumulator $\text{acc} = (\text{acc}.\text{x}, \text{acc}.\text{w})$ and a proof π_V for the accumulation verifier.
- *Accumulation verifier:* On input the accumulator verification key avk , input instances $[q_i.\text{x}]_{i=1}^n$, accumulator instances $[\text{acc}_j.\text{x}]_{j=1}^m$, a new accumulator instance $\text{acc}.\text{x}$, and a proof π_V , V outputs a bit indicating whether $\text{acc}.\text{x}$ correctly accumulates $[q_i]_{i=1}^n$ and $[\text{acc}_j]_{j=1}^m$.
- *Decider:* On input the decision key dk , and an accumulator $\text{acc} = (\text{acc}.\text{x}, \text{acc}.\text{w})$, D outputs a bit indicating whether acc is a valid accumulator.

These algorithms must satisfy two properties, *completeness* and *knowledge soundness*, defined below. We additionally define a notion of zero knowledge that we will rely on to achieve zero knowledge PCD (see Section 5).

Completeness. For every (unbounded) adversary \mathcal{A} ,

$$\Pr \left[\begin{array}{l} \forall j \in [m], D(\text{dk}, \text{acc}_j) = 1 \\ \forall i \in [n], \Phi(\text{pp}_\Phi, i_\Phi, q_i) = 1 \\ \Downarrow \\ V^\rho(\text{avk}, [q_i.\text{x}]_{i=1}^n, [\text{acc}_j.\text{x}]_{j=1}^m, \text{acc}.\text{x}, \pi_V) = 1 \\ D(\text{dk}, \text{acc}) = 1 \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow G(1^\lambda) \\ \text{pp}_\Phi \leftarrow \mathcal{H}(1^\lambda) \\ (i_\Phi, [q_i]_{i=1}^n, [\text{acc}_j]_{j=1}^m) \leftarrow \mathcal{A}^\rho(\text{pp}, \text{pp}_\Phi) \\ (\text{apk}, \text{avk}, \text{dk}) \leftarrow I^\rho(\text{pp}, \text{pp}_\Phi, i_\Phi) \\ (\text{acc}, \pi_V) \leftarrow P^\rho(\text{apk}, [q_i]_{i=1}^n, [\text{acc}_j]_{j=1}^m) \end{array} \right] = 1 .$$

Note that for $m = n = 0$, the precondition on the left-hand side holds vacuously; this is required for the completeness condition to be non-trivial.

Knowledge soundness. We say that $\text{AS} = (G, I, P, V, D)$ has *knowledge error* $k(\lambda)$ if for every (non-uniform) adversary $\tilde{\mathcal{P}}$ running in expected polynomial time there exists an extractor E running in expected

⁵We remark that in some schemes it is important, for the sake of efficiency, for the indexer I to have oracle access to the predicate parameters pp_Φ and predicate index i_Φ , rather than reading them in full. All of our constructions and statements extend, in a straightforward way, to this case.

polynomial time such that for every set Z :

$$\Pr \left[\begin{array}{l} (\text{pp}, \text{pp}_\Phi, \text{ai}, i_\Phi, \text{acc}, [\text{q}_i.\mathbb{x}]_{i=1}^n, [\text{acc}_j.\mathbb{x}]_{j=1}^m, \text{ao}) \in Z \\ \wedge \\ \forall j \in [m], D(\text{dk}, \text{acc}_j) = 1 \\ \forall i \in [n], \Phi(\text{pp}_\Phi, i_\Phi, \text{q}_i) = 1 \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow G(1^\lambda) \\ \text{pp}_\Phi \leftarrow \mathcal{H}(1^\lambda) \\ \text{ai} \leftarrow \mathcal{D}(1^\lambda) \\ \left(\begin{array}{ccc} i_\Phi & \text{acc} & \text{ao} \\ [\text{q}_i]_{i=1}^n & [\text{acc}_j]_{j=1}^m \end{array} \right) \leftarrow E_{\tilde{\mathcal{P}}}(\text{pp}, \text{pp}_\Phi, \text{ai}) \end{array} \right] \\ \geq \Pr \left[\begin{array}{l} (\text{pp}, \text{pp}_\Phi, \text{ai}, i_\Phi, \text{acc}, [\text{q}_i.\mathbb{x}]_{i=1}^n, [\text{acc}_j.\mathbb{x}]_{j=1}^m, \text{ao}) \in Z \\ \wedge \\ V^\rho(\text{avk}, [\text{q}_i.\mathbb{x}]_{i=1}^n, [\text{acc}_j.\mathbb{x}]_{j=1}^m, \text{acc}.\mathbb{x}, \pi_V) = 1 \\ D(\text{dk}, \text{acc}) = 1 \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow G(1^\lambda) \\ \text{pp}_\Phi \leftarrow \mathcal{H}^\rho(1^\lambda) \\ \text{ai} \leftarrow \mathcal{D}(1^\lambda) \\ \left(\begin{array}{ccc} i_\Phi & \text{acc} & \pi_V & \text{ao} \\ [\text{q}_i.\mathbb{x}]_{i=1}^n & [\text{acc}_j.\mathbb{x}]_{j=1}^m \end{array} \right) \leftarrow \tilde{\mathcal{P}}^\rho(\text{pp}, \text{pp}_\Phi, \text{ai}) \\ (\text{apk}, \text{avk}, \text{dk}) \leftarrow I^\rho(\text{pp}, \text{pp}_\Phi, i_\Phi) \end{array} \right] - k(\lambda) .$$

If $k(\lambda)$ is negligible then we say AS satisfies knowledge soundness.

Zero knowledge. There exists a polynomial-time simulator S such that for every polynomial-size ‘‘honest’’ adversary \mathcal{A} (see below) the following distributions are (statistically/computationally) indistinguishable:

$$\left\{ \begin{array}{l} (\rho, \text{pp}, \text{pp}_\Phi, i_\Phi, \text{acc}) \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow G(1^\lambda) \\ \text{pp}_\Phi \leftarrow \mathcal{H}(1^\lambda) \\ (i_\Phi, [\text{q}_i]_{i=1}^n, [\text{acc}_j]_{j=1}^m) \leftarrow \mathcal{A}^\rho(\text{pp}, \text{pp}_\Phi) \\ (\text{apk}, \text{avk}, \text{dk}) \leftarrow I^\rho(\text{pp}, \text{pp}_\Phi, i_\Phi) \\ (\text{acc}, \pi_V) \leftarrow P^\rho(\text{apk}, [\text{q}_i]_{i=1}^n, [\text{acc}_j]_{j=1}^m) \end{array} \right\} \\ \text{and} \\ \left\{ \begin{array}{l} (\rho[\mu], \text{pp}, \text{pp}_\Phi, i_\Phi, \text{acc}) \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ (\text{pp}, \tau) \leftarrow S^\rho(1^\lambda) \\ \text{pp}_\Phi \leftarrow \mathcal{H}(1^\lambda) \\ (i_\Phi, [\text{q}_i]_{i=1}^n, [\text{acc}_j]_{j=1}^m) \leftarrow \mathcal{A}^\rho(\text{pp}, \text{pp}_\Phi) \\ (\text{acc}, \mu) \leftarrow S^\rho(\text{pp}_\Phi, \tau, i_\Phi) \end{array} \right\} .$$

Here \mathcal{A} is *honest* if it outputs, with probability 1, a tuple $(i_\Phi, [\text{q}_i]_{i=1}^n, [\text{acc}_j]_{j=1}^m)$ such that $\Phi(\text{pp}_\Phi, i_\Phi, \text{q}_i) = 1$ and $D(\text{dk}, \text{acc}_j) = 1$ for all $i \in [n]$ and $j \in [m]$. Note that the simulator S is *not* required to simulate the accumulation verifier proof π_V .

4.1 Accumulation schemes for certain predicates

We conclude by specializing the definition of an accumulation scheme to the case of predicates induced by the verifier in a non-interactive argument (Definition 4.1) and in a polynomial commitment scheme (Definition 4.2).

Definition 4.1 (accumulation for ARG). *A non-interactive argument system $\text{ARG} = (\mathcal{G}, \mathcal{I}, \mathcal{P}, \mathcal{V})$ has an accumulation scheme if the pair $(\Phi_V, \mathcal{H}_{\text{ARG}} := \mathcal{G})$ has an accumulation scheme, where Φ_V is defined below:*

- $\Phi_V(\text{pp}_\Phi = \text{pp}, i_\Phi = \mathbf{i}, \text{q}.\mathbb{x} = (\mathbb{x}, \pi.\mathbb{x}), \text{q}.\mathbb{w} = \pi.\mathbb{w})$:
1. $(\text{ipk}, \text{ivk}) \leftarrow \mathcal{I}(\text{pp}, \mathbf{i})$.
 2. *Output* $\mathcal{V}(\text{ivk}, \mathbb{x}, (\pi.\mathbb{x}, \pi.\mathbb{w}))$.

Definition 4.2 (accumulation for PC). *A polynomial commitment scheme $PC = (\text{Setup}, \text{Trim}, \text{Commit}, \text{Open}, \text{Check})$ has an accumulation scheme if, for every $D(\lambda) = \text{poly}(\lambda)$, the pair $(\Phi_{PC}, \mathcal{H}_{PC,D})$ defined below has an accumulation scheme.*

$\Phi_{PC}(\text{pp}_\Phi = \text{pp}_{PC}, i_\Phi = d, q = ((C, z, v), \pi)):$ $\mathcal{H}_{PC,D}(1^\lambda):$
 1. $(\text{ck}, \text{rk}) \leftarrow PC.\text{Trim}(\text{pp}_{PC}, d).$ *Output* $\text{pp}_{PC} \leftarrow PC.\text{Setup}(1^\lambda, D(\lambda)).$
 2. *Output* $PC.\text{Check}(\text{rk}, C, z, v, \pi).$

5 PCD from arguments of knowledge with split accumulation

We formally restate and then prove Theorem 1, which provides a construction of proof-carrying data (PCD) from any NARK that has a split accumulation scheme with certain efficiency properties.

First, we provide some notation for these properties.

Definition 5.1. Let $AS = (G, I, P, V, D)$ be an accumulation scheme for a non-interactive argument (see Definition 4.1). We denote by $V^{(\lambda, m, N, k)}$ the circuit corresponding to the computation of the accumulation verifier V , for security parameter λ , when checking the accumulation of m instance-proof pairs and accumulators, on an index of size at most N , where each instance is of size at most k .

We denote by $v(\lambda, m, N, k)$ the size of the circuit $V^{(\lambda, m, N, k)}$, by $|\text{avk}(\lambda, m, N)|$ the size of the accumulator verification key avk , and by $|\text{acc.x}(\lambda, m, N)|$ the size of an accumulator instance.

Note that here we have specified that the size of acc.x is bounded by a function of λ, m, N ; in particular, it may not depend on the number of instances accumulated.

When we invoke the accumulation verifier in our construction of PCD, an instance will consist of an accumulator verification key, an accumulator *instance*, and some additional data of size ℓ . Thus the size of the accumulation verifier circuit used in the scheme is given by

$$v^*(\lambda, m, N, \ell) := v(\lambda, m, N, |\text{avk}(\lambda, m, N)| + |\text{acc.x}(\lambda, m, N)| + \ell) .$$

The notion of “sublinear verification” which is important here is that v^* is sublinear in N . The following theorem shows that when this is the case, this accumulation scheme can be used to construct PCD.

Theorem 5.2. *There exists a polynomial-time transformation T such that if $ARG = (G, \mathcal{I}, \mathcal{P}, \mathcal{V})$ is a NARK for circuit satisfiability and AS is a split accumulation scheme for ARG then $PCD = (G, \mathbb{I}, \mathbb{P}, \mathbb{V}) := T(ARG, AS)$ is a PCD scheme for constant-depth compliance predicates, provided*

$$\exists \epsilon \in (0, 1) \text{ and a polynomial } \alpha \text{ s.t. } v^*(\lambda, m, N, \ell) = O(N^{1-\epsilon} \cdot \alpha(\lambda, m, \ell)) .$$

Moreover:

- If ARG and AS are secure against quantum adversaries, then PCD is secure against quantum adversaries.
- If ARG and AS are (post-quantum) zero knowledge, then PCD is (post-quantum) zero knowledge.
- If the size of the predicate $\varphi: \mathbb{F}^{(m+2)\ell} \rightarrow \mathbb{F}$ is $f = \omega(\alpha(\lambda, m, \ell)^{1/\epsilon})$ then:
 - the cost of running \mathbb{I} is equal to the cost of running both \mathcal{I} and I on an index of size $f + o(f)$;
 - the cost of running \mathbb{P} is equal to the cost of accumulating m instance-proof pairs using P , and running P , on an index of size $f + o(f)$ and instance of size $o(f)$;
 - the cost of running \mathbb{V} is equal to the cost of running both \mathcal{V} and D on an index of size $f + o(f)$ and an instance of size $o(f)$.

This last point gives the conditions for a *sublinear additive* recursive overhead; i.e., when the *additional* cost of proving that φ is satisfied recursively is asymptotically smaller than the cost of proving that φ is satisfied locally. Note that the smaller the compliance predicate φ , the more efficient the accumulation scheme has to be in order to achieve this.

5.1 Construction

Let $\text{ARG} = (\mathcal{G}, \mathcal{I}, \mathcal{P}, \mathcal{V})$ be a non-interactive argument for circuit satisfiability, and let $\text{AS} = (\mathbb{G}, \mathbb{I}, \mathbb{P}, \mathbb{V}, \mathbb{D})$ be an accumulation scheme for ARG . Below we construct a PCD scheme $\text{PCD} = (\mathbb{G}, \mathbb{I}, \mathbb{P}, \mathbb{V})$.

Given a compliance predicate $\varphi: \mathbb{F}^{(m+2)\ell} \rightarrow \mathbb{F}$, the circuit that realizes the recursion is as follows.

$$R_{\mathbb{V}, \varphi}^{(\lambda, N, k)}((\text{avk}, z, \text{acc}.\mathbb{x}), (z_{\text{loc}}, [z_i, \pi_i.\mathbb{x}, \text{acc}_i.\mathbb{x}]_{i=1}^m, \pi_{\mathbb{V}})):$$

1. Check that the compliance predicate $\varphi(z, z_{\text{loc}}, z_1, \dots, z_m)$ accepts.
2. If there exists $i \in [m]$ such that $z_i \neq \perp$, check that the NARK accumulation verifier accepts:

$$\mathbb{V}^{(\lambda, m, N, k)}(\text{avk}, [(\text{avk}, z_i, \text{acc}_i.\mathbb{x}), \pi_i.\mathbb{x}]_{i=1}^m, [\text{acc}_i.\mathbb{x}]_{i=1}^m, \text{acc}.\mathbb{x}, \pi_{\mathbb{V}}) = 1 .$$

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

Below we denote by “ q_i ” the query contents $((\text{avk}, z_i, \text{acc}_i), \pi_i)$, and use “ $q_i.\mathbb{x}$ ” to denote query instances $((\text{avk}, z_i, \text{acc}_i.\mathbb{x}), \pi_i.\mathbb{x})$ and “ $q_i.\mathbb{w}$ ” to denote the corresponding query witnesses $((\text{avk}, z_i, \text{acc}_i.\mathbb{w}), \pi_i.\mathbb{w})$.

Above, $\mathbb{V}^{(\lambda, m, N, k)}$ refers to the circuit representation of \mathbb{V} with input size appropriate for security parameter λ , number of instance-proof pairs and accumulators m , index size N , and instance size k .

Next we describe the generator \mathbb{G} , indexer \mathbb{I} , prover \mathbb{P} , and verifier \mathbb{V} of the PCD scheme.

- $\mathbb{G}(1^\lambda)$: Sample $\text{pp} \leftarrow \mathcal{G}(1^\lambda)$ and $\text{pp}_{\text{AS}} \leftarrow \mathbb{G}(1^\lambda)$, and output $\mathbb{P}\mathbb{P} := (\text{pp}, \text{pp}_{\text{AS}})$.
- $\mathbb{I}(\mathbb{P}\mathbb{P}, \varphi)$:
 1. Compute the integer $N := N(\lambda, |\varphi|, m, \ell)$, where N is defined in Lemma 5.3 below.
 2. Construct the circuit $R := R_{\mathbb{V}, \varphi}^{(\lambda, N, k)}$ where $k := |\text{avk}(\lambda, N)| + \ell + |\text{acc}(\lambda, N)|$.
 3. Compute the index key pair $(\text{ipk}, \text{ivk}) \leftarrow \mathcal{I}(\text{pp}, R)$ for the circuit R for the NARK.
 4. Compute the index key triple $(\text{apk}, \text{dk}, \text{avk}) \leftarrow \mathbb{I}(\text{pp}_{\text{AS}}, \text{ipk} = (\text{pp}, R))$ for the accumulator.
 5. Output $\text{ipk} := (\text{ipk}, \text{apk})$ and $\text{ivk} := (\text{ivk}, \text{dk}, \text{avk})$.
- $\mathbb{P}(\text{ipk}, z, z_{\text{loc}}, [z_i, (\pi_i, \text{acc}_i)]_{i=1}^m)$:
 1. If $z_i = \perp$ for all $i \in [m]$ then set $(\text{acc}, \pi_{\mathbb{V}}) \leftarrow \mathbb{P}(\text{apk}, \perp)$.
 2. If $z_i \neq \perp$ for some $i \in [m]$ then compute $(\text{acc}, \pi_{\mathbb{V}}) \leftarrow \mathbb{P}(\text{apk}, [q_i]_{i=1}^m, [\text{acc}_i]_{i=1}^m)$.
 3. Compute $\pi \leftarrow \mathcal{P}(\text{ipk}, (\text{avk}, z, \text{acc}.\mathbb{x}), (z_{\text{loc}}, [z_i, \pi_i.\mathbb{x}, \text{acc}_i.\mathbb{x}]_{i=1}^m, \pi_{\mathbb{V}}))$.
 4. Output (π, acc) .
- $\mathbb{V}(\text{ivk}, z, (\pi, \text{acc}))$: Accept if both $\mathcal{V}(\text{ivk}, (\text{avk}, z, \text{acc}.\mathbb{x}), \pi)$ and $\mathbb{D}(\text{dk}, \text{acc})$ accept.

5.2 Completeness

Let \mathcal{A} be any adversary that causes the completeness condition of PCD to be satisfied with probability p . We construct an adversary \mathcal{B} , as follows, that causes the completeness condition of AS to be satisfied with probability at most p .

- $\mathcal{B}(\text{pp}, \text{pp}_{\text{AS}})$:
1. Set $\mathbb{P}\mathbb{P} := (\text{pp}, \text{pp}_{\text{AS}})$ and compute $(\varphi, z, z_{\text{loc}}, [z_i, \pi_i, \text{acc}_i]_{i=1}^m) \leftarrow \mathcal{A}(\mathbb{P}\mathbb{P})$.
 2. Run $(\text{apk}, \text{dk}, \text{avk}) \leftarrow \mathbb{I}(\text{pp}_{\text{AS}}, \text{pp}, R_{\mathbb{V}, \varphi}^{(\lambda, N, k)})$.
 3. Output $(R_{\mathbb{V}, \varphi}^{(\lambda, N, k)}, [q_i]_{i=1}^m, [\text{acc}_i]_{i=1}^m)$.

Suppose that \mathcal{A} outputs $(\varphi, z, z_{\text{loc}}, [z_i, \pi_i, \text{acc}_i]_{i=1}^m)$ such that the completeness precondition is satisfied, but $\mathbb{V}(\text{ivk}, z, (\pi, \text{acc})) = 0$. Then, by construction of \mathbb{V} , it holds that either $\mathcal{V}(\text{ivk}, (\text{avk}, z, \text{acc}.\mathbb{x}), \pi) = 0$ or $D(\text{dk}, \text{acc}) = 0$. If $z_i = \perp$ for all i , then by perfect completeness of ARG both of these algorithms output 1; hence there exists i such that $z_i \neq \perp$. Hence it holds that for all i , $\mathbb{V}(\text{ivk}, z_i, (\pi_i, \text{acc}_i)) = 1$, whence for all i , $\mathcal{V}(\text{ivk}, (\text{avk}, z_i, \text{acc}_i.\mathbb{x}), \pi_i) = \Phi_{\mathcal{V}}(\text{pp}, R_{V,\varphi}^{(\lambda,N,k)}, (\text{avk}, z_i, \text{acc}_i.\mathbb{x}), \pi_i) = 1$ and $D(\text{dk}, \text{acc}_i) = 1$.

If $\mathcal{V}(\text{ivk}, (\text{avk}, z, \text{acc}.\mathbb{x}), \pi) = 0$, then, by perfect completeness of ARG, we know that $R_{V,\varphi}^{(\lambda,N,k)}$ rejects $((\text{avk}, z, \text{acc}), (z_{\text{loc}}, [z_i, \pi_i.\mathbb{x}, \text{acc}_i.\mathbb{x}]_{i=1}^m), \pi_{\mathbb{V}})$, and so $\mathbb{V}(\text{avk}, [q_i.\mathbb{x}]_{i=1}^m, [\text{acc}_i.\mathbb{x}]_{i=1}^m, \text{acc}.\mathbb{x}) = 0$. Otherwise, $D(\text{dk}, \text{acc}) = 0$.

Now consider the completeness experiment for AS with adversary \mathcal{B} . Since $\text{pp}, \text{pp}_{\text{AS}}$ are drawn identically to the PCD experiment, the distribution of the output of \mathcal{A} is identical. Hence in particular it holds that for all i , $\Phi_{\mathcal{V}}(\text{pp}, R_{V,\varphi}^{(\lambda,N,k)}, (\text{avk}, z_i, \text{acc}_i), \pi_i) = 1$ and $D(\text{dk}, \text{acc}_i) = 1$. By the above, it holds that either $\mathbb{V}(\text{avk}, [q_i.\mathbb{x}]_{i=1}^m, [\text{acc}_i.\mathbb{x}]_{i=1}^m, \text{acc}) = 0$ or $D(\text{dk}, \text{acc}) = 0$, and so $\mathcal{B} := (\mathcal{B}_1, \mathcal{B}_2)$ causes the completeness condition for AS to be satisfied with probability at most p .

5.3 Knowledge soundness

Since the extracted transcript \mathbb{T} will be a tree, we find it convenient to associate the label $z^{(u,v)}$ of the unique outgoing edge of a node u with the node u itself, so that the node u is labelled with $(z^{(u)}, z_{\text{loc}}^{(u)})$. For the purposes of the proof we also associate with each node u a NARK proof $\pi^{(u)}$ and an accumulator $\text{acc}^{(u)}$, so that the full label for a node is $(z^{(u)}, z_{\text{loc}}^{(u)}, \pi^{(u)}, \text{acc}^{(u)})$. It is straightforward to transform such a transcript into one that satisfies Definition 3.1.

Given a malicious prover $\tilde{\mathbb{P}}$, we will define an extractor $\mathbb{E}_{\tilde{\mathbb{P}}}$ that satisfies knowledge soundness. In the process we construct a sequence of extractors $\mathbb{E}_1, \dots, \mathbb{E}_d$ for $d := d(\varphi)$ (the depth of φ); \mathbb{E}_j outputs a tree of depth $j + 1$. Let $\mathbb{E}_0(\text{pp}, \text{ai})$ run $(\varphi, \text{o}, \pi, \text{acc}) \leftarrow \tilde{\mathbb{P}}(\text{pp}, \text{ai})$ and output (φ, \mathbb{T}_0) , where \mathbb{T}_0 is a single node labeled with $(\text{o}, \pi, \text{acc})$. Let $l_{\mathbb{T}}(j)$ denote the vertices of \mathbb{T} at depth j ; $l_{\mathbb{T}}(0) := \emptyset$ and $l_{\mathbb{T}}(1)$ is the singleton containing the root.

Now we define the extractor \mathbb{E}_j inductively for each $j \in [d]$. Suppose we have already constructed \mathbb{E}_{j-1} . We construct a NARK prover $\tilde{\mathcal{P}}_j$ as follows:

$\tilde{\mathcal{P}}_j(\text{pp}, (\text{pp}_{\text{AS}}, \text{ai})):$

1. Compute $(\varphi, \mathbb{T}_{j-1}, \text{ao}) \leftarrow \mathbb{E}_{j-1}((\text{pp}, \text{pp}_{\text{AS}}), \text{ai})$.
2. For each vertex $v \in l_{\mathbb{T}_{j-1}}(j)$, denote its label by $(z^{(v)}, \pi^{(v)}, \text{acc}^{(v)})$.
3. Run the argument indexer $(\text{ipk}, \text{ivk}) \leftarrow \mathcal{I}(\text{pp}, R_{V,\varphi}^{(\lambda,N,k)})$. Run the accumulator indexer $(\text{apk}, \text{dk}, \text{avk}) \leftarrow \mathcal{I}(\text{pp}_{\text{AS}}, \text{pp}, R_{V,\varphi}^{(\lambda,N,k)})$.
4. Output

$$(\vec{\mathbf{i}}, \vec{\mathbf{x}}, \vec{\pi}, \text{ao}') := \left(\vec{R}, (\text{avk}, z^{(v)}, \text{acc}^{(v)}.\mathbb{x})_{v \in l_{\mathbb{T}_{j-1}}(j)}, (\pi^{(v)})_{v \in l_{\mathbb{T}_{j-1}}(j)}, (\varphi, \mathbb{T}_{j-1}, \text{ao}) \right)$$

where \vec{R} is the vector $(R_{V,\varphi}^{(\lambda,N,k)}, \dots, R_{V,\varphi}^{(\lambda,N,k)})$ of the appropriate length.

Next let $\mathcal{E}_{\tilde{\mathcal{P}}_j}$ be the extractor that corresponds to $\tilde{\mathcal{P}}_j$, via the knowledge soundness of the non-interactive argument ARG. We now define an accumulation scheme prover \mathbb{P}_j as follows:

$\tilde{P}_j(\text{pp}_{\text{AS}}, (\text{pp}, \text{ai})):$

1. Run the extractor $(\vec{\mathbb{i}}, \vec{\mathbb{x}}, \vec{\mathbb{w}}, \text{ao}') \leftarrow \mathcal{E}_{\tilde{P}_j}(\text{pp}, (\text{pp}_{\text{AS}}, \text{ai}))$.
2. Parse the auxiliary output ao' as (φ, T', ao) . If T' is not a transcript of depth j , abort.
3. For each vertex $v \in l_{T'}(j)$,
 - obtain $\text{acc}^{(v)}$ from T' ;
 - obtain the local data $z_{\text{loc}}^{(v)}$, input messages $(z_i^{(v)}, \pi_i^{(v)}.\mathbb{x}, \text{acc}_i^{(v)}.\mathbb{x})_{i \in [m]}$ and accumulation proof $\pi_V^{(v)}$ from $\mathbb{w}^{(v)}$;
 - append $z_{\text{loc}}^{(v)}$ to the label of v in T' ;
 - let $S_j := \{v \in l_{T'}(j) : \exists i, z_i^{(v)} \neq \perp\}$; attach m children to each $v \in S_j$, where the i -th child is labeled with $z_i^{(v)}$.
4. Output $\left((\vec{\mathbb{i}}^{(v)}, \text{acc}^{(v)}, \pi_V^{(v)}, [\mathbf{q}_i^{(v)}.\mathbb{x} = ((\text{avk}, z_i^{(v)}, \text{acc}_i^{(v)}.\mathbb{x}), \pi_i^{(v)}.\mathbb{x})]_{i=1}^m, [\text{acc}_i^{(v)}.\mathbb{x}]_{i=1}^m)_{v \in S_j}, (\varphi, T', \text{ao}) \right)$.

Let $\mathbb{E}_{\tilde{P}_j}$ be the extractor corresponding to \tilde{P}_j , by the knowledge soundness guarantee of AS. Finally, we define the extractor \mathbb{E}_j as follows:

$\mathbb{E}_j(\mathbb{PP} = (\text{pp}, \text{pp}_{\text{AS}}), \text{ai}):$

1. Run the extractor $\left((\vec{\mathbb{i}}^{(v)}, \text{acc}^{(v)}, [\mathbf{q}_i^{(v)}]_{i=1}^m, [\text{acc}_i^{(v)}]_{i=1}^m)_{v \in S_j}, \text{ao}' \right) \leftarrow \mathbb{E}_{\tilde{P}_j}(\text{pp}, \text{pp}_{\text{AS}}, \text{ai})$.
2. Parse the auxiliary output ao' as (φ, T', ao) . If T' is not a transcript of depth j , abort. Let $S_j := \{v \in l_{T'}(j) : \exists i, z_i^{(v)} \neq \perp\}$.
3. Parse each $\mathbf{q}_i^{(v)}$ as $(\text{avk}^{(v)}, z_i^{(v)}, \text{acc}_i^{(v)}.\mathbb{x}, \pi_i^{(v)})$.
4. Output $(\varphi, T_j, \text{ao})$ where T_j is the transcript constructed from T' by adding, for each vertex $v \in S_j$, $(\pi_i^{(v)}, \text{acc}_i^{(v)})$ to the label of its i -th child.

The extractor $\mathbb{E}_{\tilde{P}}$ is equal to \mathbb{E}_d .

We now show that $\mathbb{E}_{\tilde{P}}$ runs in expected polynomial time and that it outputs a transcript that is φ -compliant.

Running time of the extractor. It follows from the extraction guarantees of ARG and AS that \mathbb{E}_j runs in expected time polynomial in the expected running time of \mathbb{E}_{j-1} . Hence if $d(\varphi)$ is a constant, $\mathbb{E}_{\tilde{P}} = \mathbb{E}_{d(\varphi)}$ runs in expected polynomial time.

Correctness of the extractor. Fix a set Z , and suppose that \tilde{P} 's output falls in Z , and causes \mathbb{V} to accept, with probability μ . We show by induction that, for all $j \in \{0, \dots, d\}$, the transcript T_j output by \mathbb{E}_j is φ -compliant up to depth j , and that for all $v \in T_j$, both $\mathcal{V}(\text{ivk}, (\text{avk}, z^{(v)}, \text{acc}^{(v)}.\mathbb{x}), \pi^{(v)})$ and $\text{D}(\text{dk}, \text{acc}^{(v)})$ accept, and that $(\mathbb{PP}, \text{ai}, \varphi, \text{o}(T_j), \text{ao}) \in Z$ and $\varphi \in \text{F}$, with probability $\mu - \text{negl}(\lambda)$.

For $j = 0$ the statement holds by assumption.

Now suppose that $(\varphi, T_{j-1}) \leftarrow \mathbb{E}_{j-1}(\mathbb{PP}, \text{ai})$ is such that T_{j-1} is φ -compliant up to depth $j-1$, and that both $\mathcal{V}(\text{ivk}, (\text{avk}, z^{(v)}, \text{acc}^{(v)}.\mathbb{x}), \pi^{(v)})$ and $\text{D}(\text{dk}, \text{acc}^{(v)})$ accept for all $v \in T_{j-1}$, with probability $\mu - \text{negl}(\lambda)$.

Let $(\vec{\mathbb{i}}, (\text{avk}_v, z^{(v)}, \text{acc}^{(v)}.\mathbb{x})_v, (\pi^{(v)})_v, (\varphi, T'), \vec{\mathbb{w}}) \leftarrow \mathcal{E}_{\tilde{P}_j}(\text{pp}, (\text{pp}_{\text{AS}}, \text{ai}))$.

We let $(\text{pp}, (\text{pp}_{\text{AS}}, \text{ai}), \vec{\mathbb{i}}, (\text{avk}_v, z^{(v)}, \text{acc}^{(v)}.\mathbb{x})_v, (\varphi, T', \text{ao})) \in Z'$ if and only if, for $(\text{apk}, \text{dk}, \text{avk}) \leftarrow \text{I}(\text{pp}_{\text{AS}}, \text{pp}, R_{V, \varphi}^{(\lambda, N, k)})$:

- $((\text{pp}, \text{pp}_{\text{AS}}), \text{ai}, \varphi, \text{o}(T'), \text{ao}) \in Z$ and $\varphi \in \text{F}$,
- $\vec{\mathbb{i}}^{(v)} = R_{V, \varphi}^{(\lambda, N, k)}$ and $\text{avk}_v = \text{avk}$ for all v ,
- T' is φ -compliant up to depth $j-1$,
- $\text{D}(\text{dk}, \text{acc}^{(v)})$ accepts for all $v \in T'$, and

- for $v \in l_{T'}(j)$, v is labeled in T' with $(z^{(v)}, \pi^{(v)}, \text{acc}^{(v)})$.

By knowledge soundness, with probability $\mu - \text{negl}(\lambda)$, $(\text{pp}, (\text{pp}_{\text{AS}}, \text{ai}), \vec{\mathbf{i}}, (\text{ivk}_v, z^{(v)})_v, (\varphi, T')) \in Z'$ and for every vertex $v \in l_{T'}(j)$, $(R_{V,\varphi}^{(\lambda,N,k)}, (\text{avk}_v, z^{(v)}, \text{acc}^{(v)}), \mathbb{w}^{(v)}) \in \mathcal{R}_{\text{R1CS}}$. Here we use Z' and the auxiliary output in the knowledge soundness definition of ARG to ensure consistency between the values $z^{(v)}$ and T' , and to ensure that T' is φ -compliant and that the decider accepts.

Consider some $v \in l_{T'}(j)$. Since $(R_{V,\varphi}^{(\lambda,N,k)}, (\text{avk}^{(v)}, z^{(v)}, \text{acc}^{(v)}.\mathbb{x}), \mathbb{w}^{(v)}) \in \mathcal{R}_{\text{R1CS}}$, we obtain from $\mathbb{w}^{(v)}$ either

- local data $z_{\text{loc}}^{(v)}$, input messages $(z_i^{(v)}, \pi_i^{(v)}.\mathbb{x}, \text{acc}_i^{(v)}.\mathbb{x})_{i \in [m]}$ and proof π_V such that $\varphi(z^{(v)}, z_{\text{loc}}, z_1, \dots, z_m)$ accepts and the accumulation verifier $V^{(\lambda,N,k)}(\text{avk}^{(v)}, [q_i^{(v)}.\mathbb{x}]_{i=1}^m, [\text{acc}_i^{(v)}.\mathbb{x}]_{i=1}^m, \text{acc}^{(v)}, \pi_V^{(v)})$ accepts, where $q_i^{(v)}.\mathbb{x} := ((\text{avk}^{(v)}, z_i^{(v)}, \text{acc}_i^{(v)}.\mathbb{x}), \pi_i^{(v)}.\mathbb{x})$; or
- local data $z_{\text{loc}}^{(v)}$ such that $\varphi(z^{(v)}, z_{\text{loc}}, \perp, \dots, \perp)$ accepts.

In both cases we append $z_{\text{loc}}^{(v)}$ to the label of v . In the latter case, v has no children and so is φ -compliant by the base case condition. In the former case we label the children of v with $(z_i, \pi_i, \text{acc}_i)$, and so v is φ -compliant.

We define $(\text{pp}_{\text{AS}}, \text{pp}, \text{ai}, (\mathbf{i}^{(v)}, \text{acc}^{(v)}, [q_i^{(v)}.\mathbb{x}]_{i=1}^m, [\text{acc}_i^{(v)}.\mathbb{x}]_{i=1}^m)_v, (\varphi, T', \text{ao})) \in Z''$ if and only if

- $(\text{pp}, \text{pp}_{\text{AS}}, \text{ai}, \varphi, \text{o}(T'), \text{ao}) \in Z$ and $\varphi \in F$,
- $\mathbf{i}^{(v)} = R_{V,\varphi}^{(\lambda,N,k)}$ for all v ,
- T' is φ -compliant up to depth j ,
- for all v , $q_i^{(v)}.\mathbb{x} = ((\text{avk}, z_i^{(v)}, \text{acc}_i^{(v)}.\mathbb{x}), \pi_i^{(v)})$ where $(\text{apk}, \text{avk}, \text{dk}) \leftarrow I(\text{pp}_{\text{AS}}, \text{pp}_{\Phi}, \mathbf{i}_{\Phi})$, and
- for $u \in l_{T'}(j+1)$, where u is the i -th child of $v \in l_{T'}(j)$, u is labeled in T' with $z_i^{(u)}$.

Let $(\mathbf{i}^{(v)}, \text{acc}^{(v)}, [q_i^{(v)}.\mathbb{x}]_{i=1}^m, [\text{acc}_i^{(v)}.\mathbb{x}]_{i=1}^m)_{v \in S_j}, \text{ao}' \leftarrow E_{\tilde{P}_j}(\text{pp}_{\text{AS}}, \text{pp}, \text{ai})$. By the knowledge soundness guarantee of the accumulation scheme, $(\text{pp}, \text{pp}_{\Phi}, \text{ai}, (\mathbf{i}^{(v)}, \text{acc}^{(v)}, [q_i^{(v)}.\mathbb{x}]_{i=1}^m, [\text{acc}_i^{(v)}.\mathbb{x}]_{i=1}^m)_v, \text{ao}') \in Z''$, and it holds that for all descendants u of v in T_j , $D(\text{dk}, \text{acc}^{(u)})$ accepts and $\Phi_V(\text{pp}, R_{V,\varphi}^{(\lambda,N,k)}, (\text{avk}, z^{(u)}, \text{acc}^{(u)}.\mathbb{x}), \pi_{\text{in}}^{(u)}) = \mathcal{V}(\text{ivk}, (\text{avk}, z^{(u)}, \text{acc}^{(u)}.\mathbb{x}), \pi_{\text{in}}^{(u)})$ accepts, with probability $\mu - \text{negl}(\lambda)$; this completes the inductive step.

Hence by induction, $(\varphi, T, \text{ao}) \leftarrow \mathbb{E}(\mathbb{P}\mathbb{P}, \text{ai})$ has φ -compliant T , $(\mathbb{P}\mathbb{P}, \text{ai}, \varphi, \text{o}(T), \text{ao}) \in Z$, and $\varphi \in F$, with probability $\mu - \text{negl}(\lambda)$.

5.4 Efficiency

The efficiency argument follows from Lemma 5.3 and is essentially identical to that of [BCMS20], and so we will not repeat it. We note only that the quantity v^* (i) describes the size of the *accumulation* verifier, which in particular need not read the entire NARK proof, which may be large, and (ii) is a function of the size of the accumulator *instance* alone; the accumulator witness may be large.

Lemma 5.3. *Suppose that for every security parameter $\lambda \in \mathbb{N}$, arity m , and message size $\ell \in \mathbb{N}$ the ratio of accumulation verifier circuit size to index size $v^*(\lambda, m, N, \ell)/N$ is monotone decreasing in N . Then there exists a size function $N(\lambda, f, m, \ell)$ such that*

$$\forall \lambda, f, m, \ell \in \mathbb{N} \quad S(\lambda, f, m, \ell, N(\lambda, f, m, \ell)) \leq N(\lambda, f, m, \ell) .$$

Moreover if for some $\epsilon > 0$ and some increasing function α it holds that, for all N, λ, m, ℓ sufficiently large,

$$v^*(\lambda, m, N, \ell) \leq N^{1-\epsilon} \alpha(\lambda, m, \ell)$$

then, for all λ, m, ℓ sufficiently large, $N(\lambda, f, m, \ell) \leq O(f + \alpha(\lambda, m, \ell)^{1/\epsilon})$.

6 Accumulating Pedersen polynomial commitments

We construct a split accumulation scheme for the Pedersen polynomial commitment scheme. This polynomial commitment scheme has modest efficiency properties: commitments are succinct, but evaluation proofs (and the time to verify them) are linear in the degree of the committed polynomial.

Despite this, our split accumulation scheme achieves strong efficiency guarantees: checking the accumulation of n evaluation claims only requires the accumulation verifier V to add $O(n)$ commitments (regardless of the degree of the committed polynomials) and only requires the decider D to check one evaluation proof.

In more detail, we prove the following (more general) theorem for a polynomial commitment scheme PC_{HC} derived from any \mathbb{F}_q -linear hash function HC . The Pedersen commitment [Ped92] is an \mathbb{F}_q -linear homomorphic commitment where the commitment space is a group \mathbb{G} of order q , and is computationally binding provided the discrete logarithm problem is hard in \mathbb{G} .

Theorem 6.1. *The scheme $AS = (G, I, P, V, D)$ constructed in Section 6.1 is an accumulation scheme in the random oracle model for the polynomial commitment scheme PC_{HC} described in Fig. 4. AS achieves the following efficiency:*

- Generator: $G(1^\lambda)$ runs in time $O(\lambda)$.
- Indexer: the time of $I(pp, pp_\Phi, i_\Phi = d)$ is dominated by the time to compute $HC.Trim$ with degree bound d .
- Accumulation prover: the time of $P^p(apk, [q_i]_{i=1}^n, [acc_j]_{j=1}^m)$ is dominated by the time to commit to $n + m$ polynomials of degree d . The output accumulator acc consists of
 - an accumulator instance $acc.x$ consisting of a commitment and two field elements, and
 - an accumulator witness $acc.w$ consisting of a polynomial of degree less than d .
 The output proof π consists of n commitments and $2n + 2m$ field elements.
- Accumulation verifier: the time of $V^p(avk, [q_i.x]_{i=1}^n, [acc_j.w]_{j=1}^m, acc.x, \pi)$ is dominated by $O(n + m)$ field additions/multiplications and $O(n + m)$ group scalar multiplications.
- Decider: the time of $D(dk, acc)$ is dominated by the time to compute a commitment to a polynomial of degree less than d .

Definition 6.2. An \mathbb{F}_q -linear homomorphic commitment $HC = (Setup, Trim, Commit)$ for prime q is a binding commitment scheme (Section 3.6) with the following additional properties: the message space is \mathbb{F}_q^n , and there exists an efficient operation $+$ on the commitment space such that

$$\forall \vec{a}, \vec{b} \in \mathbb{F}_q^n, HC.Commit(\vec{a}) + HC.Commit(\vec{b}) = HC.Commit(\vec{a} + \vec{b}) .$$

PC_{HC} . We use HC to construct a homomorphic polynomial commitment scheme PC_{HC} :

- *Setup:* On input $\lambda, D \in \mathbb{N}$, output $pp_{CM} \leftarrow HC.Setup(1^\lambda, D + 1)$.
 - *Trim:* On input pp_{CM} and $d \in \mathbb{N}$, check that $d \leq D$, set $ck_{HC} \leftarrow HC.Trim(pp_{CM}, d + 1)$, and output $(ck := (ck_{HC}, d), rk := (ck_{HC}, d))$.
 - *Commit:* On input ck and $p \in \mathbb{F}_q[X]$ of degree at most $ck.d$, output $C \leftarrow HC.Commit(ck.ck_{HC}, p)$.
 - *Open:* On input (ck, p, C, z) , output $\pi := p$.
 - *Check:* On input $(rk, (C, z, v), \pi = p)$, check that $C = HC.Commit(rk.ck_{HC}, p)$, $p(z) = v$, and $\deg(p) \leq rk.d$.
- Completeness of PC_{HC} follows from that of HC , while extractability follows from the binding property of HC .

Figure 4: PC_{HC} is a homomorphic polynomial commitment scheme based on a linear homomorphic commitment.

6.1 Accumulation scheme for PC_{HC}

We describe the accumulation scheme $\text{AS} = (G, I, P, V, D)$ for the homomorphic polynomial commitment scheme PC_{HC} . First, we describe the form of inputs $q = (q.\mathbb{x}, q.\mathbb{w})$ and accumulators $\text{acc} = (\text{acc}.\mathbb{x}, \text{acc}.\mathbb{w})$. Both $q.\mathbb{x}$ and $\text{acc}.\mathbb{x}$ are tuples of the form (C, z, v) , while both $q.\mathbb{w}$ and $\text{acc}.\mathbb{w}$ consist of a polynomial p (allegedly, committed inside C and such that $p(z) = v$ and $\deg(p) < d$). Jumping ahead, we note that since the decider D is equal to the predicate Φ_{PC} , there is no distinction between inputs and prior accumulators, and so it suffices to accumulate inputs only.

Generator. The generator G receives as input $\text{pp} := 1^\lambda$ and outputs 1^λ . (In other words, G does not have to create additional public parameters beyond those used by HC .)

Indexer. On input the accumulator parameters pp , predicate parameters $\text{pp}_\Phi = \text{pp}_{\text{PC}}$, and a predicate index $i_\Phi = d$, the indexer I computes the committer and receiver keys $(\text{ck}, \text{rk}) := \text{PC}_{\text{HC}}.\text{Trim}(\text{pp}_{\text{PC}}, d)$, and then outputs the accumulator proving key $\text{apk} := (\text{ck}, d)$, the accumulator verification key $\text{avk} := d$, and the decision key $\text{dk} := \text{rk}$.

Accumulation prover. On input accumulation proving key apk and inputs $[q_i]_{i=1}^n$, P works as below.

$P^\rho(\text{apk} = (\text{ck}, d), [q_i]_{i=1}^n)$:

1. For each i in $[n]$:
 - (a) Parse the input instance $q.\mathbb{x}_i$ as (C_i, z_i, v_i) and the input witness $q.\mathbb{w}_i$ as $\pi_i = p_i \in \mathbb{F}^{\leq \text{ck}.d}[X]$.
 - (b) Compute the witness polynomial $w_i(X) := \frac{p_i(X) - v_i}{X - z_i}$.
 - (c) Compute a commitment to $w_i(X)$: $W_i := \text{PC}_{\text{HC}}.\text{Commit}(\text{ck}, w_i)$.
2. Compute the challenge $z := \rho(d, [(C_i, z_i, v_i, W_i)]_{i=1}^n)$.
3. For each i in $[n]$, compute the evaluations $y_i := p_i(z)$ and $y'_i := w_i(z)$.
4. Compute the challenge $\alpha := \rho(z, [(y_i, y'_i)]_{i=1}^n)$.
5. Compute the linear combination $p := \sum_{i=1}^n \alpha^{i-1} \cdot p_i + \sum_{i=1}^n \alpha^{n+i-1} \cdot w_i$.
6. Compute the evaluation $v := p(z)$.
7. Compute the linear combination $C := \sum_{i=1}^n \alpha^{i-1} \cdot C_i + \sum_{i=1}^n \alpha^{n+i-1} \cdot W_i$.
8. Set the accumulator $\text{acc} := (\text{acc}.\mathbb{x}, \text{acc}.\mathbb{w})$ where $\text{acc}.\mathbb{x} := (C, z, v)$ and $\text{acc}.\mathbb{w} := p$.
9. Set the proof $\pi_V := [(W_i, y_i, y'_i)]_{i=1}^n$.
10. Output (acc, π_V) .

Accumulation verifier. On input the accumulator verification key avk , input instances $[q_i.\mathbb{x}]_{i=1}^n$, a new accumulator instance $\text{acc}.\mathbb{x}$, and a proof π_V , V works as below.

$V^\rho(\text{avk} = d, [q_i.\mathbb{x}]_{i=1}^n, \text{acc}.\mathbb{x}, \pi_V)$:

1. Parse $\text{acc}.\mathbb{x}$ as (C, z, v) , and π_V as $[(W_i, y_i, y'_i)]_{i=1}^n$.
2. For each $i \in [n]$, parse $q.\mathbb{x}_i$ as (C_i, z_i, v_i) .
3. Check that $z = \rho(d, [(C_i, z_i, v_i, W_i)]_{i=1}^n)$.
4. For each $i \in [n]$, check that $y_i - v_i = y'_i \cdot (z - z_i)$.
5. Compute $\alpha := \rho(z, [(y_i, y'_i)]_{i=1}^n)$.
6. Check that $v = \sum_{i=1}^n \alpha^{i-1} \cdot y_i + \sum_{i=1}^n \alpha^{n+i-1} \cdot y'_i$.
7. Check that $C = \sum_{i=1}^n \alpha^{i-1} \cdot C_i + \sum_{i=1}^n \alpha^{n+i-1} \cdot W_i$.

Decider. On input the decision key $\text{dk} = \text{rk}$ and an accumulator $\text{acc} = (\text{acc}.\mathbb{x} = (C, z, v), \text{acc}.\mathbb{w} = \pi)$, D checks that $\text{PC}_{\text{HC}}.\text{Check}(\text{rk}, C, z, v, \pi) = 1$.

6.2 Proof of Theorem 6.1

We prove Theorem 6.1 by first proving that AS constructed above is complete and satisfies the claimed efficiency properties, and then, in Section 6.3, by proving that AS achieves knowledge soundness.

Completeness. Recalling that we need only accumulate predicate inputs, we demonstrate that the following simplified completeness property holds for every (unbounded) adversary \mathcal{A} :

$$\Pr \left[\begin{array}{l} \forall i \in [n], \Phi(\text{pp}_\Phi, i_\Phi, \mathbf{q}_i) = 1 \\ \downarrow \\ V^\rho(\text{avk}, [\mathbf{q}_i.\mathbf{x}]_{i=1}^n, \text{acc}.\mathbf{x}, \pi_V) = 1 \\ D(\text{dk}, \text{acc}) = 1 \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow G(1^\lambda) \\ \text{pp}_\Phi \leftarrow \mathcal{H}(1^\lambda) \\ (i_\Phi, [\mathbf{q}_i]_{i=1}^n) \leftarrow \mathcal{A}(\text{pp}, \text{pp}_\Phi) \\ (\text{apk}, \text{avk}, \text{dk}) \leftarrow I^\rho(\text{pp}, \text{pp}_\Phi, i_\Phi) \\ (\text{acc}, \pi_V) \leftarrow P^\rho(\text{apk}, [\mathbf{q}_i]_{i=1}^n) \end{array} \right] = 1 .$$

We prove this directly. For each $i \in [n]$, since

$$\Phi(\text{pp}_\Phi, i_\Phi, \mathbf{q}_i = (C_i, z_i, v_i), \pi_i = p_i) = \text{PC}_{\text{HC}}.\text{Check}(\text{rk}, C_i, z_i, v_i, p_i) = 1 ,$$

we know that $C_i = \text{PC}_{\text{HC}}.\text{Commit}(\text{ck}, p_i)$ and $p_i(z_i) = v_i$; this implies that each witness polynomial $w_i(X) = \frac{p_i(X) - v_i}{X - z_i}$ is indeed a polynomial of degree $d - 1$.

Together with the fact that the accumulation prover P behaves honestly, the foregoing facts imply that C is a well-formed commitment to $p = \sum_{i=1}^n \alpha^i p_i + \sum_{i=1}^n \alpha^{n+i} w_i$, and that $p(z) = v$, as required.

Efficiency. We now analyze the efficiency of our accumulation scheme.

- *Generator:* $G(1^\lambda)$ outputs 1^λ , and hence runs in time $O(\lambda)$.
- *Indexer:* $I^\rho(\text{pp}, \text{pp}_\Phi, i_\Phi)$ invokes $\text{PC}_{\text{HC}}.\text{Trim}$, and hence runs in time $O_\lambda(d)$.
- *Accumulation prover:* $P^\rho(\text{apk} = \text{ck}, [\mathbf{q}_i = (\mathbf{q}.\mathbf{x}_i, \mathbf{q}.\mathbf{w}_i)]_{i=1}^n)$ computes a commitment to the degree $\deg(p_i) - 1$ witness polynomial w_i for each input $\mathbf{q}.\mathbf{x}_i = (C_i, z_i, v_i)$. The time to generate these n commitments dominates the running time of P .
- *Accumulator size:* The accumulator instance $\text{acc}.\mathbf{x}$ consists of a polynomial commitment C , an evaluation point z and an evaluation claim v . The accumulator witness $\text{acc}.\mathbf{w}$ is a polynomial of degree d . (Separately, the accumulation proof π_V contains $O(n)$ group and field elements; this efficiency measure is not so important.)
- *Accumulation verifier:* $V^\rho(\text{avk}, [\mathbf{q}_i.\mathbf{x}]_{i=1}^n, \text{acc}.\mathbf{x}, \pi_V)$ computes a random linear combination between $2n$ commitments, and hence its running time is as claimed.
- *Decider:* $D(\text{dk}, \text{acc})$ simply invokes $\text{PC}_{\text{HC}}.\text{Check}$, which in turn invokes $\text{PC}_{\text{HC}}.\text{Commit}$, and checks that the output matches the accumulator.

6.3 Knowledge soundness

Since we do not distinguish queries and prior accumulators, and since our predicate does not access the random oracle, we can simplify the knowledge soundness condition. In particular, it suffices to demonstrate

that for every expected polynomial-time adversary $\tilde{\mathcal{P}}$ and auxiliary input distribution \mathcal{D} there exists an expected polynomial-time extractor \mathcal{E} such that for every set Z the following condition holds:

$$\Pr \left[\begin{array}{l} (\text{pp}, \text{pp}_\Phi, i_\Phi, \text{acc}, [q_i.\mathbb{x}]_{i=1}^n, \text{ai}, \text{ao}) \in Z \\ \wedge \\ \forall i \in [n], \Phi(\text{pp}_\Phi, i_\Phi, q_i) = 1 \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow G(1^\lambda) \\ \text{pp}_\Phi \leftarrow \mathcal{H}(1^\lambda) \\ \text{ai} \leftarrow \mathcal{D}(1^\lambda) \\ \left(\begin{array}{cc} i_\Phi & \text{acc} \\ [q_i]_{i=1}^n & \text{ao} \end{array} \right) \leftarrow \mathcal{E}(\text{pp}, \text{pp}_\Phi, \text{ai}) \end{array} \right] \\ \geq \Pr \left[\begin{array}{l} (\text{pp}, \text{pp}_\Phi, i_\Phi, \text{acc}, [q_i.\mathbb{x}]_{i=1}^n, \text{ai}, \text{ao}) \in Z \\ \wedge \\ \text{V}^\rho(\text{avk}, [q_i.\mathbb{x}]_{i=1}^n, \text{acc}.\mathbb{x}, \pi_V) = 1 \\ \text{D}(\text{dk}, \text{acc}) = 1 \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow G(1^\lambda) \\ \text{pp}_\Phi \leftarrow \mathcal{H}(1^\lambda) \\ \text{ai} \leftarrow \mathcal{D}(1^\lambda) \\ \left(\begin{array}{ccc} i_\Phi & \text{acc} & \text{ao} \\ [q_i.\mathbb{x}]_{i=1}^n & \pi_V & \end{array} \right) \leftarrow \tilde{\mathcal{P}}^\rho(\text{pp}, \text{pp}_\Phi, \text{ai}) \\ (\text{apk}, \text{avk}, \text{dk}) \leftarrow \text{I}^\rho(\text{pp}, \text{pp}_\Phi, i_\Phi) \end{array} \right] - \text{negl}(\lambda) .$$

The remainder of this subsection is dedicated to proving the above statement. We first establish some useful notation for algorithms with access to oracles. We then give two key technical tools: an expected-time forking lemma with negligible loss, and an expected-time variant of the zero-finding game lemma of [BCMS20]. Finally, we apply these tools to prove the knowledge soundness property.

6.3.1 Notation for oracle algorithms

Let A be a t -query oracle algorithm with access to an oracle $\rho: \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$. For $\vec{h} = (h_1, \dots, h_t) \in (\{0, 1\}^\lambda)^t$, we denote by $(q, y; \Gamma) \leftarrow A^{\vec{h}}(x; r)$ the following procedure: run A on input x and randomness r , and answer the i -th query q_i of A to its oracle with h_i for each i ; then, when A outputs (q, y) , if $q = q_k$ for some k , let k be the smallest such, and output $(q, y; \Gamma = [(q_i, h_i)]_{i=1}^k)$; otherwise, output $(q, y; \perp)$. Note that Γ is *truncated* to (and includes) the first appearance of the query q . We will write $(q, y; \Gamma) \leftarrow A^\rho(x; r)$ to denote the same procedure where h_i is defined (adaptively) to be $\rho(q_i)$ for each i .

We assume without loss of generality that A makes no oracle query more than once; in particular, we can interpret Γ as partial function $\{0, 1\}^* \rightarrow \{0, 1\}^\lambda$. For q in the support of Γ , we denote by $\Gamma[q \mapsto x]$ the partial function identical to Γ except that $\Gamma(q) = x$. If the randomness r is omitted it is assumed to be chosen uniformly.

6.3.2 An expected-time forking lemma

We give an expected-time forking lemma that is suitable for our setting. Crucially, the lemma gives a guarantee for extracting from expected-time adversaries, which will be important for extracting recursively; and it suffers only a negligible loss in success probability.

Lemma 6.3. *Let A be a t -query oracle algorithm running in expected time t_A , G be a parameter generation algorithm, and \mathfrak{p} be a t -time computable predicate. For any set Z , define*

$$\delta_Z := \Pr \left[\begin{array}{l} \Gamma \neq \perp \wedge (\text{pp}, y) \in Z \\ \wedge \mathfrak{p}(\text{pp}, q, y, \Gamma) = 1 \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow G(1^\lambda) \\ \vec{h} \leftarrow (\{0, 1\}^\lambda)^t \\ (q, y; \Gamma) \leftarrow A^{\vec{h}}(\text{pp}) \end{array} \right] .$$

Then there exists a t -query oracle algorithm B running in expected time $O(tN \cdot t_A)$ such that for all Z ,

$$\Pr \left[\begin{array}{l} \Gamma \neq \perp \wedge (\text{pp}, y_1) \in Z \\ \wedge \forall j \in [N], \text{p}(\text{pp}, q, y_j, \Gamma[q \mapsto g_j]) = 1 \\ \wedge \forall j \neq k \in [N], g_j \neq g_k \end{array} \middle| \begin{array}{l} \text{pp} \leftarrow G(1^\lambda) \\ \vec{h} \leftarrow (\{0, 1\}^\lambda)^t \\ (q, [g_i, y_i]_{i=1}^N; \Gamma) \leftarrow B^{\vec{h}}(\text{pp}, 1^N) \end{array} \right] \geq \delta_Z - \frac{2N\sqrt{t}}{2^{\lambda/2}} .$$

Proof. The algorithm $B^{\vec{h}}$ on input $(\text{pp}, 1^N)$ operates as follows.

1. Draw $r \leftarrow \{0, 1\}^R$.
2. Run $A^{\vec{h}}(x; r)$ until it terminates and outputs $(q, y; \Gamma)$. If $\Gamma = \perp$ or $p(x, q, \Gamma(q), y) = 0$, output \perp . Otherwise let $i = |\Gamma|$; in particular, $h_i = \Gamma(q)$.
3. Set $g_1 := h_i$ and $y_1 := y$.
4. Set $K := 1$ and repeat the following until $K = N$:
 - (a) Draw $h'_1, \dots, h'_t \leftarrow \{0, 1\}^\lambda$, and run $A^{h_1, \dots, h_{i-1}, h'_1, \dots, h'_t}(x; r)$ until it terminates and outputs $(q', y'; \Gamma')$.
 - (b) If $q' = q$ and $\text{p}(\text{pp}, q, y', \Gamma[q \mapsto h'_i]) = 1$, update $K \leftarrow K + 1$ and set $g_K := h'_i$ and $y_K := y'$.
5. Output $(q, g_1, y_1, \dots, g_N, y_N)$.

Let $S_i := \{(\vec{h}, r) : (q, y; D) \leftarrow A^{\vec{h}}(x; r) \wedge p(x, q, y, D) = 1\}$, and let $S := \bigcup_{i=1}^t S_i$. Define

$$p_i(h_1, \dots, h_{i-1}; r) := |\{h'_1, \dots, h'_t \in \{0, 1\}^\lambda : (h_1, \dots, h_{i-1}, h'_1, \dots, h'_t, r) \in S_i\}| / 2^{\lambda(t-i+1)} .$$

Observe that if B samples r such that $(\vec{h}, r) \in S_i$ then the probability that a single iteration of Step 4 increments K is $p_i(h_1, \dots, h_{i-1}; r)$.

We see that B halts almost surely: if B does not terminate in Step 2 then $p_i(h_1, \dots, h_{i-1}; r) \neq 0$. We also observe that the probability that B terminates in Step 2, or that $(\text{pp}, y_1) \notin Z$, is at most $1 - \delta_Z$.

We now bound the expected running time of B . Let T_A, T_B be random variables denoting the running time of A, B respectively, and let $t(h_1, \dots, h_{i-1}; r)$ denote the expected running time of $A^{\vec{h}}(x; r)$ over uniformly random h_i, \dots, h_t . For $(\vec{h}, r) \in S_i$, the number of iterations between $K \leftarrow k$ and $K \leftarrow k + 1$, which we denote $X_{\vec{h}, r}^{(k)}$, is geometrically distributed with parameter $p_i(h_1, \dots, h_{i-1}; r) \neq 0$. We denote the *time* between these increments of K by $T_{\vec{h}, r}^{(k)}$ and note that for $(\vec{h}; r) \in S_i$, $E[T_{\vec{h}, r}^{(k)} | X_{\vec{h}, r}^{(k)} = j] = j \cdot t(h_1, \dots, h_{i-1}; r)$ by linearity. Hence by law of total expectation,

$$\begin{aligned} E[T_B] &= \frac{1}{2^{\lambda t + R}} \sum_{i=1}^t \sum_{(\vec{h}, r) \in S_i} \sum_{k=1}^N \sum_{j=1}^{\infty} \Pr[X_{\vec{h}, r}^{(k)} = j] \cdot E[T_{\vec{h}, r}^{(k)} | X_{\vec{h}, r}^{(k)} = j] \\ &= \frac{N}{2^{\lambda t + R}} \sum_{i=1}^t \sum_{(\vec{h}, r) \in S_i} \sum_{j=1}^{\infty} (1 - p_i(h_1, \dots, h_{i-1}; r))^{j-1} \cdot p_i(h_1, \dots, h_{i-1}; r) \cdot j \cdot t(h_1, \dots, h_{i-1}; r) \\ &= \frac{N}{2^{\lambda t + R}} \sum_{i=1}^t \sum_{(\vec{h}, r) \in S_i} t(h_1, \dots, h_{i-1}; r) / p_i(h_1, \dots, h_{i-1}; r) \\ &\leq N \cdot \sum_{i=1}^t \sum_{h_1, \dots, h_{i-1}, r} \frac{t(h_1, \dots, h_{i-1}; r)}{2^{\lambda(i-1) + R}} = tN \cdot E[T_A] . \end{aligned}$$

where the inequality follows because for all functions $f(h_1, \dots, h_{i-1}; r)$ into \mathbb{R} ,

$$\sum_{(\vec{h}, r) \in S_i} f(h_1, \dots, h_{i-1}; r) = 2^{\lambda(t-i+1)} \cdot \sum_{\substack{h_1, \dots, h_{i-1}, r \\ p_i(h_1, \dots, h_{i-1}; r) \neq 0}} f(h_1, \dots, h_{i-1}; r) \cdot p_i(h_1, \dots, h_{i-1}; r) .$$

It can be shown similarly that the expected number of iterations is at most tN , and so the probability that B performs more than $\sqrt{t} \cdot 2^{\lambda/2}$ iterations is at most $N\sqrt{t} \cdot 2^{-\lambda/2}$. Conditioned on this, the probability that in any iteration we draw h'_i such that $h'_i = g_j$ for any $j < K$ is at most $N\sqrt{t}2^{\lambda/2} \cdot 2^{-\lambda} = N\sqrt{t} \cdot 2^{-\lambda/2}$. Applying a union bound completes the proof. \square

6.3.3 Zero-finding games

The following lemma, due to [BCMS20], bounds the probability that applying the random oracle to a commitment to a polynomial yields a zero of that polynomial. We refer to this as a *zero-finding game*. Here we have adapted the lemma to expected-time adversaries; the proof is essentially unchanged.

Lemma 6.4 ([BCMS20]). *Let $F: \mathbb{N} \rightarrow \mathbb{N}$, and $\text{CM} = (\text{Setup}, \text{Trim}, \text{Commit})$ be a commitment scheme. Fix a number of variables $M \in \mathbb{N}$ and maximum degree $N \in \mathbb{N}$. Then for every family of (possibly inefficient) functions $\{f_{\text{pp}}\}_{\text{pp}}$ and fields $\{\mathbb{F}_{\text{pp}}\}_{\text{pp}}$ where $f_{\text{pp}}: \mathcal{M}_{\text{pp}} \rightarrow \mathbb{F}_{\text{pp}}^{\leq N}[X_1, \dots, X_M]$ and $|\mathbb{F}_{\text{pp}}| \geq F(\lambda)$, and for every message format L and t -query expected polynomial time oracle algorithm \mathcal{A} , the following holds.*

$$\Pr \left[\begin{array}{l} p \neq 0 \\ \wedge \\ p(z) = 0 \end{array} \middle| \begin{array}{l} \rho \leftarrow \mathcal{U}(\lambda) \\ \text{pp} \leftarrow \text{CM.Setup}(1^\lambda, L) \\ (\ell, \mathbf{p} \in \mathcal{M}_{\text{ck}}, \omega) \leftarrow \mathcal{A}^\rho(\text{pp}) \\ \text{ck} \leftarrow \text{CM.Trim}(\text{pp}, \ell) \\ C \leftarrow \text{CM.Commit}(\text{ck}, \mathbf{p}; \omega) \\ \mathbf{z} \in \mathbb{F}_{\text{pp}}^N \leftarrow \rho(C) \\ p \leftarrow f_{\text{pp}}(\mathbf{p}) \end{array} \right] \leq \sqrt{(t+1) \cdot \frac{MN}{F(\lambda)}} + \text{negl}(\lambda) .$$

Remark 6.5. For Lemma 6.4 to hold, the algorithms of CM *must not* have access to the random oracle ρ used to generate the challenge point \mathbf{z} . The lemma is otherwise black-box with respect to CM , and so CM itself may use other oracles. The lemma continues to hold when \mathcal{A} has access to these additional oracles. We use this fact later to justify the security of domain separation.

6.3.4 Proof of knowledge soundness

We formally present the extractor below and then analyze why it satisfies the knowledge soundness definition. Let $A^\rho(\text{pp}, \text{pp}_\Phi, \text{ai})$ be the algorithm which runs $(i_\Phi, [\mathbf{q}_{i,\mathbf{x}}]_{i=1}^n, \text{acc} = (C, z, v), \pi_V = [(W_i, y_i, y'_i)]_{i=1}^n, \text{ao}) \leftarrow \tilde{P}^\rho(\text{pp}, \text{pp}_\Phi, \text{ai})$, sets $q := (z, [y_i, y'_i]_{i=1}^n)$, and outputs $(q, (i_\Phi, [\mathbf{q}_{i,\mathbf{x}}]_{i=1}^n, \text{acc}, \pi_V))$. Define the forking lemma predicate:

$\rho(\text{pp}, \text{pp}_\Phi, q, y, \Gamma)$:

1. If Γ contains a collision, output 0.
2. Parse q as $(z, [(y_i, y'_i)]_{i=1}^n)$, and y as $(i_\Phi = d, [\mathbf{q}_{i,\mathbf{x}}]_{i=1}^n, \text{acc}, \pi_V)$. If z does not match the evaluation point in $\text{acc}.\mathbf{x}$, or $[(y_i, y'_i)]_{i=1}^n$ does not match π_V , output 0.
3. Compute $(\text{apk}, \text{avk}, \text{dk}) \leftarrow \text{I}(\text{pp}, \text{pp}_\Phi, d)$.
4. Run $V(\text{avk}, [\mathbf{q}_{i,\mathbf{x}}]_{i=1}^n, \text{acc}.\mathbf{x}, \pi_V)$, answering its oracle queries according to Γ . If V asks a query outside of the support of Γ , or rejects, output 0.
5. Run $D(\text{dk}, \text{acc})$; if it rejects, output 0.
6. Otherwise, output 1.

We may assume without loss of generality that A queries the oracle at the points $(d, [(C_i, z_i, v_i, W_i)]_{i=1}^n)$ and q . The probability that A finds a collision or an inversion⁶ is $O(t^2/2^{-\lambda})$, and so the output of A satisfies ρ

⁶An *inversion* is a query whose answer is equal to the first λ bits of any previous query, including itself.

with probability $\delta - \text{negl}(\lambda)$. Let B be the algorithm given by applying Lemma 6.3 to A, p .

For the purposes of the proof, we present the extractor as an algorithm with access to a random oracle; the final extractor can then be obtained by simulating the oracle.

$E_{\bar{p}}^{\rho}(\text{pp}, \text{pp}_{\Phi})$:

1. Run $(q, \alpha_1, y_1, \dots, \alpha_{2n}, y_{2n}; \Gamma) \leftarrow B^{\rho}(\text{pp}, 1^{2n})$.
2. For $j \in [2n]$, parse y_j as $(i_{\Phi}^{(j)} = d^{(j)}, [\mathbf{q}_{i.\mathbb{X}}^{(j)}]_{i=1}^n, \pi_{\mathbb{V}}^{(j)}, \text{acc}^{(j)})$ and $\text{acc}^{(j)}$ as $((G^{(j)}, z^{(j)}, v^{(j)}), s^{(j)})$.
3. Set $\vec{s} := \begin{pmatrix} s^{(1)} \\ \vdots \\ s^{(2n)} \end{pmatrix}$, and set $A := \begin{pmatrix} 1 & \alpha_1 & \dots & \alpha_1^{2n-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha_j & \dots & \alpha_j^{2n-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha_{2n} & \dots & \alpha_{2n}^{2n-1} \end{pmatrix}$.
4. If A is invertible, compute $(\vec{p} \parallel \vec{w}) := A^{-1} \cdot \vec{s}$; otherwise, abort.
5. Output $(i_{\Phi}, \text{acc}^{(1)}, [\mathbf{q}_i = (\mathbf{q}_{i.\mathbb{X}}^{(1)}, p_i)]_{i=1}^n, \pi_{\mathbb{V}}^{(1)})$.

By Lemma 6.3, $E_{\bar{p}}$ runs in expected polynomial time and with probability $\delta - \text{negl}(\lambda)$ we have: $(\text{pp}, \text{pp}_{\Phi}, y_1) \in Z, \forall j \neq k \in [2n], \alpha_j \neq \alpha_k$ and $\forall j \in [2n], p(\text{pp}, \text{pp}_{\Phi}, q, y_j, \Gamma[q \mapsto \alpha_j]) = 1$; call this event E .

Conditioned on E , we observe the following. First, since the α_j are distinct, A is invertible. Next, let $(z, [(y_i, y'_i)]_{i=1}^n) := q$; note that $z^{(j)} = z$ and $(d^{(j)}, [\mathbf{q}_{i.\mathbb{X}}^{(j)}, W_i^{(j)}]_{i=1}^n) = \Gamma^{-1}(z) = (d^{(1)}, [\mathbf{q}_{i.\mathbb{X}}^{(1)}, W_i^{(1)}]_{i=1}^n)$ for all j , whence also $\pi_{\mathbb{V}}^{(j)} = \pi_{\mathbb{V}}^{(1)}$ for all j . For the remainder of the proof we therefore omit the superscripts on $d, z, \mathbf{q}_{i.\mathbb{X}} = (C_i, z_i, v_i)$, and $\pi_{\mathbb{V}} = [(W_i, y_i, y'_i)]_{i=1}^n$. It is immediate that the knowledge soundness predicate p is satisfied. We now show a sequence of two claims. The first shows that the extracted polynomials \vec{p}, \vec{w} are openings of the corresponding commitments, and that their evaluations at z are as claimed.

Claim 6.6. *Conditioned on E , for each $i \in [n]$, $C_i = \text{PC}_{\text{HC}}.\text{Commit}(\text{ck}, p_i)$, $W_i = \text{PC}_{\text{HC}}.\text{Commit}(\text{ck}, w_i)$, $\deg(p_i) \leq d$, $\deg(w_i) \leq d$, $p_i(z) = y_i$ and $w_i(z) = y'_i$.*

Proof. Let $\vec{C} := (C_1, \dots, C_n)$ and $\vec{W} := (W_1, \dots, W_n)$. Let $\vec{y} := (y_1, \dots, y_n)$ and $\vec{y}' := (y'_1, \dots, y'_n)$. Then for each j , let $((G^{(j)}, z, v^{(j)}), p^{(j)}) := \text{acc}^{(j)}$, and let $G := (G^{(1)}, \dots, G^{(2n)})$ and $\vec{v} := (v^{(1)}, \dots, v^{(2n)})$.

Since the verifier accepts $(\text{avk}, [\mathbf{q}_{i.\mathbb{X}}]_{i=1}^n, \text{acc}^{(j)_{.\mathbb{X}}}, \pi_{\mathbb{V}})$ for all j , it holds that $\vec{v} = A(\vec{y} \parallel \vec{y}')$, and $\vec{G} = A(\vec{C} \parallel \vec{W})$. Moreover, since the decider accepts $(\text{avk}, \text{acc}^{(j)})$ for all j , it holds for all j that $G^{(j)} = \text{PC}_{\text{HC}}.\text{Commit}(\text{ck}, s^{(j)})$, $s^{(j)}(z) = v^{(j)}$ and $\deg(s^{(j)}) \leq d$, from which the degree bounds on p_i, w_i follow directly.

Using the homomorphic property of PC_{HC} , and because $(\vec{C} \parallel \vec{W}) = A^{-1}\vec{G}$, it holds that $C_i = \sum_j A_{i,j}^{-1} G^{(j)} = \text{PC}_{\text{HC}}.\text{Commit}(\text{ck}, \sum_j A_{i,j}^{-1} s^{(j)}) = \text{PC}_{\text{HC}}.\text{Commit}(\text{ck}, p_i)$. Similarly, $W_i = \text{PC}_{\text{HC}}.\text{Commit}(\text{ck}, w_i)$. In addition, since $(\vec{y}, \vec{y}') = A^{-1}\vec{v}$, and $p^{(j)}(z) = v^{(j)}$, we have that $p_i(z) = \sum_j A_{i,j}^{-1} v^{(j)} = y_i$, and $w_i(z) = \sum_j A_{n+i,j}^{-1} v^{(j)} = y'_i$. \square

The second claim shows that the evaluations of the p_i on the original query points z_i are as claimed.

Claim 6.7. *With probability at least $\delta - \text{negl}(\lambda)$, it holds that for all $i \in [n]$, $p_i(z_i) = v_i$.*

Proof. Consider a slight modification to $E_{\bar{p}}$ which also outputs \vec{w} . From the previous claim, if E occurs, then for each i , the tuple (C_i, z_i, v_i, W_i) constitutes a binding commitment to the polynomial $p_i(X) - v_i - w_i(X)(X - z_i)$ of degree at most $d + 1$, and $p_i(z) = y_i, w_i(z) = y'_i$. Since the verifier accepts the output

of $E_{\bar{p}}$, we have that $z = \Gamma(d, [(C_i, z_i, v_i, W_i)]_{i=1}^n)$, and that $y_i - v_i = y'_i(z - z_i)$. By Lemma 6.4, with probability $\delta - \text{negl}(\lambda)$, $p_i(X) - v_i - w_i(X)(X - z_i)$ is the zero polynomial, and so $p_i(z_i) = v_i$. \square

Together these claims establish that with probability at least $\delta - \text{negl}(\lambda)$, for all $i \in [n]$ it holds that $\text{PC}_{\text{HC}}.\text{Check}(\text{rk}, (C_i, z_i, v_i), p_i) = 1$, which completes the proof of knowledge soundness.

7 Implementation

We contribute a generic and modular implementation of proof-carrying data based on accumulation schemes. Our implementation includes several components of independent interest.

Framework for accumulation. We design and implement a generic framework for accumulation schemes that supports arbitrary predicates/relations. The main interface is a Rust trait that defines the behavior of any (atomic or split) accumulation scheme. We implement this trait for accumulation schemes for popular polynomial commitments (and thus for popular SNARKs that rely on PC schemes, such as [CHMMVW20]). In more detail, we implemented:

- the atomic accumulation scheme AS_{AGM} in [BCMS20] for the PC scheme PC_{AGM} ;
- the atomic accumulation scheme AS_{IPA} in [BCMS20] for the PC scheme PC_{IPA} ;
- the split accumulation scheme AS_{HC} in Section 6 for PC_{HC} .

Our framework also provides a generic trait for defining RICS constraints for the verifier of an accumulation scheme. We use this trait to implement RICS constraints for all these accumulation schemes.

PCD from accumulation. We provide a generic construction of PCD from accumulation, which simultaneously supports the case of atomic accumulation from [BCMS20] and the case of split accumulation from Section 5. Our code builds on and extends an existing PCD library that offers a generic “PCD” trait.⁷ We instantiate this PCD trait via a modular construction, which takes as ingredients any NARK (as defined by an appropriate trait), accumulation scheme for that NARK that implements the accumulation trait (from above), and constraints for the accumulation verifier. We use our concrete instantiations of these ingredients to construct PCD that is based on atomic accumulation of PC_{AGM} and PC_{IPA} , and split accumulation of PC_{HC} .

Cycles of elliptic curves. All PCD constructions in our implementation rely on the technique of cycles of elliptic curves [BCTV14]: PCD based on PC_{AGM} uses cycles of pairing-friendly curves, while PCD based on PC_{IPA} and PC_{HC} uses cycles of standard curves. For all of these, we rely on existing implementations from the `arkworks` ecosystem:⁸ for pairing-friendly cycles we use the MNT cycle of curves (low security and high security variants), while for standard cycles we use the Pasta cycle of curves [Hop20].

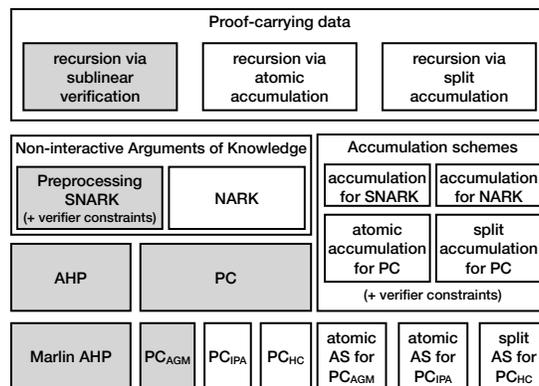


Figure 5: Diagram illustrating components in our implementation. The gray boxes denote components that exist in prior libraries, while the white boxes denote components contributed in this work.

⁷<https://github.com/arkworks-rs/pcd>

⁸<https://github.com/arkworks-rs/curves>

8 Evaluation

We perform an evaluation focused on understanding the tradeoffs of atomic versus split accumulation in the discrete logarithm setting.⁹ We compare two routes for recursion:

- recursion based on the atomic accumulation scheme AS_{IPA} in [BCMS20] for the PC scheme PC_{IPA} ;
- recursion based on the split accumulation scheme AS_{HC} in Section 6 for PC_{HC} .

We proceed as follows. In Section 8.1 we compare the two polynomial commitment schemes PC_{IPA} and PC_{HC} , and in Section 8.2 we compare the two corresponding accumulation schemes AS_{IPA} and AS_{HC} .

Experimental setup. All experiments are performed on a machine with an Intel Xeon 6136 CPU at 3.0 GHz. The reported numbers are for schemes instantiated over the 255-bit prime-order Pallas curve in the Pasta cycle [Hop20]; results for the Vesta curve in that cycle would be similar.

8.1 Comparing polynomial commitments based on DLs

We compare the performance of PC_{IPA} and PC_{HC} in Table 2, reporting experiments for an illustrative choice of polynomial degree d . In both PC schemes all operations (commit, open, check) are linear in the degree d , though for PC_{HC} opening is concretely much cheaper than PC_{IPA} (primarily because PC_{HC} has a trivial opening procedure). The main difference between the two PC schemes is that an evaluation proof in PC_{HC} is $O(d)$ field elements while an evaluation proof in PC_{IPA} is $O(\log d)$ group elements; this asymptotic difference is apparent in the reported numbers (the proof size for PC_{HC} is significantly larger than for PC_{IPA}). We also report lines of code to realize the same abstract PC scheme trait, to support the (intuitive) claim that PC_{HC} is a much simpler primitive than PC_{IPA} .

PC scheme	Commit	Open	Check	$ C $	$ \pi $	LoC
PC_{IPA}	8.0 s	106.6 s	8.2 s	33 B	1.4 kB $O(\log d) \mathbb{G}$	1120
PC_{HC}	8.1 s	0.43 s	8.3 s	33 B	33.5 MB $O(d) \mathbb{F}$	608

Table 2: Comparison between the PC schemes PC_{IPA} and PC_{HC} for polynomials of degree $d = 2^{20}$.

8.2 Comparing accumulation schemes based on DLs

We compare the performance of AS_{IPA} and AS_{HC} in Table 3, reporting experiments for an illustrative choice of polynomial degree d . We focus on the special case where the accumulation scheme is used to accumulate one new polynomial evaluation claim into one old accumulator to obtain a new accumulator. Our experiments indicate that AS_{HC} is cheaper than AS_{IPA} across all metrics *except for accumulator size*, and more generally that performance is consistent with the asymptotic comparison from Table 1. In more detail:

- While prover time (per claim) in both AS_{IPA} and AS_{HC} are linear in the degree d , our experiments show that AS_{IPA} is concretely much more expensive than AS_{HC} .
- Decider time in both AS_{IPA} and AS_{HC} are linear in the degree d , and our experiments show that the two schemes have similar concrete performance.
- Verifier time (per claim) in AS_{IPA} is logarithmic while in AS_{HC} it is constant, and our experiments confirm that AS_{IPA} is concretely significantly more expensive than AS_{HC} .
- Verifier constraint cost is *much* higher for AS_{IPA} , even though both schemes use the same underlying constraint gadget libraries.

⁹The pairing setting is also part of our implementation, as described in Section 7, but we do not include an evaluation for it here.

- The size of an atomic accumulator for AS_{IPA} is logarithmic, and amounts to a few kilobytes; in contrast an accumulator for AS_{HC} is much larger, but is split into a short instance part (106 bytes) and a long witness part (33.5 megabytes).

Overall the expensive parts of AS_{HC} are exactly where intended (a large accumulation witness part) in exchange for a very cheap verifier and a very short accumulation instance part; all other metrics are comparable to (and concretely better than for) AS_{IPA} .

scheme	P	V	D	acc		V	LoC	
				x	w		native	constraints
AS_{IPA}	117.6 s	14 ms	8.3 s	1.58 kB	0	1006×10^3	664	1232
AS_{HC}	25.2 s	2 ms	8.1 s	106 B	33.5 MB	61×10^3	571	395

Table 3: Comparison between the accumulation schemes AS_{IPA} and AS_{HC} for polynomials of degree $d = 2^{20}$, when accumulating one old accumulator and one evaluation claim into a new accumulator.

In Figure 6, we also compare $|V|$ (the constraint cost of V) in both AS_{HC} and AS_{IPA} as we accumulate polynomial evaluation claims of degree d in the range 2^{10} to 2^{20} . As expected, the cost for AS_{HC} is a small constant, whereas the cost of AS_{IPA} grows logarithmically (and is concretely much larger).

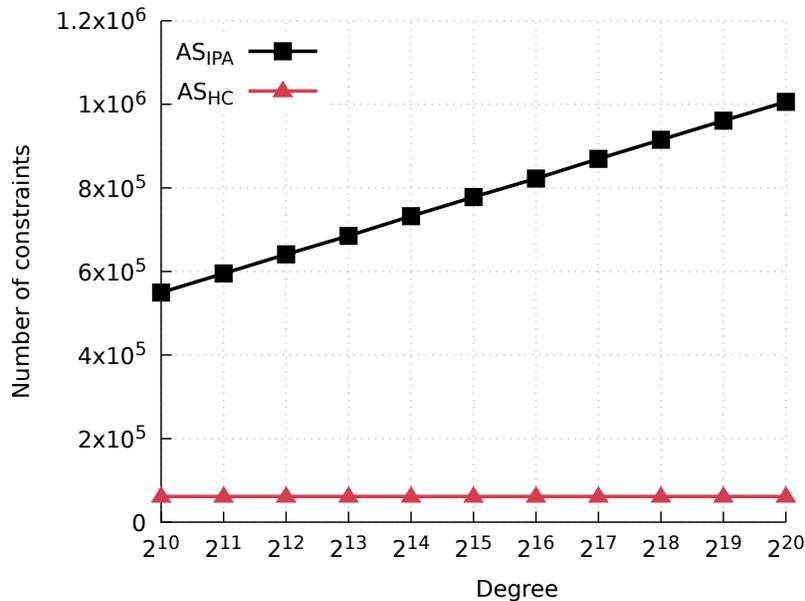


Figure 6: Comparison of the constraint cost of the accumulation verifier V in AS_{IPA} and AS_{HC} when varying the degree of the accumulated polynomial from 2^{10} to 2^{20} .

A Split accumulation scheme for R1CS

We describe a split accumulation scheme $AS = (G, I, P, V, D)$ for the NARK for R1CS in Figure 3. As a subroutine we use an accumulation scheme $AS_{\text{HC}} = (G_{\text{HC}}, I_{\text{HC}}, P_{\text{HC}}, V_{\text{HC}}, D_{\text{HC}})$ for the Pedersen PC scheme PC_{HC} (e.g., the one we construct in Section 6). We use domain separation on the given random oracle ρ for different tasks: we use ρ_{HC} to denote the oracle used for one invocation of AS_{HC} ; ρ_{NARK} to denote the oracle used to run the NARK for R1CS; and ρ_{AS} to denote the random oracle used by AS for other tasks.

Predicate inputs. The predicate to accumulate is the NARK verifier, with the following split in a predicate input q obtained from an R1CS instance x and proof π :

- The instance part of q consists of the R1CS input x , the first prover message $(C_w, C_A, C_B, C_C, C_q, C_w)$, and the claimed evaluations sent in the second prover message $(v_A, v_B, v_C, v_q, v_H)$. This amounts to 6 group elements and $|x| + 5$ field elements (which is short).
- The witness part of q consists of the openings sent in the second prover message (w, z_A, z_B, z_C, q) . This amounts to $O(n + m)$ field elements (which is proportional to the dimensions of the R1CS matrices).

Accumulator. The format of an accumulator acc is as follows:

- The instance part of acc consists of $\text{acc}.\mathbb{x} = (\text{acc}_{\text{HC}}.\mathbb{x}, x, C_w, C_A, C_B, C_C)$.
- The witness part of acc consists of $\text{acc}.\mathbb{w} = (\text{acc}_{\text{HC}}.\mathbb{w}, w, z_A, z_B, z_C)$.

Generator. The generator G runs G_{HC} as a subroutine and outputs its output.

Indexer. The indexer I receives as input accumulation public parameters pp (output by G), predicate public parameters $\text{pp}_{\Phi} = \text{pp}_{\text{NARK}}$ (they are the public parameters of the NARK per Definition 4.1), predicate index $i_{\Phi} = (A, B, C)$ (it is the index of the relation verified by the NARK per Definition 4.1), and works as follows:

- Hash the coefficient matrices: $\sigma := \rho_{\text{AS}}(A, B, C)$ and $\tau := \rho_{\text{NARK}}(A, B, C)$.
- Set the degree bound to be $d := m$, the number of rows in each R1CS coefficient matrix.
- Trim the public parameters of PC_{HC} for degree d : $\text{ck} \leftarrow \text{CM.Trim}(\text{pp}_{\text{NARK}}, d + 1)$. (In the NARK the public parameters are the public parameters for the commitment scheme.)
- Invokes $I_{\text{HC}}(\text{pp}, \text{pp}_{\Phi}, d)$ to obtain $(\text{apk}_{\text{HC}}, \text{avk}_{\text{HC}}, \text{dk}_{\text{HC}})$.
- Output $(\text{apk}, \text{avk}, \text{dk}) := ((\sigma, \tau, \text{apk}_{\text{HC}}), (\sigma, \tau, \text{avk}_{\text{HC}}), ((A, B, C), \text{ck}, \text{dk}_{\text{HC}}))$.

Decider. The decider D receives as input the decision key $\text{dk} = ((A, B, C), \text{ck}, \text{dk}_{\text{HC}})$ and an accumulator acc (both instance and witness parts) and determines the validity of the accumulator as follows.

1. Set $z := (\text{acc}.x, \text{acc}.w)$.
2. Check that $\text{acc}.C_w = \text{CM.Commit}(\text{ck}, \text{acc}.w)$.
3. Check that $\text{acc}.C_A = \text{CM.Commit}(\text{ck}, \text{acc}.z_A)$, $\text{acc}.C_B = \text{CM.Commit}(\text{ck}, \text{acc}.z_B)$, $\text{acc}.C_C = \text{CM.Commit}(\text{ck}, \text{acc}.z_C)$.
4. Check that $\text{acc}.z_A = A \cdot z$, $\text{acc}.z_B = B \cdot z$, $\text{acc}.z_C = Cz$.
5. Check that $D_{\text{HC}}(\text{dk}_{\text{HC}}, \text{acc}.\text{acc}_{\text{HC}}) = 1$.

Prover. The prover P , given a random oracle ρ , receives as input the accumulator proving key $\text{apk} = (\sigma, \tau, \text{apk}_{\text{HC}})$, inputs $[q_i = (q.\mathbb{x}_i, q.\mathbb{w}_i)]_{i=1}^n$, and old accumulators $[\text{acc}_j = (\text{acc}.\mathbb{x}_j, \text{acc}.\mathbb{w}_j)]_{j=1}^m$, and outputs a new accumulator $\text{acc} = (\text{acc}.\mathbb{x}, \text{acc}.\mathbb{w})$ and a proof π_V that are computed as follows.

1. Compute $\alpha := \rho_{\text{AS}}(\sigma, [\text{acc}_j.\mathbb{x}]_{j=1}^m, [q_i.\mathbb{x}]_{i=1}^n)$.
2. For each $i \in [n]$, compute a combined polynomial evaluation instance $q_{\text{HC},i}.\mathbb{x} := (C, \gamma, v)$ and witness $q_{\text{HC},i}.\mathbb{w} := f$ where

$$q_{\text{HC},i}.\mathbb{x}.C := q_i.C_A + \alpha \cdot q_i.C_B + \alpha^2 \cdot q_i.C_C + \alpha^3 \cdot q_i.C_q + \alpha^4 \cdot q_i.C_H ,$$

$$\begin{aligned}
\mathbf{q}_{\text{HC},i,\mathbb{X}}.\gamma &:= \rho_{\text{NARK}}(\tau, \mathbf{q}_i.x, \mathbf{q}_i.\{C_w, C_A, C_B, C_C, C_q, C_H\}) , \\
\mathbf{q}_{\text{HC},i,\mathbb{X}}.v &:= \mathbf{q}_i.v_A + \alpha \cdot \mathbf{q}_i.v_B + \alpha^2 \cdot \mathbf{q}_i.v_C + \alpha^3 \cdot \mathbf{q}_i.v_q + \alpha^4 \cdot \mathbf{q}_i.v_H , \\
\mathbf{q}_{\text{HC},i,\mathbb{W}}.f &:= \mathbf{q}_i.z_A + \alpha \cdot \mathbf{q}_i.z_B + \alpha^2 \cdot \mathbf{q}_i.z_C + \alpha^3 \cdot \mathbf{q}_i.q + \alpha^4 \cdot \mathbf{q}_i.V_H .
\end{aligned}$$

3. Accumulate all polynomial evaluations: $(\text{acc}_{\text{HC}}, \pi_{\text{V,HC}}) := \text{P}^{\rho_{\text{HC}}}(\text{apk}_{\text{HC}}, [\text{acc}_j.\text{acc}_{\text{HC}}]_{j=1}^m, [\mathbf{q}_{\text{HC},i}]_{i=1}^n)$.
4. Accumulate the linear computation for the instance part:

$$\begin{aligned}
\text{acc}.x &:= \sum_{j=1}^m \alpha^j \cdot \text{acc}_j.x + \sum_{i=1}^n \alpha^{m+i} \cdot \mathbf{q}_i.x , \\
\text{acc}.C_w &:= \sum_{j=1}^m \alpha^j \cdot \text{acc}_j.C_w + \sum_{i=1}^n \alpha^{m+i} \cdot \mathbf{q}_i.C_w , \\
\text{acc}.C_A &:= \sum_{j=1}^m \alpha^j \cdot \text{acc}_j.C_A + \sum_{i=1}^n \alpha^{m+i} \cdot \mathbf{q}_i.C_A , \\
\text{acc}.C_B &:= \sum_{j=1}^m \alpha^j \cdot \text{acc}_j.C_B + \sum_{i=1}^n \alpha^{m+i} \cdot \mathbf{q}_i.C_B , \\
\text{acc}.C_C &:= \sum_{j=1}^m \alpha^j \cdot \text{acc}_j.C_C + \sum_{i=1}^n \alpha^{m+i} \cdot \mathbf{q}_i.C_C .
\end{aligned}$$

5. Accumulate the linear computation for the witness part:

$$\begin{aligned}
\text{acc}.w &:= \sum_{j=1}^m \alpha^j \cdot \text{acc}_j.w + \sum_{i=1}^n \alpha^{m+i} \cdot \mathbf{q}_i.w , \\
\text{acc}.z_A &:= \sum_{j=1}^m \alpha^j \cdot \text{acc}_j.z_A + \sum_{i=1}^n \alpha^{m+i} \cdot \mathbf{q}_i.z_A , \\
\text{acc}.z_B &:= \sum_{j=1}^m \alpha^j \cdot \text{acc}_j.z_B + \sum_{i=1}^n \alpha^{m+i} \cdot \mathbf{q}_i.z_B , \\
\text{acc}.z_C &:= \sum_{j=1}^m \alpha^j \cdot \text{acc}_j.z_C + \sum_{i=1}^n \alpha^{m+i} \cdot \mathbf{q}_i.z_C .
\end{aligned}$$

6. Set $\text{acc}.\mathbb{X} := (\text{acc}_{\text{HC}}.\mathbb{X}, x, C_w, C_A, C_B, C_C)$ and $\text{acc}.\mathbb{W} := (\text{acc}_{\text{HC}}.\mathbb{W}, w, z_A, z_B, z_C)$.
7. Output the new accumulator $\text{acc} := (\text{acc}.\mathbb{X}, \text{acc}.\mathbb{W})$ and accumulation proof $\pi_{\text{V}} := \pi_{\text{V,HC}}$.

Verify. The verifier V , given a random oracle ρ , receives as input the accumulator verification key $\text{avk} = (\sigma, \tau, \text{avk}_{\text{HC}})$, input instances $[\mathbf{q}_i.\mathbb{X}]_{i=1}^n$, accumulator instances $[\text{acc}_j.\mathbb{X}]_{j=1}^m$, a new accumulator instance $\text{acc}.\mathbb{X}$, and a proof π_{V} , and determines whether to accept or reject as follows.

1. Compute $\alpha \in \mathbb{F}$ as in Item 1 of the accumulation prover.
2. For each $i \in [n]$, compute the polynomial evaluation instance $\mathbf{q}_{\text{HC},i,\mathbb{X}}$ as in Item 2 of the accumulation prover (excluding the witness computations) and check that

$$\mathbf{q}_i.v_q \cdot \mathbf{q}_i.v_H = \mathbf{q}_i.v_A \cdot \mathbf{q}_i.v_B - \mathbf{q}_i.v_C .$$

3. Verify accumulation of polynomial evaluations: $\text{V}^{\rho_{\text{HC}}}(\text{avk}_{\text{HC}}, [\text{acc}_j.\text{acc}_{\text{HC}}]_{j=1}^m, [\mathbf{q}_{\text{HC},i}]_{i=1}^n, \text{acc}_{\text{HC}}, \pi_{\text{V,HC}}) = 1$.
4. Check the accumulation of linear computations by performing the assignments in Item 4 of the accumulation prover as equality checks (between the new accumulator instance and the input instances and old accumulator instances).

Acknowledgements

This research was supported in part by the Ethereum Foundation, NSF, DARPA, a grant from ONR, and the Simons Foundation. Nicholas Spooner was supported by DARPA under Agreement No. HR00112020023.

References

- [BBBPWM18] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. “Bulletproofs: Short Proofs for Confidential Transactions and More”. In: *Proceedings of the 39th IEEE Symposium on Security and Privacy*. S&P ’18. 2018, pp. 315–334.
- [BCCGP16] J. Bootle, A. Cerulli, P. Chaidos, J. Groth, and C. Petit. “Efficient Zero-Knowledge Arguments for Arithmetic Circuits in the Discrete Log Setting”. In: *Proceedings of the 35th Annual International Conference on Theory and Application of Cryptographic Techniques*. EUROCRYPT ’16. 2016, pp. 327–357.
- [BCCT13] N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer. “Recursive Composition and Bootstrapping for SNARKs and Proof-Carrying Data”. In: *Proceedings of the 45th ACM Symposium on the Theory of Computing*. STOC ’13. 2013, pp. 111–120.
- [BCMS20] B. Bünz, A. Chiesa, P. Mishra, and N. Spooner. “Proof-Carrying Data from Accumulation Schemes”. In: TCC ’20 (2020).
- [BCS16] E. Ben-Sasson, A. Chiesa, and N. Spooner. “Interactive Oracle Proofs”. In: *Proceedings of the 14th Theory of Cryptography Conference*. TCC ’16-B. 2016, pp. 31–60.
- [BCTV14] E. Ben-Sasson, A. Chiesa, E. Tromer, and M. Virza. “Scalable Zero Knowledge via Cycles of Elliptic Curves”. In: *Proceedings of the 34th Annual International Cryptology Conference*. CRYPTO ’14. 2014, pp. 276–294.
- [BDFG20] D. Boneh, J. Drake, B. Fisch, and A. Gabizon. *Halo Infinite: Recursive zk-SNARKs from any Additive Polynomial Commitment Scheme*. Cryptology ePrint Archive, Report 2020/1536. 2020.
- [BGH19] S. Bowe, J. Grigg, and D. Hopwood. *Halo: Recursive Proof Composition without a Trusted Setup*. Cryptology ePrint Archive, Report 2019/1021. 2019.
- [BMMTV19] B. Bünz, M. Maller, P. Mishra, N. Tyagi, and P. Vesely. *Proofs for Inner Pairing Products and Applications*. Cryptology ePrint Archive, Report 2019/1177. 2019.
- [BMRS20] J. Bonneau, I. Meckler, V. Rao, and E. Shapiro. *Coda: Decentralized Cryptocurrency at Scale*. Cryptology ePrint Archive, Report 2020/352. 2020.
- [BN06] M. Bellare and G. Neven. “Multi-signatures in the plain public-key model and a general forking lemma”. In: *Proceedings of the 13th ACM Conference on Computer and Communications Security*. CCS ’06. 2006, pp. 390–399.
- [BR93] M. Bellare and P. Rogaway. “Random Oracles Are Practical: A Paradigm for Designing Efficient Protocols”. In: *Proceedings of the 1st ACM Conference on Computer and Communications Security*. CCS ’93. 1993, pp. 62–73.
- [CCDW20] W. Chen, A. Chiesa, E. Dauterman, and N. P. Ward. *Reducing Participation Costs via Incremental Verification for Ledger Systems*. Cryptology ePrint Archive, Report 2020/1522. 2020.
- [CHMMVW20] A. Chiesa, Y. Hu, M. Maller, P. Mishra, N. Vesely, and N. Ward. “Marlin: Preprocessing zkSNARKs with Universal and Updatable SRS”. In: *Proceedings of the 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’20. 2020, pp. 738–768.
- [COS20] A. Chiesa, D. Ojha, and N. Spooner. “Fractal: Post-Quantum and Transparent Recursive Proofs from Holography”. In: *Proceedings of the 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’20. 2020, pp. 769–793.

- [CT10] A. Chiesa and E. Tromer. “Proof-Carrying Data and Hearsay Arguments from Signature Cards”. In: *Proceedings of the 1st Symposium on Innovations in Computer Science*. ICS ’10. 2010, pp. 310–331.
- [CTV13] S. Chong, E. Tromer, and J. A. Vaughan. *Enforcing Language Semantics Using Proof-Carrying Data*. Cryptology ePrint Archive, Report 2013/513. 2013.
- [CTV15] A. Chiesa, E. Tromer, and M. Virza. “Cluster Computing in Zero Knowledge”. In: *Proceedings of the 34th Annual International Conference on Theory and Application of Cryptographic Techniques*. EUROCRYPT ’15. 2015, pp. 371–403.
- [GT20] A. Ghoshal and S. Tessaro. *Tight State-Restoration Soundness in the Algebraic Group Model*. Cryptology ePrint Archive, Report 2020/1351. 2020.
- [Halo20] S. Bowe, J. Grigg, and D. Hopwood. *Halo2*. 2020. URL: <https://github.com/zcash/halo2>.
- [Hop20] D. Hopwood. *The Pasta Curves for Halo 2 and Beyond*. <https://electriccoin.co/blog/the-pasta-curves-for-halo-2-and-beyond/>. 2020.
- [KB20] A. Kattis and J. Bonneau. *Proof of Necessary Work: Succinct State Verification with Fairness Guarantees*. Cryptology ePrint Archive, Report 2020/190. 2020.
- [KZG10] A. Kate, G. M. Zaverucha, and I. Goldberg. “Constant-Size Commitments to Polynomials and Their Applications”. In: *Proceedings of the 16th International Conference on the Theory and Application of Cryptology and Information Security*. ASIACRYPT ’10. 2010, pp. 177–194.
- [Lin03] Y. Lindell. “Parallel Coin-Tossing and Constant-Round Secure Two-Party Computation”. In: *Journal of Cryptology* 16.3 (2003), pp. 143–184.
- [Mina] O(1) Labs. *Mina Cryptocurrency*. <https://minaprotocol.com/>. 2017.
- [NT16] A. Naveh and E. Tromer. “PhotoProof: Cryptographic Image Authentication for Any Set of Permissible Transformations”. In: *Proceedings of the 37th IEEE Symposium on Security and Privacy*. S&P ’16. 2016, pp. 255–271.
- [Pas03] R. Pass. “On Deniability in the Common Reference String and Random Oracle Model”. In: *Proceedings of the 23rd Annual International Cryptology Conference*. CRYPTO ’03. 2003, pp. 316–337.
- [Ped92] T. P. Pedersen. “Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing”. In: *Proceedings of the 11th Annual International Cryptology Conference*. CRYPTO ’91. 1992, pp. 129–140.
- [Pickles20] O(1) Labs. *Pickles*. URL: <https://github.com/o1-labs/marlin>.
- [Val08] P. Valiant. “Incrementally Verifiable Computation or Proofs of Knowledge Imply Time/Space Efficiency”. In: *Proceedings of the 5th Theory of Cryptography Conference*. TCC ’08. 2008, pp. 1–18.
- [WTSTW18] R. S. Wahby, I. Tzialla, A. Shelat, J. Thaler, and M. Walfish. “Doubly-Efficient zkSNARKs Without Trusted Setup”. In: *Proceedings of the 39th IEEE Symposium on Security and Privacy*. S&P ’18. 2018, pp. 926–943.