

On Knowledge-Soundness of Plonk in ROM from Falsifiable Assumptions*

October 10, 2024

Helger Lipmaa ¹, Roberto Parisella ², and Janno Siim ²

¹ University of Tartu, Tartu, Estonia

² Simula UiB, Bergen, Norway

Abstract. Lipmaa, Parisella, and Siim [Eurocrypt, 2024] proved the extractability of the KZG polynomial commitment scheme under the falsifiable assumption ARSDH. They also showed that variants of real-world zk-SNARKs like Plonk can be made knowledge-sound in the random oracle model (ROM) under the ARSDH assumption. However, their approach did not consider various batching optimizations, resulting in their variant of Plonk having approximately 3.5 times longer argument. Our contributions are: (1) We prove that several batch-opening protocols for KZG, used in modern zk-SNARKs, have computational special-soundness under the ARSDH assumption. (2) We prove that interactive Plonk has computational special-soundness under the ARSDH assumption and a new falsifiable assumption SplitRSDH. We also prove that two minor modifications of the interactive Plonk have computational special-soundness under only the ARSDH and a simpler variant of SplitRSDH. We define a new type-safe oracle framework of the AGMOS (AGM with oblivious sampling) and prove SplitRSDH is secure in it. The Fiat-Shamir transform can be applied to obtain non-interactive versions, which are secure in the ROM under the same assumptions.

Keywords: Batching · KZG · Plonk · special-soundness · zk-SNARKs

1 Introduction

As shown by Gentry and Wichs [GW11], the soundness of SNARKs necessarily relies on non-falsifiable assumptions. However, their result applies only in the standard model, where one does not use random oracles. In the current landscape of universal zk-SNARKs, where the trusted setup does not depend on the computation, we can observe three ways (a trichotomy) to consider Gentry-Wichs. First, the GKMMM zk-SNARK of Groth et al. [GKM⁺18] uses non-falsifiable knowledge assumptions but does not use random oracles. However, GKMMM has a quadratic-length Structured Reference String (SRS) and is thus only of theoretical interest. Second, many well-known zk-SNARKs like Bulletproofs [BBB⁺18],

* Second eprint. Substantial rewrite of the previous eprint version, with several new results, including remodeling AGMOS, a more standard security assumption, a new analysis of Lin (including the impossibility result), and the analysis of SmallPlonk.

Brakedown [GLS⁺23], and FRI-based [BBHR18] STARKs work in the random oracle model (ROM) and only use falsifiable assumptions. While zk-SNARKs of this category excel in other aspects (e.g., they offer excellent prover time [DP23] or a transparent SRS), they have a longer argument and higher verification costs.

Until recently, constant-argument and linear-SRS length zk-SNARKs relied on both the ROM and non-falsifiable assumptions (including an idealized group model like the Generic Group Model (GGM) [Sho97,Mau05] or the Algebraic Group Model (AGM) [FKL18]). This includes well-known updatable zk-SNARKs [GWC19,CHM⁺20,RZ21,CFF⁺21,LSZ22]. In particular, Plonk [GWC19] is widely deployed in practice due to its general efficiency and the use of the Plonkish arithmetization that allows for custom gates [GW19] and lookup arguments [GW20,EFG22,STW24,CFF⁺24]. Thus, this paper focuses on Plonk, though our techniques are broadly applicable.

Both the ROM [CGH98] and idealized group models [Den02] are known to be non-instantiable. Using only a single idealized model in a security proof is preferable, whether it be the ROM, the AGM, or the GGM. The ROM is more established and better understood, making it desirable to avoid idealized group models as much as possible. In addition, the understanding of the AGM is lacking, with several recent papers finding bugs and proposing fixes [Zha22,ZZK22,LPS23,BFHK23,JM24]. Using the terminology of Jaeger and Mohan [JM24], most of the found bugs were illusory or ambiguous: they relied on one possible interpretation of AGM and disappeared in a different interpretation of AGM. Following Zhandry [Zha22], Jaeger and Mohan proposed a coherent interpretation of the AGM, in particular, proving that under reasonable assumptions, AGM security implies GGM security, [JM24].

Separately, Lipmaa et al. [LPS23] proposed AGM with Oblivious Sampling (AGMOS), a more realistic variant of the AGM where the adversary has an additional capacity to sample group elements obliviously.³ They noted that a common use of the seminal KZG polynomial commitment scheme [KZG10], secure in the AGM, is not secure in the AGMOS and thus not in the standard model. While [JM24]’s framework for AGM is more rigorous, they also do not consider oblivious sampling, which can significantly affect the security of KZG-based zk-SNARKs. One also faces the tedious but necessary job of rewriting many existing AGM proofs. Only a few AGMOS proofs exist (e.g., for Plonk [FFR24] and Poly-math [Lip24]); the security of other KZG-based zk-SNARKs in the AGMOS is still unproven. Because idealized group models frequently change, using them as minimally as possible is desirable. For example, such models are a valuable tool for sanity-checking new standard-looking falsifiable assumptions, which has been their primary (and significantly less controversial) use case in the past.

The trichotomic situation changed in 2024 when Lipmaa et al. (LPS, [LPS24]) proved that the KZG polynomial commitment scheme [KZG10] is sound under a new falsifiable assumption ARSDH (Adaptive Rational

³ Oblivious sampling means sampling a group element without knowing the respective discrete logarithm. In elliptic curve groups, it is possible with hash-and-increment or admissible encodings [Ica09,LPS23].

Table 1. Comparison of different versions of **Plonk**, secure under falsifiable assumptions. The number of bits are given for the BLS381-12 pairing $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$. Additionally, n denotes the number of gates, and ℓ is the number of elements in the public input of the circuit whose satisfiability is being proven.

Argument	Proof size (bits)	Prover	Verifier	Assumptions
Plonk’s polynomial IOP compiled with [LPS24]	$23 \mathbb{F} + 30 \mathbb{G}_1 $ (17408)	$30n \mathbb{G}_1$ Exp., $\mathcal{O}(n \log n) \mathbb{F}$ Ops.	46 Pair., $24 \mathbb{G}_1$ Exp., $\mathcal{O}(\ell + \log n) \mathbb{F}$ Ops.	ARSDH
SanPlonk (current work)	$7 \mathbb{F} + 9 \mathbb{G}_1 $ (5248)	$9n \mathbb{G}_1$ Exp., $\mathcal{O}(n \log n) \mathbb{F}$ Ops.	2 Pair., 19 \mathbb{G}_1 Exp., $\mathcal{O}(\ell + \log n) \mathbb{F}$ Ops.	ARSDH
Plonk (current work)	$6 \mathbb{F} + 9 \mathbb{G}_1 $ (4992)	$9n \mathbb{G}_1$ Exp., $\mathcal{O}(n \log n) \mathbb{F}$ Ops.	2 Pair., 18 \mathbb{G}_1 Exp., $\mathcal{O}(\ell + \log n) \mathbb{F}$ Ops.	ARSDH, SplitRSDH
SmallPlonk (current work)	$6 \mathbb{F} + 7 \mathbb{G}_1 $ (4224)	$11n \mathbb{G}_1$ Exp., $\mathcal{O}(n \log n) \mathbb{F}$ Ops.	2 Pair., 16 \mathbb{G}_1 Exp., $\mathcal{O}(\ell + \log n) \mathbb{F}$ Ops.	ARSDH, SplitRSDH

Strong Diffie-Hellman assumption). Importantly, they only use AGM(OS) to build confidence that ARSDH is secure. LPS proposed a compiler that uses KZG to convert an efficient polynomial IOP (PIOP, [BFS20]) like Plonk’s to a constant-communication zk-SNARK. Since Plonk and the zk-SNARKs [GWC19, CHM⁺20, RZ21, CFF⁺21, LSZ22] combine polynomial IOPs and KZG, the LPS compiler results in zk-SNARKs that have both constant-length argument and linear-length SRS. Importantly, the security of the resulting zk-SNARK relies only on the ROM and the ARSDH assumption.

Unfortunately, the LPS compiler has serious drawbacks. First, it works only with abstract PIOPs, meaning one cannot operate on polynomial commitments without opening them. Thus, one cannot use batching techniques like the classical “linearization trick” [GWC19, CHM⁺20]. (We will explain this trick in Sections 1.1 and 4.2.) Plonk, being highly optimized, uses batching techniques very aggressively. One reason why [LPS24] does not handle batching techniques is that [LPS23] observed that the linearization trick is sometimes not secure in the AGMOS. Very recently, Faonio et al. [FFR24] defined a precise condition under which the linearization trick is secure in the AGMOS; they used this to prove that Plonk is knowledge-sound in the AGMOS. Their proof techniques heavily rely on the properties of the AGMOS and seem not to work in the standard model. It is possible that the interactive Plonk is insecure in the standard model.

The LPS compiler introduces even more overhead: compared to the original PIOP, the prover has to open each polynomial commitment at a new random point. Put together, [LPS24] proves that a *variant* of (interactive) Plonk is knowledge-sound under the ARSDH assumption. However, the variant is much less efficient than the “real-world” Plonk. The most striking is the verifier’s slowdown, who needs to implement 46 pairings instead of 2. See Table 1 for a com-

parison. We believe practitioners are unwilling to sacrifice efficiency by such a factor or change their battle-tested codebases.

An additional weakness of [LPS24] is that similarly to [GWC19], it proves *knowledge-soundness* of the interactive Plonk. As well-known [AFK22], knowledge-sound interactive protocols can suffer a significant security loss after Fiat-Shamir compilation to zk-SNARKs. Practitioners often assume tight security for Fiat-Shamir zk-SNARKs in their implementations. Given that [GWC19,LPS24] prove that the interactive Plonk is knowledge-sound, this holds only heuristically. Achieving tight security with Fiat-Shamir requires more robust soundness notions, such as round-by-round soundness, state-restoration soundness, or special soundness.

To sum it up, Plonk is known to be knowledge-sound, assuming both the ROM and the AGMOS [FFR24]. Lipmaa et al. [LPS24] proved that a variant of Plonk is knowledge-sound in the ROM under falsifiable assumptions. However, their variant of Plonk is too inefficient to be used in practice. Given the importance of Plonk in industry, this situation is highly non-satisfactory. A proof that works in the ROM and the AGM(OS) but not in the ROM under falsifiable assumptions leaves it open to unknown problems of the AGM(OS).

Finally, Plonk’s paper [GWC19] has been modified repeatedly to correct bugs and improve efficiency. A knowledge-soundness proof of Plonk under reasonable falsifiable assumptions in the ROM can be seen as an independent security audit, showing that the current variant of Plonk is secure and does not have to be changed anymore (except to improve on efficiency). It also ascertains Plonk will stay secure even if more problems are found in the AGM(OS) or GGM, as long as one does not break the concrete falsifiable assumptions.

Our Contributions. We prove that interactive Plonk has computational special-soundness under falsifiable assumptions ARSDH [LPS24] and SplitRSDH. While SplitRSDH is novel, it is not too different from ARSDH. After applying the Fiat-Shamir transformation, Plonk is tightly [AFK22,DG23,AFKR23] knowledge-sound in the ROM under the same assumptions. Towards this end, we first study KZG batching protocols that are not covered by the LPS compiler. We prove the well-known KZG batching protocol (**Batch**) has computational special-soundness under the ARSDH assumption. We prove that the linearization trick protocol (**Lin**) never has computational special-soundness in the standard model. ([FFR24] showed that **Lin** is secure in the AGMOS in many situations, relevant in practice.) Since Plonk uses the linearization trick, this seems to contradict our main positive result about Plonk. To overcome this, we introduce two new techniques, sanitization and Rhino (reduction to a hard assumption if not polynomial). By using sanitization, we show that slightly less efficient “sanitized” versions **SanLin** and **SanPlonk** of the linearization trick and Plonk are special-sound under the ARSDH assumption alone. By using Rhino, we show that the “real-world” Plonk is special-sound under ARSDH and SplitRSDH, and its well-known variant **SmallPlonk** (also described in [GWC19]) is special-sound under ARSDH alone.

As a contribution of independent interest, we propose a rigorous framework to define AGMOS so that the reductions in the AGMOS can be lifted to reductions in the GGM. This framework relies heavily on [JM24]. We show that SplitRSDH is secure in the resulting framework of AGMOS.

1.1 Our Techniques And Results

We use the bracket notation for additive groups, with say $[f]_1 = f \cdot [1]_1 \in \mathbb{G}_1$, see Section 2. KZG’s [KZG10] commitment to a polynomial $f(X)$ is $[f]_1 = [f(x)]_1$, and its opening proof of $[f]_1$ at point \mathfrak{z} is $[t]_1 = [t(x)]_1$, where x is a trapdoor. Crucially, $t(X) = (f(X) - f(\mathfrak{z})) / (X - \mathfrak{z})$ is a polynomial iff the prover is honest.

Differently from previous work on Plonk and zk-SNARKs in general ([LPS24] proved computational special-soundness only for KZG, but not for their interactive argument), we will prove computational κ -special-soundness for the arguments we consider. (See Section 2.2 for the definition of special-soundness and related notions like transcript trees.)

Case 1: Batch. To demonstrate related proof techniques, we start by proving that the standard KZG batching protocol **Batch** (see Section 4.1) is computational $(n + 1, m)$ -special-sound under the assumption that KZG is computational $(n + 1)$ -special-sound. Here, m is the batching factor. Like all subsequent proofs, **Batch**’s proof consists of an information-theoretical tree extraction lemma (Lemma 1) and the main theorem (Theorem 2). In the lemma, we construct a tree extractor that, given as an input an accepting transcript tree \mathcal{T} , outputs a tuple $\mathbf{k.tr}$ of m accepting KZG transcripts $([f_t]_1, \mathfrak{z}_i, \dots)$ for $t \in [1, m]$ and $i \in [1, n + 1]$, where $[f_t]_1$ are polynomial commitments to be batch-opened and \mathfrak{z}_i are mutually different opening points. The lemma uses standard linear algebraic techniques and relies on the fact that KZG is homomorphic. In Theorem 2, we call Lemma 1 to obtain the accepting tuple $\mathbf{k.tr}$. Given a promised $(n + 1)$ -special-soundness extractor of KZG and an $(n + 1, m)$ -special-soundness adversary for **Batch**, we construct an $(n + 1, m)$ -special-soundness extractor $\text{Ext}_{\text{ss}}^{\text{batch}}$ for **Batch** and an $(n + 1)$ -special-soundness adversary $\mathcal{B}_{\text{ss}}^{\text{kzg}}$ for KZG, such that either $\text{Ext}_{\text{ss}}^{\text{batch}}$ succeeds in extracting the valid witness for **Batch** (thus **Batch** is special-sound) or $\mathcal{B}_{\text{ss}}^{\text{kzg}}$ succeeds in breaking the $(n + 1)$ -special-soundness of KZG.

The rest of the proofs use a similar structure and will only describe their additional complications compared to **Batch**.

Case 2: Lin. We next consider the linearization trick **Lin** [GWC19, CHM⁺20] as formalized in [FFR24]. Assume that $\mathbf{a}(X) = (a_j(X))_{j=1}^{n_a}$ and $\mathbf{d}(X) = (d_t(X))_{t=1}^{n_b}$ are committed polynomials and $(\mathbf{g}_t(\mathbf{X}))_{t=1}^{n_b}$ are publicly known (“core”) polynomials. The goal of the **Lin** prover is to prove that $\sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}(X)) \mathbf{d}_t(X) = 0$; for example, $\mathbf{a}_1(X) \mathbf{d}_1(X) + (-1) \cdot \mathbf{d}_2(X)$ in the case of R1CS. In **Lin**, the verifier sends a random point \mathfrak{z} . The prover batch-opens $n_a + 1$ polynomials at \mathfrak{z} : (1) each $a_j(X)$ to $\bar{a}_j := a_j(\mathfrak{z})$, and (2) $\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}) \mathbf{d}_t(X)$ to 0. **Lin**’s communication is much better than that of **Batch** (see Table 2), and thus it and its variants are widely used in contemporary zk-SNARKS like Plonk.

Lin is well-known to be knowledge-sound in the AGM, independently of the choice of the core polynomials $\mathbf{g}_t(\mathbf{X})$. Lipmaa et al. [LPS23] described a simple attack, showing that Lin is not always knowledge-sound in the AGMOS. Their attack was based on the fact that the AGMOS adversary can do oblivious sampling. This result raised the question of the security of all zk-SNARKs that use Lin and have security proof in the AGM (or GGM) without oblivious sampling. Faonio et al. [FFR24] solved this issue by showing that Lin is computationally special-sound iff the public core polynomials are linearly independent. They then showed that the corresponding polynomials in Plonk are linearly independent and thus proved that Plonk is knowledge-sound in the AGMOS.

Adapting an idea of [GKP22], we prove that if DL is hard, Lin does not have computational special-soundness. Thus, the result of [FFR24] does not hold in the standard model. While our proof is for a simple choice of linearly independent core polynomials $\mathbf{g}_t(\mathbf{X})$, it can be generalized to *any* core polynomials. Hence, [FFR24]’s knowledge-soundness proof of Plonk cannot be used in the standard model since it relies on knowledge-soundness of some instance of Lin .

We introduce a new proof technique *sanitization* to solve this issue. We note that the AGMOS attacks of [LPS23,FFR24] are possible since one can obviously sample the polynomials $d_t(X)$ from some joint distribution. To protect against this, we ask the prover of SanLin (sanitized Lin , see Section 4.3) to batch-open all $d_t(X)$ at a single location. SanLin is a minimal modification of Lin , with the prover’s computation having just a single new field element. We then prove SanLin is computational special-sound (see Theorem 4) assuming that KZG has computational special-soundness and evaluation-binding, that is, under ARSDH.

Case 3: Plonk . Next, we refer always to *interactive* Plonk and its variants. Plonk is knowledge-sound in the AGM [GWC19]. Using Lin ’s special-soundness in the AGMOS, [FFR24] recently proved Plonk is knowledge-sound in the AGMOS. However, since (under the DL assumption) Lin is not special-sound in the standard model, their proof for Plonk does also not work in the standard model. To solve this issue, we propose three different solutions.

Case 3a: SanPlonk . When we replace Lin ’s use in Plonk with SanLin , we obtain SanPlonk , a variant of Plonk (see Table 1) with a communication of one additional field element. Since SanLin is computationally special-sound, assuming KZG is special-sound and evaluation-binding, so is SanPlonk . (The combined full computational special-soundness proof of our three Plonk variants is ten pages; we refer to Section 6 for details.) Since we use SanLin , this proof relies on our proof technique of sanitization.

Case 3b: Plonk . Since it is unlikely practitioners want to change their implementations, we also prove the actual real-world Plonk (as in the current version of [GWC19]) has computational special-soundness. This proof relies on Rhino (reduction to a hard assumption if not polynomial), a different proof technique. We use the same notation as in Plonk ’s paper [GWC19] to simplify reading. To explain this proof, we have to go into more detail about how Plonk works (see Section 6.2) and also explain some philosophy of KZG versus AGM.

In the security proof, we reach the following point. The Plonk’s prover has sent a few polynomial commitments like $[a, b, c]_1$ and a few other group elements like $[t_{lo}, t_{mid}, t_{hi}]_1$. For the random opening point \mathfrak{z} chosen by the verifier, the latter group elements satisfy $[t]_1 = [t_{lo} + \mathfrak{z}^n t_{mid} + \mathfrak{z}^{2n} t_{hi}]_1$ for some n . Here, $[t]_1$ is a KZG’s opening proof. Now, since $[t_{lo}, t_{mid}, t_{hi}]_1$ are group elements, in the AGM proof, one can extract polynomials $t_{lo}(X)$, $t_{mid}(X)$, and $t_{hi}(X)$ from $[t_{lo}, t_{mid}, t_{hi}]_1$. One can then compute the polynomial $t(X) = t_{lo}(X) + \mathfrak{z}^n t_{mid}(X) + \mathfrak{z}^{2n} t_{hi}(X)$ corresponding to the group element $[t]_1$.

Unfortunately, there is insufficient information in the standard model proof to extract $t(X)$. In our security proof of Plonk (see Theorem 6), we thus do not extract $t(X)$ from $[t_{lo}, t_{mid}, t_{hi}]_1$. Instead, we compute a candidate *rational function* $t(X)$ from other polynomials extracted in the security proof thus far. If $t(X)$ is a polynomial, then it is a valid opening of $[t]_1$, and then we can use it to finish the special-soundness proof of Plonk. If $t(X)$ is *not* a polynomial, we break a novel but standard-looking assumption SplitRSDH (Split RSDH, see Definition 4). (Thus, the name *Rhino* for the proof technique.) This completes the computational special-soundness proof of Plonk under KZG’s computational special-soundness and evaluation-binding and SplitRSDH.

Case 3c: SmallPlonk. We also consider SmallPlonk, a well-known variant of Plonk where $[t]_1$ is not split into $[t_{lo}, t_{mid}, t_{hi}]_1$. Thus, SmallPlonk has a shorter argument than Plonk (see Table 1). Using Rhino again, we prove that SmallPlonk has computational special-soundness under KZG’s computational special-soundness and evaluation-binding and SplitRSDH. However, we have a simpler variant of SplitRSDH, and the proof itself is slightly simpler.

On the New Assumption. SplitRSDH is a novel but falsifiable and standard-looking assumption (see Definition 4). We prove that SplitRSDH is secure in the AGMOS [LPS23], probably the most realistic known variant of AGM. This proof can be lifted to the GGM, [JM24]. To show that the latter result holds, we describe AGMOS by using the oracle framework of AGM from [JM24] but adding extra oracles for oblivious sampling as was done in [LPS23]. (See Section 3.) Thus, the lifting result of [JM24] holds for the AGMOS. The new AGMOS framework is simpler to use and understand than the one in [LPS23]. This forms an independent contribution of the current paper.

We emphasize that Plonk does not rely on AGMOS or other ideal models for groups. The use of AGMOS to prove the security of SplitRSDH should be seen as an independent element of cryptanalysis for the new assumption.

Generalizations. Our impossibility proof for Lin shows that the special-soundness of zk-SNARKs cannot be solely based on the special-soundness of Lin. Our two techniques, sanitization, and Rhino, can also be used in other zk-SNARKs to overcome this issue. We leave it to the future work.

Fiat-Shamir. We proved that *interactive* Plonk, SanPlonk, and SmallPlonk have computational κ -special-soundness (for a slightly different κ) under falsifiable assumptions. Under the same assumptions, the Fiat-Shamir transformation can be applied to obtain a zk-SNARK that is knowledge-sound in the ROM. The tight-

ness of Fiat-Shamir when applied to computationally κ -special-sound arguments was analyzed in [DG23,AFKR23]. The tightness of Fiat-Shamir is significantly better for special-sound interactive arguments than for knowledge-sound interactive arguments. Notably, this results in tighter security in the ROM, compared to previous proofs, which had less tight security and relied on both ROM and AGM/AGMOS. See Appendix D.6 for a brief discussion.

Zero-Knowledge. Up to now, we have ignored the issue of zero knowledge. Since sanitization implies batch-opening certain polynomials at one extra point, one sometimes has to add another randomizer to one of these polynomials to obtain zero knowledge. Since the needed change is usually straightforward (instead, the major innovation of the current work is in the analysis of special-soundness), we ignore the issue everywhere except for SanPlonk. SanPlonk’s description includes a new randomizer. In Appendix E, we prove that SanPlonk has zero knowledge. Zero-knowledge of Plonk (and SmallPlonk) was proven in [Sef24].

2 Preliminaries

Let λ denote the security parameter. By $f(\lambda) \approx_\lambda 0$, we mean that f is a negligible function. PPT (resp. DPT) stands for probabilistic (resp. deterministic) polynomial time. We denote the concatenation of vectors \mathbf{u} and \mathbf{v} as $\mathbf{u} \parallel \mathbf{v}$. \mathbb{F} is a finite field of prime order p . $\mathbb{F}[X]$ is the polynomial ring in variable X over the field \mathbb{F} and $\mathbb{F}_{\leq n}[X] \subset \mathbb{F}[X]$ is the set of polynomials of at most degree n . We denote $[a, b] := \{a, a + 1, \dots, b\}$, where $a \leq b$ are integers. Our notation is inspired by Plonk [GWC19] (for example, we denote polynomials by using SansSerif), but we do not follow it universally. For any set \mathcal{S} , $Z_{\mathcal{S}}(X) = \prod_{s \in \mathcal{S}} (X - s)$ denotes the vanishing polynomial over a set \mathcal{S} .

Bilinear Groups. A bilinear group generator $\text{Pgen}(1^\lambda)$ returns $\mathbf{p} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \hat{e}, [1]_1, [1]_2)$, where $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T are additive cyclic (thus, abelian) groups of prime order p , $\hat{e} : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is a non-degenerate efficiently computable bilinear pairing, and $[1]_\iota$ is a fixed generator of \mathbb{G}_ι . While $[1]_\iota$ is part of \mathbf{p} , we often give it as an explicit input to different algorithms for clarity. The bilinear pairing is of Type-3, that is, there is no efficient isomorphism between \mathbb{G}_1 and \mathbb{G}_2 . We use the common bracket notation, that is, for $\iota \in \{1, 2, T\}$ and $a \in \mathbb{Z}_p$, we write $[a]_\iota$ to denote $a[1]_\iota$. We denote $\hat{e}([a]_1, [b]_2)$ by $[a]_1 \bullet [b]_2$ and assume $[1]_T = [1]_1 \bullet [1]_2$. Thus, $[a]_1 \bullet [b]_2 = [ab]_T$ for any $a, b \in \mathbb{F}$, where $\mathbb{F} = \mathbb{Z}_p$.

We recall the following falsifiable assumption ARSDH [LPS24].

Definition 1. *The $(n + 1)$ -ARSDH (Adaptive Rational Strong Diffie-Hellman) assumption holds for Pgen in \mathbb{G}_1 if for any PPT \mathcal{A} , $\text{Adv}_{\text{Pgen}, n, \mathbb{G}_1, \mathcal{A}}^{\text{arsdh}}(\lambda) :=$*

$$\Pr \left[\begin{array}{l} \mathcal{S} \subset \mathbb{F} \wedge |\mathcal{S}| = n + 1 \wedge \\ [g]_1 \neq [0]_1 \wedge [\varphi]_1 = [\frac{g}{Z_{\mathcal{S}}(x)}]_1 \end{array} \middle| \begin{array}{l} \mathbf{p} \leftarrow \text{Pgen}(1^\lambda); x \leftarrow_{\$} \mathbb{F}; \\ \mathbf{ck} \leftarrow ([(x^i)_{i=1}^n]_1, [1, x]_2); \\ (\mathcal{S}, [g, \varphi]_1) \leftarrow \mathcal{A}(\mathbf{ck}) \end{array} \right] \approx_\lambda 0, \quad ,$$

where $Z_{\mathcal{S}}(X) := \prod_{s \in \mathcal{S}} (X - s)$. (The condition $[\varphi]_1 = [\frac{g}{Z_{\mathcal{S}}(x)}]_1$ is equivalent to $[g]_1 \bullet [1]_2 = [\varphi]_1 \bullet [Z_{\mathcal{S}}(x)]_2$.)

2.1 Polynomial Commitment Schemes

In a (univariate) polynomial commitment scheme (PCS, [KZG10]), the prover commits to a polynomial $f \in \mathbb{F}_{\leq n}[X]$ and later opens it to $f(\mathfrak{z})$ for $\mathfrak{z} \in \mathbb{F}$ chosen by the verifier. A *non-interactive* polynomial commitment scheme [KZG10] consists of the following algorithms:

Setup $\text{Pgen}(1^\lambda) \mapsto \mathfrak{p}$: Given 1^λ , return system parameters \mathfrak{p} .

Commitment key generation $\text{KGen}(\mathfrak{p}, n) \mapsto (\text{ck}, \text{tk})$: Given a system parameter \mathfrak{p} and an upperbound n on the polynomial degree, return (ck, tk) , where ck is the commitment key and tk is the trapdoor. We assume ck implicitly contains \mathfrak{p} . In the current paper, we do not use the trapdoor.

Commitment $\text{Com}(\text{ck}, f) \mapsto C$: Given a commitment key ck and a polynomial $f \in \mathbb{F}_{\leq n}[X]$, return a commitment C to f .

Opening $\text{Open}(\text{ck}, C, \mathfrak{z}, f) \mapsto (\bar{f}, \pi)$: Given a commitment key ck , a commitment C , an evaluation point $\mathfrak{z} \in \mathbb{F}$, and a polynomial $f \in \mathbb{F}_{\leq n}[X]$, return (\bar{f}, π) , where $\bar{f} \leftarrow f(\mathfrak{z})$ and π is an opening proof.

Verification $\text{V}(\text{ck}, C, \mathfrak{z}, \bar{f}, \pi) \mapsto \{0, 1\}$: Given a commitment key ck , a commitment C , an evaluation point \mathfrak{z} , a purported evaluation $\bar{f} \stackrel{?}{=} f(\mathfrak{z})$, and an opening proof π , return 1 (accept) or 0 (reject).

The KZG commitment scheme is a well-known non-interactive PCS [KZG10]; another such scheme is PST (multilinear KZG) [PST13]. Many PCSs have either an interactive opening or verification phase [BBHR18, BBB⁺18, BFS20].

A non-interactive PCS PC is *complete*, if for any λ , $\mathfrak{p} \leftarrow \text{Pgen}(1^\lambda)$, $n \in \text{poly}(\lambda)$, $\mathfrak{z} \in \mathbb{F}$, $f \in \mathbb{F}_{\leq n}[X]$,

$$\Pr \left[\text{V}(\text{ck}, C, \mathfrak{z}, \bar{f}, \pi) = 1 \mid \begin{array}{l} (\text{ck}, \text{tk}) \leftarrow \text{KGen}(\mathfrak{p}, n); C \leftarrow \text{Com}(\text{ck}, f); \\ (\bar{f}, \pi) \leftarrow \text{Open}(\text{ck}, C, \mathfrak{z}, f) \end{array} \right] = 1.$$

A non-interactive PCS PC is *binding*, if for any PPT \mathcal{A} , $\text{Adv}_{\text{Pgen}, \text{PC}, n, \mathcal{A}}^{\text{bind}}(\lambda) :=$

$$\Pr \left[C = \text{Com}(\text{ck}, f) = \text{Com}(\text{ck}, g) \wedge \begin{array}{l} \mathfrak{p} \leftarrow \text{Pgen}(1^\lambda); (\text{ck}, \text{tk}) \leftarrow \text{KGen}(\mathfrak{p}, n); \\ (C, f, g) \leftarrow \mathcal{A}(\text{ck}) \end{array} \mid f \neq g \wedge \deg(f) \leq n, \deg(g) \leq n \right] \approx_\lambda 0.$$

A PCS is *evaluation-binding* [KZG10] if it is hard to open the same evaluation point to different evaluations: PC is *evaluation-binding* for Pgen , if for any $n \in \text{poly}(\lambda)$, and PPT adversary \mathcal{A} , $\text{Adv}_{\text{Pgen}, \text{PC}, n, \mathcal{A}}^{\text{evb}}(\lambda) :=$

$$\Pr \left[\begin{array}{l} \text{V}(\text{ck}, C, \mathfrak{z}, \bar{f}, \pi) = 1 \wedge \\ \text{V}(\text{ck}, C, \mathfrak{z}, \bar{f}', \pi') = 1 \wedge \bar{f} \neq \bar{f}' \end{array} \mid \begin{array}{l} \mathfrak{p} \leftarrow \text{Pgen}(1^\lambda); (\text{ck}, \text{tk}) \leftarrow \text{KGen}(\mathfrak{p}, n); \\ (C, \mathfrak{z}, \bar{f}, \pi, \bar{f}', \pi') \leftarrow \mathcal{A}(\text{ck}) \end{array} \right] \approx_\lambda 0.$$

Evaluation-binding implies binding. Really, suppose $\mathcal{A}_{\text{bind}}$ succeeded in breaking binding, outputting $([c]_1, f(X), f'(X))$ such that $c = f(x) = f'(x)$ and $f(X) \neq f'(X)$. Then, we can find a point \mathfrak{z} , such that $f(\mathfrak{z}) \neq f'(\mathfrak{z})$, open $[c]_1$ at $f(\alpha)$, and $f'(\alpha)$, and break evaluation-binding.

We rely on the following terminology from [LPS24]. We call $\text{tr} = (C, \mathfrak{z}, \bar{f}, \pi)$ a *transcript* of the PCS. We say that a commitment key ck and a transcript

\mathbf{tr} is *accepting* when $V(\mathbf{ck}, \mathbf{tr}) = 1$. For any $n \geq 1$ and any commitment key \mathbf{ck} outputted by $\text{KGen}(\mathbf{p}, n)$, we define the following relations.

$$\begin{aligned} \mathcal{R}_{\mathbf{ck}} &:= \{(C, f) : C = \text{PC.Com}(\mathbf{ck}, f) \wedge \deg(f) \leq n\} , \\ \mathcal{R}_{\mathbf{ck}, \mathbf{tr}} &:= \{(C, f) : (C, f) \in \mathcal{R}_{\mathbf{ck}} \wedge \forall j \in [0, n]. f(\mathfrak{z}_j) = \bar{f}_j\} , \end{aligned} \quad (1)$$

where $\mathbf{tr} = (\mathbf{tr}_0, \dots, \mathbf{tr}_n)$ contains $n + 1$ accepting transcripts $\mathbf{tr}_j = (C, \mathfrak{z}_j, \bar{f}_j, \pi_j)$ such that C is the same in all transcripts, but \mathfrak{z}_j -s are pairwise distinct.

Let $n \in \text{poly}(\lambda)$ with $n \geq 1$. A non-interactive polynomial commitment scheme PC is *computationally* $(n+1)$ -*special-sound* for Pgen , if there exists a DPT extractor Ext_{ss} , such that for any PPT adversary \mathcal{A}_{ss} , $\text{Adv}_{\text{Pgen}, \text{PC}, \text{Ext}_{\text{ss}}, n, \mathcal{A}_{\text{ss}}}^{\text{ss}}(\lambda) :=$

$$\Pr \left[\begin{array}{l} \mathbf{tr} = (\mathbf{tr}_j)_{j=0}^n \wedge \\ \forall j \in [0, n]. \left(\begin{array}{l} \mathbf{tr}_j = (C, \mathfrak{z}_j, \bar{f}_j, \pi_j) \\ \wedge V(\mathbf{ck}, \mathbf{tr}_j) = 1 \end{array} \right) \\ \wedge (\forall i \neq j. \mathfrak{z}_i \neq \mathfrak{z}_j) \wedge (C, f) \notin \mathcal{R}_{\mathbf{ck}, \mathbf{tr}} \end{array} \middle| \begin{array}{l} \mathbf{p} \leftarrow \text{Pgen}(1^\lambda); \\ (\mathbf{ck}, \mathbf{tk}) \leftarrow \text{KGen}(\mathbf{p}, n); \\ \mathbf{tr} \leftarrow \mathcal{A}_{\text{ss}}(\mathbf{ck}); \\ f \leftarrow \text{Ext}_{\text{ss}}(\mathbf{ck}, \mathbf{tr}) \end{array} \right] \approx_\lambda 0 .$$

The KZG [KZG10] polynomial commitment scheme is defined as follows:

$\text{KZG.Pgen}(\lambda)$: return $\mathbf{p} \leftarrow \text{Pgen}(1^\lambda)$.

$\text{KZG.KGen}(\mathbf{p}, n)$: $\mathbf{tk} = x \leftarrow_{\$} \mathbb{Z}_p^*$; $\mathbf{ck} \leftarrow (\mathbf{p}, [(x^i)_{i=0}^n]_1, [1, x]_2)$; return $(\mathbf{ck}, \mathbf{tk})$.

$\text{KZG.Com}(\mathbf{ck}, f)$: return $C \leftarrow [f(x)]_1 = \sum_{j=0}^n f_j [x^j]_1$.

$\text{KZG.Open}(\mathbf{ck}, C, \mathfrak{z}, f)$: $\bar{f} \leftarrow f(\mathfrak{z})$; $\varphi(X) \leftarrow (f(X) - \bar{f}) / (X - \mathfrak{z})$; $\pi \leftarrow [\varphi(x)]_1$; return (\bar{f}, π) .

$\text{KZG.V}(\mathbf{ck}, C, \mathfrak{z}, \bar{f}, \pi)$: Return 1 iff $(C - \bar{f}[1]_1) \bullet [1]_2 = \pi \bullet [x - \mathfrak{z}]_2$.

KZG is evaluation-binding under the n -SDH assumption [KZG10] and non-black-box extractable in the AGM [FKL18] under the PDL assumption [Lip12, CHM⁺20] and in AGMOS [LPS23] under the PDL and TOFR assumptions. All of these are falsifiable assumptions. We refer to the respective papers for the definition of the assumptions. Lipmaa et al. [LPS24] proved the following result.

Theorem 1. *If the $(n+1)$ -ARSDH assumption holds, then KZG for degree $\leq n$ polynomials is computationally $(n+1)$ -special-sound: There exists a DPT KZG special-soundness extractor $\text{Ext}_{\text{ss}}^{\text{kzgg}}$, such that for any PPT \mathcal{A}_{ss} , there exists a PPT \mathcal{B} , such that $\text{Adv}_{\text{Pgen}, \text{PC}, \text{Ext}_{\text{ss}}^{\text{kzgg}}, n, \mathcal{A}_{\text{ss}}}^{\text{ss}}(\lambda) \leq \text{Adv}_{\text{Pgen}, n, \mathbb{G}_1, \mathcal{B}}^{\text{arsdh}}(\lambda)$.*

We say that a non-interactive polynomial commitment scheme is *triply homomorphic*, if: if $(C_j, \mathfrak{z}, \bar{f}_j, \pi_j)$ is an accepting transcript for every j , then so is $(\sum s_j C_j, \mathfrak{z}, \sum s_j \bar{f}_j, \sum s_j \pi_j)$ for any s_j . Clearly, KZG is triply homomorphic.

2.2 Interactive Arguments

Let $\mathcal{R} \subseteq \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^*$ be a ternary relation. \mathcal{R} contains triples $(\mathbf{srs}, \mathbf{x}, \mathbf{w}) \in \mathcal{R}$ where \mathbf{srs} is a public common reference string, \mathbf{x} is a public statement, and \mathbf{w} is a private witness. We only consider NP-*relations* relations \mathcal{R} , where the validity of a witness \mathbf{w} can be verified in time polynomial in $|\mathbf{x}| + |\mathbf{srs}|$.

Let $\text{Pgen}(1^\lambda)$ generate system parameters \mathbf{p} that are available to all algorithms. We do not always explicitly write \mathbf{p} as an input.

An *interactive argument* $\Pi = (\text{KGen}, \text{P}, \text{V})$ for a relation \mathcal{R} is an interactive protocol between a probabilistic polynomial time prover P and verifier V . The key generator KGen generates a common reference string \mathbf{srs} at the beginning of the protocol. Both P and V take as public input \mathbf{srs} and a statement \mathbf{x} and, additionally, P takes as private input a witness \mathbf{w} . The verifier V either accepts or rejects the transcript (all messages exchanged in the protocol execution). Accordingly, we say the transcript is accepting or rejecting.

Let $\boldsymbol{\kappa} = (\kappa_1, \dots, \kappa_\mu) \in \mathbb{N}^\mu$. A $\boldsymbol{\kappa}$ -tree of transcripts for a $(2\mu+1)$ -move public-coin interactive argument $\Pi = (\text{KGen}, \text{P}, \text{V})$ is a set of $K = \prod_{i=1}^\mu \kappa_i$ transcripts arranged in the following tree structure. The nodes in this tree correspond to the prover's messages and the edges to the verifier's challenges. Every node at depth i has precisely κ_i children corresponding to κ_i pairwise distinct challenges. Every transcript corresponds to exactly one path from the root node to a leaf.

Definition 2. Let $\boldsymbol{\kappa} = (\kappa_1, \dots, \kappa_\mu)$. A $(2\mu+1)$ -move public-coin interactive argument $\Pi = (\text{KGen}, \text{P}, \text{V})$ for relation \mathcal{R} is computational $\boldsymbol{\kappa}$ -special-sound if there exists a DPT extractor Ext_{ss} such that for any PPT \mathcal{A}_{ss} , $\text{Adv}_{\text{Pgen}, \Pi, \text{Ext}_{\text{ss}}, \boldsymbol{\kappa}, \mathcal{A}}^{\text{ss}}(\lambda) :=$

$$\Pr \left[\begin{array}{l} \mathcal{T} \text{ is a } \boldsymbol{\kappa}\text{-tree of} \\ \text{accepting transcripts} \\ \wedge (\mathbf{srs}, \mathbf{x}, \mathbf{w}) \notin \mathcal{R} \end{array} \middle| \begin{array}{l} \mathbf{p} \leftarrow \text{Pgen}(1^\lambda); (\mathbf{srs}, \text{tk}) \leftarrow \text{KGen}(\mathbf{p}); \\ (\mathbf{x}, \mathcal{T}) \leftarrow \mathcal{A}_{\text{ss}}(\mathbf{srs}); \mathbf{w} \leftarrow \text{Ext}_{\text{ss}}(\mathbf{srs}, \mathbf{x}, \mathcal{T}) \end{array} \right] \approx_\lambda 0 .$$

3 An Oracle Framework for AGMOS

Later, we propose a new assumption (SplitRSDH, Definition 4) and prove it is secure in the AGMOS (AGM with oblivious sampling, [LPS23]). We also compare our work to the work of Faonio et al. [FFR24]. We emphasize that we only use AGMOS as a sanity check that the new standard-looking assumption is likely to be secure. This is a standard and least contraversial use of the ideal group models. However, to have this part of paper in a more solid foundation, we recast AGMOS in the type-safe oracle framework of Jaeger and Mohan (Crypto 2024, [JM24]). Importantly, [JM24] proved an AGM to GGM lifting result that also carries over to our framework for AGMOS.

The resulting variant of AGMOS is more natural than AGMOS's original description in [LPS23]. Namely [LPS23] added sampling oracles on top of the AGM framework of [FKL18] which restricted the adversaries (by requiring them to output explanations) but did not have any other oracles. On the other hand, the oracle AGM framework of [JM24] restricts only the group access by providing oracles to group operations and other interesting functionalities. To this framework, our AGMOS framework just adds two more sampling oracles. We hope this makes AGMOS simpler to understand and use.

Algebraic group model (AGM, [FKL18]) is a very popular idealized model to prove the security of various assumptions and protocols. Compared to GGM,

AGM offers two major advantages. First, it is more realistic, allowing the adversaries to see the bit-representations of group elements. (Shoup’s GGM [Sho97] assumes the bit-representations are randomized and Maurer’s GGM [Mau05] does not reveal bit-representations.) Second, AGM proofs are shorter than GGM proofs: GGM proofs usually start and end with somewhat boilerplate steps, common to all GGM proofs. AGM allows to abstract these steps away and thus directly concentrate on the meat of the proof.

Fuchsbauer et al. [FKL18] left certain aspects of the AGM unformalized. This gave rise to a number of criticisms [Zha22,ZZK22] and fixes [Zha22,JM24]. In Crypto 2024, Jaeger and Mohan proposed [JM24] a rigorous interpretation of the AGM that answers to the criticisms. In particular, they prove that under mild assumptions, an AGM proof of security (in their interpretation) implies a GGM proof of security; see [JM24, Theorem 2] for the concrete statement. We will follow [JM24] when describing AGM and constructing AGM proofs. More precisely, we will use their oracle framework to AGM that is similar to Maurer’s type-safe oracle GGM framework [Mau05]. (In addition, [JM24] considered a pseudocode and a circuit framework.)

Fix a pairing description $\mathfrak{p} \leftarrow \mathbf{Pgen}(1^\lambda)$ and let σ be a fixed injective encoding of elements of any \mathbb{G}_ι as a bit-representation. In the oracle framework of bilinear AGM, there are two tables, **Val** (values) and **Expl** (explanations). The adversary is given handles to **Val** and access to a fixed family $\Pi_{\mathcal{O}}$ of oracles. We will denote the adversary’s initial input (e.g., input from the challenger) in \mathbb{G}_ι by $[\mathfrak{x}_\iota]_\iota$. We assume $[\mathfrak{x}_\iota]_\iota$ always includes $[1]_\iota$. Let $\mathfrak{x} = ([\mathfrak{x}_1]_1, [\mathfrak{x}_2]_2)$. **Val** is initialized with the elements of \mathfrak{x} and the corresponding entries of **Expl** are initialized with unit vectors \mathbf{e}_i (corresponding to the explanations of type $\mathfrak{x}_\iota[i] = \mathbf{e}_i^\top \mathfrak{x}_\iota$). Most oracles only operate on group elements based on handles. All oracles can output \perp (for example, when one is asked to add together elements from different groups).

In the bilinear setting, the two most important oracles are \mathcal{O}_p (takes two elements of the same group, referenced by handles i and j and stores their sum to a location k) and $\mathcal{P}air$ (takes elements of groups \mathbb{G}_1 and \mathbb{G}_2 , referenced by handles i and j , and stores their sum to a location k). The oracles also compute explanations of the output. The explanation of the output of \mathcal{O}_p is the sum of explanations of its inputs. The output explanation of $\mathcal{P}air$ is a tensor product of the input explanations. Two oracles that output non-group elements are $\mathcal{E}q$ (takes two elements of the same group, referenced by handles i and j and returns 1 if they are equal) and $\mathcal{E}ncode$ (takes a group element, referenced by a handle i and returns its bit representation). For the sake of reference, we depict all mentioned oracles in Fig. 1.

Additional oracles are possible as long as they work on handles and are efficient to implement. For example, [JM24] considers a scalar multiplication oracle; however, it can be implemented by using \mathcal{O}_p . Similarly, $\mathcal{E}q$ can be implemented by using $\mathcal{E}ncode$. If not being careful, adding or removing oracles can change the model drastically. The presence of $\mathcal{E}ncode$ is the only difference between the oracle framework of AGM and Maurer’s GGM; in the latter, only $\mathcal{E}q$ can be used to obtain any information about the bit-representation of group elements. Accord-

$\mathcal{Op}(i, j, k)$	$\mathcal{Paix}(i, j, k)$
if $\neg \exists \iota \in \{1, 2, T\}. (\mathbf{Val}_i \in \mathbb{G}_\iota \wedge \mathbf{Val}_j \in \mathbb{G}_\iota)$ then $\mathbf{Val}_k \leftarrow \perp$; else $\mathbf{Val}[k] \leftarrow \mathbf{Val}[i] + \mathbf{Val}[j]$; $\mathbf{Expl}[k] \leftarrow \mathbf{Expl}[i] + \mathbf{Expl}[j]$;	if $\neg (\mathbf{Val}[i] \in \mathbb{G}_1 \wedge \mathbf{Val}[j] \in \mathbb{G}_2)$ then $\mathbf{Val}[k] \leftarrow \perp$; else $\mathbf{Val}[k] \leftarrow \mathbf{Val}[i] \bullet \mathbf{Val}[j]$; $\mathbf{Expl}[k] \leftarrow \mathbf{Expl}[i] \otimes \mathbf{Expl}[j]$;
$\mathcal{Eq}(i, j)$	$\mathcal{Encode}(i)$
if $\neg \exists \iota \in \{1, 2, T\}. (\mathbf{Val}[i] \in \mathbb{G}_\iota \wedge \mathbf{Val}[j] \in \mathbb{G}_\iota)$ then return \perp ; else return $\mathbf{Val}[i] =? \mathbf{Val}[j]$	if $\neg \exists \iota \in \{i, j\}. (\mathbf{Val}[i] \in \mathbb{G}_\iota)$ then return \perp ; else return $\sigma(\mathbf{Val}[i])$
$\mathcal{Samp}_\iota(i, E, D)$	
if $E \notin \mathcal{EF}_{p,\iota} \vee D \notin \mathcal{DF}_p$ then return \perp ; fi $s \leftarrow s D$; $[\mathbf{q}]_\iota \leftarrow E(s)$; $\text{il}_\iota \leftarrow \text{il}_\iota + 1$; $i \leftarrow \mathbf{x}_\iota + \text{il}_\iota$; $\mathbf{Val}[i] \leftarrow [\mathbf{q}]_\iota$; $\mathbf{Expl}[i] \leftarrow (\mathbf{0}_{i-1}, 1)$; return s ;	

Fig. 1. The description of all considered oracles.

ing to [JM24], as the only difference with the AGM, in the *standard model*, the adversary has also access to a *Decode* oracle that takes as an input a bit-string (a bit-representation of a group element) and outputs the group element.

Lipmaa et al. [LPS23] recently defined AGMOS (AGM with oblivious sampling). AGMOS is more realistic than AGM since AGMOS adversaries are given an additional power of sampling group elements without knowing their discrete logarithms. Importantly, as shown in [LPS23], certain uses of KZG [KZG10] are secure in AGM but not in AGMOS.

Let $\mathcal{EF}_{p,\iota}$ be a set of (polynomially many) functions $\mathbb{F} \rightarrow \mathbb{G}_\iota$. This can include elliptic-curve hashing [Ica09]. Let \mathcal{DF}_p be a family of distributions over \mathbb{F} . Lipmaa et al. [LPS23] introduced two oracles \mathcal{Samp}_1 and \mathcal{Samp}_2 , one for each \mathbb{G}_1 and \mathbb{G}_2 . The i th query (E, D) to \mathcal{Samp}_ι consists of a function $E \in \mathcal{EF}_{p,\iota}$ and a distribution $D \in \mathcal{DF}_p$. The \mathcal{Samp}_ι oracle samples a random field element $s_i \leftarrow s D$, computes $[\mathbf{q}_{\iota_i}]_\iota \leftarrow E(s_i)$, stores $[\mathbf{q}_{\iota_i}]_\iota$ in the table (with a self-referential explanation $(0, \dots, 0, 1)$), and returns a field element s_i . One also keeps track on the number of \mathcal{Samp}_ι queries il_ι at any moment.⁴

Let $\Pi_\Theta = \{\mathcal{Op}, \mathcal{Paix}, \mathcal{Eq}, \mathcal{Encode}, \mathcal{Samp}_1, \mathcal{Samp}_2\}$ be the collection of available oracles. We require that for any PPT oracle adversary \mathcal{A}^{Π_Θ} , there exists a (non-uniform) PPT extractor $\text{Ext}_{\mathcal{A}^{\Pi_\Theta}}$, such that: if $\mathcal{A}^{\Pi_\Theta}(\mathbf{x})$ outputs a vector of group elements $[\mathbf{y}]_\iota$, on input $\mathbf{x} = ([\mathbf{x}_1]_1, [\mathbf{x}_2]_2)$, then with an overwhelming probability, $\text{Ext}_{\mathcal{A}^{\Pi_\Theta}}$ outputs field-element matrices $\boldsymbol{\gamma}$, $\boldsymbol{\delta}$, and $[\mathbf{q}_\iota]_\iota$ (\mathcal{Samp}_ι 's

⁴ AGMOS suits better to the oracle framework of the AGM than to the original FKL [FKL18] framework of the AGM. In particular, [LPS23] defined sampling by using oracles \mathcal{Samp}_ι . Since in the FKL framework, one does not define any other oracles (instead, one restricts the adversaries), \mathcal{Samp}_ι were the only oracles in [LPS23]. In the current paper, \mathcal{Samp}_ι are part of a larger family of oracles.

answer vector), such that

$$\mathbf{y} = \boldsymbol{\gamma}^\top \mathbf{x}_\iota + \boldsymbol{\delta}^\top \mathbf{q}_\iota . \quad (2)$$

Definition 3 (AGMOS). Let $\mathcal{EF} = \{\mathcal{EF}_{\mathbf{p},\iota}\}$ be a collection of functions. Let $\mathcal{DF} = \{\mathcal{DF}_{\mathbf{p}}\}$ be a family of distributions. A PPT algorithm \mathcal{A} is an $(\mathcal{EF}, \mathcal{DF})$ -AGMOS adversary for Pgen if there exists a PPT extractor $\text{Ext}_{\mathcal{A}}$, such that for any $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2)$, $\text{Adv}_{\text{Pgen}, \mathcal{EF}, \mathcal{DF}, \mathcal{A}, \text{Ext}_{\mathcal{A}}}^{\text{agmos}}(\lambda) :=$

$$\Pr \left[\begin{array}{l} y_1 \neq \gamma_1^\top \mathbf{x}_1 + \boldsymbol{\delta}_1^\top \mathbf{q}_1 \vee \\ y_2 \neq \gamma_2^\top \mathbf{x}_2 + \boldsymbol{\delta}_2^\top \mathbf{q}_2 \end{array} \middle| \begin{array}{l} \mathbf{p} \leftarrow \text{Pgen}(1^\lambda); r \leftarrow \text{RND}_\lambda(\mathcal{A}); \\ ([y_1]_1, [y_2]_2) \leftarrow \mathcal{A}^{H^\Theta}(\mathbf{p}, \mathbf{x}; r); \\ (\gamma_\iota, \boldsymbol{\delta}_\iota, [\mathbf{q}_\iota]_\iota)_{\iota=1}^2 \leftarrow \text{Ext}_{\mathcal{A}}^{H^\Theta}(\mathbf{p}, \mathbf{x}; r) \end{array} \right] = 0 .$$

Samp_ι is the non-programmable oracle depicted in Fig. 1. Here, $[\mathbf{q}_\iota]_\iota$ is required to be the tuple of elements output by Samp_ι . We denote by il_ι the number of Samp_ι calls.

In the Jaeger-Mohan oracle framework of AGM(OS), such extractor can be constructed efficiently since one is given access to all oracle queries and answers and by keeping track of the explanations of all already computed group elements. (Jaeger and Mohan [JM24] describe it in their Figure 8.) The extractor $\text{Ext}_{\mathcal{A}}$ just runs \mathcal{A} and then returns the outputs of \mathcal{A} together with the corresponding entries of the table **Expl**. Notably, $\text{Ext}_{\mathcal{A}}$ depends only on \mathcal{A} being algebraic (i.e., having to access the oracles to perform group operations) and no other aspects of \mathcal{A} 's code. Both in this case and the original AGM framework of [FKL18], $\text{Adv}_{\text{Pgen}, \mathcal{EF}, \mathcal{DF}, \mathcal{A}, \text{Ext}_{\mathcal{A}}}^{\text{agmos}}(\lambda) = 0$.

Typically, AGMOS proofs rely on the following two assumptions, also used in the current work. (Their use does not depend on the details of [JM24].)

Let $d_1(\lambda), d_2(\lambda) \in \text{poly}(\lambda)$. Pgen is $(d_1(\lambda), d_2(\lambda))$ -PDL (*Power Discrete Logarithm*, [Lip12]) secure if for any non-uniform PPT \mathcal{A} , $\text{Adv}_{d_1, d_2, \text{Pgen}, \mathcal{A}}^{\text{ddl}}(\lambda) :=$

$$\Pr \left[\mathcal{A}(\mathbf{p}, [(x^i)_{i=0}^{d_1}]_1, [(x^i)_{i=0}^{d_2}]_2) = x \mid \mathbf{p} \leftarrow \text{Pgen}(1^\lambda), x \leftarrow \mathfrak{s}\mathbb{F} \right] = \text{negl}(\lambda) .$$

Let \mathcal{EF} be some family of function and \mathcal{DF} a family of distributions. We say that Pgen is $(\mathcal{EF}, \mathcal{DF})$ -TOFR (*Tensor Oracle FindRep*, [LPS23]) secure if for any PPT \mathcal{A} , $\text{Adv}_{\text{Pgen}, \mathcal{A}}^{\text{tofr}}(\lambda) :=$

$$\Pr \left[\mathbf{v} \neq \mathbf{0} \wedge \mathbf{v}^\top \cdot \begin{pmatrix} 1 \\ \mathbf{q}_1 \\ \mathbf{q}_2 \\ \mathbf{q}_1 \otimes \mathbf{q}_2 \end{pmatrix} = 0 \mid \mathbf{p} \leftarrow \text{Pgen}(1^\lambda); \mathbf{v} \leftarrow \mathcal{A}^{\text{Samp}_1, \text{Samp}_2}(\mathbf{p}) \right] \approx_\lambda 0 .$$

Here, Samp_ι , \mathbf{q}_1 , and \mathbf{q}_2 are as in Definition 3. \mathcal{A} is a standard model adversary, it just does not see the discrete logarithms of elements created by Samp_ι .

Depending on the application, one might want to restrict the family of distributions to say uniform distributions. The smaller \mathcal{DF} is, the weaker TOFR will become. See [LPS23] for discussions.

A *reduction* \mathcal{B} from an assumption A to an assumption B is a map that maps an A -adversary \mathcal{A} to a B -adversary $\mathcal{B}^\mathcal{A}$. \mathcal{B} is *model-preserving* [JM24,

KGen(p): $x \leftarrow_{\$} \mathbb{F}^*$; return $\mathbf{srs} \leftarrow ((x^s)_{s=0}^n)_1, [1, x]_2$ and $\mathbf{td}_{\mathbf{srs}} \leftarrow x$;
<hr/> P(srs, w = (f_t(X))_{t=1}^m): for $t \in [1, m]$, $[f_t]_1 \leftarrow [f_t(x)]_1$; return $[(f_t)_{t=1}^m]_1$;
V: return $\mathfrak{z} \leftarrow_{\$} \mathbb{F}$; // Evaluation point
P: for $t \in [1, m]$, $\bar{f}_t \leftarrow f_t(\mathfrak{z})$; return $(\bar{f}_t)_{t=1}^m$;
V: return $v \leftarrow_{\$} \mathbb{F}$; // Batch coefficient
P: $h(X) \leftarrow (\sum_{t=1}^m v^{t-1} (f_t(X) - \bar{f}_t)) / (X - \mathfrak{z})$; return $[h]_1 \leftarrow [h(X)]_1$;
V: check $[\sum_{t=1}^m v^{t-1} (f_t - \bar{f}_t)]_1 \bullet [1]_2 = [h]_1 \bullet [x - \mathfrak{z}]_2$;

Fig. 2. The protocol Batch.

Definition 9], if $\mathcal{B}^{\mathcal{A}}$ is a type-safe (i.e. Maurer) generic adversary whenever \mathcal{A} is one. That is, \mathcal{B} does not use *Encode* (or *Decode*). \mathcal{B} is *efficient* if (1) it runs in time polynomial in the time of \mathcal{A} , and (2) $\mathcal{B}^{\mathcal{A}}$'s success in breaking B is at most a constant-time different from the success of \mathcal{A} in breaking A . Jaeger and Mohan [JM24, Theorem 9] proved the following result that we state informally.

Fact 1 *If there exists a model-preserving and efficient algebraic reduction from assumption A to assumption B and B is generically hard, then A is generically hard.*

4 Special Soundness of KZG Batching Protocols

It is a common practice to prove the knowledge-soundness of KZG-based interactive arguments in idealized group models (GGM, AGM, AGMOS). Stretching the techniques of [LPS24], we prove that some of these arguments are computationally special-sound under falsifiable assumptions. We describe an argument Batch for KZG-batching, and the linearization trick Lin. The former, we prove to be computationally special-sound as it is. However, Lin we show is not special-sound. We apply a novel technique of sanitization to obtain a computationally special-sound version of Lin.

4.1 Special-Soundness of Batch-KZG

Assume the usual setting of KZG with trapdoor x and degree bound n . Consider the well-known interactive protocol Batch from Fig. 2, where the prover has committed to m polynomials and then engages in a single batch proof that opens all m polynomials simultaneously at the same point \mathfrak{z} . Here, $[f_t]_1$ are commitments to some polynomials, \mathfrak{z} is the common evaluation point, f_t are purported evaluations of $[f_t]_1$ at \mathfrak{z} , and $[h]_1$ is the batched opening of all m commitments. Note that Batch's verifier essentially checks that $\mathbf{k.tr} = (\sum_{t=1}^m v^{t-1} [f_t]_1, \mathfrak{z}, \sum_{t=1}^m v^{t-1} f_t, [h]_1)$ is an accepting KZG transcript for a randomly chosen v .

In Lemma 1, we show how to extract from a transcript tree a tuple of admissible transcripts. In Theorem 2, we use the constructed extractor to establish

```

TEbatch(ck,  $\mathcal{T}$ )


---


Parse  $\mathcal{T} = (\text{tr}_{ij})_{i \in [1, n+1], j \in [1, m]}$ ; //  $\text{tr}_{ij} = (([f_t]_{t=1}^m]_1, \mathfrak{z}_i, (\bar{f}_{ti})_{t=1}^m, v_{ij}, [h_{ij}]_1)$ 
 $[h'_{i1}, \dots, h'_{im}]_1^\top \leftarrow \mathbf{V}_i^{-1} [h_{i1}, \dots, h_{im}]_1^\top$ ; (*) // See  $\mathbf{V}_i$  in Eq. (3)
for  $t \in [1, m]$  do
  for  $i \in [1, n+1]$  do  $\mathbf{k}.\text{tr}'_{it} \leftarrow ([f_t]_1, \mathfrak{z}_i, \bar{f}_{ti}, [h'_{ti}]_1)$ ; endfor (**)
   $\mathbf{k}.\text{tr}'_t \leftarrow (\mathbf{k}.\text{tr}'_{1t}, \dots, \mathbf{k}.\text{tr}'_{n+1,t})$ ; endfor
return  $(\mathbf{k}.\text{tr}'_{t=1}^m)$ ;

```

Fig. 3. The subroutine TE_{batch} .

Batch's special-soundness. Thus, Batch can be used without modification when one moves away from the proofs in idealized group models. This is important since Batch and its variants are ubiquitous in modern updatable zk-SNARKs like Plonk [GWC19], Marlin [CHM⁺20], and others [CFF⁺21, RZ21, LSZ22].

Lemma 1. *Let $\mathcal{T} = (\text{tr}_{ij})$ be an $(n+1, m)$ -tree of Batch's accepting transcripts, where tr_{ij} are as in Fig. 3. The DPT algorithm $\text{TE}_{\text{batch}}(\text{ck}, \mathcal{T})$ in Fig. 3 computes a tuple of accepting KZG transcripts $(\mathbf{k}.\text{tr}'_{t=1}^m)$, such that $\mathbf{k}.\text{tr}'_{it} = ([f_t]_1, \mathfrak{z}_i, \dots)$, with mutually different \mathfrak{z}_i for $i \in [1, n+1]$.*

Proof. Let \mathcal{T} be the given accepting tree of transcripts and

$$\mathbf{V}_i = \begin{pmatrix} 1 & v_{i1} & v_{i1}^2 & \cdots & v_{i1}^{m-1} \\ 1 & v_{i2} & v_{i2}^2 & \cdots & v_{i2}^{m-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & v_{im} & v_{im}^2 & \cdots & v_{im}^{m-1} \end{pmatrix} \quad (3)$$

be a Vandermonde matrix. Given \mathcal{T} 's structure, \mathfrak{z}_i -s are distinct and v_{ij} -s are distinct for each \mathfrak{z}_i , rendering \mathbf{V}_i non-singular for every $i \in [1, n+1]$.

Define $[h'_{i1}, \dots, h'_{im}]_1^\top$ as in (*) in Fig. 3. As noted above, by the definition of Batch (see Fig. 2), for any i and j , since tr_{ij} is accepted by the Batch verifier, $\mathbf{k}.\text{tr}_{ij} := ([\varphi_{ij}]_1, \mathfrak{z}_i, \Phi_{ij}, [h_{ij}]_1)$ is accepted by the KZG verifier, where $\varphi_{ij} := \sum_{t=1}^m v_{ij}^{t-1} f_t$, $\Phi_{ij} := \sum_{t=1}^m v_{ij}^{t-1} \bar{f}_{ti}$, and $h_{ij} = \sum_{t=1}^m v_{ij}^{t-1} h'_{it}$. Namely,

$$\begin{bmatrix} \varphi_{i1} - \Phi_{i1} \\ \vdots \\ \varphi_{im} - \Phi_{im} \end{bmatrix}_1 \bullet [1]_2 = \mathbf{V}_i \begin{bmatrix} f_1 - \bar{f}_{i1} \\ \vdots \\ f_m - \bar{f}_{im} \end{bmatrix}_1 \bullet [1]_2 = \begin{bmatrix} h_{i1} \\ \vdots \\ h_{im} \end{bmatrix}_1 \bullet [x - \mathfrak{z}_i]_2,$$

for any $i \in [1, n+1]$. But then

$$\begin{bmatrix} f_1 \\ \vdots \\ f_m \end{bmatrix}_1 = \mathbf{V}_i^{-1} \begin{bmatrix} \varphi_{i1} \\ \vdots \\ \varphi_{im} \end{bmatrix}_1, \quad \begin{pmatrix} \bar{f}_{i1} \\ \vdots \\ \bar{f}_{im} \end{pmatrix} = \mathbf{V}_i^{-1} \begin{pmatrix} \Phi_{i1} \\ \vdots \\ \Phi_{im} \end{pmatrix}, \quad \begin{bmatrix} h'_{i1} \\ \vdots \\ h'_{im} \end{bmatrix}_1 := \mathbf{V}_i^{-1} \begin{bmatrix} h_{i1} \\ \vdots \\ h_{im} \end{bmatrix}_1.$$

KZG's triple homomorphism ensures $\mathbf{k}.\text{tr}'_{it}$ (see (**) in Fig. 3) is an accepting KZG transcript. Thus, TE_{batch} returns accepting KZG transcripts $\mathbf{k}.\text{tr}'_t = (\mathbf{k}.\text{tr}'_{1t}, \dots, \mathbf{k}.\text{tr}'_{n+1,t})$ for $t \in [1, m]$, of the claimed form. \square

$\text{Ext}_{\text{SS}}^{\text{batch}}(\text{ck}, \mathcal{T})$	$\mathcal{B}_{\text{SS}}^{\text{kzg}}(\text{ck})$
$(\mathbf{k}, \text{tr}'_{t=1}^m) \leftarrow \text{TE}_{\text{batch}}(\text{ck}, \mathcal{T});$ for $t \in [1, m]$ do $f_t^*(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k}, \text{tr}'_t);$ endfor return $(f_t^*(X))_{t=1}^m;$	$\mathcal{T} \leftarrow \mathcal{A}_{\text{SS}}^{\text{batch}}(\text{ck});$ $(\mathbf{k}, \text{tr}'_{t=1}^m) \leftarrow \text{TE}_{\text{batch}}(\text{ck}, \mathcal{T});$ for $t \in [1, m]$ do $f_t^*(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k}, \text{tr}'_t);$ if $([f_t]_1, f_t^*(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k}, \text{tr}'_t}$ then return $\mathbf{k}, \text{tr}'_t;$ fi endfor return $\perp;$

Fig. 4. The κ -special-soundness extractor $\text{Ext}_{\text{SS}}^{\text{batch}}$ and the KZG $(n+1)$ -special-soundness adversary $\mathcal{B}_{\text{SS}}^{\text{kzg}}$ in Theorem 2.

Theorem 2. *Let $n, m \in \text{poly}(\lambda)$ and $\kappa = (n+1, m)$. If KZG is computational $(n+1)$ -special-sound, then Batch is computational κ -special-sound.*

Proof. Let $\text{Ext}_{\text{SS}}^{\text{kzg}}$ be the promised $(n+1)$ -special-soundness extractor of KZG and let $\mathcal{A}_{\text{SS}}^{\text{batch}}$ be any Batch κ -special-soundness adversary. In Fig. 4, we depict a κ -special-soundness extractor $\text{Ext}_{\text{SS}}^{\text{batch}}$ for Batch and an $(n+1)$ -special-soundness adversary $\mathcal{B}_{\text{SS}}^{\text{kzg}}$ for KZG. Here, $\text{Ext}_{\text{SS}}^{\text{batch}}$ has an oracle access to $\text{Ext}_{\text{SS}}^{\text{kzg}}$ and $\mathcal{B}_{\text{SS}}^{\text{kzg}}$ has an oracle access to $\mathcal{A}_{\text{SS}}^{\text{batch}}$ and $\text{Ext}_{\text{SS}}^{\text{kzg}}$.

$\text{Ext}_{\text{SS}}^{\text{batch}}$ inputs ck and a κ -tree $\mathcal{T} = (\text{tr}_{ij})_{i \in [1, n+1], j \in [1, m]}$ of accepting Batch transcripts, where tr_{ij} is defined as in Fig. 3. $\text{Ext}_{\text{SS}}^{\text{batch}}$ calls the (deterministic) algorithm TE_{batch} (see Fig. 3) to compute $n+1$ valid transcripts $\mathbf{k}, \text{tr}'_{it} = ([f_t]_1, \mathfrak{z}_i, \dots)$ for each polynomial that has to be extracted. Then, $\text{Ext}_{\text{SS}}^{\text{batch}}$ calls m times the $(n+1)$ -special-soundness extractor $\text{Ext}_{\text{SS}}^{\text{kzg}}$ to compute the witness.

Let us bound the advantage of Batch's adversary $\mathcal{A}_{\text{SS}}^{\text{batch}}$ against the extractor $\text{Ext}_{\text{SS}}^{\text{batch}}$. Assume that \mathcal{T} is an $(n+1, m)$ -tree of Batch transcripts, each accepted by the Batch verifier. Let bad be the event that for at least one t , $([f_t]_1, f_t^*(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k}, \text{tr}'_t}$. If bad does not occur, then $\text{Ext}_{\text{SS}}^{\text{batch}}$ has computed a valid witness for Batch. Since \mathfrak{z}_i are all distinct, $\mathcal{B}_{\text{SS}}^{\text{kzg}}$ wins the $(n+1)$ -special-soundness game if and only if bad happened. Thus, $\text{Adv}_{\text{Pgen}, \text{KZG}, \text{Ext}_{\text{SS}}^{\text{kzg}}, n+1, \mathcal{B}_{\text{SS}}^{\text{kzg}}}^{\text{SS}}(\lambda) = \Pr[\text{bad}] = \text{Adv}_{\text{Pgen}, \text{Batch}, \text{Ext}_{\text{SS}}^{\text{batch}}, n+1, m, \mathcal{A}_{\text{SS}}^{\text{batch}}}^{\text{SS}}(\lambda)$. This concludes the proof. \square

4.2 Lin: Linearization Trick

Definition of Lin. In many zk-SNARKs, a natural subtask⁵ is for the prover to prove that $\sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}(X)) \mathbf{d}_t(X) = 0$, where $\mathbf{g}(X) = (\mathbf{g}_1(X), \dots, \mathbf{g}_{n_b}(X))$ are publicly known polynomials, $\mathbf{a}(X) = (\mathbf{a}_1(X), \dots, \mathbf{a}_{n_a}(X))$, and $\mathbf{a}_\mathfrak{s}(X)$ (for $\mathfrak{s} \in [1, n_a]$) and $\mathbf{d}_t(X)$ (for $t \in [1, n_b]$) are committed polynomials.

⁵ Here, we follow closely the formalism of [FFR24]. One can easily generalize this framework to allow for more general expressions.

Table 2. Comparison of different protocols for the relation $\mathcal{R}_{\text{ck},\mathbf{g}}^{\text{Lin}}$. The communication does not include the commitment costs. KS stands for knowledge-soundness, and SS stands for special-soundness. The number of bits are given for the BLS381-12 pairing.

Method	$ \pi $ (bits)	KS in the AGM	KS and SS in the standard model
Opening $\mathbf{a}_s, \mathbf{d}_t$ separately	$(n_a + n_b)(\mathbb{F} + \mathbb{G}_t)$ $(640(n_a + n_b))$	✓	✓
Batch-opening a_s, d_t	$(n_a + n_b) \mathbb{F} + \mathbb{G}_t $ $(256(n_a + n_b) + 384)$	✓	✓
Lin	$n_a \mathbb{F} + \mathbb{G}_t $ $(256n_a + 384)$	✓	✗
SanLin (our work)	$(n_a + 1) \mathbb{F} + \mathbb{G}_t $ $(256n_a + 640)$	✓	✓

Example 1. In the simple quadratic check (e.g., in R1CS), one has $n_b = 2$, $n_a = 1$, $\mathbf{g}_1(\mathbf{a}(X)) = \mathbf{a}_1(X)$, and $\mathbf{g}_2(\mathbf{a}(X)) = -1$. One checks that $\mathbf{a}_1(X)\mathbf{d}_1(X) - \mathbf{d}_2(X) = 0$.

We define the underlying relation as

$$\mathcal{R}_{\text{ck},\mathbf{g}}^{\text{Lin}} = \left(\begin{array}{l} ([a_s]_{s=1}^{n_a}, [d_t]_{t=1}^{n_b}]_1, \\ ((\mathbf{a}_s(X))_{s=1}^{n_a}, (\mathbf{d}_t(X))_{t=1}^{n_b}) \end{array} \middle| \begin{array}{l} \forall s. ([a_s]_1, \mathbf{a}_s(X)) \in \mathcal{R}_{\text{ck}} \wedge \\ \forall t. ([d_t]_1, \mathbf{d}_t(X)) \in \mathcal{R}_{\text{ck}} \wedge \\ \sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}(X))\mathbf{d}_t(X) = 0 \end{array} \right).$$

We define $\mathcal{R}_{\text{ck},\mathbf{g},\text{tr}}^{\text{Lin}}$ as usual, requiring the committed polynomials to be consistent with the provided transcripts. Following [FFR24], we say $\mathbf{a}_s(X)$ are *left* polynomials, $\mathbf{g}_t(\mathbf{a}(X))$ are *core* polynomials, and $\mathbf{d}_t(X)$ are *right* polynomials.

For this task, several natural solutions exist. First, one can open all $n_a + n_b$ polynomials separately at a random point \mathfrak{z} and test that $\sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}(\mathfrak{z}))\mathbf{d}_t(\mathfrak{z}) = 0$. Second, one can do the same, but batch the opening proofs. Third, one can use the linearization trick, known at least from [GWC19, CHM⁺20]. Using the linearization trick, one performs a single batch-opening at a random point \mathfrak{z} , opening (1) all $\mathbf{a}_s(X)$, $s \in [1, n_a]$, to $\bar{a}_s := \mathbf{a}_s(\mathfrak{z})$, and (2) the *linearization polynomial* $\Lambda(X) := \sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}})\mathbf{d}_t(X)$ to 0. (See Fig. 5 for a detailed description.)

Insecurity of Lin. Lin is knowledge-sound in the AGM, [CHM⁺20] since one can extract all polynomials $\mathbf{a}_s(X)$ and $\mathbf{d}_t(X)$ from their commitments. However, Lipmaa et al. [LPS23] showed that (a simple case of) Lin is not knowledge-sound in the AGMOS. Faonio et al. [FFR24] generalized the attack of [LPS23], which we recall in Appendix A.1.

Faonio et al. [FFR24] also proved that Lin is knowledge-sound in the AGMOS iff $\mathbf{g}_t(\mathbf{a}(X))$ are linearly independent. Furthermore, they pointed out that in the case of Plonk, the core polynomials are linearly independent, thus justifying the use of Lin in Plonk. However, this still leaves it open whether Lin is special-sound in the standard model, assuming linear independence of core polynomials. We give a proof of the following impossibility result in Appendix A.2.

Theorem 3. *Let $\kappa_1, \kappa_2 \in \text{poly}(\lambda)$. Assuming DL holds in \mathbb{G}_1 , Lin is not computational (κ_1, κ_2) -special-sound even when the core polynomials $\mathbf{g}_t(\mathbf{a}(X))$ are linearly independent.*

<p>KGen: $x \leftarrow_{\\$} \mathbb{F}^*$; return $\mathbf{srs} \leftarrow ([x^t]_{t=0}^n, [1, x]_2)$ and $\mathbf{td}_{\mathbf{srs}} \leftarrow x$;</p> <hr/> <p>P($\mathbf{srs}, \mathbb{w} = ((\mathbf{a}_s(X))_{s=1}^{n_a}, (\mathbf{d}_t(X))_{t=1}^{n_b})$): for $s \in [1, n_a]$: $[a_s]_1 \leftarrow [\mathbf{a}_s(x)]_1$; for $t \in [1, n_b]$: $[d_t]_1 \leftarrow [\mathbf{d}_t(x)]_1$; return $[(a_s)_{s=1}^{n_a}, (d_t)_{t=1}^{n_b}]_1$;</p> <p>V: return $\mathfrak{z} \leftarrow_{\\$} \mathbb{F}$;</p> <p>P: for $s \in [1, n_a]$: $\bar{a}_s \leftarrow \mathbf{a}_s(\mathfrak{z})$; return $(\bar{a}_s)_{s=1}^{n_a}$;</p> <p>V: return $\gamma \leftarrow_{\\$} \mathbb{F}$;</p> <p>P: return $\bar{d} \leftarrow \sum_{t=1}^{n_b} \gamma^{t-1} \mathbf{d}_t(\mathfrak{z})$;</p> <p>V: return $\beta \leftarrow_{\\$} \mathbb{F}$;</p> <p>P: $\mathbf{H}(X) \leftarrow \sum_{s=1}^{n_a} \beta^{s-1} (\mathbf{a}_s(X) - \bar{a}_s) + \beta^{n_a} \cdot \sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}) \mathbf{d}_t(X) + \beta^{n_a+1} \cdot (\sum_{t=1}^{n_b} \gamma^{t-1} \mathbf{d}_t(X) - \bar{d})$; $h(X) \leftarrow \mathbf{H}(X)/(X - \mathfrak{z})$; return $[h]_1 \leftarrow [h(x)]_1$;</p> <p>V: check $\left[\sum_{s=1}^{n_a} \beta^{s-1} (a_s - \bar{a}_s) + \beta^{n_a} \sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}) d_t + \beta^{n_a+1} (\sum_{t=1}^{n_b} \gamma^{t-1} d_t - \bar{d}) \right]_1 \bullet [1]_2 = [h]_1 \bullet [x - \mathfrak{z}]_2$;</p>
--

Fig. 5. Lin (without highlighted parts) and SanLin (with highlighted parts).

To give some intuition of the above result, consider the simple case when one uses Lin to show that $1 \cdot \mathbf{d}_1(X) + X \cdot \mathbf{d}_2(X) = 0$. In this example, 1 and X are linearly independent polynomials. We can construct the following DL adversary that gets a challenge $[y]_1$ as an input. The adversary can pick a KZG public key \mathbf{ck} and store the trapdoor x . In such case, the adversary can commit to polynomials $\mathbf{d}_1(X) = yX$ as $[d_1]_1 = x[y]_1$ and to $\mathbf{d}_2(X) = -y$ as $[d_2]_1 = -[y]_1$, which satisfy the requirement $1 \cdot \mathbf{d}_1(X) + X \cdot \mathbf{d}_2(X) = 0$. We show in the proof of Theorem 3 that one can also simulate the rest of (κ_1, κ_2) -tree of accepting transcripts (for any $\kappa_1, \kappa_2 \in \text{poly}(\lambda)$) of Lin. If an efficient computational special-soundness extractor exists, the reduction can run it with an input \mathbf{ck} and the transcript tree to recover the polynomial $\mathbf{d}_2(X) = -y$. Thus, the adversary can return y and break the DL assumption in \mathbb{G}_1 .

4.3 SanLin: Sanitized Lin

Next, we will modify Lin so that it will become secure for any choice of $\mathbf{a}_s(X)$. The attack of [LPS23,FFR24] is possible since Lin’s prover never opens any polynomials $\mathbf{d}_t(X)$, allowing it to “obliviously sample” $[d_t]_1$. However, $[a_s]_1$ (that are already opened) cannot be sampled obliviously. We counter this attack introducing a technique called *sanitization*. Sanitization involves batch-opening all polynomial commitments (e.g., that were not batch-opened otherwise). In SanLin (*sanitized* Lin), this means batch-opening all polynomials $\mathbf{d}_t(X)$ at a random point. Crucially, we do not need the actual evaluations $\mathbf{d}_t(\mathfrak{z})$. Thus, it suffices for the prover to send $\bar{d} \leftarrow \sum_{t=1}^{n_b} \gamma^{t-1} \mathbf{d}_t(\mathfrak{z})$, for a new batching coefficient γ , adding a single field element \bar{d} (and one round) to Lin’s communication. We depict SanLin in Fig. 5 and compare it to more simplistic protocols in Table 2.

Theorem 4. *Let $n, n_a, n_b \in \text{poly}(\lambda)$, $n_g := \max_t \deg g_t(\mathbf{a}_t(X)) \in \text{poly}(\lambda)$, $\kappa_1 := \max(n_g, n) + 1$, and $\kappa = (\kappa_1, n_b, n_a + 2)$. If KZG for degree $\leq n$ polynomials has computational $(n+1)$ -special-soundness and evaluation-binding, then SanLin has computational κ -special-soundness.*

We postpone SanLin’s computational special-soundness proof to Appendix B.

5 New Assumption: SplitRSDH

Our proofs for Plonk and SmallPlonk rely on the following novel falsifiable assumption SplitRSDH.

Definition 4 (par-SplitRSDH). *Let $\text{par} = (m, n_\psi, n_1, n_S, (\psi_k(X))_{k=1}^m)$. Assume $m, n_\psi, n_1, n_S \in \text{poly}(\lambda)$ are positive integers so that $n_S > n_1 + n_\psi + 1$. Assume $\psi_k(X) \in \mathbb{F}_{\leq n_\psi}[X]$. For any PPT \mathcal{A} , $\text{Adv}_{\text{Gen, par}, \mathcal{A}}^{\text{splitrsdh}}(\lambda) :=$*

$$\Pr \left[\begin{array}{l} \mathcal{S} \subset \mathbb{F} \wedge |\mathcal{S}| = n_S \wedge L(X) \in \mathbb{F}[X] \wedge \\ \deg L(X) \in [n_1 + n_\psi + 1, n_S - 1] \wedge \\ \left(\sum_{k=1}^m [\tau_k \psi_k(x)]_1 = [L(x)]_1 + [\varphi Z_S(x)]_1 \right) \end{array} \middle| \begin{array}{l} x \leftarrow_{\$} \mathbb{F}; \\ \text{ck} \leftarrow ([1, x, \dots, x^{n_1}]_1, [1, x]_2); \\ (\mathcal{S}, L(X), [(\tau_k)_{k=1}^m, \varphi]_1) \leftarrow \mathcal{A}(\text{p}, \text{ck}) \end{array} \right]$$

is negligible. Here, $Z_S(X) := \prod_i (X - \mathfrak{z}_i)$ is the vanishing polynomial. (The condition $\sum_{k=1}^m [\tau_k \psi_k(x)]_1 = [L(x)]_1 + [\varphi Z_S(x)]_1$ is equivalent to $\sum_{k=1}^m ([\tau_k]_1 \bullet [\psi_k(x)]_2) = [1]_1 \bullet [L(x)]_2 + [\varphi]_1 \bullet [Z_S(x)]_2$.)

SplitRSDH is falsifiable since the challenger, knowing x , can verify the adversary’s success. In Appendix C, we prove SplitRSDH’s security in the new type-safe framework of AGMOS. The intuition behind SplitRSDH and its AGMOS proof is simple. Let $\tau_k(X)$ be the explanations of $[\tau_k]_1$ provided by the AGMOS adversary. Let $T(X) := \sum_{k=1}^m \tau_k(X) \psi_k(X) \in \mathbb{F}_{\leq n_1 + n_\psi}[X]$. If the challenger accepts, then $T(X) - L(X)$ has a degree smaller than n_S and thus $\varphi(X) = 0$. But then $T(X) = L(X)$, which means that $\deg L(X) \leq n_1 + n_\psi$, a contradiction.

We define the parameters

$$\text{par}_n^{\text{plonk}} = (m, n_\psi, n_1, n_S, (\psi_k(X))_{k=1}^m),$$

with $m = 3$, $n_\psi = 2n$, $n_1 = \kappa_{\text{kzgz}} = n + 5$, $n_S = \kappa_3 = 4\kappa_{\text{kzgz}} + 1$, and $\psi_k(X) = X^{(k-1)n}$, and

$$\text{par}_n^{\text{smallplonk}} = (m, n_\psi, n_1, n_S, (\psi_k(X))_{k=1}^m),$$

with $m = 1$, $n_\psi = 0$, $n_1 = \kappa_{\text{kzgz}} = 3n + 5$, $n_S = \kappa_3 = 4\kappa_{\text{kzgz}} + 1$, and $\psi_1(X) = 1$. Thus, $n_S > n_1 + n_\psi + n$ (we need the latter in Theorem 8). Here, $\text{par}_n^{\text{plonk}}$ is only useful for Plonk and $\text{par}_n^{\text{smallplonk}}$ is only useful for SmallPlonk.

For $m \geq 1$, we can reduce SplitRSDH to the following nicer-looking but purportedly stronger assumption TIDRSDH (target intermediate degree RSDH), where we allow the adversary to output $[g]_T \leftarrow \sum_{k=1}^m ([\tau_k]_1 \bullet [\psi_k(x)]_2)$. Then, it suffices for the following probability to be negligible (under the conditions of Definition 4):

$$\Pr \left[\begin{array}{l} \mathcal{S} \subset \mathbb{F} \wedge |\mathcal{S}| = n_S \wedge L(X) \in \mathbb{F}[X] \wedge \\ \deg L(X) \in [n_1 + n_\psi + 1, n_S - 1] \wedge \\ [g - L(x)]_T = [\varphi Z_S(x)]_T \end{array} \middle| \begin{array}{l} x \leftarrow_{\$} \mathbb{F}; \\ \text{ck} \leftarrow ([1, x, \dots, x^{n_1}]_1, [1, x, \dots, x^{n_\psi}]_2); \\ (\mathcal{S}, L(X), [\varphi]_1, [g]_T) \leftarrow \mathcal{A}(\text{p}, \text{ck}); \end{array} \right].$$

The AGM(OS) proof of TIDRSDH follows from noticing that if the adversary succeeds, then $g(X) - L(X)$ has degree $< n_S$ and thus $\varphi(X) = 0$. But then $g(X) = L(X)$, which means $\deg L(X) \leq n_1 + n_\psi$ and the adversary failed. TIDRSDH is not weaker and (purportedly) stronger than SplitRSDH for the same par. First, it allows the adversary to output \mathbb{G}_T elements⁶. Second, a SplitRSDH adversary has less freedom to choose $[g]_T$. Third, ck has more elements so the AGMOS reduction is to a stronger PDL. We will explicitly use SplitRSDH within the current paper but one can instead use TIDRSDH.

SplitRSDH can also be reduced to two somewhat more standard assumptions. Assume a par-SplitRSDH adversary succeeds. Consider two cases:

- $\varphi \neq 0$. Then we construct the following reduction to a “target ARSDH” assumption. Given $([1, \dots, x^{n_S - 1 - n_\psi}]_1, [1, \dots, x^{n_\psi}]_2)$, it outputs $(\mathcal{S}, [g]_T, [\varphi]_1)$, where $[g]_T = \sum_{k=1}^m ([\tau_k]_1 \bullet [\psi_k(x)]_2) - [L(x)]_T$, and checks that $[g]_T = [\varphi]_1 \bullet [Z_S(x)]_2$ and $\varphi \neq 0$.
- $\varphi = 0$. Then we construct the following reduction to a “target BDE assumption”. Given $([1, \dots, x^{n_1}]_1, [1, \dots, x^{n_\psi}]_2)$, it outputs $(L(X), [g]_T)$, where $[g]_T = \sum_{k=1}^m ([\tau_k]_1 \bullet [\psi_k(x)]_2)$, and once checks that $\deg L(X) > n_1 + n_\psi$ and $[g]_T = [L(x)]_T$.

6 Special-Soundness of Plonk And Variants

In this section, we prove that interactive Plonk [GWC19] has special-soundness, assuming that KZG is evaluation-binding and specially sound and a new, falsifiable assumption SplitRSDH holds. Recall that KZG is evaluation-binding and specially sound under the ARSDH assumption. In addition, we prove that both SanPlonk (a sanitized variant of Plonk with one field element of extra communication) and SmallPlonk (a well-known variant of Plonk, where one does not split $[t]_1$ to three group elements) have special-soundness under ARSDH alone. By applying the Fiat-Shamir transform to any of the three constructions, one can obtain a zk-SNARKs secure in the ROM under the same assumptions.

6.1 Preliminaries For Plonk

We recall Plonk [GWC19], a popular zk-SNARK for proving satisfiability of arbitrary arithmetic circuits. We closely follow the notation of [GWC19].

Let \mathbb{H} be a multiplicative subgroup of \mathbb{F} containing the n th roots of unity. Let ω be a primitive n th root of unity and a generator of \mathbb{H} , $\mathbb{H} = \{1, \omega, \dots, \omega^{n-1}\}$. Let $Z_{\mathbb{H}}(X) := X^n - 1$ be the vanishing polynomial on \mathbb{H} . For $i \in [1, n]$, $L_i(X)$ denotes the i th Lagrange polynomial on \mathbb{H} . Namely, $L_i(X)$ is the unique polynomial of at most degree $n - 1$ such that $L_i(\omega^i) = 1$ and $L_i(\omega^j) = 0$ for all $j \in [1, n] \setminus \{i\}$. We assume that the number of constraints is upper bounded by n .

⁶ Since \mathbb{G}_T (a subgroup of the multiplicative group of a finite field) is not a generic group [JR10], we prefer not to handle adversaries who output \mathbb{G}_T elements. See [JR10] for a discussion.

The SNARK proof relation. Let $\mathcal{P} = \{\mathfrak{q}_M(X), \mathfrak{q}_L(X), \mathfrak{q}_R(X), \mathfrak{q}_O(X), \mathfrak{q}_C(X), \mathfrak{S}_{\sigma_1}(X), \mathfrak{S}_{\sigma_2}(X), \mathfrak{S}_{\sigma_3}(X)\}$ be a set of polynomials that defines a given circuit. See Appendix D.1 for the exact conditions they satisfy. They are exactly the same as in [GWC19]. We use the notation $\mathfrak{q}_{X_i} := \mathfrak{q}_X(\omega^i)$.

Given $\ell \leq n$ and \mathcal{P} , we wish to prove statements of knowledge for the relation $\mathcal{R}_{\mathcal{P}} \subset \mathbb{F}^\ell \times \mathbb{F}^{3n-\ell}$ containing all pairs $\mathfrak{x} = (w_i)_{i=1}^\ell$, $\mathfrak{w} = (w_i)_{i=\ell+1}^{3n}$ such that

1. For $i \in [1, \ell]$: $\mathfrak{q}_{M_i} = \mathfrak{q}_{R_i} = \mathfrak{q}_{O_i} = \mathfrak{q}_{C_i} = 0$ and $\mathfrak{q}_{L_i} = -1$, which guarantees

$$\mathfrak{q}_{M_i} w_i w_{n+i} + \mathfrak{q}_{L_i} w_i + \mathfrak{q}_{R_i} w_{n+i} + \mathfrak{q}_{O_i} w_{2n+i} + \mathfrak{q}_{C_i} = -w_i . \quad (4)$$

We see later that this is needed to force the prover to use the correct \mathfrak{x} .

2. For all $i \in [\ell + 1, n]$:

$$\mathfrak{q}_{M_i} w_i w_{n+i} + \mathfrak{q}_{L_i} w_i + \mathfrak{q}_{R_i} w_{n+i} + \mathfrak{q}_{O_i} w_{2n+i} + \mathfrak{q}_{C_i} = 0 , \quad (5)$$

3. For all $i \in [1, 3n]$:

$$w_i = w_{\sigma(i)} . \quad (6)$$

We refer to [GWC19] for the explanation how these constraints are related to arithmetic circuits.

6.2 Plonk And Variants

We present Plonk, SmallPlonk, and Plonk’s sanitized variant SanPlonk. While we describe their interactive versions, to save space, we will omit the adjective “interactive”. We describe them in parallel, highlighting changes in SanPlonk and SmallPlonk compared to Plonk. Compared to Plonk, SanPlonk batch opens $[t_{lo}]_1$, $[t_{mid}]_1$, and $[t_{hi}]_1$ (three group elements sent in Plonk), applying the sanitization technique from Section 4.2. This results in an interactive argument with two additional rounds and one more field element sent by the prover, compared to Plonk. On the other hand, SmallPlonk sends just a single group element $[t]_1$ instead of three elements $[t_{lo}, t_{mid}, t_{hi}]_1$. We prove the computational special-soundness of Plonk and SmallPlonk assuming that (1) KZG is evaluation-binding and special sound, and (2) a new falsifiable assumption SplitRSDH holds. We prove the computational special-soundness of SanPlonk solely under (1).

Common preprocessed input: n , $[x, \dots, x^{n+5}]_1$ plus additional elements $[x^{n+6}, \dots, x^{3n+5}]_1$ in SmallPlonk), $(\mathfrak{q}_{M_i}, \mathfrak{q}_{L_i}, \mathfrak{q}_{R_i}, \mathfrak{q}_{O_i}, \mathfrak{q}_{C_i})_{i=1}^n$, σ^* , $\mathfrak{q}_M(X) = \sum_{i=1}^n \mathfrak{q}_{M_i} L_i(X)$, $\mathfrak{q}_L(X) = \sum_{i=1}^n \mathfrak{q}_{L_i} L_i(X)$, $\mathfrak{q}_R(X) = \sum_{i=1}^n \mathfrak{q}_{R_i} L_i(X)$, $\mathfrak{q}_O(X) = \sum_{i=1}^n \mathfrak{q}_{O_i} L_i(X)$, $\mathfrak{q}_C(X) = \sum_{i=1}^n \mathfrak{q}_{C_i} L_i(X)$, $\mathfrak{S}_{\sigma_1}(X) = \sum_{i=1}^n \sigma^*(i) L_i(X)$, $\mathfrak{S}_{\sigma_2}(X) = \sum_{i=1}^n \sigma^*(n+i) L_i(X)$, $\mathfrak{S}_{\sigma_3}(X) = \sum_{i=1}^n \sigma^*(2n+i) L_i(X)$.

Verifier preprocessed input: $[\mathfrak{q}_M]_1 := \mathfrak{q}_M(x) \cdot [1]_1$, $[\mathfrak{q}_L]_1 := \mathfrak{q}_L(x) \cdot [1]_1$, $[\mathfrak{q}_R]_1 := \mathfrak{q}_R(x) \cdot [1]_1$, $[\mathfrak{q}_O]_1 := \mathfrak{q}_O(x) \cdot [1]_1$, $[\mathfrak{q}_C]_1 := \mathfrak{q}_C(x) \cdot [1]_1$, $[s_{\sigma_1}]_1 := \mathfrak{S}_{\sigma_1}(x) \cdot [1]_1$, $[s_{\sigma_2}]_1 := \mathfrak{S}_{\sigma_2}(x) \cdot [1]_1$, $[s_{\sigma_3}]_1 := \mathfrak{S}_{\sigma_3}(x) \cdot [1]_1$, $x \cdot [1]_2$,

Public input: $(\ell, (w_i)_{i=1}^\ell)$.

First round. On input $(w_i)_{i=1}^{3n}$, the prover does the following. Sample $(b_1, \dots, b_9) \leftarrow_{\mathfrak{s}} \mathbb{F}$. Compute wire polynomials $\mathfrak{a}(X) \leftarrow \sum_{i=1}^n w_i L_i(X) + (b_1 X + b_2) Z_{\mathbb{H}}(X) \in$

$\mathbb{F}_{\leq n+1}[X]$, $\mathbf{b}(X) \leftarrow \sum_{i=1}^n w_{n+i} L_i(X) + (b_3 X + b_4) \mathbf{Z}_{\mathbb{H}}(X) \in \mathbb{F}_{\leq n+1}[X]$, $\mathbf{c}(X) \leftarrow \sum_{i=1}^n w_{2n+i} L_i(X) + (b_5 X + b_6) \mathbf{Z}_{\mathbb{H}}(X) \in \mathbb{F}_{\leq n+1}[X]$. Send $[\mathbf{a}(x), \mathbf{b}(x), \mathbf{c}(x)]_1$ to \mathbf{V} . \mathbf{V} replies with $\beta, \gamma \leftarrow_{\$} \mathbb{F}$.

Second round. The prover computes polynomial $\mathbf{z}(X) \leftarrow L_1(X) + \sum_{i=1}^{n-1} \left(\prod_{j=1}^i \frac{(w_j + \beta \omega^j + \gamma)(w_{n+j} + \beta k_1 \omega^j + \gamma)(w_{2n+j} + \beta k_2 \omega^j + \gamma)}{(w_j + \sigma^*(j) \omega^j + \gamma)(w_{n+j} + \sigma^*(n+j) \omega^j + \gamma)(w_{2n+j} + \sigma^*(2n+j) \omega^j + \gamma)} \right) L_{i+1}(X) + (b_7 X^2 + b_8 X + b_9) \mathbf{Z}_{\mathbb{H}}(X) \in \mathbb{F}_{\leq n+2}[X]$ and sends $[\mathbf{z}(x)]_1$ to \mathbf{V} . \mathbf{V} replies with $\alpha \leftarrow_{\$} \mathbb{F}$.

Third round. The prover does the following. Compute

$$\begin{aligned} \mathbf{F}_0(X) &:= \mathbf{a}(X) \mathbf{b}(X) \mathbf{q}_{\mathbb{M}}(X) + \mathbf{a}(X) \mathbf{q}_{\mathbb{L}}(X) + \mathbf{b}(X) \mathbf{q}_{\mathbb{R}}(X) + \mathbf{c}(X) \mathbf{q}_{\mathbb{O}}(X) + \text{PI}(X) + \mathbf{q}_{\mathbb{C}}(X) \\ \mathbf{F}_1(X) &:= (\mathbf{a}(X) + \beta X + \gamma) (\mathbf{b}(X) + \beta k_1 X + \gamma) (\mathbf{c}(X) + \beta k_2 X + \gamma) \mathbf{z}(X) \\ &\quad - (\mathbf{a}(X) + \beta S_{\sigma_1}(X) + \gamma) (\mathbf{b}(X) + \beta S_{\sigma_2}(X) + \gamma) (\mathbf{c}(X) + \beta S_{\sigma_3}(X) + \gamma) \mathbf{z}(X \omega) \\ \mathbf{F}_2(X) &:= (\mathbf{z}(X) - 1) L_1(X) \\ \mathbf{F}(X) &:= \mathbf{F}_0(X) + \alpha \mathbf{F}_1(X) + \alpha^2 \mathbf{F}_2(X) , \\ \mathbf{t}(X) &:= \frac{\mathbf{F}(X)}{\mathbf{Z}_{\mathbb{H}}(X)} \in \mathbb{F}_{\leq 3n+5}[X] . \end{aligned} \tag{7}$$

In **SmallPlonk**, the prover just sends $[\mathbf{t}(x)]_1$ to the verifier. In **Plonk** and **SanPlonk**, we split $\mathbf{t}(X)$ into polynomials $\mathbf{t}_{\text{lo}}'(X), \mathbf{t}_{\text{mid}}'(X)$ (both of degree less than n) and $\mathbf{t}_{\text{hi}}'(X)$ (of degree at most $n + 5$), such that $\mathbf{t}(X) = \mathbf{t}_{\text{lo}}'(X) + X^n \mathbf{t}_{\text{mid}}'(X) + X^{2n} \mathbf{t}_{\text{hi}}'(X)$. Sample $b_{10}, b_{11}, b_{12} \leftarrow_{\$} \mathbb{F}$ and define $\mathbf{t}_{\text{lo}}(X) := \mathbf{t}_{\text{lo}}'(X) + b_{10} X^n + b_{12} X^{n+1}$, $\mathbf{t}_{\text{mid}}(X) := \mathbf{t}_{\text{mid}}'(X) - b_{10} - b_{12} X + b_{11} X^n$, and $\mathbf{t}_{\text{hi}}(X) := \mathbf{t}_{\text{hi}}'(X) - b_{11}$. (b_{12} is required for the zero-knowledge proof of **SanPlonk**.) Note that $\mathbf{t}(X) = \mathbf{t}_{\text{lo}}(X) + X^n \mathbf{t}_{\text{mid}}(X) + X^{2n} \mathbf{t}_{\text{hi}}(X)$. Send $[\mathbf{t}_{\text{lo}}(x), \mathbf{t}_{\text{mid}}(x), \mathbf{t}_{\text{hi}}(x)]_1$ to the verifier. The verifier replies with the evaluation randomness $\mathfrak{z} \leftarrow_{\$} \mathbb{F}$.

Fourth round. The prover does the following. Set $\bar{a} \leftarrow \mathbf{a}(\mathfrak{z})$, $\bar{b} \leftarrow \mathbf{b}(\mathfrak{z})$, $\bar{c} \leftarrow \mathbf{c}(\mathfrak{z})$, $\bar{s}_{\sigma_1} \leftarrow S_{\sigma_1}(\mathfrak{z})$, $\bar{s}_{\sigma_2} \leftarrow S_{\sigma_2}(\mathfrak{z})$, $\bar{z}_{\omega} \leftarrow \mathbf{z}(\omega \mathfrak{z})$. Send $(\bar{a}, \bar{b}, \bar{c}, \bar{s}_{\sigma_1}, \bar{s}_{\sigma_2}, \bar{z}_{\omega})$ to the verifier. The verifier replies with the sanitization randomness $\delta \leftarrow_{\$} \mathbb{F}$.

Fourth (Plonk/SmallPlonk) or fifth (SanPlonk) round. The prover sends $[\bar{t}_{\mathfrak{z}} \leftarrow \mathbf{t}_{\text{lo}}(\mathfrak{z}) + \delta \mathbf{t}_{\text{mid}}(\mathfrak{z}) + \delta^2 \mathbf{t}_{\text{hi}}(\mathfrak{z})]$. The verifier replies with $v \leftarrow_{\$} \mathbb{F}$.

Fifth (Plonk/SmallPlonk) or sixth (SanPlonk) round. Prover does the following. Compute the linearization polynomial $\mathbf{r}(X)$:

$$\begin{aligned} \mathbf{A}_0(X) &= \bar{a} \bar{b} \cdot \mathbf{q}_{\mathbb{M}}(X) + \bar{a} \cdot \mathbf{q}_{\mathbb{L}}(X) + \bar{b} \cdot \mathbf{q}_{\mathbb{R}}(X) + \bar{c} \cdot \mathbf{q}_{\mathbb{O}}(X) + \text{PI}(\mathfrak{z}) + \mathbf{q}_{\mathbb{C}}(X) , \\ \mathbf{A}_1(X) &= (\bar{a} + \beta \mathfrak{z} + \gamma) (\bar{b} + \beta k_1 \mathfrak{z} + \gamma) (\bar{c} + \beta k_2 \mathfrak{z} + \gamma) \cdot \mathbf{z}(X) \\ &\quad - (\bar{a} + \beta \bar{s}_{\sigma_1} + \gamma) (\bar{b} + \beta \bar{s}_{\sigma_2} + \gamma) (\bar{c} + \beta \cdot S_{\sigma_3}(X) + \gamma) \bar{z}_{\omega} , \\ \mathbf{A}_2(X) &= (\mathbf{z}(X) - 1) L_1(\mathfrak{z}) , \\ \mathbf{r}(X) &= \mathbf{A}_0(X) + \alpha \mathbf{A}_1(X) + \alpha^2 \mathbf{A}_2(X) \\ &\quad - \begin{cases} \mathbf{Z}_{\mathbb{H}}(\mathfrak{z}) \cdot (\mathbf{t}_{\text{lo}}(X) + \mathfrak{z}^n \mathbf{t}_{\text{mid}}(X) + \mathfrak{z}^{2n} \mathbf{t}_{\text{hi}}(X)) & \text{(in Plonk and SanPlonk)} \\ \mathbf{Z}_{\mathbb{H}}(\mathfrak{z}) \cdot \mathbf{t}(X) & \text{(in SmallPlonk)} \end{cases} . \end{aligned} \tag{8}$$

Let

$$\begin{aligned} H(X) &:= r(X) + va(X) + v^2b(X) + v^3c(X) + v^4S_{\sigma_1}(X) + v^5S_{\sigma_2}(X) \\ &\quad + v^6(t_{lo}(X) + \delta t_{mid}(X) + \delta^2 t_{hi}(X)) , \\ \bar{H} &:= v\bar{a} + v^2\bar{b} + v^3\bar{c} + v^4\bar{s}_{\sigma_1} + v^5\bar{s}_{\sigma_2} + v^6\bar{t}_3 . \end{aligned} \quad (9)$$

Compute the opening proof polynomials $W_{\mathfrak{z}}(X) := (H(X) - \bar{H})/(X - \mathfrak{z})$ and

$$W_{\mathfrak{z}\omega}(X) = \frac{z(X) - \bar{z}_\omega}{X - \mathfrak{z}\omega} .$$

$$[W_{\mathfrak{z}}]_1 := [W_{\mathfrak{z}}(x)]_1; [W_{\mathfrak{z}\omega}]_1 := [W_{\mathfrak{z}\omega}(x)]_1; \text{ Send } [W_{\mathfrak{z}}, W_{\mathfrak{z}\omega}]_1;$$

Verification algorithm.

1. Validate $(w_i)_{i=1}^\ell \in \mathbb{F}^\ell$ and $(\bar{a}, \bar{b}, \bar{c}, \bar{s}_{\sigma_1}, \bar{s}_{\sigma_2}, \bar{z}_\omega, \bar{t}_3) \in \mathbb{F}^7$.
2. Validate that $[a, b, c, z]_1, [t_{lo}, t_{mid}, t_{hi}]_1$ ($[t]_1$ in **SmallPlonk**), $[W_{\mathfrak{z}}, W_{\mathfrak{z}\omega}]_1 \in \mathbb{G}_1$.
3. Compute $Z_{\mathbb{H}}(\mathfrak{z}) = \mathfrak{z}^n - 1$, $L_1(\mathfrak{z}) = \frac{\omega(\mathfrak{z}^n - 1)}{n(\mathfrak{z} - \omega)}$, and $\text{PI}(\mathfrak{z}) = \sum_{i \in [\ell]} w_i L_i(\mathfrak{z})$.
4. Split r into its constant and non-constant terms. Compute r 's constant term: $r_0 := \text{PI}(\mathfrak{z}) - L_1(\mathfrak{z})\alpha^2 - \alpha(\bar{a} + \beta\bar{s}_{\sigma_1} + \gamma)(\bar{b} + \beta\bar{s}_{\sigma_2} + \gamma)(\bar{c} + \gamma)\bar{z}_\omega$, and let $r'(X) := r(X) - r_0$.
5. Sample $u \leftarrow \mathbb{F}$ and compute the first part of the batched polynomial commitment $[D]_1 := [r'(x)]_1 + u \cdot [z]_1$:

$$\begin{aligned} [D]_1 &:= \bar{a}\bar{b} \cdot [\mathbf{q}_M]_1 + \bar{a} \cdot [\mathbf{q}_L]_1 + \bar{b} \cdot [\mathbf{q}_R]_1 + \bar{c} \cdot [\mathbf{q}_O]_1 + [\mathbf{q}_C]_1 \\ &\quad + ((\bar{a} + \beta\mathfrak{z} + \gamma)(\bar{b} + \beta k_1\mathfrak{z} + \gamma)(\bar{c} + \beta k_2\mathfrak{z} + \gamma)\alpha + L_1(\mathfrak{z})\alpha^2 + u) \cdot [z]_1 \\ &\quad - (\bar{a} + \beta\bar{s}_{\sigma_1} + \gamma)(\bar{b} + \beta\bar{s}_{\sigma_2} + \gamma)\alpha\beta\bar{z}_\omega \cdot [s_{\sigma_3}]_1 \\ &\quad - \begin{cases} Z_{\mathbb{H}}(\mathfrak{z}) \cdot ([t_{lo}]_1 + \mathfrak{z}^n \cdot [t_{mid}]_1 + \mathfrak{z}^{2n} \cdot [t_{hi}]_1) & \text{(in Plonk and SanPlonk)} \\ Z_{\mathbb{H}}(\mathfrak{z}) \cdot [t]_1 & \text{(in SmallPlonk)} \end{cases} . \end{aligned}$$

6. Compute full batched polynomial commitment $[F]_1 := [D]_1 + v \cdot [a]_1 + v^2 \cdot [b]_1 + v^3 \cdot [c]_1 + v^4 \cdot [s_{\sigma_1}]_1 + v^5 \cdot [s_{\sigma_2}]_1 + v^6 [t_{lo} + \delta t_{mid} + \delta^2 t_{hi}]_1$.
7. Compute the batch evaluation $E_0 := -r_0 + v\bar{a} + v^2\bar{b} + v^3\bar{c} + v^4\bar{s}_{\sigma_1} + v^5\bar{s}_{\sigma_2} + v^6\bar{t}_3$, $E_1 := \bar{z}_\omega$, and $E := E_0 + uE_1$.
8. Batch validate all evaluations:

$$([W_{\mathfrak{z}}]_1 + u \cdot [W_{\mathfrak{z}\omega}]_1) \bullet [x]_2 \stackrel{?}{=} (\mathfrak{z} \cdot [W_{\mathfrak{z}}]_1 + u\mathfrak{z}\omega \cdot [W_{\mathfrak{z}\omega}]_1 + [F]_1 - [E]_1) \bullet [1]_2 . \quad (10)$$

Clearly, **SanPlonk** remains complete after the highlighted changes. In Appendix E, we prove that **SanPlonk** has zero knowledge.

6.3 Special-Soundness Proof of (San)Plonk's IP

We prove that **Plonk**, **SanPlonk**, and **SmallPlonk** (all described in the previous section) have computational special-soundness.

Before going on, we note that the **Plonk** verifier performs batch verification, using a batching coefficient u created after the prover's last message. That is,

after unrolling all optimizations, Eq. (10) can be replaced by two checks,

$$\begin{aligned} \left[\begin{array}{l} r+v(a-\bar{a})+v^2(b-\bar{b})+v^3(c-\bar{c})+v^4(s_{\sigma 1}-\bar{s}_{\sigma 1}) \\ +v^5(s_{\sigma 2}-\bar{s}_{\sigma 2})+v^6(t_{lo}+\delta t_{mid}+\delta^2 t_{hi}-\bar{t}_3) \end{array} \right]_1 \bullet [1]_2 = [W_3]_1 \bullet [x-3]_2, \\ [z-\bar{z}_\omega]_1 \bullet [1]_2 = [W_{3\omega}]_1 \bullet [x-3\omega]_2, \end{aligned} \quad (11)$$

where $[r]_1 = [D]_1 - u[z]_1 + r_0[1]_1$. Batching with u just introduces a soundness error $1/|\mathbb{F}|$. We say that a Plonk transcript is accepting when Eq. (11) are satisfied.

We divide the proof into several smaller lemmas. In Section 6.4, we will analyze a subtree of Plonk's, SanPlonk's, and SmallPlonk's accepting transcripts for fixed β , γ , and α , showing that from it, one can extract certain polynomials. In Section 6.5, we use that result to prove the special-soundness of Plonk, SanPlonk, and SmallPlonk.

In the following, $n \in \text{poly}(\lambda)$ and,

$$\begin{aligned} \kappa_{\text{kzg}} &= \begin{cases} n+5 & (\text{in Plonk and SanPlonk}) \\ 3n+5 & (\text{in SmallPlonk}) \end{cases}, \\ \kappa_{\text{Plonk}} &= (\kappa_\beta = 3n+1, \kappa_\gamma = 3n+1, \kappa_\alpha = 3, \kappa_3 = 4\kappa_{\text{kzg}}+1, \kappa_v^{\text{Plonk}} = 6), \\ \kappa_{\text{san}} &= (\kappa_\beta = 3n+1, \kappa_\gamma = 3n+1, \kappa_\alpha = 3, \kappa_3 = 4\kappa_{\text{kzg}}+1, \kappa_\delta = 3, \kappa_v^{\text{san}} = 7) \end{aligned} \quad (12)$$

corresponding to the branching factors of KZG, Plonk, SanPlonk, and SmallPlonk. Moreover, SmallPlonk uses $\kappa_{\text{SmallPlonk}} = \kappa_{\text{Plonk}}$, except for a different value of $\kappa_{\text{kzg}} = 3n+5$. Moreover, define $\kappa_v = \kappa_v^{\text{Plonk}}$ in the case of Plonk and $\kappa_v = \kappa_v^{\text{san}}$ in the case of SanPlonk.

6.4 Subtree Analysis

In this subsection, we analyze a subtree of Plonk's transcripts that results from fixing β , γ , and α . As usual, we start with a tree extractor lemma that gets a tree of accepting Plonk transcripts as input and outputs many accepting KZG transcripts that open relevant polynomial commitments at many different locations.

Lemma 2 (Plonk/SanPlonk/SmallPlonk transcript tree to KZG transcripts). *Let \mathcal{T} be a κ_{Plonk} -tree of Plonk's (resp., κ_{san} -tree of SanPlonk's and $\kappa_{\text{SmallPlonk}}$ -tree of SmallPlonk's) accepting transcripts. Let $\hat{\mathcal{T}}_{\beta\gamma\alpha}$ be a $(1, 1, 1, \kappa_3, \kappa_\delta, \kappa_v)$ -subtree of \mathcal{T} for any fixed β , γ , and α , with the transcripts in this subtree denoted as*

$$\text{tr}_{ijk} = \left(\begin{array}{l} [a, b, c]_1, \beta, \gamma, [z]_1, \alpha, [t_{lo}, t_{mid}, t_{hi}]_1 \text{ ([}t_1\text{ in SmallPlonk)}, \\ \mathfrak{z}_i, \bar{a}_i, \bar{b}_i, \bar{c}_i, \bar{s}_{\sigma 1i}, \bar{s}_{\sigma 2i}, \bar{z}_{\omega i}, \delta_{ij}, \bar{t}_{3ij}, v_{ijk}, [W_{3ijk}, W_{3\omega ij}]_1 \end{array} \right). \quad (13)$$

The DPT algorithm $\text{TE}_{*\text{plonk}}(\text{ck}, \hat{\mathcal{T}}_{\beta\gamma\alpha})$ in Fig. 6 computes a tuple $((\mathbf{k.tr}_k)_k, \mathbf{k.tr}^\omega)$, where $k \in [1, \kappa_v^{\text{Plonk}}] = [1, 6]$ in Plonk and

$\text{TE}_{*\text{plonk}}(\text{ck}, \hat{\mathcal{T}}_{\beta\gamma\alpha})$	
1 :	Parse $\hat{\mathcal{T}}_{\beta\gamma\alpha} = (\text{tr}_{ijk})_{i \in [1, \kappa_3], j \in [1, \kappa_\delta], k \in [1, \kappa_v]}$; // tr_{ijk} as in Eq. (13); $\kappa_\delta = 1$ in Plonk
2 :	for $i \in [1, \kappa_3]$ do
3 :	$[A_{0i}]_1 \leftarrow \bar{a}_i \bar{b}_i [\mathbf{q}_M]_1 + \bar{a}_i [\mathbf{q}_L]_1 + \bar{b}_i [\mathbf{q}_R]_1 + \bar{c}_i [\mathbf{q}_O]_1 + \text{Pl}(\mathfrak{z})[1]_1 + [\mathbf{q}_c]_1$;
4 :	$[A_{1i}]_1 \leftarrow (\bar{a}_i + \beta \mathfrak{z}_i + \gamma)(\bar{b}_i + \beta k_1 \mathfrak{z}_i + \gamma)(\bar{c}_i + \beta k_2 \mathfrak{z}_i + \gamma)[z]_1$
5 :	$\quad \quad \quad - (\bar{a}_i + \beta \bar{s}_{\sigma 1} + \gamma)(\bar{b}_i + \beta \bar{s}_{\sigma 2} + \gamma)(\bar{c}_i[1]_1 + \beta[\mathbf{S}_{\sigma 3}(x)]_1 + \gamma[1]_1)\bar{z}_{\omega, i}$;
6 :	$[A_{2i}]_1 \leftarrow [z - 1]_1 L_1(\mathfrak{z}_i)$;
7 :	$[r_i]_1 \leftarrow [A_{0i}]_1 + \alpha[A_{1i}]_1 + \alpha^2[A_{2i}]_1 - \begin{cases} Z_{\mathbb{H}}(\mathfrak{z}_i) \cdot ([t_{l0}]_1 + \mathfrak{z}_i^n [t_{mid}]_1 + \mathfrak{z}_i^{2n} [t_{hi}]_1) \\ Z_{\mathbb{H}}(\mathfrak{z}_i) \cdot [t]_1 \end{cases}$;
8 :	$[W'_{\mathfrak{z}_i 11}, \dots, W'_{\mathfrak{z}_i 1\kappa_v}]_1^\top \leftarrow \mathbf{V}_{\mathfrak{z}_i 1}^{-1} [W_{\mathfrak{z}_i 11}, \dots, W_{\mathfrak{z}_i 1\kappa_v}]_1^\top$; // $\mathbf{V}_{\mathfrak{z}_i 1}$ as in 15, 16
9 :	$\mathbf{k.tr}_{1i} \leftarrow ([r]_1, \mathfrak{z}_i, 0, [W'_{\mathfrak{z}_i 11}]_1)$; $\mathbf{k.tr}_{2i} \leftarrow ([a]_1, \mathfrak{z}_i, \bar{a}_i, [W'_{\mathfrak{z}_i 12}]_1)$;
10 :	$\mathbf{k.tr}_{3i} \leftarrow ([b]_1, \mathfrak{z}_i, \bar{b}_i, [W'_{\mathfrak{z}_i 13}]_1)$; $\mathbf{k.tr}_{4i} \leftarrow ([c]_1, \mathfrak{z}_i, \bar{c}_i, [W'_{\mathfrak{z}_i 14}]_1)$;
11 :	$\mathbf{k.tr}_{5i} \leftarrow ([s_{\sigma 1}]_1, \mathfrak{z}_i, \bar{s}_{\sigma 1}, [W'_{\mathfrak{z}_i 15}]_1)$; $\mathbf{k.tr}_{6i} \leftarrow ([s_{\sigma 2}]_1, \mathfrak{z}_i, \bar{s}_{\sigma 2}, [W'_{\mathfrak{z}_i 16}]_1)$;
12 :	for $j \in \{1, 2, 3\}$ do $\bar{t}_{\mathfrak{z}_i j} \leftarrow (\mathbf{V}_{ij}^{-1})_7 (\bar{H}_{ij1}, \dots, \bar{H}_{ij7})^\top$; endfor
13 :	$(\bar{t}_{\mathfrak{z}_i lo, i}, \bar{t}_{\mathfrak{z}_i mid, i}, \bar{t}_{\mathfrak{z}_i hi, i})^\top := \mathbf{C}_i^{-1} (\bar{t}_{\mathfrak{z}_i 1}, \bar{t}_{\mathfrak{z}_i 2}, \bar{t}_{\mathfrak{z}_i 3})^\top$;
14 :	$[W_{lo, i}, W_{mid, i}, W_{hi, i}]_1 \leftarrow \mathbf{C}_i^{-1} [W'_{\mathfrak{z}_i 17}, W'_{\mathfrak{z}_i 27}, W'_{\mathfrak{z}_i 37}]_1^\top$;
15 :	$\mathbf{k.tr}_{7i} \leftarrow ([t_{l0}]_1, \mathfrak{z}_i, \bar{t}_{\mathfrak{z}_i lo, i}, [W_{lo, i}]_1)$; $\mathbf{k.tr}_{8i} \leftarrow ([t_{mid}]_1, \mathfrak{z}_i, \bar{t}_{\mathfrak{z}_i mid, i}, [W_{mid, i}]_1)$;
16 :	$\mathbf{k.tr}_{9i} \leftarrow ([t_{hi}]_1, \mathfrak{z}_i, \bar{t}_{\mathfrak{z}_i hi, i}, [W_{hi, i}]_1)$;
17 :	$\mathbf{k.tr}_i^\omega \leftarrow ([z]_1, \mathfrak{z}_i \omega, \bar{z}_{\omega i}, [W_{\mathfrak{z}_i \omega i 11}]_1)$; endfor
18 :	$\mathbf{k.tr}^\omega \leftarrow (\mathbf{k.tr}_i^\omega)_{i \in [1, \kappa_3]}$; for $k \in [1, \kappa_v + 2]$ do $\mathbf{k.tr}_k \leftarrow (\mathbf{k.tr}_{ki})_{i \in [1, \kappa_3]}$; endfor
19 :	return $((\mathbf{k.tr}_k)_{k \in [1, \kappa_v + 2]}, \mathbf{k.tr}^\omega)$;

Fig. 6. The subroutine $\text{TE}_{*\text{plonk}}$.

$k \in [1, \kappa_v^{\text{san}} + 2] = [1, 9]$ in SanPlonk , of KZG accepting transcripts, such that (1) $\mathbf{k.tr}_{1i}$, $\mathbf{k.tr}_{2i}$, $\mathbf{k.tr}_{3i}$, $\mathbf{k.tr}_{4i}$, $\mathbf{k.tr}_{5i}$, and $\mathbf{k.tr}_{6i}$ open (respectively) $[r_i]_1$, $[a]_1$, $[b]_1$, $[c]_1$, $[s_{\sigma 1}]_1$ and $[s_{\sigma 2}]_1$ to 0, \bar{a}_i , \bar{b}_i , \bar{c}_i , $\bar{s}_{\sigma 1}$, and $\bar{s}_{\sigma 2}$ at \mathfrak{z}_i , and (2) $\mathbf{k.tr}_i^\omega$ opens $[z]_1$ to $\bar{z}_{\omega i}$ at $\mathfrak{z}_i \omega$. In addition, in SanPlonk , $\mathbf{k.tr}_{7i}$, $\mathbf{k.tr}_{8i}$, and $\mathbf{k.tr}_{9i}$ open (respectively) $[t_{l0}, t_{mid}, t_{hi}]_1$ to some values $\bar{t}_{\mathfrak{z}_i lo, i}$, $\bar{t}_{\mathfrak{z}_i mid, i}$, and $\bar{t}_{\mathfrak{z}_i hi, i}$ at \mathfrak{z}_i . Moreover, \mathfrak{z}_i are mutually different.

Proof. Let \mathcal{T} be a $(\kappa_\beta, \kappa_\gamma, \kappa_\alpha, \kappa_3, \dots)$ -tree of accepting transcripts. We fix a $(1, 1, 1, \kappa_3, \kappa_\delta, \kappa_v)$ -subtree $\hat{\mathcal{T}}_{\beta\gamma\alpha}$ of \mathcal{T} for some β, γ , and α . Then, $\hat{\mathcal{T}}_{\beta\gamma\alpha} = \{\text{tr}_{ijk}\}$ contains accepting transcripts given in Eq. (13) with mutually different \mathfrak{z}_i .

Let us unload some of the formulas in Fig. 6. First, $[A_{0i}]_1$, $[A_{1i}]_1$, $[A_{2i}]_1$, $[r_i]_1$ (lines 3, 4, 6, and 7 in Fig. 6) are commitments to $(\mathfrak{z}_i$ -dependent) polynomials A_{0i} , A_{1i} , A_{2i} , and r , defined as in Eq. (8).

Recall that we analyze in the case the verifier individually tests the two verification equations in Eq. (11), ignoring the optimization induced by using the batching variable u . Let

$$\begin{aligned}
 \bar{H}_{ijk} &\leftarrow v_{ijk}^0 \cdot 0 + v_{ijk}^1 \bar{a}_i + v_{ijk}^2 \bar{b}_i + v_{ijk}^3 \bar{c}_i + v_{ijk}^4 \bar{s}_{\sigma 1} + v_{ijk}^5 \bar{s}_{\sigma 2} + v_{ijk}^6 \bar{t}_{\mathfrak{z}_i j} \quad , \\
 [\mathbf{t}_{\delta_{ij}}]_1 &\leftarrow [t_{l0} + \delta_{ij} t_{mid} + \delta_{ij}^2 t_{hi}]_1 \quad , \\
 [\mathbf{H}_{ijk}]_1 &\leftarrow v_{ijk}^0 [r]_1 + v_{ijk}^1 [a]_1 + v_{ijk}^2 [b]_1 + v_{ijk}^3 [c]_1 + v_{ijk}^4 [s_{\sigma 1}]_1 + v_{ijk}^5 [s_{\sigma 2}]_1 + v_{ijk}^6 [\mathbf{t}_{\delta_{ij}}]_1 .
 \end{aligned} \tag{14}$$

Since KZG is triply homomorphic, tr_{ijk} is an accepting Plonk transcript iff $([H_{ijk}]_1, \mathfrak{z}_i, \bar{H}_{ijk}, [W_{3ijk}]_1)$ and k.tr_i^ω (line 17 in Fig. 6) are accepting KZG transcripts. We get this by slight rewriting of Plonk's verification equation. In particular, k.tr_i^ω are accepting transcripts, with different values of \mathfrak{z}_i .

We will separate the rest of the proof to the case of (1) Plonk and SmallPlonk, and (2) SanPlonk. However, the subroutine on Fig. 6 corresponds to both.

Plonk And SmallPlonk. Here,

$$\mathbf{V}_i = \begin{pmatrix} 1 & v_{i1} & \dots & v_{i1}^5 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & v_{i6} & \dots & v_{i6}^5 \end{pmatrix}, \quad (15)$$

is an invertible Vandermonde matrix. According to Eq. (14), $(\bar{H}_{i1}, \dots, \bar{H}_{i6})^\top = \mathbf{V}_i \cdot (0, \bar{a}_i, \bar{b}_i, \bar{c}_i, \bar{s}_{\sigma 1i}, \bar{s}_{\sigma 2i})^\top$ and $[H_{i1}, \dots, H_{i6}]_1^\top = \mathbf{V}_i \cdot [r_i, a, b, c, s_{\sigma 1}, s_{\sigma 2}]_1^\top$. Let $[W'_{3i1}, \dots, W'_{3i6}]_1^\top$ be as on line 8 of Fig. 6. By the triple homomorphism of KZG, k.tr_{ki} are accepting KZG transcripts for every i and k .

SanPlonk. Here,

$$\mathbf{V}_{ij} = \begin{pmatrix} 1 & v_{ij1} & \dots & v_{ij1}^6 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & v_{ij7} & \dots & v_{ij7}^6 \end{pmatrix} \quad \text{and} \quad \mathbf{C}_i := \begin{pmatrix} 1 & \delta_{i1} & \delta_{i1}^2 \\ 1 & \delta_{i2} & \delta_{i2}^2 \\ 1 & \delta_{i3} & \delta_{i3}^2 \end{pmatrix} \quad (16)$$

are invertible Vandermonde matrices. According to Eq. (14), $[H_{ij1}, \dots, H_{ij7}]_1 = \mathbf{V}_{ij} \cdot [r_i, a, b, c, s_{\sigma 1}, s_{\sigma 2}, \mathfrak{t}_{\delta_{ij}}]_1^\top$ and $(\bar{H}_{ij1}, \dots, \bar{H}_{ij7})^\top = \mathbf{V}_{ij} \cdot (0, \bar{a}_i, \bar{b}_i, \bar{c}_i, \bar{s}_{\sigma 1i}, \bar{s}_{\sigma 2i}, \bar{t}_{3ij})^\top$. Let $[W'_{3ij1}, \dots, W'_{3ij7}]_1^\top$ be defined as on line 8 of Fig. 6. By triple homomorphism, $\text{k.tr}_{1i}, \dots, \text{k.tr}_{6i}$ and k.tr_{7i}^* are accepting KZG transcripts, where k.tr_{1i} to k.tr_{6i} are as in Fig. 6 and $\text{k.tr}_{7i}^* := ([\mathfrak{t}_{\delta_{ij}}]_1, \mathfrak{z}_i, \bar{t}_{3ij}, [W'_{3ij7}]_1)$.

Next, $\mathbf{C}_i \cdot [t_{lo}, t_{mid}, t_{hi}]_1^\top = [\mathfrak{t}_{\delta_{i1}}, \mathfrak{t}_{\delta_{i2}}, \mathfrak{t}_{\delta_{i3}}]_1^\top$. Define $(\bar{t}_{3lo,i}, \bar{t}_{3mid,i}, \bar{t}_{3hi,i})^\top := \mathbf{C}_i^{-1} \cdot (\bar{t}_{3i1}, \bar{t}_{3i2}, \bar{t}_{3i3})^\top$ and $[W_{lo,i}, W_{mid,i}, W_{hi,i}]_1 := \mathbf{C}_i^{-1} [W'_{3i17}, W'_{3i27}, W'_{3i37}]_1^\top$. Thus, k.tr_{7i} , k.tr_{8i} , and k.tr_{9i} are accepting KZG transcripts for every i . \square

Constructing the Subtree Extractor.

Theorem 5 (Subtree extractor). *Let \mathcal{T} be a κ_{Plonk} -tree of Plonk's (resp., κ_{San} -tree of SanPlonk's or $\kappa_{\text{SmallPlonk}}$ -tree of SmallPlonk's) accepting transcripts.*

Let $\hat{\mathcal{T}}_{\beta\gamma\alpha} = (\text{tr}_{ijk})$ be a subtree of \mathcal{T} for any fixed β, γ , and α , where tr_{ijk} are as in Eq. (13). Assume that KZG is evaluation-binding and computational $(\kappa_{\text{kzg}} + 1)$ -special-sound. In the case of Plonk (resp., SmallPlonk), assume the $\text{par}_n^{\text{plonk}}$ -SplitRSDH (resp., $\text{par}_n^{\text{smallplonk}}$ -SplitRSDH) assumption holds. There exists a DPT extractor $\text{Ext}_{\text{ss}}^{\text{sub}}$ that, given $\hat{\mathcal{T}}_{\beta\gamma\alpha}$, outputs $(z(X), a(X), b(X), c(X))$, where $z(X), a(X), b(X)$, and $c(X)$ are consistent with the commitments and all κ_3 openings of $[z, a, b, c]_1$. Moreover, $\mathfrak{t}(X)$ (defined as in Eq. (7)) is a polynomial.

We postpone the proof of this result to Appendix D.2.

6.5 Full Special-Soundness Proof

Finally, we are ready to prove the special-soundness of Plonk, SanPlonk, and SmallPlonk. For this, we combine the subtree analysis of Section 6.4, KZG's

binding (required to guarantee that extracted polynomials like $a(X)$ are the same in all subtrees), and an analysis of the permutation argument.

Theorem 6. *Let $n \in \text{poly}(\lambda)$ and κ_{kzg} , κ_{Plonk} , and κ_{san} be as in Eq. (12).*

1. *If KZG is computational $(\kappa_{\text{kzg}} + 1)$ -special-sound and evaluation-binding, and $\text{par}_n^{\text{smallplonk}}$ -SplitRSDH holds, then SmallPlonk is computational κ_{Plonk} -special-sound.*
2. *If KZG is computational $(\kappa_{\text{kzg}} + 1)$ -special-sound and evaluation-binding, and $\text{par}_n^{\text{plonk}}$ -SplitRSDH holds, then Plonk is computational κ_{Plonk} -special-sound.*
3. *If KZG is computational $(\kappa_{\text{kzg}} + 1)$ -special-sound and evaluation-binding, then SanPlonk is computational κ_{san} -special-sound.*

Due to the page limit, we postpone the proof to Appendix D.5. Note that we prove special-soundness for the information-theoretic argument underlying Plonk as an implicit sub-claim in the proof of the previous theorem. To our knowledge, this is the first time that such a strong security property has been proven for Plonk’s information-theoretic component.

Acknowledgments. We thank Antonio Faonio for useful discussions about [FFR24]. The first author was co-funded by the European Union and Estonian Research Council via project TEM-TA119.

References

- AFK22. Thomas Attema, Serge Fehr, and Michael Kloof. Fiat-shamir transformation of multi-round interactive proofs. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 113–142. Springer, Cham, November 2022. doi:10.1007/978-3-031-22318-1_5. 1, D.6, 10
- AFKR23. Thomas Attema, Serge Fehr, Michael Kloof, and Nicolas Resch. The fiat–shamir transformation of $(\gamma_1, \dots, \gamma_\mu)$ -special-sound interactive proofs. Technical Report 2023/1945, IACR, December 22, 2023. URL: <https://eprint.iacr.org/2023/1945>. 1, 1.1, D.6
- BBB⁺18. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018. doi:10.1109/SP.2018.00020. 1, 2.1
- BBHR18. Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast reed-solomon interactive oracle proofs of proximity. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *ICALP 2018*, volume 107 of *LIPICs*, pages 14:1–14:17. Schloss Dagstuhl, July 2018. doi:10.4230/LIPICs.ICALP.2018.14. 1, 2.1
- BFHK23. Balthazar Bauer, Pooya Farshim, Patrick Harasser, and Markulf Kohlweiss. The Uber-Knowledge Assumption: A Bridge to the AGM. Technical Report 2023/1601, IACR, November 9, 2023. URL: <https://eprint.iacr.org/2023/1601>. 1

- BFS20. Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 677–706. Springer, Cham, May 2020. doi:10.1007/978-3-030-45721-1_24. 1, 2.1
- CFF⁺21. Matteo Campanelli, Antonio Faonio, Dario Fiore, Anaïs Querol, and Hadrián Rodríguez. Lunar: A toolbox for more efficient universal and updatable zkSNARKs and commit-and-prove extensions. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 3–33. Springer, Cham, December 2021. doi:10.1007/978-3-030-92078-4_1. 1, 4.1
- CFF⁺24. Matteo Campanelli, Antonio Faonio, Dario Fiore, Tianyu Li, and Helger Lipmaa. Lookup arguments: Improvements, extensions and applications to zero-knowledge decision trees. In Qiang Tang and Vanessa Teague, editors, *PKC 2024, Part II*, volume 14602 of *LNCS*, pages 337–369. Springer, Cham, April 2024. doi:10.1007/978-3-031-57722-2_11. 1
- CGH98. Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th ACM STOC*, pages 209–218. ACM Press, May 1998. doi:10.1145/276698.276741. 1
- CHM⁺20. Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Psi Veseley, and Nicholas P. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 738–768. Springer, Cham, May 2020. doi:10.1007/978-3-030-45721-1_26. 1, 1.1, 2.1, 4.1, 4.2
- Den02. Alexander W. Dent. Adapting the weaknesses of the random oracle model to the generic group model. In Yuliang Zheng, editor, *ASIACRYPT 2002*, volume 2501 of *LNCS*, pages 100–109. Springer, Berlin, Heidelberg, December 2002. doi:10.1007/3-540-36178-2_6. 1
- DG23. Quang Dao and Paul Grubbs. Spartan and bulletproofs are simulation-extractable (for free!). In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part II*, volume 14005 of *LNCS*, pages 531–562. Springer, Cham, April 2023. doi:10.1007/978-3-031-30617-4_18. 1, 1.1
- DP23. Benjamin E. Diamond and Jim Posen. Succinct arguments over towers of binary fields. Technical Report 2023/1784, IACR, November 17, 2023. URL: <https://eprint.iacr.org/2023/1784>. 1
- EFG22. Liam Eagen, Dario Fiore, and Ariel Gabizon. cq: Cached quotients for fast lookups. Cryptology ePrint Archive, Report 2022/1763, 2022. URL: <https://eprint.iacr.org/2022/1763>. 1
- FFR24. Antonio Faonio, Dario Fiore, and Luigi Russo. Real-world Universal zk-SNARKs Are Non-Malleable. In *ACM CCS 2024*, pages ?–?, Salt Lake City, USA, October 14–18, 2024. ACM. Accepted. 1, 1, 1.1, 3, 5, 4.2, 4.2, 4.2, 4.3, 6.5, A.1, A.2
- FKL18. Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Cham, August 2018. doi:10.1007/978-3-319-96881-0_2. 1, 2.1, 3, 4, 3
- GKM⁺18. Jens Groth, Markulf Kohlweiss, Mary Maller, Sarah Meiklejohn, and Ian Miers. Updatable and universal common reference strings with applications to zk-SNARKs. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 698–728. Springer, Cham, August 2018. doi:10.1007/978-3-319-96878-0_24. 1

- GKP22. Chaya Ganesh, Hamidreza Khoshakhlagh, and Roberto Parisella. Niwi and new notions of extraction for algebraic languages. In Clemente Galdi and Stanislaw Jarecki, editors, *Security and Cryptography for Networks*, pages 687–710, Cham, 2022. Springer International Publishing. 1.1
- GLS⁺23. Alexander Golovnev, Jonathan Lee, Srinath T. V. Setty, Justin Thaler, and Riad S. Wahby. Brakedown: Linear-time and field-agnostic SNARKs for R1CS. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 193–226. Springer, Cham, August 2023. doi:10.1007/978-3-031-38545-2_7. 1
- GW11. Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108. ACM Press, June 2011. doi:10.1145/1993636.1993651. 1
- GW19. Ariel Gabizon and Zachary J. Williamson. The turbo-plonk program syntax for specifying snark programs, 2019. URL: https://docs.zkproof.org/pages/standards/accepted-workshop3/proposal-turbo_plonk.pdf. 1
- GW20. Ariel Gabizon and Zachary J. Williamson. plookup: A simplified polynomial protocol for lookup tables. Cryptology ePrint Archive, Report 2020/315, 2020. URL: <https://eprint.iacr.org/2020/315>. 1
- GWC19. Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. URL: <https://eprint.iacr.org/2019/953>. 1, 1, 1.1, 2, 4.1, 4.2, 6, 6.1, 6.1, D.1, E
- Ica09. Thomas Icart. How to hash into elliptic curves. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 303–316. Springer, Berlin, Heidelberg, August 2009. doi:10.1007/978-3-642-03356-8_18. 3, 3
- JM24. Joseph Jaeger and Deep Inder Mohan. Generic and algebraic computation models: When AGM proofs transfer to the GGM. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part V*, volume 14924 of *LNCS*, pages 14–45. Springer, Cham, August 2024. doi:10.1007/978-3-031-68388-6_2. 1, 1, 1.1, 3, 3, 3, C
- JR10. Tibor Jager and Andy Rupp. The semi-generic group model and applications to pairing-based cryptography. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 539–556. Springer, Berlin, Heidelberg, December 2010. doi:10.1007/978-3-642-17373-8_31. 6
- KZG10. Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Berlin, Heidelberg, December 2010. doi:10.1007/978-3-642-17373-8_11. 1, 1.1, 2.1, 2.1, 3
- Lip12. Helger Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 169–189. Springer, Berlin, Heidelberg, March 2012. doi:10.1007/978-3-642-28914-9_10. 2.1, 3, C
- Lip24. Helger Lipmaa. Polymath: Groth16 is not the limit. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part X*, volume 14929 of *LNCS*, pages 170–206. Springer, Cham, August 2024. doi:10.1007/978-3-031-68403-6_6. 1

- LPS23. Helger Lipmaa, Roberto Parisella, and Janno Siim. Algebraic group model with oblivious sampling. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part IV*, volume 14372 of *LNCS*, pages 363–392. Springer, Cham, November / December 2023. doi:10.1007/978-3-031-48624-1_14. 1, 3, 1.1, 2.1, 3, 3, 4, 3, 4.2, 4.3, A.1, C, C
- LPS24. Helger Lipmaa, Roberto Parisella, and Janno Siim. Constant-size zk-SNARKs in ROM from falsifiable assumptions. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 34–64. Springer, Cham, May 2024. doi:10.1007/978-3-031-58751-1_2. 1, 1, 1.1, 2, 2.1, 2.1, 4
- LSZ22. Helger Lipmaa, Janno Siim, and Michal Zajac. Counting vampires: From univariate sumcheck to updatable ZK-SNARK. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part II*, volume 13792 of *LNCS*, pages 249–278. Springer, Cham, December 2022. doi:10.1007/978-3-031-22966-4_9. 1, 4.1
- Mau05. Ueli M. Maurer. Abstract models of computation in cryptography (invited paper). In Nigel P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12. Springer, Berlin, Heidelberg, December 2005. doi:10.1007/11586821_1. 1, 3
- PST13. Charalampos Papamanthou, Elaine Shi, and Roberto Tamassia. Signatures of correct computation. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 222–242. Springer, Berlin, Heidelberg, March 2013. doi:10.1007/978-3-642-36594-2_13. 2.1
- RZ21. Carla Ràfols and Arantxa Zapico. An algebraic framework for universal and updatable SNARKs. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 774–804, Virtual Event, August 2021. Springer, Cham. doi:10.1007/978-3-030-84242-0_27. 1, 4.1
- Sef24. Marek Sefranek. How (not) to simulate PLONK. Technical Report 2024/848, IACR, May 31, 2024. URL: <https://eprint.iacr.org/2024/848>. 1.1, E, 7, E
- Sho97. Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Berlin, Heidelberg, May 1997. doi:10.1007/3-540-69053-0_18. 1, 3
- STW24. Srinath T. V. Setty, Justin Thaler, and Riad S. Wahby. Unlocking the lookup singularity with Lasso. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 180–209. Springer, Cham, May 2024. doi:10.1007/978-3-031-58751-1_7. 1
- Zha22. Mark Zhandry. To label, or not to label (in generic groups). In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part III*, volume 13509 of *LNCS*, pages 66–96. Springer, Cham, August 2022. doi:10.1007/978-3-031-15982-4_3. 1, 3
- ZZK22. Cong Zhang, Hong-Sheng Zhou, and Jonathan Katz. An analysis of the algebraic group model. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 310–322. Springer, Cham, December 2022. doi:10.1007/978-3-031-22972-5_11. 1, 3

A Insecurity of Lin

A.1 Attack Against Knowledge-Soundness of Lin.

For the sake of completeness, we reproduce the attack of [LPS23,FFR24] against knowledge-soundness of Lin.

The adversary \mathcal{A} chooses some tuple $\{\mathbf{a}_s(X)\}$ of polynomials, such that $V := \{\mathbf{g}_t(\mathbf{a}(X))\}_{t=1}^{n_b}$ is linearly dependent. \mathcal{A} chooses a non-zero vector $\bar{\mathbf{d}}^*$ that is orthogonal to V , i.e., $\sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}(X))\bar{d}_t^* = 0$. \mathcal{A} samples $[D]_1 \leftarrow_{\$} \mathbb{G}_1$ obliviously, and sets $[d_t]_1 \leftarrow \bar{d}_t^*[D]_1$. \mathcal{A} sets $\bar{a}_s \leftarrow \mathbf{a}_s(\mathfrak{z})$ and $[h]_1 \leftarrow [(\sum_{s=1}^{n_a} \beta^{s-1}(\mathbf{a}_s(x) - \bar{a}_s)) / (x - \mathfrak{z})]_1$. Clearly,

$$\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}})d_t = \sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}(\mathfrak{z}))\bar{d}_t^* D = (\sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}(\mathfrak{z}))\bar{d}_t^*) D = 0 \quad ,$$

and thus

$$\begin{aligned} [\sum_{s=1}^{n_a} \beta^{s-1}(\mathbf{a}_s(x) - \bar{a}_s) + \beta^{n_a} \sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}})d_t]_1 \bullet [1]_2 &= [\sum_{s=1}^{n_a} \beta^{s-1}(\mathbf{a}_s(x) - \bar{a}_s)]_1 \\ &= [h]_1 \bullet [x - \mathfrak{z}]_2 \quad . \end{aligned}$$

Thus, the Lin verifier accepts, but it is impossible to extract from obliviously sampled $[D]_1$ and thus from any of $[d_t]_1$.

A.2 Proof of Theorem 3.

Finally, we give the proof of Theorem 3, which shows that linear independence is insufficient for proving computational special-soundness of Lin.

Proof. We prove impossibility for the next special case. The general case follows. For the rest of the proof, let us fix some arbitrary polynomials $\mathbf{a}_1, \mathbf{a}_2, \mathbf{d}'_1, \mathbf{d}'_2 \in \mathbb{F}_{\leq n}[X]$, and $\mathbf{g}_1, \mathbf{g}_2 \in \mathbb{F}[X]$, which satisfy the equation $\mathbf{g}_1(\mathbf{a}_1(X))\mathbf{d}'_1(X) + \mathbf{g}_2(\mathbf{a}_2(X))\mathbf{d}'_2(X) = 0$. In particular, $\mathbf{g}_1(X)$ and $\mathbf{g}_2(X)$ can be linearly independent, as was required in [FFR24]. For instance, one can choose $\mathbf{g}_1(X) = X$, and $\mathbf{g}_2 = 1$.

We also define $\mathbf{d}_1(X) = y\mathbf{d}'_1(X)$ and $\mathbf{d}_2(X) = y\mathbf{d}'_2(X)$ for a randomly chosen $y \leftarrow_{\$} \mathbb{F}$ (y will later act as a DL challenge). Observe that $\mathbf{g}_1(\mathbf{a}_1(X))\mathbf{d}_1(X) + \mathbf{g}_2(\mathbf{a}_2(X))\mathbf{d}_2(X) = 0$. We show that if Lin has computational (κ_1, κ_2) -special-soundness (for any $\kappa_1, \kappa_2 \in \text{poly}(\lambda)$) for committed polynomials $\mathbf{a}_1, \mathbf{a}_2, \mathbf{d}_1, \mathbf{d}_2$ and public polynomials $\mathbf{g}_1, \mathbf{g}_2$ as defined above, then it is possible to break the discrete logarithm assumption in \mathbb{G}_1 .

Let Ext_{ss} be a PPT extractor that gets as an input $\text{ck} = ([x^i]_{i=1}^n]_1, [x]_2)$ and $\text{tr}_{ij} = ([a_1, d_1, a_2, d_2]_1, \mathfrak{z}_i, A_{1i}, A_{2i}, \beta_{ij}, [h_{ij}]_1)$, for $i, j \in [1, \kappa_1] \times [1, \kappa_2]$, such that

$$[(a_1 - A_{1i}) + \beta_{ij}(a_2 - A_{2i}) + \beta_{ij}^2(\mathbf{g}_1(A_{1i})d_1 + \mathbf{g}_1(A_{2i})d_2)]_1 \bullet [1]_2 = [h_{ij}]_1 \bullet [x - \mathfrak{z}_i]_2 \quad .$$

Assume that for any computational special-soundness adversary \mathcal{B} , Ext_{ss} outputs a correct tuple $(\mathbf{a}_1(X), \mathbf{d}_1(X), \mathbf{a}_2(X), \mathbf{d}_2(X)) \in \mathcal{R}_{\text{ck}, \mathbf{g}, \text{tr}}^{\text{Lin}}$ with probability at least $1 - \varepsilon_{\mathcal{A}}$.

$\mathcal{A}(\mathbf{p}, [y]_1) :$ <hr/> $x \leftarrow_{\mathfrak{s}} \mathbb{F}; \mathbf{ck} \leftarrow ([x^i]_{i=1}^n)_1, [x]_2;$ $[a_1]_1 \leftarrow \mathbf{a}_1(x)[1]_1; [a_2]_1 \leftarrow \mathbf{a}_2(x)[1]_1; [d_1]_1 \leftarrow \mathbf{d}'_1(x)[y]_1; [d_2]_1 \leftarrow \mathbf{d}'_2(x)[y]_1;$ for $i \in [1, \kappa_1]$ do $\mathfrak{z}_i \leftarrow i; A_{1i} \leftarrow a_1(\mathfrak{z}_i); A_{2i} \leftarrow a_2(\mathfrak{z}_i);$ for $j \in [1, \kappa_2]$ do $\beta_{ij} \leftarrow j;$ $[h_{ij}]_1 \leftarrow \frac{[(a_1 - A_{1i}) + \beta_{ij}(a_2 - A_{2i})]_1}{x - \mathfrak{z}_i} + \frac{\beta_{ij}^2(\mathbf{g}_1(A_{1i})[d_1]_1 + \mathbf{g}_2(A_{2i})[d_2]_1)}{x - \mathfrak{z}_i};$ $\mathbf{tr}_{ij} \leftarrow ([a_1, d_1, a_2, d_2]_1, \mathfrak{z}_i, A_{1i}, A_{2i}, \beta_{ij}, [h_{ij}]_1);$ $(\tilde{a}_1(X), \tilde{d}_1(X), \tilde{a}_2(X), \tilde{d}_2(X)) \leftarrow \text{Ext}_{\text{ss}}(\mathbf{p}, \mathbf{ck}, \mathbf{tr});$ return $\tilde{d}_1(x)/d'_1(x);$
$\mathcal{B}(\mathbf{p}, \mathbf{ck}) :$ <hr/> $y \leftarrow_{\mathfrak{s}} \mathbb{F};$ $[a_1]_1 \leftarrow [a_1(x)]_1; [a_2]_1 \leftarrow [a_2(x)]_1; [d_1]_1 \leftarrow y[d'_1(x)]_1; [d_2]_1 \leftarrow y[d'_2(x)]_1;$ for $i \in [1, \kappa_1]$ do $\mathfrak{z}_i \leftarrow i; A_{1i} \leftarrow \mathbf{a}_1(\mathfrak{z}_i); A_{2i} \leftarrow \mathbf{a}_2(\mathfrak{z}_i);$ for $j \in [1, \kappa_2]$ do $\beta_{ij} \leftarrow j;$ $h_{ij}(X) \leftarrow \frac{a_1(X) - A_{1i} + \beta_{ij}(a_2(X) - A_{2i})}{X - \mathfrak{z}_i} + \frac{\beta_{ij}^2(\mathbf{g}_1(A_{1i})y d'_1(X) - \mathbf{g}_2(A_{2i})y d'_2(X))}{X - \mathfrak{z}_i};$ $[h_{ij}]_1 = [h_{ij}(x)]_1;$ $\mathbf{tr}_{ij} \leftarrow ([a_1, d_1, a_2, d_2]_1, \mathfrak{z}_i, A_{1i}, A_{2i}, \beta_{ij}, [h_{ij}]_1);$ return $\mathbf{tr};$

Fig. 7. Discrete logarithm adversary \mathcal{A} and computational special-soundness adversary \mathcal{B} for any fixed polynomials $\mathbf{a}_1, \mathbf{a}_2, \mathbf{d}'_1, \mathbf{d}'_2 \in \mathbb{F}_{\leq n}[X]$, and $\mathbf{g}_1, \mathbf{g}_2 \in \mathbb{F}[X]$ that satisfy $\mathbf{g}_1(\mathbf{a}_1(X))\mathbf{d}'_1(X) + \mathbf{g}_2(\mathbf{a}_2(X))\mathbf{d}'_2(X) = 0$.

In Fig. 7, we depict a PPT adversary \mathcal{A} that breaks DL with probability $\geq 1 - \varepsilon$. \mathcal{A} gets as an input a DL challenge $[y]_1$ for $y \leftarrow_{\mathfrak{s}} \mathbb{F}$. \mathcal{A} samples a new trapdoor x and generates \mathbf{ck} from x . \mathcal{A} computes commitments $[a_1]_1 \leftarrow [\mathbf{a}_1(x)]_1$, $[a_2]_1 \leftarrow [\mathbf{a}_2(x)]_1$, $[d_1]_1 \leftarrow \mathbf{d}'_1(x)[y]_1$, and $[d_2]_1 \leftarrow \mathbf{d}'_2(x)[y]_1$. (\mathcal{A} can compute $[d_1, d_2]_1$ only since it knows x .) \mathcal{A} then generates transcripts \mathbf{tr}_{ij} for $\mathfrak{z}_i = i$ and $\beta_{ij} = j$, with the opening proofs

$$[h_{ij}]_1 = \frac{[(a_1 - A_{1i}) + \beta_{ij}(a_2 - A_{2i}) + \beta_{ij}^2(\mathbf{g}_1(A_{1i})\tilde{d}_1(X) - \mathbf{g}_2(A_{2i})\tilde{d}_2(X))]_1}{x - \mathfrak{z}_i}.$$

Finally, it runs $\text{Ext}_{\text{ss}}(\mathbf{p}, \mathbf{ck}, \mathbf{tr})$ to obtain polynomials $(\tilde{a}_1(X), \tilde{d}_1(X), \tilde{a}_2(X), \tilde{d}_2(X))$ and outputs $\tilde{d}_1(x)/d'_1(x)$ as a candidate for the DL y .

Suppose \mathcal{A} succeeds in breaking DL with probability $\varepsilon_{\mathcal{A}}$. We present a sequence of games to show that $\varepsilon_{\mathcal{A}} \geq 1 - \varepsilon$, which implies that ε cannot be negligible if the discrete logarithm assumption holds.

Game₁: The first game in Fig. 8 samples \mathbf{ck} in the first step and generates a discrete logarithm challenge $[y]_1 \leftarrow_{\mathfrak{s}} \mathbb{G}_1$. The rest of the steps are identical to \mathcal{A} . The game outputs 1 when $(\tilde{d}_1(x)/d'_1(x))[1]_1 = [y]_1$. It is easy to see that the

<p>Game₁ :</p> <hr/> <pre> 1 : $x \leftarrow_{\\$} \mathbb{F}; \text{ck} \leftarrow ((x^i)_{i=1}^n)_1, [x]_2$; 2 : $[y]_1 \leftarrow_{\\$} \mathbb{G}_1$; 3 : $[a_1]_1 \leftarrow \mathbf{a}_1(x)[1]_1; [a_2]_1 \leftarrow \mathbf{a}_2(x)[1]_1; [d_1]_1 \leftarrow \mathbf{d}'_1(x)[y]_1; [d_2]_1 \leftarrow \mathbf{d}'_2(x)[y]_1$; 4 : for $i \in [1, \kappa_1]$ do 5 : $\mathfrak{z}_i \leftarrow i; A_{1i} \leftarrow \mathbf{a}_1(\mathfrak{z}_i); A_{2i} \leftarrow \mathbf{a}_2(\mathfrak{z}_i)$; 6 : for $j \in [1, \kappa_2]$ do 7 : $\beta_{ij} \leftarrow j$; 8 : $[h_{ij}]_1 \leftarrow \frac{[(a_x - A_{1i}) + \beta_{ij}(a_2(x) - A_{2i})]_1}{x - \mathfrak{z}_i} + \frac{\beta_{ij}^2(\mathbf{g}_1(A_{1i})[d_1]_1 + \mathbf{g}_2(A_{2i})[d_2]_1)}{x - \mathfrak{z}_i}$; 9 : 10 : $\text{tr}_{ij} \leftarrow ([a_1, d_1, a_2, d_2]_1, \mathfrak{z}_i, A_{1i}, A_{2i}, \beta_{ij}, [h_{ij}]_1)$; 11 : $(\tilde{\mathbf{a}}_1(X), \tilde{\mathbf{d}}_1(X), \tilde{\mathbf{a}}_2(X), \tilde{\mathbf{d}}_2(X)) \leftarrow \text{Ext}_{\text{ss}}(\mathbf{p}, \text{ck}, \text{tr})$; 12 : return $(\tilde{\mathbf{d}}_1(x)/\tilde{\mathbf{d}}_1(x))[1]_1 = [y]_1$; </pre> <p>Game₂ :</p> <hr/> <pre> 1 : $x \leftarrow_{\\$} \mathbb{F}; \text{ck} \leftarrow ((x^i)_{i=1}^n)_1, [x]_2$; 2 : $y \leftarrow_{\\$} \mathbb{F}$; 3 : $[a_1]_1 \leftarrow [a_1(x)]_1; [a_2]_1 \leftarrow [a_2(x)]_1$; 4 : $[d_1]_1 \leftarrow y[\mathbf{d}'_1(x)]_1; [d_2]_1 \leftarrow y[\mathbf{d}'_2(x)]_1$; 5 : for $i \in [1, \kappa_1]$ do 6 : $\mathfrak{z}_i \leftarrow i$; 7 : $A_{1i} \leftarrow \mathbf{a}_1(\mathfrak{z}_i); A_{2i} \leftarrow \mathbf{a}_2(\mathfrak{z}_i)$; 8 : for $j \in [1, \kappa_2]$ do 9 : $\beta_{ij} \leftarrow j$; 10 : $h_{ij}(X) \leftarrow \frac{a_1(X) - A_{1i} + \beta_{ij}(a_2(X) - A_{2i})}{X - \mathfrak{z}_i} + \frac{\beta_{ij}^2(\mathbf{g}_1(A_{1i})y\mathbf{d}'_1(X) - \mathbf{g}_2(A_{2i})y\mathbf{d}'_2(X))}{X - \mathfrak{z}_i}$; 11 : $[h_{ij}]_1 = [h_{ij}(x)]_1$; 12 : $\text{tr}_{ij} \leftarrow ([a_1, d_1, a_2, d_2]_1, \mathfrak{z}_i, A_{1i}, A_{2i}, \beta_{ij}, [h_{ij}]_1)$; 13 : $(\tilde{\mathbf{a}}_1(X), \tilde{\mathbf{d}}_1(X), \tilde{\mathbf{a}}_2(X), \tilde{\mathbf{d}}_2(X)) \leftarrow \text{Ext}_{\text{ss}}(\mathbf{p}, \text{ck}, \text{tr})$; 14 : return $(\tilde{\mathbf{d}}_1(x)/\tilde{\mathbf{d}}_1(x))[1]_1 = [y]_1$; </pre>
--

Fig. 8. Games in the proof of Theorem 3. Here, $\mathbf{a}_1, \mathbf{a}_2, \mathbf{d}'_1, \mathbf{d}'_2 \in \mathbb{F}_{\leq n}[X]$, and $\mathbf{g}_1, \mathbf{g}_2 \in \mathbb{F}[X]$ are arbitrary fixed polynomials that satisfy $\mathbf{g}_1(\mathbf{a}_1(X))\mathbf{d}'_1(X) + \mathbf{g}_2(\mathbf{a}_2(X))\mathbf{d}'_2(X) = 0$

probability that this game outputs 1 is $\varepsilon_{\mathcal{A}}$, the probability that \mathcal{A} breaks the DL assumption in the DL game.

Game₂: The second game in Fig. 8 is a very slight modification of **Game₁**. The game samples $y \leftarrow_{\$} \mathbb{F}$ instead of directly sampling the group element $[y]_1$. This game uses the knowledge of y (instead of knowledge of x) to simulate identically distributed transcripts. Namely, on step 4 it computes now $[d_1]_1 \leftarrow y[\mathbf{d}'_1(x)]_1$, $[d_2]_1 \leftarrow y[\mathbf{d}'_2(x)]_1$, and it uses the fact that $h_{ij}(X)$ on step 10 is a polynomial to compute $[h_{ij}]_1 = [h_{ij}(x)]_1$. The latter is true since, $\mathbf{g}_1(\mathbf{a}_1(\mathfrak{z}_i))\mathbf{d}_1(\mathfrak{z}_i) + \mathbf{g}_2(\mathbf{a}_2(\mathfrak{z}_i))\mathbf{d}_2(\mathfrak{z}_i) = 0$ for any $\mathfrak{z}_i = i$, which means that $\mathbf{g}_1(\mathbf{a}_1(\mathfrak{z}_i))\mathbf{d}_1(X) + \mathbf{g}_2(\mathbf{a}_2(\mathfrak{z}_i))\mathbf{d}_2(X)$ is divisible by $X - \mathfrak{z}_i$. Thus, we can efficiently compute $[h_{ij}(x)]_1$ by only using ck (and not x). In fact, **Game₂** uses the discrete logarithm of $[x]_1$ only on the very first step when it computes ck . Since inputs to Ext_{ss} are identically distributed to **Game₁**, the probability that **Game₂** outputs 1 is also $\varepsilon_{\mathcal{A}}$.

Now we can take the steps 2 to 12 in Game_2 and write a computational special-soundness adversary \mathcal{B} based on it as shown in Fig. 7. This is possible because these steps do not require knowledge of x . By inlining \mathcal{B} to the special-soundness game, it is clear that $\Pr[\text{Game}_2 \neq 1]$ is bounded by the advantage of \mathcal{B} in the special-soundness game. Namely, $\text{Game}_2 \neq 1$ happens only if $\tilde{d}_1(x) \neq d_1$. If this is the case, then $(\tilde{a}_1(X), \tilde{d}_1(X), \tilde{a}_2(X), \tilde{d}_2(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{g}}^{\text{Lin}}$, which means that \mathcal{B} wins the special-soundness game. This happens at most with probability ε as we assumed Ext_{ss} to fail at most with probability ε in the beginning of the proof. Therefore, the probability that $\tilde{d}_1(x)/d_d(x) = y$ in Game_2 must be greater than $1 - \varepsilon$.

We conclude that $\varepsilon_{\mathcal{A}} \geq 1 - \varepsilon$, which implies that if ε is negligible (and thus special soundness holds) then we break the DL assumption. \square

B Proof of SanLin

Before proving Theorem 4, we state the following lemma.

Lemma 3. *Assume that $\mathcal{T} = (\text{tr}_{ijk})$ is a κ -tree of SanLin's accepting transcripts, where tr_{ijk} are as in step 2 in Fig. 9. The PPT algorithm $\text{TE}_{\text{lin}}(\text{ck}, \mathcal{T})$ in Fig. 9 computes accepting KZG transcripts $(\text{tr}_{a_s})_{s=1}^{n_a}$ and $(\text{tr}_{d_t})_{t=1}^{n_b}$, s.t. $\text{tr}_{a_s, i} = ([a_s]_1, \mathfrak{z}_i, \dots)$ and $\text{tr}_{d_t, i} = ([d_t]_1, \mathfrak{z}_i, \dots)$ for $i \in [1, n+1]$, and \mathfrak{z}_i are mutually distinct.*

Proof. Let \mathcal{T} be the given accepting tree of transcripts. For $i \in [1, \kappa_1]$ and $j \in [1, n_b]$, let

$$\mathbf{C}_i = \begin{pmatrix} 1 & \gamma_{i1} & \gamma_{i1}^2 & \cdots & \gamma_{i1}^{n_b-1} \\ 1 & \gamma_{i2} & \gamma_{i2}^2 & \cdots & \gamma_{i2}^{n_b-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \gamma_{in_b} & \gamma_{in_b}^2 & \cdots & \gamma_{in_b}^{n_b-1} \end{pmatrix} \quad \text{and} \quad \mathbf{B}_{ij} = \begin{pmatrix} 1 & \beta_{ij1} & \cdots & \beta_{ij1}^{n_a+1} \\ 1 & \beta_{ij2} & \cdots & \beta_{ij2}^{n_a+1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \beta_{ij, n_a+2} & \cdots & \beta_{ij, n_a+2}^{n_a+1} \end{pmatrix} \quad (17)$$

be Vandermonde matrices. Since \mathcal{T} is a tree of transcripts, \mathfrak{z}_i are all distinct, and for each \mathfrak{z}_i , all γ_{ij} are distinct, and for each γ_{ij} , all β_{ijk} are distinct. Thus, for each $i \in [1, \kappa_1]$ and $j \in [1, n_b]$, \mathbf{C}_i and \mathbf{B}_{ij} are non-singular.

By the construction of SanLin, tr_{ijk} is an accepting SanLin transcript iff

$$\mathbf{k} \cdot \text{tr}_{ijk} = ([\varphi_{ijk}]_1, \mathfrak{z}_i, \Phi_{ijk}, [h_{ijk}]_1)$$

is an accepting KZG transcript, where

$$\varphi_{ijk} := \sum_{s=1}^{n_a} \beta_{ijk}^{s-1} a_s + \beta_{ijk}^{n_a} \sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) d_t + \beta_{ijk}^{n_a+1} \sum_{t=1}^{n_b} \gamma_{ij}^{t-1} d_t$$

and

$$\Phi_{ijk} := \sum_{s=1}^{n_a} \beta_{ijk}^{s-1} \bar{a}_{is} + \beta_{ijk}^{n_a+1} \bar{d}_{ij}.$$

Thus,

$$\begin{aligned} (a_1, \dots, a_{n_a}, \sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) d_t, \sum_{t=1}^{n_b} \gamma_{ij}^{t-1} d_t)^\top &= \mathbf{B}_{ij}^{-1} \cdot (\varphi_{ij1}, \dots, \varphi_{ij, n_a+2})^\top, \\ (\bar{a}_{i1}, \dots, \bar{a}_{in_a}, 0, \bar{d}_{ij})^\top &= \mathbf{B}_{ij}^{-1} \cdot (\Phi_{ij1}, \dots, \Phi_{ij, n_a+2})^\top. \end{aligned}$$

$\text{TE}_{\text{lin}}(\text{ck}, \mathcal{J})$
1: Parse $\mathcal{J} = (\text{tr}_{ijk})_{i \in [1, \kappa_1], j \in [1, n_b], k \in [1, n_a + 2]}$;
2: Parse $\text{tr}_{ijk} = ([(\mathbf{a}_s)_{s=1}^{n_a}, (\mathbf{d}_t)_{t=1}^{n_b}]_1, \mathfrak{z}_i, (\bar{\mathbf{a}}_{it})_{t=1}^{n_b}, \gamma_{ij}, \bar{\mathbf{d}}_{ij}, \beta_{ijk}, [h_{ijk}]_1)$;
3: for $i \in [1, \kappa_1]$ do for $j \in [1, n_b]$ do
4: $[h_{ij1}^*, \dots, h_{ij, n_a+2}^*]_1^\top \leftarrow \mathbf{B}_{ij}^{-1} [h_{ij1}, \dots, h_{ij, n_a+2}]_1^\top$;
5: for $s \in [1, n_a]$ do $\text{k.tr}_{ijs}^* \leftarrow ([a_s]_1, \mathfrak{z}_i, \bar{\mathbf{a}}_{is}, [h_{ijs}^*]_1)$;
6: $\text{k.tr}_{ij, n_a+1}^* \leftarrow ([\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) \mathbf{d}_t]_1, \mathfrak{z}_i, 0, [h_{ij, n_a+1}^*]_1)$;
7: $\text{k.tr}_{ij, n_a+2}^* \leftarrow ([\sum_{t=1}^{n_b} \gamma_{ij}^{t-1} \mathbf{d}_t]_1, \mathfrak{z}_i, \bar{\mathbf{d}}_{ij}, [h_{ij, n_a+2}^*]_1)$;
8: for $s \in [1, n_a]$ do
9: for $i \in [1, \kappa_1]$ do $\text{k.tr}_{a_s, i} \leftarrow \text{k.tr}_{i1s}^*$;
10: $\text{tr}_{a_s} \leftarrow (\text{k.tr}_{a_s, 1}, \dots, \text{k.tr}_{a_s, n+1})$;
11: for $i \in [1, \kappa_1]$ do
12: $\begin{pmatrix} \bar{\mathbf{d}}'_{i1} \\ \vdots \\ \bar{\mathbf{d}}'_{in_b} \end{pmatrix} \leftarrow \mathbf{C}_i^{-1} \begin{pmatrix} \bar{\mathbf{d}}_{i1} \\ \vdots \\ \bar{\mathbf{d}}_{in_b} \end{pmatrix}$; $\begin{bmatrix} h'_{i1} \\ \vdots \\ h'_{in_b} \end{bmatrix}_1 \leftarrow \mathbf{C}_i^{-1} \begin{bmatrix} h_{i1, n_a+2}^* \\ \vdots \\ h_{in_b, n_a+2}^* \end{bmatrix}_1$;
13: for $t \in [1, n_b]$ do
14: $\text{k.tr}_{d_t, i} \leftarrow ([d_t]_1, \mathfrak{z}_i, \bar{\mathbf{d}}'_{it}, [h'_{it}]_1)$;
15: $\text{tr}_{d_t} \leftarrow (\text{k.tr}_{d_t, 1}, \dots, \text{k.tr}_{d_t, n+1})$;
16: return $(\text{tr}_{a_s})_{s=1}^{n_a}, (\text{tr}_{d_t})_{t=1}^{n_b}$;

Fig. 9. The TE_{lin} subroutine.

Since KZG is triply homomorphic, for all i, j, s , k.tr_{ijs}^* (see Steps 5, 6, 7 in Fig. 9) are accepting KZG transcripts. Thus, one can define $\text{tr}_{a_s, i} = \text{k.tr}_{ijs}^*$ for $j = 1$ (one can choose any value of j).

Let $h_{i1, n_a+2}^*, \dots, h_{in_b, n_a+2}^*$ be as in Step 4 of Fig. 9. We define $(\bar{\mathbf{d}}'_{i1}, \dots, \bar{\mathbf{d}}'_{in_b})^\top \leftarrow \mathbf{C}_i^{-1} \cdot (\bar{\mathbf{d}}_{i1}, \dots, \bar{\mathbf{d}}_{in_b})^\top$ and $[h'_{i1}, \dots, h'_{in_b}]_1^\top \leftarrow \mathbf{C}_i^{-1} \cdot [h_{i1, n_a+2}^*, \dots, h_{in_b, n_a+2}^*]_1^\top$. Moreover, the vector of first elements $[\sum_{t=1}^{n_b} \gamma_{ij}^{t-1} \mathbf{d}_t]_1$ of $\text{k.tr}_{ij, n_a+2}^*$ is equal to $\mathbf{C}_i \cdot [d_1, \dots, d_{n_b}]_1^\top$. Since KZG is triple homomorphic and $\text{k.tr}_{ij, n_a+2}^*$ are accepting, $\text{k.tr}_{d_t, i} \leftarrow ([d_t]_1, \mathfrak{z}_i, \bar{\mathbf{d}}'_{it}, [h'_{it}]_1)$ is an accepting KZG transcript. \square

Next, we prove Theorem 4.

Proof (Theorem 4). Recall that a valid witness contains $n_a + n_b$ polynomials $(\mathbf{a}_s^*(X))_{s=1}^{n_a}$ and $(\mathbf{d}_t^*(X))_{t=1}^{n_b}$ of degree at most n , such that $\sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}^*(X)) \mathbf{d}_t^*(X) = 0$ and $\mathbf{a}_s^*(x) = a_s, \mathbf{d}_t^*(x) = d_t$ for $s \in [1, n_a], t \in [1, n_b]$.

Recall $\kappa = (\kappa_1, n_b, n_a + 2)$. Let $\text{Ext}_{\text{ss}}^{\text{kzg}}$ be an $(n + 1)$ -special-soundness extractor $\text{Ext}_{\text{ss}}^{\text{kzg}}$ of KZG. We depict the κ -special-soundness extractor $\text{Ext}_{\text{ss}}^{\text{sanlin}}$ for SanLin in Fig. 10. It has blackbox access to $\text{Ext}_{\text{ss}}^{\text{kzg}}$. On input a κ -tree of SanLin accepting transcripts, $\text{Ext}_{\text{ss}}^{\text{sanlin}}$ calls TE_{lin} from Fig. 9 that returns accepting KZG transcripts $\text{tr}_{a_s}, \text{tr}_{d_t}$ for $s \in [1, n_a], t \in [1, n_b]$. As stated in Lemma 3, each $\text{tr}_{a_s, i}$ (resp., $\text{tr}_{d_t, i}$) contains a transcript for the commitment $[a_s]_1$ (resp., $[d_t]_1$) with a distinct evaluation point \mathfrak{z}_i . We feed them separately into KZG's $(n + 1)$ -special-soundness extractor $\text{Ext}_{\text{ss}}^{\text{kzg}}$ to extract all $\mathbf{a}_s^*(X)$ and $\mathbf{d}_t^*(X)$.

$\text{Ext}_{\text{SS}}^{\text{sanlin}}(\text{ck}, \mathcal{T})$	$\mathcal{B}_{\text{SS}}^{\text{kzG}}(\text{ck})$
$((\mathbf{tr}_{a_s})_{s=1}^{n_a}, (\mathbf{tr}_{d_t})_{t=1}^{n_b}) \leftarrow \text{TE}_{\text{lin}}(\text{ck}, \mathcal{T});$ for $s \in [1, n_a]$ do $\mathbf{a}_s^*(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzG}}(\text{ck}, \mathbf{tr}_{a_s});$ for $t \in [1, n_b]$ do $\mathbf{d}_t^*(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzG}}(\text{ck}, \mathbf{tr}_{d_t});$ return $((\mathbf{a}_s^*(X))_{s=1}^{n_a}, (\mathbf{d}_t^*(X))_{t=1}^{n_b});$	$\mathcal{T} \leftarrow \mathcal{A}(\text{ck});$ $((\mathbf{tr}_{a_s})_{s=1}^{n_a}, (\mathbf{tr}_{d_t})_{t=1}^{n_b}) \leftarrow \text{TE}_{\text{lin}}(\text{ck}, \mathcal{T});$ for $s \in [1, n_a]$ do $\mathbf{a}_s^*(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzG}}(\text{ck}, \mathbf{tr}_{a_s});$ if $([a_s]_1, \mathbf{a}_s^*(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{tr}_{a_s}}$ then return $\mathbf{tr}_{a_s};$ for $t \in [1, n_b]$ do $\mathbf{d}_t^*(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzG}}(\text{ck}, \mathbf{tr}_{d_t});$ if $([d_t]_1, \mathbf{d}_t^*(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{tr}_{d_t}}$ then return $\mathbf{tr}_{d_t};$ return $\perp;$
$\mathcal{C}_{\text{evb}}(\text{ck})$	
$\mathcal{T} \leftarrow \mathcal{A}(\text{ck}); ((\mathbf{tr}_{a_s})_{s=1}^{n_a}, (\mathbf{tr}_{d_t})_{t=1}^{n_b}) \leftarrow \text{TE}_{\text{lin}}(\text{ck}, \mathcal{T});$ for $s \in [1, n_a]$ do $\mathbf{a}_s^*(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzG}}(\text{ck}, \mathbf{tr}_{a_s});$ for $t \in [1, n_b]$ do $\mathbf{d}_t^*(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzG}}(\text{ck}, \mathbf{tr}_{d_t});$ Parse $((\mathbf{tr}_{a_s})_{s=1}^{n_a}, (\mathbf{tr}_{d_t})_{t=1}^{n_b})$ as in Fig. 9; for $i \in [n+2, \kappa_1]$ do for $s \in [1, n_a]$ do if $\bar{a}_{is} \neq \mathbf{a}_s^*(\mathfrak{z}_i)$ then return $([a_s]_1, \mathfrak{z}_i, \bar{a}_{is}, [h'_{i1}]_1, \mathbf{a}_s^*(\mathfrak{z}_i), \left[\frac{\mathbf{a}_s^*(x) - \mathbf{a}_s^*(\mathfrak{z}_i)}{x - \mathfrak{z}_i} \right]_1);$ for $t \in [1, n_b]$ do if $\bar{d}'_{it} \neq \mathbf{d}_t^*(\mathfrak{z}_i)$ then return $([d_t]_1, \mathfrak{z}_i, \bar{d}'_{it}, [h'_{it}]_1, \mathbf{d}_t^*(\mathfrak{z}_i), \left[\frac{\mathbf{d}_t^*(x) - \mathbf{d}_t^*(\mathfrak{z}_i)}{x - \mathfrak{z}_i} \right]_1);$ for $i \in [1, \kappa_1]$ do if $\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) \bar{d}'_{it} \neq 0$ then return $([\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) d_t]_1, \mathfrak{z}_i, \sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) \bar{d}'_{it}, [\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) h'_{it}]_1, 0, [h'_{i1, n_a+1}]_1);$ return $\perp;$	

Fig. 10. The extractor $\text{Ext}_{\text{SS}}^{\text{sanlin}}$, the KZG special soundness-adversary $\mathcal{B}_{\text{SS}}^{\text{kzG}}$, and the KZG evaluation-binding adversary \mathcal{C}_{evb} .

Next, we show that $\varepsilon = \text{Adv}_{\text{Pgen, SanLin, Ext}_{\text{SS}}^{\text{sanlin}}, \kappa, \mathcal{A}}^{\text{SS}}(\lambda)$ (see Definition 2) is negligible for any PPT adversary \mathcal{A} . That is, if \mathcal{A} outputs an accepting tree of transcripts, then $\text{Ext}_{\text{SS}}^{\text{sanlin}}$ fails to extract $(\mathbf{a}_s^*(X))_{s=1}^{n_a}, (\mathbf{d}_t^*(X))_{t=1}^{n_b}$, such that

$$((([a_s]_1, \bar{a}_s)_{s=1}^{n_a}, ([d_t]_1, \bar{d}_t)_{t=1}^{n_b}), ((\mathbf{a}_s^*(X))_{s=1}^{n_a}, (\mathbf{d}_t^*(X))_{t=1}^{n_b})) \in \mathcal{R}_{\text{ck}, \mathbf{g}, \mathbf{tr}}^{\text{Lin}},$$

(for naturally defined \mathbf{tr}) with probability $\varepsilon = \text{negl}(\lambda)$.

We consider the following two failure events for $\text{Ext}_{\text{SS}}^{\text{sanlin}}$:

1. The event bad_{ext} happens when $([a_s]_1, \mathbf{a}_s^*(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{tr}_{a_s}}$ or $([d_t]_1, \mathbf{d}_t^*(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{tr}_{d_t}}$ for some $s \in [1, n_a], t \in [1, n_b]$.
2. The event bad_{evb} happens when bad_{ext} did not happen, but one of the following conditions hold:
 - (a) For any $s \in [1, n_a], t \in [1, n_b]$, either $\mathbf{a}_s^*(\mathfrak{z}_i) \neq \bar{a}_{is}$ or $\mathbf{d}_t^*(\mathfrak{z}_i) \neq \bar{d}'_{it}$ for some $i \in [n+2, \kappa_1]$.

(b) $\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) \bar{d}'_{it} \neq 0$ for any $i \in [1, \kappa_1]$.

Consider the scenario where neither $\mathbf{bad}_{\text{ext}}$ nor $\mathbf{bad}_{\text{evb}}$ occurs. Then, we have extracted polynomials $\mathbf{a}_s^*(X)$ and $\mathbf{d}_t^*(X)$ for $s \in [1, n_a]$, $t \in [1, n_b]$ of degree at most n which are consistent with the commitments. Denote $\mathbf{g}^*(X) := \sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}^*(X)) \mathbf{d}_t^*(X)$. For any $i \in [1, \kappa_1]$,

$$\mathbf{g}^*(\mathfrak{z}_i) = \sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}^*(\mathfrak{z}_i)) \mathbf{d}_t^*(\mathfrak{z}_i) = \sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) \bar{d}'_{it} = 0 .$$

Since $\mathbf{g}^*(X)$ is at most of degree $n_{\mathbf{g}} < \kappa_1$, it follows that $\mathbf{g}^*(X) = 0$ and consequently $\sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}^*(X)) \mathbf{d}_t^*(X) = 0$. Hence, $\text{Ext}_{\text{ss}}^{\text{sanlin}}$ extracts a valid witness.

Next, we bound the probabilities $\Pr[\mathbf{bad}_{\text{ext}}]$ and $\Pr[\mathbf{bad}_{\text{evb}}]$. To bound $\Pr[\mathbf{bad}_{\text{ext}}]$, we construct a PPT adversary $\mathcal{B}_{\text{ss}}^{\text{kzg}}$ (see Fig. 10) that breaks the special-soundness of KZG when the event $\mathbf{bad}_{\text{ext}}$ happens. $\mathcal{B}_{\text{ss}}^{\text{kzg}}$ runs $\mathcal{A}(\text{ck})$ to recover the tree \mathcal{T} . Then, it obtains transcript vectors $((\mathbf{tr}_{a_s})_{s=1}^{n_a}, (\mathbf{tr}_{d_t})_{t=1}^{n_b}) \leftarrow \text{TE}_{\text{lin}}(\text{ck}, \mathcal{T})$. For each transcript vector \mathbf{tr}_f , where $f \in \{a_s\}_{s=1}^{n_a} \cup \{d_t\}_{t=1}^{n_b}$, $\mathcal{B}_{\text{ss}}^{\text{kzg}}$ runs the deterministic extractor $\text{Ext}_{\text{ss}}^{\text{kzg}}(\text{ck}, \mathbf{tr}_f)$ to extract the polynomial $\mathbf{f}^*(X) \in \{\mathbf{a}_s^*\}_{s=1}^{n_a} \cup \{\mathbf{d}_t^*\}_{t=1}^{n_b}$. $\mathcal{B}_{\text{ss}}^{\text{kzg}}$ returns one transcript vector \mathbf{tr}_f , which satisfies $([f]_1, \mathbf{f}^*(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{tr}_f}$. When the event $\mathbf{bad}_{\text{ext}}$ happens, $([f]_1, \mathbf{f}^*(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{tr}_f}$ for some f , and hence $\mathcal{B}_{\text{ss}}^{\text{kzg}}$ breaks the special soundness of KZG. Thus,

$$\Pr[\mathbf{bad}_{\text{ext}}] = \text{Adv}_{\text{Pgen, KZG, Ext}_{\text{ss}}^{\text{kzg}}, n+1, \mathcal{B}_{\text{ss}}^{\text{kzg}}}^{\text{ss}}(\lambda) .$$

Second, we bind the probability $\Pr[\mathbf{bad}_{\text{evb}}]$. We construct a PPT KZG's evaluation-binding adversary \mathcal{C}_{evb} , depicted in Fig. 10. \mathcal{C}_{evb} runs \mathcal{A} , TE_{lin} and $\text{Ext}_{\text{ss}}^{\text{sanlin}}$ to obtain \mathcal{T} and corresponding $((\mathbf{a}_s^*(X))_{s=1}^{n_a}, (\mathbf{d}_t^*(X))_{t=1}^{n_b})$. If the event $\mathbf{bad}_{\text{evb}}$ happens, then for all s and t , $([a_s]_1, \mathbf{a}_s^*(X)) \in \mathcal{R}_{\text{ck}, \mathbf{tr}_{a_s}}$ and $([d_t]_1, \mathbf{d}_t^*(X)) \in \mathcal{R}_{\text{ck}, \mathbf{tr}_{d_t}}$. Thus, $\mathbf{a}_s^*(x) = a_s$, and $\mathbf{d}_t^*(x) = d_t$. Note that

$$h_{a_s}(X) := (\mathbf{a}_s^*(X) - \mathbf{a}_s^*(\mathfrak{z}_i)) / (X - \mathfrak{z}_i)$$

is always a polynomial. Hence, \mathcal{C}_{evb} can compute $[h_{a_s}(x)]_1$ that satisfies

$$[h_{a_s}(x)]_1 \bullet [x - \mathfrak{z}_i]_2 = [\mathbf{a}_s^*(x) - \mathbf{a}_s^*(\mathfrak{z}_i)]_T = [a_s - \mathbf{a}_s^*(\mathfrak{z}_i)]_T .$$

Thus,

$$([a_s]_1, \mathfrak{z}_i, \mathbf{a}_s^*(\mathfrak{z}_i), [h_{a_s}(x)]_1)$$

is an accepting transcript. However, $\mathbf{k.tr}_{i,j_s}^* = ([a_s]_1, \mathfrak{z}_i, \bar{a}_{i,s}, [h_{i,j_s}^*]_1)$ is also an accepting transcript. If $\mathbf{a}_s^*(\mathfrak{z}_i) \neq \bar{a}_{i,s}$, \mathcal{C}_{evb} has broken KZG's evaluation-binding by finding a collision. Analogously, one can show that $\mathbf{d}_t^*(\mathfrak{z}_i) \neq \bar{d}'_{it}$ implies that \mathcal{C}_{evb} broke KZG's evaluation-binding.

We also need that $\mathbf{g}^*(\mathfrak{z}_i) = \sum_{t=1}^{n_b} \mathbf{g}_t(\mathbf{a}^*(\mathfrak{z}_i)) \mathbf{d}_t^*(\mathfrak{z}_i) = 0$. From Lemma 3,

$$\mathbf{k.tr}_{d_t, i} = ([d_t]_1, \mathfrak{z}_i, \bar{d}'_{it}, [h'_{it}]_1)$$

and

$$\left(\left[\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) d_t \right]_1, \mathfrak{z}_i, 0, [h_{i,j,n_a+1}^*]_1 \right)$$

are accepting transcripts for all $i \in [1, \kappa_1]$. Since KZG is triply homomorphic, from all $\mathbf{k.tr}_{d_t, i}$ being accepting, we get

$$([\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) d_t]_1, \mathfrak{z}_i, \sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) \bar{d}'_{it}, [\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) h'_{it}]_1)$$

is an accepting transcript. If $\sum_{t=1}^{n_b} \mathbf{g}_t(\bar{\mathbf{a}}_i) \bar{d}'_{it} \neq 0$, \mathcal{C}_{evb} has found a collision and thus broken KZG's evaluation-binding. Therefore,

$$\Pr[\text{bad}_{\text{evb}}] = \text{Adv}_{\text{Pgen, KZG, } n, \mathcal{C}_{\text{evb}}}^{\text{evb}}(\lambda)$$

and we can conclude that

$$\text{Adv}_{\text{Pgen, SanLin, Ext}_{\text{ss}}^{\text{sanlin}}, \kappa, \mathcal{A}}^{\text{ss}}(\lambda) = \text{Adv}_{\text{Pgen, KZG, Ext}_{\text{ss}}^{\text{kzg}}, n+1, \mathcal{D}_{\text{ss}}^{\text{kzg}}}^{\text{ss}}(\lambda) + \text{Adv}_{\text{Pgen, KZG, } n, \mathcal{C}_{\text{evb}}}^{\text{evb}}(\lambda) .$$

□

C Security of SplitRSDH in the AGMOS

Next, we prove that par-SplitRSDH is secure in the AGMOS (AGM with oblivious sampling, [LPS23]). We give the proof in the new oracle framework of Section 3. Because of Fact 1 in Section 3, the reduction implies that SplitRSDH is generically hard.

Theorem 7. *As in Definition 4, let $\text{par} = (m, n_\psi, n_1, n_S, (\psi_k(X))_{k=1}^m)$. Assume $m, n_\psi, n_1, n_S \in \text{poly}(\lambda)$ are positive integers so that $n_S > n_1 + n_\psi + 1$. Assume $\psi_k(X) \in \mathbb{F}_{\leq n_\psi}[X]$. Fix \mathcal{EF} and \mathcal{DF} . If the $(n_1, 1)$ -PDL and $(\mathcal{EF}, \mathcal{DF})$ -TOFR assumptions hold, then par-SplitRSDH holds in the AGMOS. More precisely, for every PPT \mathcal{A} , there exist PPT \mathcal{B}_{PDL} and PPT $\mathcal{C}_{\text{TOFR}}$, such that*

$$\text{Adv}_{\text{Pgen, par, } \mathcal{A}}^{\text{splitsrdh}}(\lambda) \leq \text{Adv}_{n_1, 1, \text{Pgen, } \mathcal{B}_{\text{PDL}}}^{\text{pdl}}(\lambda) + \text{Adv}_{\text{Pgen, } \mathcal{C}_{\text{TOFR}}}^{\text{tofr}}(\lambda) .$$

As usually, AGM and AGMOS proofs have a large number of boilerplate steps, needed to construct formal reductions. For the ease of parsing, we have marked one paragraph of the following proof as “meat of the analysis”: this is essentially the only part that deeply depends on the concrete assumption. (Alternatively, it is a more formal rewording of the intuition we gave just after Definition 4.) Readers familiar with AGM(OS) proofs may skip the rest of the proof.

Proof. Let \mathcal{A} be an $(\mathcal{EF}, \mathcal{DF})$ -AGMOS adversary against the SplitRSDH assumption that succeeds with some non-negligible probability ε . In Fig. 11, we depict a PDL adversary \mathcal{B}_{PDL} and a TOFR adversary $\mathcal{C}_{\text{TOFR}}$, such that $\Pr[\mathcal{A} \text{ is successful}] \leq \Pr[\mathcal{B}_{\text{PDL}} \text{ is successful} \mid E] + \Pr[\mathcal{C}_{\text{TOFR}} \text{ is successful} \mid \neg E]$ for some event E .

\mathcal{B}_{PDL} is given an input ck while $\mathcal{C}_{\text{TOFR}}$ samples ck itself. (In the oracle framework of [JM24], it means \mathcal{B}_{PDL} can only use oracles to access ck while $\mathcal{C}_{\text{TOFR}}$ knows x and can use it in computation.) After that, both reductions call \mathcal{A} with ck , obtaining the output of a SplitRSDH adversary. According to Definition 3, there exists an extractor $\text{Ext}_{\mathcal{A}}$ who (with a high probability) outputs the

$\mathcal{B}_{\text{PDL}}(\mathbf{p}, \mathbf{ck} = ([1, x, \dots, x^{n_1}]_1, [1, x]_2)) : \mathcal{C}_{\text{TOFR}}(\mathbf{p}) :$
<pre> 1 : $x \leftarrow_{\\$} \mathbb{F}; \mathbf{ck} \leftarrow ([1, x, \dots, x^{n_1}]_1, [1, x]_2);$ // Only in the TOFR reduction: 2 : $(\mathcal{S} = \{\mathfrak{z}_i\}_{i=1}^{n_{\mathcal{S}}}, L(X), [(\tau_k)_{k=1}^m, \varphi]_1) \leftarrow \mathcal{A}^{\Pi_{\mathcal{O}}}(\mathbf{p}, \mathbf{ck});$ // Calling \mathcal{A} 3 : $((\{\tau'_k(X), \hat{\tau}_k\}_{k=1}^m, (\varphi'(X), \hat{\varphi}))) \leftarrow \text{Ext}_{\mathcal{A}}^{\Pi_{\mathcal{O}}}(\mathbf{p}, \mathbf{ck});$ // Calling $\text{Ext}_{\mathcal{A}}$ 4 : // Defining verification polynomial based on extractor's output 5 : for $k \in [1, m]$ do $\tau_k(X, \mathbf{Q}) := \tau'_k(X) + \hat{\tau}_k^{\top} \mathbf{Q};$ 6 : $\varphi(X, \mathbf{Q}) := \varphi'(X) + \hat{\varphi}^{\top} \mathbf{Q};$ 7 : $V_0(X) := \sum_{k=1}^m \tau'_k(X) \psi_k(X) - L(X) - \varphi'(X) \mathbf{Z}_{\mathcal{S}}(X);$ 8 : for $j \in [1, \text{il}]$ do $V_j(X) := \sum_{k=1}^m \hat{\tau}_{kj} \psi_k(X) - \hat{\varphi}_j(X) \mathbf{Z}_{\mathcal{S}}(X);$ 9 : $V(X, \mathbf{Q}) := V_0(X) + \sum_{j=1}^{\text{il}} V_j(X) \mathbf{Q}_j;$ 10 : if $V(X, \mathbf{Q}) = 0$ then return $\perp;$ 11 : // Finishing off either PDL or TOFR reduction 12 : if $\forall j \in [0, \text{il}]. [V_j(x)]_1 = [0]_1$ then // PDL Reduction 13 : Let $j_0 \in [0, \text{il}]$ be s.t. $(V_{j_0}(X) \neq 0);$ 14 : $(x_1, \dots, x_{n_1}) \leftarrow \text{Solve}(V_{j_0}(X) = 0);$ 15 : for $i \in [1, n_1]$ do if $[x]_1 = x_i [1]_1$ then return $x_i;$ 16 : else // TOFR Reduction 17 : for $j \in [1, \text{il}]$ do $v_j \leftarrow V_j(x);$ 18 : return $v;$ </pre>

Fig. 11. The PDL and TOFR reductions in the proof of Theorem 7. We have boxed the lines where the reduction handles group elements.

explanations behind \mathcal{A} 's output group elements. We use the explanations to define a verification polynomial $V(X, \mathbf{Q})$, where \mathbf{Q} is the vector of indeterminates corresponding to sampling oracle queries. This polynomial is defined so that $V(x, \mathbf{q}) = 0$ iff the SplitRSDH adversary succeeds in satisfying the last verification in Definition 4, i.e., $\sum_{k=1}^m ([\tau_k]_1 \bullet [\psi_k(x)]_2) = [1]_1 \bullet [L(x)]_2 + [\varphi]_1 \bullet [\mathbf{Z}_{\mathcal{S}}(x)]_2$. After that, as usually in AGMOS proofs [LPS23], we define an event E , such that if $E = \text{true}$, we can and will construct a successful PDL adversary and otherwise a successful TOFR adversary.

Let us now give a detailed description of the reductions and their success probability. Consider the reduction \mathcal{B}_{PDL} . (The analysis of $\mathcal{C}_{\text{TOFR}}$ differs minimally.) \mathcal{B}_{PDL} (line 2) first calls \mathcal{A} with correct input. Recall that $\psi_k(X) \in \mathbb{F}_{\leq n_{\psi}}[X]$, and $n_{\mathcal{S}} > n_1 + n_{\psi} + 1$. The AGMOS adversary \mathcal{A} has access to $\Pi_{\mathcal{O}}$ that includes the sampling oracles. Let $\text{bad}_{\mathcal{A}}$ be the event that \mathcal{A} succeeds. Thus, $\Pr[\text{bad}_{\mathcal{A}}] = \varepsilon \neq \text{negl}(\lambda)$. That is, according to Definition 4, the following holds with probability ε :

$$\begin{aligned}
& (\mathcal{S} \subset \mathbb{F}) \wedge (|\mathcal{S}| = n_{\mathcal{S}}) \wedge (L(X) \in \mathbb{F}[X]) \wedge (L(X) \in [n_1 + n_{\psi} + 1, n_{\mathcal{S}} - 1]) \wedge \\
& \left(\sum_{k=1}^m [\tau_k \psi_k(x)]_1 = [L(x)]_1 + [\varphi \mathbf{Z}_{\mathcal{S}}(x)]_1 \right)
\end{aligned}$$

We note that since SplitRSDH is not publicly-verifiable, the (model-preserving) reduction cannot test whether $\text{bad}_{\mathcal{A}}$ holds (more precisely, it cannot test that

the boxed equality holds). This is however not important: we only show that either \mathcal{B}_{PDL} or $\mathcal{C}_{\text{TOFR}}$ succeeds with high probability whenever $\text{bad}_{\mathcal{A}} = \text{true}$.

Assume now $\text{bad}_{\mathcal{A}} = \text{true}$. (This holds with probability ε .) According to Definition 3, there exists an extractor $\text{Ext}_{\mathcal{A}}^{H_{\mathcal{O}}}$ that, on input (\mathbf{p}, \mathbf{x}) and access to \mathcal{A} 's oracle queries, outputs the values on line 3, such that $\tau'_k(X), \varphi'(X) \in \mathbb{F}_{\leq n_1}[X]$ and $\hat{\tau}_{kj}, \hat{\varphi}_j \in \mathbb{F}$, and \mathbf{q} is a vector of discrete logarithms of the answers of the oblivious sampling oracle Samp_1 . (Since \mathcal{A} does not output \mathbb{G}_2 elements, we ignore the oracle Samp_2 .) According to Definition 3, if $\text{bad}_{\mathcal{A}} = \text{true}$, the reduction always succeeds, that is, for all k , $[\tau_k]_1 = [\tau'_k(x)]_1 + \hat{\tau}_k^{\top}[\mathbf{q}]_1$ and $[\varphi]_1 = [\varphi'(x)]_1 + \hat{\varphi}^{\top}[\mathbf{q}]_1$. As explained earlier, this extractor can be constructed by observing the inputs and outputs of all oracle queries. Moreover, if \mathcal{A} is model-preserving, then $\text{Ext}_{\mathcal{A}}$ is model-preserving: $\text{Ext}_{\mathcal{A}}$ never applies Encode or Decode to any group elements unless \mathcal{A} does.

Next, given the extractor's outputs, \mathcal{B}_{PDL} defines verification polynomials $V(X, \mathbf{Q})$ and $V_j(X)$, such that $V(X, \mathbf{Q}) := V_0(X) + \sum_{j=1}^{\text{il}} V_j(X)\mathbf{Q}_j$. Since the verifier accepts, $V(x, \mathbf{q}) = 0$.

“Meat of the analysis.” Let us show that if $\text{bad}_{\mathcal{A}} = \text{true}$, then \mathcal{B}_{PDL} aborts on line 10 (i.e., $V(X, \mathbf{Q}) = 0$) with probability 0. Really, assume $V(X, \mathbf{Q}) = 0$ as a polynomial. Since $\deg L(X) < n_{\mathcal{S}}$, $\varphi'(X) = 0$. Let $T(X) := \sum_{k=1}^m \tau'_k(X)\psi_k(X) \in \mathbb{F}_{\leq n_1+n_{\psi}}[X]$. Then, $T(X) = L(X)$. Since $\deg T(X) \leq n_1 + n_{\psi}$, we have $\deg L(X) = n_1 + n_{\psi}$, and thus this case is impossible. Thus, the reduction never aborts on line 10.

Now, know that $V(x, \mathbf{q}) = 0$ but $V(X, \mathbf{Q}) \neq 0$. Let E be the event that the check on line 12 succeeds, i.e., $V_j(x) = 0$ for all $j \in [0, \text{il}]$. If $E = \text{true}$, we finish the PDL reduction. Otherwise, we finish the TOFR reduction.

$E = \text{true}$. Since $E = \text{true}$, $V_j(x) = 0$ for all $j \in [0, \text{il}]$. Since $V(X, \mathbf{Q}) \neq 0$, $V_{j_0}(X) \neq 0$ for some $j_0 \in [0, \text{il}]$. Thus, $V_{j_0}(X) \neq 0$, but $V_{j_0}(x) = 0$. Clearly, $V_{j_0}(X)$ is univariate and of degree at most n_1 . Thus, it has at most n_1 roots. The PDL adversary computes x as one of the roots of $V_{j_0}(X)$.

$E = \text{false}$. Since $E = \text{false}$, $\exists j. (V_j(x) \neq 0)$ but the verifier accepts ($V(x, \mathbf{q}) = 0$). In this case, we construct an adversary $\mathcal{C}_{\text{TOFR}}$ that breaks the TOFR assumption. $\mathcal{C}_{\text{TOFR}}$ starts exactly as \mathcal{B}_{PDL} , only sampling ck itself. Knowing the trapdoor x , $\mathcal{C}_{\text{TOFR}}$ outputs the vector \mathbf{v} defined by setting $v_j = V_j(x)$ for $j \geq 0$. Thus, this adversary is successful in breaking the TOFR assumption.

To sum it up,

$$\Pr[\mathcal{A} \text{ is successful}] \leq \Pr[\mathcal{B}_{\text{PDL}} \text{ fails} \mid E] + \Pr[\mathcal{C}_{\text{TOFR}} \text{ fails} \mid \neg E] .$$

□

AGM to GGM Lifting. To be able to prove the lifting result of [JM24], we need to show that both \mathcal{B}_{PDL} and $\mathcal{C}_{\text{TOFR}}$ are model-preserving and efficient. (We already know that PDL [Lip12] and TOFR [LPS23] are generically hard.) Really, assume \mathcal{A} is an generic adversary. Thus, $\text{Ext}_{\mathcal{A}}$ is also generic. The rest of the reductions touches the adversary's group element outputs (and, in the case of \mathcal{B}_{PDL} , the input ck) in the following steps:

- Line 12 uses \mathcal{O}_p in \mathbb{G}_2 , $\mathcal{P}air$, and $\mathcal{E}q$ in \mathbb{G}_T to check if $E = \text{true}$.
- Line 15 uses $\mathcal{E}q$ in \mathbb{G}_1 to find correct root.

Thus, both reductions are model-preserving. Since PDL and TOFR are generically hard, we obtain that SplitRSDH is secure in Maurer’s GGM.

We remark that first use of oracles is actually not needed. Really, there is no need to check if $E = \text{true}$ or not (without this check, the success probability of the reduction can only increase). The second use is needed to solve PDL with a high probability.

D Postponed Material from Section 6

D.1 Plonk’s Polynomials That Define A Specific Circuit

The following polynomials, along with the integer n , uniquely define our circuit:

- $q_M(X), q_L(X), q_R(X), q_O(X), q_C(X)$ are selector polynomials that define the circuit’s arithmetization. We refer to [GWC19] for the explanation how they are defined based on an arithmetic circuit.
- $S_{ID_1}(X) = X, S_{ID_2}(X) = k_1X, S_{ID_3}(X) = k_2X$ encode an identity permutation respectively on groups $k_0 \cdot \mathbb{H}, k_1 \cdot \mathbb{H}$, and $k_2 \cdot \mathbb{H}$. Here, $k_0 := 1$ and $k_1, k_2 \in \mathbb{F}$ are chosen such that $\mathbb{H}, k_1 \cdot \mathbb{H}, k_2 \cdot \mathbb{H}$ are distinct cosets of \mathbb{H} in \mathbb{F}^* , and thus consist of $3n$ distinct elements. For example, one can take ω to be a quadratic residue in \mathbb{F} , k_1 to be any quadratic non-residue, and k_2 to be a quadratic non-residue not contained in $k_1 \cdot \mathbb{H}$.
- Let us denote $\mathbb{H}' := \mathbb{H} \cup (k_1 \cdot \mathbb{H}) \cup (k_2 \cdot \mathbb{H})$. Let $\sigma : [1, 3n] \rightarrow [1, 3n]$ be a permutation. We encode an element $i \in [1, 3n]$ in \mathbb{H}' such that if we express $i = \vartheta n + j$ for the unique $\vartheta \in [0, 2]$ and $0 \leq j < n$, then $\mathbb{H}'[i] = k_{\vartheta}\omega^j$. Finally, define $\sigma^*(i) := \mathbb{H}'[\sigma(i)]$, which is an injective map on \mathbb{H}' . We encode σ^* by the three permutation polynomials $S_{\sigma_1}(X) := \sum_{i=1}^n \sigma^*(i)L_i(X), S_{\sigma_2}(X) := \sum_{i=1}^n \sigma^*(n+i)L_i(X)$, and $S_{\sigma_3}(X) := \sum_{i=1}^n \sigma^*(2n+i)L_i(X)$.

D.2 Proof of Theorem 5

Proof. Let \mathcal{T} be a $(\kappa_\beta, \kappa_\gamma, \kappa_\alpha, \kappa_3, \kappa_\delta, \kappa_v)$ -tree of accepting transcripts and let $\hat{\mathcal{T}}_{\beta\gamma\alpha}$ be its $(1, 1, 1, \kappa_3, \kappa_\delta, \kappa_v)$ -subtree for any fixed α, β, γ . By Lemma 2, $\text{TE}_{*\text{plonk}}(\text{ck}, \hat{\mathcal{T}}_{\beta\gamma\alpha})$ extracts accepting KZG transcripts $\mathbf{k.tr}_{ki}$ and $\mathbf{k.tr}_i^\omega$ for every i and k . We will consider separately the cases of (1) Plonk and SmallPlonk and (2) SanPlonk. However, the extractors and adversaries on most figures (say, Fig. 12) correspond to both.

Case of Plonk And SmallPlonk. In Fig. 12, we depict an extractor $\text{Ext}_{\text{ss}}^{\text{sub}}$. $\text{Ext}_{\text{ss}}^{\text{sub}}$ invokes $\text{Ext}_{\text{ss}}^{\text{kzG}}(\text{ck}, \mathbf{k.tr}^\omega)$ and $\text{Ext}_{\text{ss}}^{\text{kzG}}(\text{ck}, \mathbf{k.tr}_k)$ for $k \in [2, 4]$, extracting polynomials $z(X), a(X), b(X)$, and $c(X)$ of at most degree κ_{kzG} (see Eq. (12)). After executing $\text{Ext}_{\text{ss}}^{\text{sub}}$, we use the following procedure to possibly set one of the “bad” flags:

$\text{Ext}_{\text{SS}}^{\text{sub}}(\text{ck}, \hat{\mathcal{J}}_{\beta\gamma\alpha})$ <hr/> $\begin{aligned} & ((\mathbf{k.tr}_k)_{k \in [1, \kappa_v + 2]}, \mathbf{k.tr}^\omega) \leftarrow \text{TE}_{*\text{plonk}}(\text{ck}, \hat{\mathcal{J}}_{\beta\gamma\alpha}); \\ & z(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k.tr}^\omega); a(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k.tr}_2); \\ & b(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k.tr}_3); c(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k.tr}_4); \\ & t_{\text{lo}}(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k.tr}_7); t_{\text{mid}}(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k.tr}_8); t_{\text{hi}}(X) \leftarrow \text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k.tr}_9); \\ & \text{return } (z(X), a(X), b(X), c(X), t_{\text{lo}}(X), t_{\text{mid}}(X), t_{\text{hi}}(X)); \end{aligned}$
--

Fig. 12. Plonk's/SanPlonk's/SmallPlonk's special-soundness extractor $\text{Ext}_{\text{SS}}^{\text{sub}}$.

$\mathcal{B}_{\text{SS}}^{\text{kzg}}(\text{ck})$ <hr/> $\begin{aligned} & \hat{\mathcal{J}}_{\beta\gamma\alpha} \leftarrow \mathcal{A}_{\text{SS}}^{\text{plonk}}(\text{ck}); \\ & ((\mathbf{k.tr}_k)_{k \in [1, \kappa_v + 2]}, \mathbf{k.tr}^\omega) \leftarrow \text{TE}_{*\text{plonk}}(\text{ck}, \hat{\mathcal{J}}_{\beta\gamma\alpha}); \\ & (z(X), a(X), b(X), c(X), t_{\text{lo}}(X), t_{\text{mid}}(X), t_{\text{hi}}(X)) \leftarrow \text{Ext}_{\text{SS}}^{\text{sub}}(\text{ck}, \hat{\mathcal{J}}_{\beta\gamma\alpha}); \\ & \text{if } ([z]_1, z(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}^\omega} \text{ then return } \mathbf{k.tr}^\omega; \text{ fi} \\ & \text{if } ([a]_1, a(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_2} \text{ then return } \mathbf{k.tr}_2; \text{ fi} \\ & \text{if } ([b]_1, b(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_3} \text{ then return } \mathbf{k.tr}_3; \text{ fi} \\ & \text{if } ([c]_1, c(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_4} \text{ then return } \mathbf{k.tr}_4; \text{ fi} \\ & \text{if } ([t_{\text{lo}}]_1, t_{\text{lo}}(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_7} \text{ then return } \mathbf{k.tr}_7; \text{ fi} \\ & \text{if } ([t_{\text{mid}}]_1, t_{\text{mid}}(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_8} \text{ then return } \mathbf{k.tr}_8; \text{ fi} \\ & \text{if } ([t_{\text{hi}}]_1, t_{\text{hi}}(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_9} \text{ then return } \mathbf{k.tr}_9; \text{ fi} \\ & \text{return } \perp; \end{aligned}$

Fig. 13. Plonk's/SanPlonk's KZG's special-soundness adversary $\mathcal{B}_{\text{SS}}^{\text{kzg}}$.

- (i) $\text{bad}_{\text{ext}} \leftarrow \text{false}$; $\text{bad}_{\text{evb}} \leftarrow \text{false}$; $\text{bad}_{\text{rhino}} \leftarrow \text{false}$;
- (ii) if $([z]_1, z(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_1} \vee ([a]_1, a(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_2} \vee ([b]_1, b(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_3} \vee ([c]_1, c(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k.tr}_4}$ (see Eq. (1)) then $\text{bad}_{\text{ext}} \leftarrow \text{true}$; abort;
- (iii) for $i \in [\kappa_{\text{kzg}} + 2, \kappa_3]$: if $z(\mathfrak{z}_i\omega) \neq \bar{z}_{\omega i} \vee a(\mathfrak{z}_i) \neq \bar{a}_i \vee b(\mathfrak{z}_i) \neq \bar{b}_i \vee c(\mathfrak{z}_i) \neq \bar{c}_i$ then $\text{bad}_{\text{evb}} \leftarrow \text{true}$; abort;
- (iv) for $i \in [1, \kappa_3]$: if $S_{\sigma 1}(\mathfrak{z}_i) \neq \bar{s}_{\sigma 1 i} \vee S_{\sigma 2}(\mathfrak{z}_i) \neq \bar{s}_{\sigma 2 i}$ then $\text{bad}_{\text{evb}} \leftarrow \text{true}$; abort;
- (v) if $Z_{\mathbb{H}}(X) \nmid F(X)$, where $F(X) = F_0(X) + \alpha F_1(X) + \alpha^2 F_2(X)$, then $\text{bad}_{\text{rhino}} \leftarrow \text{true}$;

Importantly, only one of the “bad” flags is set at a time. Thus, for example, $\text{bad}_{\text{evb}} = \text{true}$ implies that $\text{bad}_{\text{ext}} = \text{false}$. Let \mathcal{E} be the event $\text{Ext}_{\text{SS}}^{\text{sub}}$ succeeds. Thus, \mathcal{E} is the event that $a(X)$, $b(X)$, $c(X)$, and $z(X)$ are consistent with the commitments and all openings, and $t(X) = (F_0(X) + \alpha F_1(X) + \alpha^2 F_2(X)) / Z_{\mathbb{H}}(X)$ is a polynomial. Let $\overline{\text{bad}}$ be the event none of the bad flags was set. We analyze the success probability of $\text{Ext}_{\text{SS}}^{\text{sub}}$. For this, we make the following claims.

1. **Claim 1.** $\Pr[\mathcal{E} \mid \overline{\text{bad}}] = 1$.

Really, assume that $\overline{\text{bad}}$ holds. Since $\text{bad}_{\text{ext}} = \text{bad}_{\text{evb}} = \text{false}$, we get that $a(X)$, $b(X)$, $c(X)$, $z(X)$ are consistent with the commitments and all open-

```

 $C_{\text{evb}}(\mathbf{p}, \text{ck})$ 


---


 $\hat{\mathcal{J}}_{\beta\gamma\alpha} \leftarrow \mathcal{A}_{\text{ss}}^{\text{plonk}}(\text{ck}); ((\mathbf{k}, \text{tr}_k)_{k \in [1, \kappa_v + 2]}, \mathbf{k}, \text{tr}^\omega) \leftarrow \text{TE}_{*\text{plonk}}(\text{ck}, \hat{\mathcal{J}}_{\beta\gamma\alpha});$ 
 $(z(X), \mathbf{a}(X), \mathbf{b}(X), \mathbf{c}(X), \mathbf{t}_{\text{lo}}(X), \mathbf{t}_{\text{mid}}(X), \mathbf{t}_{\text{hi}}(X)) \leftarrow \text{Ext}_{\text{ss}}^{\text{sub}}(\text{ck}, \hat{\mathcal{J}}_{\beta\gamma\alpha});$ 
for  $i \in [\kappa_{\text{kzg}} + 2, \kappa_3]$  do
  if  $z(\mathfrak{z}_i\omega) \neq \bar{z}_{\omega_i}$  then
    return  $([\bar{z}]_1, \mathfrak{z}_i\omega, \bar{z}_{\omega_i}, [W'_{\omega_i \mathfrak{j}1}]_1, z(\mathfrak{z}_i\omega), [(z(x) - z(\mathfrak{z}_i\omega))/(x - \mathfrak{z}_i\omega)]_1);$  fi
  if  $\mathbf{a}(\mathfrak{z}_i) \neq \bar{a}_i$  then return  $([a]_1, \mathfrak{z}_i, \bar{a}_i, [W'_{\omega_i \mathfrak{j}2}]_1, \mathbf{a}(\mathfrak{z}_i), [(a(x) - a(\mathfrak{z}_i))/(x - \mathfrak{z}_i)]_1);$  fi
  if  $\mathbf{b}(\mathfrak{z}_i) \neq \bar{b}_i$  then return  $([b]_1, \mathfrak{z}_i, \bar{b}_i, [W'_{\omega_i \mathfrak{j}3}]_1, \mathbf{b}(\mathfrak{z}_i), [(b(x) - b(\mathfrak{z}_i))/(x - \mathfrak{z}_i)]_1);$  fi
  if  $\mathbf{c}(\mathfrak{z}_i) \neq \bar{c}_i$  then return  $([c]_1, \mathfrak{z}_i, \bar{c}_i, [W'_{\omega_i \mathfrak{j}4}]_1, \mathbf{c}(\mathfrak{z}_i), [(c(x) - c(\mathfrak{z}_i))/(x - \mathfrak{z}_i)]_1);$  fi
  if  $\mathbf{t}_{\text{lo}}(\mathfrak{z}_i) \neq \bar{t}_{3\text{lo},i}$  then return  $([t_{\text{lo}}]_1, \mathfrak{z}_i, \bar{t}_{3\text{lo},i}, [W'_{\omega_i \mathfrak{j}7}]_1, \mathbf{t}_{\text{lo}}(\mathfrak{z}_i), [(t_{\text{lo}}(x) - t_{\text{lo}}(\mathfrak{z}_i))/(x - \mathfrak{z}_i)]_1);$  fi
  if  $\mathbf{t}_{\text{mid}}(\mathfrak{z}_i) \neq \bar{t}_{3\text{mid},i}$  then return  $([t_{\text{mid}}]_1, \mathfrak{z}_i, \bar{t}_{3\text{mid},i}, [W'_{\omega_i \mathfrak{j}8}]_1, \mathbf{t}_{\text{mid}}(\mathfrak{z}_i), [(t_{\text{mid}}(x) - t_{\text{mid}}(\mathfrak{z}_i))/(x - \mathfrak{z}_i)]_1);$  fi
  if  $\mathbf{t}_{\text{hi}}(\mathfrak{z}_i) \neq \bar{t}_{3\text{hi},i}$  then return  $([t_{\text{hi}}]_1, \mathfrak{z}_i, \bar{t}_{3\text{hi},i}, [W'_{\omega_i \mathfrak{j}9}]_1, \mathbf{t}_{\text{hi}}(\mathfrak{z}_i), [(t_{\text{hi}}(x) - t_{\text{hi}}(\mathfrak{z}_i))/(x - \mathfrak{z}_i)]_1);$  fi
endfor ;
for  $i \in [1, \kappa_3]$  do
   $r_i(X) \leftarrow A_{0i}(X) + \alpha A_{1i}(X) + \alpha^2 A_{2i}(X) - Z_{\mathbb{H}}(\mathfrak{z}_i) \cdot (\mathbf{t}_{\text{lo}}(X) + \mathfrak{z}_i^n \mathbf{t}_{\text{mid}}(X) + \mathfrak{z}_i^{2n} \mathbf{t}_{\text{hi}}(X));$ 
  if  $r_i(\mathfrak{z}_i) \neq 0$  then return  $([r_i]_1, \mathfrak{z}_i, 0, [W'_{\mathfrak{z}_i 11}]_1, r_i(\mathfrak{z}_i), [r_i(x) - r_i(\mathfrak{z}_i)/(x - \mathfrak{z}_i)]_1);$ 
endfor
for  $i \in [1, \kappa_3]$  do
  if  $S_{\sigma 1}(\mathfrak{z}_i) \neq \bar{s}_{\sigma 1i}$  then
    return  $([s_{\sigma 1}]_1, \mathfrak{z}_i, \bar{s}_{\sigma 1i}, [W'_{\omega_i \mathfrak{j}5}]_1, S_{\sigma 1}(\mathfrak{z}_i), [(S_{\sigma 1}(x) - S_{\sigma 1}(\mathfrak{z}_i))/(x - \mathfrak{z}_i)]_1);$  fi
  if  $S_{\sigma 2}(\mathfrak{z}_i) \neq \bar{s}_{\sigma 2i}$  then
    return  $([s_{\sigma 2}]_1, \mathfrak{z}_i, \bar{s}_{\sigma 2i}, [W'_{\omega_i \mathfrak{j}6}]_1, S_{\sigma 2}(\mathfrak{z}_i), [(S_{\sigma 2}(x) - S_{\sigma 2}(\mathfrak{z}_i))/(x - \mathfrak{z}_i)]_1);$  fi
endfor ;
return  $\perp$ ;

```

Fig. 14. KZG's evaluation-binding adversary C_{evb} .

ings. Since $\text{bad}_{\text{rhino}} = \text{false}$, $\mathbf{t}(X) = (\mathbf{F}_0(X) + \alpha \mathbf{F}_1(X) + \alpha^2 \mathbf{F}_2(X))/Z_{\mathbb{H}}(X)$ is a polynomial.

2. **Claim 2.** There exists a KZG's $(\kappa_{\text{kzg}} + 1)$ -special-soundness adversary $\mathcal{B}_{\text{ss}}^{\text{kzg}}$ (see Fig. 13), such that $\Pr[\mathcal{B}_{\text{ss}}^{\text{kzg}} \text{ succeeds} \mid \text{bad}_{\text{ext}}] = 1$. Recall from Item ii that bad_{ext} is set if one of the four bad events happens. $\mathcal{B}_{\text{ss}}^{\text{kzg}}$ just tests which of the cases is true and returns the corresponding transcript. Clearly, $\Pr[\mathcal{B}_{\text{ss}}^{\text{kzg}} \text{ succeeds} \mid \text{bad}_{\text{ext}}] = 1$.

3. **Claim 3.** There exists an evaluation-binding adversary C_{evb} (see Fig. 14) for KZG, such that $\Pr[C_{\text{evb}} \text{ succeeds} \mid \text{bad}_{\text{evb}}] = 1$.

Assume that $\text{bad}_{\text{evb}} = \text{true}$. (Note that $\text{bad}_{\text{evb}} = \text{true}$ means that $\text{bad}_{\text{ext}} = \text{false}$, that is, $\text{Ext}_{\text{ss}}^{\text{sub}}$ managed to extract all polynomials.) Then, one of the bad cases in Item iii or Item iv happens. C_{evb} just finds out which of these events happens, and depending on the case, returns a collision. By the correctness of the extraction, and the completeness property of KZG, any of the returned values in Fig. 14 is a collision. Thus, $\Pr[C_{\text{evb}} \text{ succeeds} \mid \text{bad}_{\text{evb}}] = 1$.

4. **Claim 4.** In the case of Plonk or SmallPlonk, there exists a SplitRSDH adversary $\mathcal{D}_{\text{splitrshd}}$ (see Fig. 15), such that $\Pr[\mathcal{D}_{\text{splitrshd}} \text{ succeeds} \mid \text{bad}_{\text{rhino}}] = 1$. In Appendix D.3, we prove this claim as a separate theorem, Theorem 8. Recalling all three bad events are disjoint,

$$\begin{aligned} \Pr[\bar{\mathcal{E}}] &= \Pr[\bar{\mathcal{E}} \mid \overline{\text{bad}}] \Pr[\overline{\text{bad}}] + \Pr[\bar{\mathcal{E}} \mid \text{bad}_{\text{ext}}] \Pr[\text{bad}_{\text{ext}}] + \Pr[\bar{\mathcal{E}} \mid \text{bad}_{\text{evb}}] \Pr[\text{bad}_{\text{evb}}] \\ &\quad + \Pr[\bar{\mathcal{E}} \mid \text{bad}_{\text{rhino}}] \Pr[\text{bad}_{\text{rhino}}] \\ &\leq 0 + \Pr[\text{bad}_{\text{ext}}] + \Pr[\text{bad}_{\text{evb}}] + \Pr[\text{bad}_{\text{rhino}}]. \end{aligned}$$

Since \mathcal{C}_{evb} succeeds whenever bad_{evb} is set and KZG is evaluation-binding, $\Pr[\text{bad}_{\text{evb}}] = \text{negl}(\lambda)$. Similarly, $\Pr[\text{bad}_{\text{ext}}] = \Pr[\text{bad}_{\text{rhino}}] = \text{negl}(\lambda)$. Thus, $\Pr[\bar{\mathcal{E}}] \leq \text{negl}(\lambda)$. This proves the claim.

Case of SanPlonk. We postpone the proof of this case to Appendix D.4. \square

D.3 Rhino for Plonk

In this section, we prove Theorem 8. Recall that SplitRSDH is defined in Definition 4. Let $L_i^{\mathcal{S}}(X) = \prod_{j \neq i} \frac{X - \mathfrak{z}_j}{\mathfrak{z}_i - \mathfrak{z}_j}$ be Lagrange polynomials of $\mathcal{S} = \{\mathfrak{z}_i\}_{i=1}^{n_{\mathcal{S}}}$

Theorem 8. Consider Plonk (with $\text{par} = \text{par}_n^{\text{plonk}}$) or SmallPlonk (with $\text{par} = \text{par}_n^{\text{smallplonk}}$). There exists a par-SplitRSDH adversary $\mathcal{D}_{\text{splitrshd}}$ (see Fig. 15), such that $\Pr[\mathcal{D}_{\text{splitrshd}} \text{ succeeds} \mid \text{bad}_{\text{rhino}}] = 1$.

Proof (Proof of Theorem 8). Assume $\text{bad}_{\text{rhino}} = \text{true}$. Let $\mathcal{D}_{\text{splitrshd}}$ be the par-SplitRSDH adversary in Fig. 15. $\mathcal{D}_{\text{splitrshd}}$ uses $\mathcal{A}_{\text{ss}}^{\text{plonk}}$ to obtain a subtree $\hat{\mathcal{T}}_{\beta\gamma\alpha} = (\text{tr}_{ijk})$ (see Eq. (13)); note that $\hat{\mathcal{T}}_{\beta\gamma\alpha}$ contains, say, $[t_{lo}, t_{mid}, t_{hi}]_1$ in the case of Plonk or $[t]_1$ in the case of SmallPlonk). Then, $\mathcal{D}_{\text{splitrshd}}$ runs $\text{TE}_{*\text{plonk}}$ on $\hat{\mathcal{T}}_{\beta\gamma\alpha}$ to obtain several accepting KZG transcripts, and then runs $\text{Ext}_{\text{ss}}^{\text{sub}}$ on $\hat{\mathcal{T}}_{\beta\gamma\alpha}$ to obtain polynomials $(\mathbf{z}(X), \mathbf{a}(X), \mathbf{b}(X), \mathbf{c}(X))$. After that, $\mathcal{D}_{\text{splitrshd}}$ computes the polynomials $F_s(X)$, $F(X)$ and $\mathbf{t}(X)$ (as defined in Eq. (7)). $\mathcal{D}_{\text{splitrshd}}$ also defines linearization polynomials $A_{0i}(X)$, $A_{1i}(X)$, $A_{2i}(X)$ (see Eq. (8)) of $F_0(X)$, $F_1(X)$, and $F_2(X)$, and their batched sum $A_i(X)$ corresponding to $F(X)$. Clearly, $[A_{si}(x)]_1 = [A_{si}]_1$ (from Fig. 6) for $s \in [0, 2]$ and $i \in [1, \kappa_3]$. If the SplitRSDH's requirement $Z_{\mathbb{H}}(X) \dagger F(X)$ is satisfied (this holds always since $\text{bad}_{\text{rhino}} = \text{true}$, see Item v), $\mathcal{D}_{\text{splitrshd}}$ outputs a SplitRSDH adversary's output.

Let us argue that $\mathcal{D}_{\text{splitrshd}}$ is successful, i.e., the output satisfies par-SplitRSDH's conditions. First, it is easy to see that \mathcal{S} is of size κ_3 since the values \mathfrak{z}_i output by $\mathcal{A}_{\text{ss}}^{\text{plonk}}$ are mutually different by the definition of special-soundness. Second, $L(X) = \sum_{i=1}^{\kappa_3} \mathbf{t}(\mathfrak{z}_i) L_i^{\mathcal{S}}(X)$ as defined in Fig. 15 has degree $\leq \kappa_3 - 1$. It remains to show that $\deg(L) \geq \kappa_{\text{kzg}} + n_{\psi} + 1$ and that

$$[t_{lo}]_1 \bullet [1]_2 + [t_{mid}]_1 \bullet [x^n]_2 + [t_{hi}]_1 \bullet [x^{2n}]_2 = [1]_1 \bullet [L(x)]_2 + [\varphi]_1 \bullet [Z_{\mathcal{S}}(x)]_2 \quad (18)$$

in the case of Plonk or

$$[t]_1 \bullet [1]_2 = [1]_1 \bullet [L(x)]_2 + [\varphi]_1 \bullet [Z_{\mathcal{S}}(x)]_2 \quad (19)$$

$\mathcal{D}_{\text{splitrsdh}}(\mathbf{p}, \mathbf{ck})$	
1 :	$\hat{\mathcal{T}}_{\beta\gamma\alpha} \leftarrow \mathcal{A}_{\text{ss}}^{\text{plonk}}(\mathbf{ck}); ((\mathbf{k}, \mathbf{tr}_k)_{k \in [1, \kappa_v + 2]}, \mathbf{k}, \mathbf{tr}^\omega) \leftarrow \mathbf{TE}_{*\text{plonk}}(\mathbf{ck}, \hat{\mathcal{T}}_{\beta\gamma\alpha});$
2 :	$(\mathbf{z}(X), \mathbf{a}(X), \mathbf{b}(X), \mathbf{c}(X)) \leftarrow \mathbf{Ext}_{\text{ss}}^{\text{sub}}(\mathbf{ck}, \hat{\mathcal{T}}_{\beta\gamma\alpha});$
3 :	$\mathbf{F}_0(X) \leftarrow \mathbf{a}(X)\mathbf{b}(X)\mathbf{q}_M(X) + \mathbf{a}(X)\mathbf{q}_L(X) + \mathbf{b}(X)\mathbf{q}_R(X) + \mathbf{c}(X)\mathbf{q}_O(X) + \mathbf{PI}(X) + \mathbf{q}_C(X);$
4 :	$\mathbf{F}_1(X) \leftarrow (\mathbf{a}(X) + \beta X + \gamma)(\mathbf{b}(X) + \beta k_1 X + \gamma)(\mathbf{c}(X) + \beta k_2 X + \gamma)\mathbf{z}(X)$
5 :	$\quad \quad \quad - (\mathbf{a}(X) + \beta S_{\sigma_1}(X) + \gamma)(\mathbf{b}(X) + \beta S_{\sigma_2}(X) + \gamma)(\mathbf{c}(X) + \beta S_{\sigma_3}(X) + \gamma)\mathbf{z}(\omega X);$
6 :	$\mathbf{F}_2(X) \leftarrow (\mathbf{z}(X) - 1)L_1(X);$
7 :	$\mathbf{F}(X) \leftarrow \mathbf{F}_0(X) + \alpha \mathbf{F}_1(X) + \alpha^2 \mathbf{F}_2(X); (**)$
8 :	$\mathbf{t}(X) \leftarrow \mathbf{F}(X)/\mathbf{Z}_{\mathbb{H}}(X);$
9 :	if $\mathbf{Z}_{\mathbb{H}}(X) \mid \mathbf{F}(X)$ then return \perp ; fi
10 :	for $i \in [1, \kappa_3]$ do
11 :	$\Lambda_{0i}(X) \leftarrow \bar{a}_i \bar{b}_i \mathbf{q}_M(X) + \bar{a}_i \mathbf{q}_L(X) + \bar{b}_i \mathbf{q}_R(X) + \bar{c}_i \mathbf{q}_O(X) + \mathbf{PI}(\mathfrak{z}) + \mathbf{q}_C(X);$
12 :	$\Lambda_{1i}(X) \leftarrow (\bar{a}_i + \beta \mathfrak{z}_i + \gamma)(\bar{b}_i + \beta k_1 \mathfrak{z}_i + \gamma)(\bar{c}_i + \beta k_2 \mathfrak{z}_i + \gamma)\mathbf{z}(X)$
13 :	$\quad \quad \quad - (\bar{a}_i + \beta \bar{s}_{\sigma_1} + \gamma)(\bar{b}_i + \beta \bar{s}_{\sigma_2} + \gamma)(\bar{c}_i + \beta S_{\sigma_3}(X) + \gamma)\bar{z}_{\omega, i};$
14 :	$\Lambda_{2i}(X) \leftarrow (\mathbf{z}(X) - 1)L_1(\mathfrak{z}_i);$
15 :	$\Lambda_i(X) \leftarrow \Lambda_{0i}(X) + \alpha \Lambda_{1i}(X) + \alpha^2 \Lambda_{2i}(X);$
16 :	$\chi'_i(X) \leftarrow \frac{\Lambda_i(X) - \Lambda_i(\mathfrak{z}_i)}{X - \mathfrak{z}_i};$
17 :	$[\mathbf{W}'_{\mathfrak{z}_i 1}, \dots, \mathbf{W}'_{\mathfrak{z}_i \kappa_v}]_1^\top \leftarrow \mathbf{V}_i^{-1}[\mathbf{W}_{\mathfrak{z}_i 1}, \dots, \mathbf{W}_{\mathfrak{z}_i \kappa_v}];$ // \mathbf{V}_i as in Eq. (15)
18 :	$[\chi_i]_1 \leftarrow \frac{1}{\mathbf{Z}_{\mathbb{H}}(\mathfrak{z}_i)}[\chi'_i(x) - \mathbf{W}'_{\mathfrak{z}_i 1}]_1;$ // $\mathbf{W}'_{\mathfrak{z}_i 1}$ is as in Line 8 in Fig. 6
19 :	$\Delta_i \leftarrow 1 / \prod_{j \neq i} (\mathfrak{z}_i - \mathfrak{z}_j);$ endfor
20 :	$[\varphi]_1 \leftarrow \sum_{i=1}^{\kappa_3} \Delta_i [\chi_i]_1;$
21 :	$L(X) \leftarrow \sum_{i=1}^{\kappa_3} \mathbf{t}(\mathfrak{z}_i) L_i^S(X);$ // $L_i^S(X)$ are Lagrange basis polynomials
22 :	// Only the return value depends on whether we have Plonk or SmallPlonk
23 :	return $(\mathcal{S} = \{\mathfrak{z}_i\}_{i=1}^{\kappa_3}, L(X), [t, \varphi]_1);$ // In SmallPlonk
24 :	return $(\mathcal{S} = \{\mathfrak{z}_i\}_{i=1}^{\kappa_3}, L(X), [t_{lo}, t_{mid}, t_{hi}, \varphi]_1);$ // In Plonk

Fig. 15. The SplitRSDH adversary $\mathcal{D}_{\text{splitrsdh}}$.

in the case of SmallPlonk. (here, we took into account the concrete par).

Recall from the proof of Theorem 5 that $\mathbf{bad}_{\text{rhino}} = \mathbf{true}$ implies $\mathbf{bad}_{\text{evb}} = \mathbf{bad}_{\text{ext}} = \mathbf{false}$. Thus,

1. $[z]_1 = [z(x)]_1$, $[a]_1 = [a(x)]_1$, $[b]_1 = [b(x)]_1$, $[c]_1 = [c(x)]_1$, and
2. $\mathbf{z}(\mathfrak{z}_i \omega) = \bar{z}_{\omega, i}$, $\mathbf{a}(\mathfrak{z}_i) = \bar{a}_i$, $\mathbf{b}(\mathfrak{z}_i) = \bar{b}_i$, and $\mathbf{c}(\mathfrak{z}_i) = \bar{c}_i$ for $i \in [1, \kappa_3]$.

Thus, $\mathbf{F}_s(\mathfrak{z}_i) = \Lambda_{si}(\mathfrak{z}_i)$, $\mathbf{F}(\mathfrak{z}_i) = \Lambda_i(\mathfrak{z}_i)$, and $\mathbf{t}(\mathfrak{z}_i) = \Lambda_i(\mathfrak{z}_i)/\mathbf{Z}_{\mathbb{H}}(\mathfrak{z}_i)$ for $s \in \{0, 1, 2\}$ and $i \in [1, \kappa_3]$.

Denote $y_i := t_{lo} + \mathfrak{z}_i^n t_{mid} + \mathfrak{z}_i^{2n} t_{hi}$ in the case of Plonk and $y_i := t$ in the case of SmallPlonk. Recall

$$[r_i]_1 = [\Lambda_i(x) - \mathbf{Z}_{\mathbb{H}}(\mathfrak{z}_i)y_i]_1 \quad (20)$$

from the lines 3, 4, and 6 in Fig. 6. Next, we use the fact that $\mathcal{D}_{\text{splitrsdh}}$ knows $\Lambda_i(X)$ as a polynomial. Since it can open both $[r_i]_1$ and $[\Lambda_i(x)]_1$ at \mathfrak{z}_i , it can also open $[y_i]_1$. More precisely, let

$$\chi'_i(X) = (\Lambda_i(X) - \Lambda_i(\mathfrak{z}_i))/(X - \mathfrak{z}_i) = (\Lambda_i(X) - \mathbf{F}(\mathfrak{z}_i))/(X - \mathfrak{z}_i)$$

be the KZG opening polynomial of $A_i(X)$, $i \in [1, \kappa_3]$, at point \mathfrak{z}_i . The second equality follows from $\text{bad}_{\text{evb}} = \text{bad}_{\text{ext}} = \text{false}$.

Recall that as part of k.tr_{1i} (see line 9 in Fig. 6), $[W'_{\mathfrak{z}_{i1}}]_1$ is the opening proof of $[r_i]_1$ at \mathfrak{z}_i to 0. Since k.tr_{1i} is accepting, Eq. (20) implies that $W'_{\mathfrak{z}_{i1}} \cdot (x - \mathfrak{z}_i) = r_i$. It follows from this, $\text{bad}_{\text{evb}} = \text{bad}_{\text{ext}} = \text{false}$, and the definition of $\chi'_i(X)$ that

$$W'_{\mathfrak{z}_{i1}} \cdot (x - \mathfrak{z}_i) = r_i = \chi'_i(x)(x - \mathfrak{z}_i) + A_i(\mathfrak{z}_i) - Z_{\mathbb{H}}(\mathfrak{z}_i)y_i$$

and thus

$$\chi'_i(x) - W'_{\mathfrak{z}_{i1}} = \frac{1}{x - \mathfrak{z}_i} (Z_{\mathbb{H}}(\mathfrak{z}_i)y_i - A_i(\mathfrak{z}_i))$$

for $i \in [1, \kappa_3]$. Since $A_i(\mathfrak{z}_i) = F(\mathfrak{z}_i)$ and $\mathfrak{t}(\mathfrak{z}_i) = F(\mathfrak{z}_i)/Z_{\mathbb{H}}(\mathfrak{z}_i)$, we get $\chi'_i(x) - W'_{\mathfrak{z}_{i1}} = \frac{1}{x - \mathfrak{z}_i} \cdot (Z_{\mathbb{H}}(\mathfrak{z}_i)y_i - F(\mathfrak{z}_i)) = \frac{Z_{\mathbb{H}}(\mathfrak{z}_i)}{x - \mathfrak{z}_i} (y_i - \mathfrak{t}(\mathfrak{z}_i))$. Defining $\chi_i \leftarrow \frac{1}{Z_{\mathbb{H}}(\mathfrak{z}_i)} (\chi'_i(x) - W'_{\mathfrak{z}_{i1}})$ as in Fig. 15 for $i \in [1, \kappa_3]$, we get

$$y_i - \mathfrak{t}(\mathfrak{z}_i) = \chi_i(x - \mathfrak{z}_i) . \quad (21)$$

We multiply κ_3 equations Eq. (21) individually by $L_i^{\mathcal{S}}(X)$ and then sum the results, getting

$$\sum_{i=1}^{\kappa_3} (y_i - \mathfrak{t}(\mathfrak{z}_i)) L_i^{\mathcal{S}}(x) = \sum_{i=1}^{\kappa_3} \chi_i(x - \mathfrak{z}_i) L_i^{\mathcal{S}}(x) . \quad (22)$$

Since $L_i^{\mathcal{S}}(X) = \prod_{j \neq i} \frac{X - \mathfrak{z}_j}{\mathfrak{z}_i - \mathfrak{z}_j} = \Delta_i \prod_{j \neq i} (X - \mathfrak{z}_j) = \Delta_i \frac{Z_{\mathbb{S}}(X)}{X - \mathfrak{z}_i}$, $(X - \mathfrak{z}_i) L_i^{\mathcal{S}}(X) = \Delta_i Z_{\mathbb{S}}(X)$. Let $L(X) = \sum_{i=1}^{\kappa_3} \mathfrak{t}(\mathfrak{z}_i) L_i^{\mathcal{S}}(X)$ be the interpolating polynomial of $\{(\mathfrak{z}_i, \mathfrak{t}(\mathfrak{z}_i))\}_{i=1}^{\kappa_3}$. Denote $T(x) := t_{lo} + t_{mid}x^n + t_{hi}x^{2n}$ in the case of Plonk and $T(x) := t$ in the case of SmallPlonk. In the case of Plonk,

$$\begin{aligned} T(x) - L(x) &= t_{lo} + t_{mid}x^n + t_{hi}x^{2n} - L(x) \\ &\stackrel{(*)}{=} t_{lo} + t_{mid} \sum_{i=1}^{\kappa_3} \mathfrak{z}_i^n L_i^{\mathcal{S}}(x) + t_{hi} \sum_{i=1}^{\kappa_3} \mathfrak{z}_i^{2n} L_i^{\mathcal{S}}(x) - L(x) \\ &= \sum_{i=1}^{\kappa_3} (y_i - \mathfrak{t}(\mathfrak{z}_i)) L_i^{\mathcal{S}}(x) , \end{aligned}$$

where $(*)$ follows from (say) $X^n = \sum_{i=1}^{\kappa_3} \mathfrak{z}_i^n L_i^{\mathcal{S}}(X)$. In the case of SmallPlonk,

$$T(x) - L(x) = t - L(x) = \sum_{i=1}^{\kappa_3} (y_i - \mathfrak{t}(\mathfrak{z}_i)) L_i^{\mathcal{S}}(x) .$$

Thus,

$$T(x) - L(x) \stackrel{22}{=} \sum_{i=1}^{\kappa_3} \chi_i(x - \mathfrak{z}_i) L_i^{\mathcal{S}}(x) = Z_{\mathbb{S}}(x) \sum_{i=1}^{\kappa_3} \Delta_i \chi_i = Z_{\mathbb{S}}(x) \varphi .$$

Thus, Eq. (18) (in the case of Plonk) or Eq. (19) (in the case of SmallPlonk) holds.

Finally, clearly $\deg L(X) \leq \kappa_3 - 1$. On the other hand, assume $\deg L(X) \leq \kappa_{\text{kzg}} + n_{\psi} \leq \kappa_3 - n - 1$. In the case of Plonk, the last inequality holds due to the choice of $\text{par}_n^{\text{plonk}}$, since $\kappa_3 = 4\kappa_{\text{kzg}} + 1 = 4n + 21 \geq (n + 5) + 2n + n + 1 = \kappa_{\text{kzg}} + n_{\psi} + n + 1$. In SmallPlonk, it is even more obvious since $\kappa_3 = 4\kappa_{\text{kzg}} + 1 = 12n + 21 \geq (3n + 5) + 0 + n + 1 = \kappa_{\text{kzg}} + n_{\psi} + n + 1$. Since $L(X)$ and $F(X)/Z_{\mathbb{H}}(X)$ match on κ_3 points \mathfrak{z}_i and both $L(X)Z_{\mathbb{H}}(X)$ and $F(X)$ have degree $\leq \kappa_3 - 1$, $L(X) = F(X)/Z_{\mathbb{H}}(X)$, which means $Z_{\mathbb{H}}(X) \mid F(X)$, contradiction with $\text{bad}_{\text{rhino}} = \text{true}$. Thus, if $Z_{\mathbb{H}}(X) \nmid F(X)$, then $\deg L(X) \in [\kappa_{\text{kzg}} + n_{\psi} + 1, \kappa_3 - 1]$. \square

D.4 SanPlonk’s Case from Theorem 5

Proof (Finishing the proof of Theorem 5).

Case of SanPlonk. The case of SanPlonk is similar to Plonk but differs in quite many details. For the sake of clarity, we **highlight** additional text, but we also removed some text that is not relevant for SanPlonk. (Intuitively, the removed text corresponds to the part in Plonk’s proof where one handles the reduction to SplitRSDH.)

In Fig. 12, we depict an extractor $\text{Ext}_{\text{SS}}^{\text{sub}}$. $\text{Ext}_{\text{SS}}^{\text{sub}}$ invokes $\text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k}, \mathbf{tr}^\omega)$ and $\text{Ext}_{\text{SS}}^{\text{kzg}}(\text{ck}, \mathbf{k}, \mathbf{tr}_k)$ for $k \in [2, 4] \cup [7, 9]$, extracting polynomials $z(X)$, $a(X)$, $b(X)$, $c(X)$, $\mathbf{t}_{\text{lo}}(X)$, $\mathbf{t}_{\text{mid}}(X)$, and $\mathbf{t}_{\text{hi}}(X)$ of at most degree $\kappa_{\text{kzg}} = n + 5$. After executing $\text{Ext}_{\text{SS}}^{\text{sub}}$, we use the following procedure to possibly set one of the “bad” flags:

- (i’) $\text{bad}_{\text{ext}} \leftarrow \text{false}$; $\text{bad}_{\text{evb}} \leftarrow \text{false}$;
- (ii’) if $([z]_1, z(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k}, \text{tr}_1} \vee ([a]_1, a(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k}, \text{tr}_2} \vee ([b]_1, b(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k}, \text{tr}_3} \vee ([c]_1, c(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k}, \text{tr}_4} \vee ([\mathbf{t}_{\text{lo}}]_1, \mathbf{t}_{\text{lo}}(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k}, \text{tr}_7} \vee ([\mathbf{t}_{\text{mid}}]_1, \mathbf{t}_{\text{mid}}(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k}, \text{tr}_8} \vee ([\mathbf{t}_{\text{lo}}]_1, \mathbf{t}_{\text{hi}}(X)) \notin \mathcal{R}_{\text{ck}, \mathbf{k}, \text{tr}_9}$ (see Eq. (1)) then $\text{bad}_{\text{ext}} \leftarrow \text{true}$; abort;
- (iii’) for $i \in [\kappa_{\text{kzg}} + 2, \kappa_3]$: if $z(\mathfrak{z}_i \omega) \neq \bar{z}_{\omega i} \vee a(\mathfrak{z}_i) \neq \bar{a}_i \vee b(\mathfrak{z}_i) \neq \bar{b}_i \vee c(\mathfrak{z}_i) \neq \bar{c}_i \vee \mathbf{t}_{\text{lo}}(\mathfrak{z}_i) \neq \bar{t}_{3\text{lo}, i} \vee \mathbf{t}_{\text{mid}}(\mathfrak{z}_i) \neq \bar{t}_{3\text{mid}, i} \vee \mathbf{t}_{\text{hi}}(\mathfrak{z}_i) \neq \bar{t}_{3\text{hi}, i}$ then $\text{bad}_{\text{evb}} \leftarrow \text{true}$; abort;
- (iv’) for $i \in [1, \kappa_3]$:
 - if $S_{\sigma 1}(\mathfrak{z}_i) \neq \bar{s}_{\sigma 1 i} \vee S_{\sigma 2}(\mathfrak{z}_i) \neq \bar{s}_{\sigma 2 i}$ then $\text{bad}_{\text{evb}} \leftarrow \text{true}$; abort;
 - $r_i(X) \leftarrow A_{0i}(X) + \alpha A_{1i}(X) + \alpha^2 A_{2i}(X) - Z_{\mathbb{H}}(\mathfrak{z}_i) \cdot (\mathbf{t}_{\text{lo}}(X) + \mathfrak{z}_i^n \mathbf{t}_{\text{mid}}(X) + \mathfrak{z}_i^{2n} \mathbf{t}_{\text{hi}}(X))$;
 - if $r_i(\mathfrak{z}_i) \neq 0$ then $\text{bad}_{\text{evb}} \leftarrow \text{true}$;

Importantly, only one of the “bad” flags is set at a time. Thus, for example, $\text{bad}_{\text{evb}} = \text{true}$ implies that $\text{bad}_{\text{ext}} = \text{false}$. Let \mathcal{E} be the event $\text{Ext}_{\text{SS}}^{\text{sub}}$ succeeds and $\overline{\text{bad}}$ be the event none of the bad flags was set. Thus, \mathcal{E} is the event that $a(X)$, $b(X)$, $c(X)$, $z(X)$, $\mathbf{t}_{\text{lo}}(X)$, $\mathbf{t}_{\text{mid}}(X)$, and $\mathbf{t}_{\text{hi}}(X)$ are consistent with the commitments and all openings, and $\mathbf{t}(X) = (\mathbf{F}_0(X) + \alpha \mathbf{F}_1(X) + \alpha^2 \mathbf{F}_2(X)) / Z_{\mathbb{H}}(X)$ is a polynomial. We analyze the success probability of $\text{Ext}_{\text{SS}}^{\text{sub}}$. For this, we make the following claims.

1. **Claim 1.** $\Pr[\mathcal{E} | \overline{\text{bad}}] = 1$.

Really, assume that $\overline{\text{bad}}$ holds. Since $\text{bad}_{\text{ext}} = \text{bad}_{\text{evb}} = \text{false}$, we get that $a(X)$, $b(X)$, $c(X)$, $z(X)$, $\mathbf{t}_{\text{lo}}(X)$, $\mathbf{t}_{\text{mid}}(X)$, $\mathbf{t}_{\text{hi}}(X)$ are consistent with the commitments and all openings.

Recall that in the case of Plonk, we deduced here that since $\text{bad}_{\text{rhino}} = \text{false}$, $\mathbf{t}(X) = (\mathbf{F}_0(X) + \alpha \mathbf{F}_1(X) + \alpha^2 \mathbf{F}_2(X)) / Z_{\mathbb{H}}(X)$ is a polynomial. In the case of SanPlonk, we handle this differently. Since $\text{bad}_{\text{evb}} = \text{false}$, $r_i(X) = 0$ for all $i \in [1, \kappa_3]$. Let $g(X) := \mathbf{F}(X) - Z_{\mathbb{H}}(X) \mathbf{t}^*(X)$, where $\mathbf{t}^*(X) := \mathbf{t}_{\text{lo}}(X) + X^n \mathbf{t}_{\text{mid}}(X) + X^{2n} \mathbf{t}_{\text{hi}}(X)$. Since for $i \in [1, \kappa_3]$, $\mathbf{F}_s(\mathfrak{z}_i) = A_{si}(\mathfrak{z}_i)$, we have $g(\mathfrak{z}_i) = 0$. Since this holds for $\kappa_3 = 4\kappa_{\text{kzg}} + 1$ evaluation points and $\deg g(X) \leq 4\kappa_{\text{kzg}}$, $g(X) = 0$. Thus, $\mathbf{t}(X)$ is a polynomial.

2. **Claim 2.** There exists a KZG's $(\kappa_{\text{kzg}} + 1)$ -special-soundness adversary $\mathcal{B}_{\text{SS}}^{\text{kzg}}$ (see Fig. 13), such that $\Pr[\mathcal{B}_{\text{SS}}^{\text{kzg}} \text{ succeeds} \mid \text{bad}_{\text{ext}}] = 1$.

Recall from Item ii that bad_{ext} is set if one of the **seven** bad events happens. $\mathcal{B}_{\text{SS}}^{\text{kzg}}$ just tests which of the cases is true and returns the corresponding transcript. Clearly, $\Pr[\mathcal{B}_{\text{SS}}^{\text{kzg}} \text{ succeeds} \mid \text{bad}_{\text{ext}}] = 1$.

3. **Claim 3.** There exists an evaluation-binding adversary \mathcal{C}_{evb} (see Fig. 14) for KZG, such that $\Pr[\mathcal{C}_{\text{evb}} \text{ succeeds} \mid \text{bad}_{\text{evb}}] = 1$.

Assume that $\text{bad}_{\text{evb}} = \text{true}$. (Note that $\text{bad}_{\text{evb}} = \text{true}$ means that $\text{bad}_{\text{ext}} = \text{false}$, that is, $\text{Ext}_{\text{SS}}^{\text{sub}}$ managed to extract all polynomials.) Then, one of the bad cases in Item iii or Item iv happens. \mathcal{C}_{evb} just finds out which of these events happens, and depending on the case, returns a collision. By the correctness of the extraction, and the completeness property of KZG, any of the returned values in Fig. 14 is a collision. Thus, $\Pr[\mathcal{C}_{\text{evb}} \text{ succeeds} \mid \text{bad}_{\text{evb}}] = 1$.

Thus,

$$\begin{aligned} \Pr[\bar{\mathcal{E}}] &= \Pr[\bar{\mathcal{E}} \mid \overline{\text{bad}}] \Pr[\overline{\text{bad}}] + \Pr[\bar{\mathcal{E}} \mid \text{bad}_{\text{ext}}] \Pr[\text{bad}_{\text{ext}}] + \Pr[\bar{\mathcal{E}} \mid \text{bad}_{\text{evb}}] \Pr[\text{bad}_{\text{evb}}] \\ &\leq 0 + \Pr[\text{bad}_{\text{ext}}] + \Pr[\text{bad}_{\text{evb}}] \end{aligned}$$

Since \mathcal{C}_{evb} succeeds whenever bad_{evb} is set and KZG is evaluation-binding, $\Pr[\text{bad}_{\text{evb}}] = \text{negl}(\lambda)$. Similarly, $\Pr[\text{bad}_{\text{ext}}] = \text{negl}(\lambda)$. Thus, $\Pr[\bar{\mathcal{E}}] \leq \text{negl}(\lambda)$. This proves the claim. \square

D.5 Proof of Theorem 6

In this section, we restate and then prove the theorem Theorem 6 from Section 6.

Theorem 9. *Let $n \in \text{poly}(\lambda)$ and κ_{kzg} , κ_{Plonk} , and κ_{san} be as in Eq. (12).*

1. *If KZG is computational $(\kappa_{\text{kzg}} + 1)$ -special-sound and evaluation-binding, and $\text{par}_n^{\text{smallplonk}}$ -SplitRSDH holds, then SmallPlonk is computational κ_{Plonk} -special-sound.*
2. *If KZG is computational $(\kappa_{\text{kzg}} + 1)$ -special-sound and evaluation-binding, and $\text{par}_n^{\text{plonk}}$ -SplitRSDH holds, then Plonk is computational κ_{Plonk} -special-sound.*
3. *If KZG is computational $(\kappa_{\text{kzg}} + 1)$ -special-sound and evaluation-binding, then SanPlonk is computational κ_{san} -special-sound.*

Proof. Fix $\mathcal{P} = \{\text{q}_M(X), \text{q}_L(X), \text{q}_R(X), \text{q}_O(X), \text{q}_C(X), \text{S}_{\sigma 1}(X), \text{S}_{\sigma 2}(X), \text{S}_{\sigma 3}(X)\}$ for a relation defined by encoding a circuit as in Section 6.1. Let \mathcal{A}_{SS} be a PPT adversary in the computational special-soundness game that outputs a $(\kappa_\beta, \kappa_\gamma, \kappa_\alpha, \kappa_3, \kappa_\delta, \kappa_v)$ -tree \mathcal{T} . We describe the special-soundness extractor $\text{Ext}_{\text{SS}}^{\text{plonk}}$ for Plonk in Fig. 16. The extractor picks an arbitrary $(1, 1, 1, \kappa_3, \kappa_\delta, \kappa_v)$ -subtree $\hat{\mathcal{T}} = \hat{\mathcal{T}}_{\beta\gamma\alpha}$ of \mathcal{T} and runs the subtree extractor $\text{Ext}_{\text{SS}}^{\text{sub}}$ from Theorem 5 on it to obtain the polynomials $\text{z}(X)$, $\text{a}(X)$, $\text{b}(X)$, $\text{c}(X)$, $\text{t}_{\text{lo}}(X)$, $\text{t}_{\text{mid}}(X)$, and $\text{t}_{\text{hi}}(X)$. In the honest protocol, the witness is encoded in $\text{a}(X)$, $\text{b}(X)$, and $\text{c}(X)$. Namely, $\bar{w}_i \leftarrow \text{a}(\omega^i)$, $\bar{w}_{n+i} \leftarrow \text{b}(\omega^i)$, $\bar{w}_{2n+i} \leftarrow \text{c}(\omega^i)$ for $i \in [1, n]$, and $\mathbf{w} = (\bar{w}_i)_{i=\ell+1}^{3n}$. The rest of the proof shows that $(\mathbf{x} = (\bar{w}_i)_{i=1}^\ell, \mathbf{w}) \in \mathcal{R}_{\mathcal{P}}$.

$\text{Ext}_{\text{ss}}^{*\text{plonk}}(\text{ck}, \mathfrak{x}, \mathcal{T})$ <hr style="border: 0.5px solid black;"/> <p style="margin: 0;">Pick an arbitrary $(1, 1, 1, \kappa_\beta, \kappa_\delta, \kappa_\nu)$-subtree $\hat{\mathcal{T}}$ of \mathcal{T}. $(z(X), a(X), b(X), c(X), \mathfrak{t}_{\text{lo}}(X), \mathfrak{t}_{\text{mid}}(X), \mathfrak{t}_{\text{hi}}(X)) \leftarrow \text{Ext}_{\text{ss}}^{\text{sub}}(\text{ck}, \hat{\mathcal{T}})$; for $i \in [1, n]$ do $\bar{w}_i \leftarrow a(\omega^i); \bar{w}_{n+i} \leftarrow b(\omega^i); \bar{w}_{2n+i} \leftarrow c(\omega^i)$; endfor return $w \leftarrow (\bar{w}_i)_{i=\ell+1}^{3n}$</p>
--

Fig. 16. The special-soundness extractor $\text{Ext}_{\text{ss}}^{*\text{plonk}}$ for Plonk/SanPlonk.

To be sure we compute a correct witness, we must assume that for each β_i , we have κ_γ mutually different values γ_{ij} . Whence the double index on challenges γ_{ij} , despite they are sent in the same round by Plonk and its variants' prover. One can implement this by rewinding the protocol with $\kappa_\beta \kappa_\gamma$ different challenges $(\beta_i \gamma_{ij})$. Alternatively, one can consider Plonk as the interactive argument where the prover first receives the challenge β from the verifier, then replies with an empty message, then receives the challenge γ , and goes on with the execution.⁷

Let $\text{SubTrees}_{\mathcal{T}} := \{\hat{\mathcal{T}}_{\beta_i \gamma_{ij} \alpha_{ijk}} : i \in [1, \kappa_\beta], j \in [1, \kappa_\gamma], k \in [1, \kappa_\alpha]\}$ be the set of all $(1, 1, 1, \kappa_\beta, \kappa_\delta, \kappa_\nu)$ -subtrees of \mathcal{T} . For each $\hat{\mathcal{T}} \in \text{SubTrees}_{\mathcal{T}}$, we can apply the extractor from Theorem 5 and obtain the subtree transcript $\text{s.tr}_{\hat{\mathcal{T}}} = (z_{\hat{\mathcal{T}}}(X), a_{\hat{\mathcal{T}}}(X), b_{\hat{\mathcal{T}}}(X), c_{\hat{\mathcal{T}}}(X))$, where $\hat{\mathcal{T}} = \hat{\mathcal{T}}_{\beta_i \gamma_{ij} \alpha_{ijk}}$ and $i \in [1, \kappa_\beta]$, $j \in [1, \kappa_\gamma]$, and $k \in [1, \kappa_\alpha]$. Observe that the extracted polynomials may depend on the specific subtree $\hat{\mathcal{T}}$.

Next, we argue that the extractor $\text{Ext}_{\text{ss}}^{*\text{plonk}}$ can fail only if one of the following events happens.

bad_{sub}: For some $\hat{\mathcal{T}} \in \text{SubTrees}_{\mathcal{T}}$, $\text{Ext}_{\text{ss}}^{\text{sub}}(\text{ck}, \hat{\mathcal{T}})$ outputs $\text{s.tr}_{\hat{\mathcal{T}}}$ such that either (1) $z_{\hat{\mathcal{T}}}(X)$, $a_{\hat{\mathcal{T}}}(X)$, $b_{\hat{\mathcal{T}}}(X)$, and $c_{\hat{\mathcal{T}}}(X)$ are inconsistent with the commitments $[z_{ij}, a, b, c]_1$, or (2) $\mathfrak{t}_{\hat{\mathcal{T}}_{ijk}}(X)$ (defined as in Eq. (7)) is not a polynomial.

bad_{bind}: (1) The event **bad_{sub}** does not happen. (2) Define $W_{\hat{\mathcal{T}}} := \{(a(X), b(X), c(X)) : (z(X), a(X), b(X), c(X)) \leftarrow \text{Ext}_{\text{ss}}^{\text{sub}}(\text{ck}, \hat{\mathcal{T}})\}$. There exist two subtrees $\hat{\mathcal{T}} \neq \hat{\mathcal{T}}' \in \text{SubTrees}_{\mathcal{T}}$, such that $W_{\hat{\mathcal{T}}} \neq W_{\hat{\mathcal{T}'}}$.

Theorem 5 implies that if KZG is evaluation-binding and computationally special-sound (and the additional assumption SplitRSDH holds in the case of Plonk or SmallPlonk), then $\Pr[\text{bad}_{\text{sub}}]$ is negligible. The fact that $\Pr[\text{bad}_{\text{bind}}]$ is negligible follows straightforwardly from the binding property of KZG. For the sake of completeness we present the binding adversary $\mathcal{A}_{\text{bind}}$ in Fig. 17.

In the following, suppose that neither **bad_{sub}** or **bad_{bind}** happened. Thus, $W_{\hat{\mathcal{T}}}(X) = W_{\hat{\mathcal{T}}'}(X)$ for any $\hat{\mathcal{T}}, \hat{\mathcal{T}}' \in \text{SubTrees}_{\mathcal{T}}$. This justifies the notation $\text{s.tr}_{ijk} = (z_{ij}(X), a(X), b(X), c(X))$, where $a(X)$, $b(X)$, and $c(X)$ do not depend on the specific subtree $\hat{\mathcal{T}}$ while $z_{ij}(X)$ depends on β_i and γ_{ij} . Further-

⁷ We note that this does not change actual Plonk/SanPlonk: when applying Fiat-Shamir, one defines $\beta = H(\text{view}, 0)$ and $\gamma = H(\text{view}, 1)$. To rewind only γ and not β , one can reprogram the random oracle at input $(\text{view}, 1)$ but not input $(\text{view}, 0)$.

$\mathcal{A}_{\text{bind}}(\text{ck})$ <hr/> $(\mathbf{x}, \mathcal{J}) \leftarrow \text{Ext}_{\text{ss}}^{\text{kzg}}(\text{ck});$ for $\hat{\mathcal{J}} \in \text{SubTrees}_{\mathcal{J}}$ do $\text{s.tr}_{\hat{\mathcal{J}}} \leftarrow \text{Ext}_{\text{ss}}^{\text{sub}}(\text{ck}, \hat{\mathcal{J}});$ endfor // Extracts polynomials from each subtree for distinct $\hat{\mathcal{J}} \neq \hat{\mathcal{J}}' \in \text{SubTrees}_{\mathcal{J}}$ do $(\mathbf{a}(X), \mathbf{b}(X), \mathbf{c}(X)) \leftarrow W_{\hat{\mathcal{J}}}; (\mathbf{a}'(X), \mathbf{b}'(X), \mathbf{c}'(X)) \leftarrow W_{\hat{\mathcal{J}}'};$ if $\mathbf{a}(X) \neq \mathbf{a}'(X)$ then return $([a]_1, \mathbf{a}(X), \mathbf{a}'(X));$ if $\mathbf{b}(X) \neq \mathbf{b}'(X)$ then return $([b]_1, \mathbf{b}(X), \mathbf{b}'(X));$ if $\mathbf{c}(X) \neq \mathbf{c}'(X)$ then return $([c]_1, \mathbf{c}(X), \mathbf{c}'(X));$ endfor return $\perp;$

Fig. 17. The binding adversary $\mathcal{A}_{\text{bind}}$ in Theorem 6.

more, each $\mathbf{t}_{ijk}(X) := F_{ijk}(X)/Z_{\mathbb{H}}(X)$ is a polynomial, where $F_{ijk}(X) := F_0(X) + \alpha_{ijk}F_{1ij}(X) + \alpha_{ijk}^2F_{2ij}(X)$, and $F_0(X)$, $F_{1ij}(X)$ and $F_{2ij}(X)$ are defined as in Eq. (7) (but they may depend on β_i and γ_{ij}). But then $F_0(X) + \alpha_{ijk}F_{1ij}(X) + \alpha_{ijk}^2F_{2ij}(X) = Z_{\mathbb{H}}(X)\mathbf{t}_{ijk}(X)$. Thus, for every $s \in [1, n]$, $F_0(\omega^s) + \alpha_{ijk}F_{1ij}(\omega^s) + \alpha_{ijk}^2F_{2ij}(\omega^s) = 0$. Let

$$\mathbf{A}_{ij} = \begin{pmatrix} 1 & \alpha_{ij1} & \alpha_{ij1}^2 \\ 1 & \alpha_{ij2} & \alpha_{ij2}^2 \\ 1 & \alpha_{ij3} & \alpha_{ij3}^2 \end{pmatrix}.$$

be a Vandermonde matrix. Then, $\mathbf{A}_{ij} \cdot (F_0(\omega^s), F_{1ij}(\omega^s), F_{2ij}(\omega^s))^{\top} = 0$. Since α_{ij1} , α_{ij2} , and α_{ij3} are distinct, \mathbf{A}_{ij} is invertible. Thus, $F_0(\omega^s) = F_{1ij}(\omega^s) = F_{2ij}(\omega^s) = 0$. We analyze these three equalities individually.

$F_0(\omega^s) = 0$. Let $\bar{w}_s := \mathbf{a}(\omega^s)$, $\bar{w}_{n+s} := \mathbf{b}(\omega^s)$, and $\bar{w}_{2n+s} := \mathbf{c}(\omega^s)$. Then, $F_0(\omega^s) = 0$ iff $\mathbf{q}_{M_s}\bar{w}_s\bar{w}_{n+s} + \mathbf{q}_{L_s}\bar{w}_s + \mathbf{q}_{R_s}\bar{w}_{n+s} + \mathbf{q}_{O_s}\bar{w}_{2n+s} + \mathbf{q}_{C_s} + \text{Pl}(\omega^s) = 0$. For $s > \ell$, $\text{Pl}(\omega^s) = 0$ and we obtain the constraint in Eq. (5). Recall that for $s \leq \ell$, $\mathbf{q}_{M_s} = \mathbf{q}_{R_s} = \mathbf{q}_{O_s} = \mathbf{q}_{C_s} = 0$ and $\mathbf{q}_{L_s} = -1$ (see Eq. (4)). Thus, $-\bar{w}_s + \text{Pl}(\omega^s) = -\bar{w}_s + w_s = 0$. It follows that $w_s = \bar{w}_s$ for $1 \leq s \leq \ell$. Hence, $\mathbf{a}(X)$ encoded the public statement \mathbf{x} correctly.

$F_{2ij}(\omega^s) = 0$. We get $F_{2ij}(\omega^s) = (\mathbf{z}_{ij}(\omega^s) - 1)L_1(\omega^s) = 0$ for all $\omega^s \in \mathbb{H}$. Thus, $\mathbf{z}_{ij}(\omega) = 1$.

$F_{1ij}(\omega^s) = 0$. We will conclude from $F_{1ij}(\omega^s) = 0$ that $w_j = w_{\sigma(j)}$ for all $j \in [1, 3n]$. We will first prove a warm-up lemma (Lemma 4), which we will later expand to a more technical result Lemma 5 that better suits our needs.

Lemma 4. *Let σ be a permutation on $[1, n]$ and $a_1, \dots, a_n, b_1, \dots, b_n \in \mathbb{F}$. Let $\beta_1, \dots, \beta_{n+1} \in \mathbb{F}$ be mutually distinct and $\gamma_1, \dots, \gamma_{n+1} \in \mathbb{F}$ be mutually distinct. If $\prod_{s=1}^n (a_s + \beta_i \omega^s + \gamma_j) = \prod_{s=1}^n (b_s + \beta_i \omega^{\sigma(s)} + \gamma_j)$ for all $i, j \in [1, n+1]$, then $b_s = a_{\sigma(s)}$ for all $s \in [1, n]$.*

Proof. Consider the polynomials

$$f(Y, Z) := \prod_{s=1}^n (a_s + \omega^s Y + Z)$$

and

$$g(Y, Z) := \prod_{s=1}^n (b_s + \omega^{\sigma(s)} Y + Z) ;$$

the degree of Y or Z in both f and g is at most n . Denote $B := \{\beta_s\}_{s=1}^{n+1}$ and $\Gamma := \{\gamma_s\}_{s=1}^{t+1}$.

Define the Lagrange polynomials $L_s^B(Y) := \prod_{i \neq s} \frac{Y - \beta_i}{\beta_s - \beta_i}$ and $L_s^\Gamma(Z) := \prod_{k \neq s} \frac{Z - \gamma_k}{\gamma_s - \gamma_k}$ for $s = 1, \dots, n+1$. Clearly, $\{L_i^B(Y) \cdot L_j^\Gamma(Z)\}_{ij}$ is a basis of bi-variate polynomials where each variable has at most degree n . Thus, we can express f and g uniquely as

$$\begin{aligned} f(Y, Z) &:= \sum_{i=1}^{n+1} \sum_{j=1}^{n+1} f_{ij} L_i^B(Y) L_j^\Gamma(Z) , \\ g(Y, Z) &:= \sum_{i=1}^{n+1} \sum_{j=1}^{n+1} g_{ij} L_i^B(Y) L_j^\Gamma(Z) , \end{aligned}$$

for some $f_{ij}, g_{ij} \in \mathbb{F}$. Since by the hypothesis of the lemma, $f(\beta_i, \gamma_j) = f_{ij} = g(\beta_i, \gamma_j) = g_{ij}$ for all $i, j \in [1, n+1]$, it follows that $f(X, Y) = g(X, Y)$.

Observe that the polynomials $a_i + \omega^i Y + Z$ and $b_s + \omega^{\sigma(s)} Y + Z$ are irreducible. Thus, $f(X, Y) = g(X, Y)$ implies that for every s there exists exactly one i such that

$$a_s + \omega^s Y + Z = b_i + \omega^{\sigma(i)} Y + Z .$$

Thus, $\omega^i = \omega^{\sigma(s)}$, which implies that $i = \sigma(s)$, which in turn implies that $a_i = a_{\sigma(s)} = b_s$. The result follows since the above holds for all $s \in [1, n]$. \square

Next, we prove a more involved version of Lemma 4, that directly applies to Plonk. Let us first define the polynomials $f_\vartheta(Y, Z) := \prod_{s=1}^n (w_{\vartheta n+s} + k_\vartheta \omega^s Y + Z)$ and $g_\vartheta(Y, Z) := \prod_{s=1}^n (w_{\vartheta n+s} + S_{\sigma,(\vartheta+1)}(\omega^s) Y + Z)$. for $\vartheta \in \{0, 1, 2\}$, where $k_0 := 1$ and k_1, k_2 are defined as in Section 6.1.

Lemma 5. *If $\prod_{\vartheta=0}^2 f_\vartheta(\beta_i, \gamma_{ij}) = \prod_{\vartheta=0}^2 g_\vartheta(\beta_i, \gamma_{ij})$ for all $i, j \in [1, 3n+1]$, then $w_s = w_{\sigma(s)}$ for all $s \in [1, 3n]$.*

Proof. Let $f(Y, Z) := \prod_{\vartheta=0}^2 f_\vartheta(Y, Z)$ and $g(Y, Z) := \prod_{\vartheta=0}^2 g_\vartheta(Y, Z)$. Both f and g have at most degree $3n$ in both Y and Z . Using the same reasoning as in Lemma 4, we conclude that $f(Y, Z) = g(Y, Z)$. The polynomials $w_{\vartheta n+s} + k_\vartheta \omega^s Y + Z$ are irreducible and pairwise distinct for all $\vartheta \in \{0, 1, 2\}$ and $s \in [1, n]$. The same holds for the polynomials $w_{\vartheta n+s} + S_{\sigma,(\vartheta+1)}(\omega^s) Y + Z$.

Recall that $\sigma^*(i) = \mathbb{H}'[\sigma(i)]$ for all $i \in [1, 3n]$ and $S_{\sigma, (\vartheta+1)}(\omega^s) = \sigma^*(\vartheta n + s) = \mathbb{H}'[\sigma(\vartheta n + s)]$. Therefore, we can express

$$f(Y, Z) = \prod_{j=1}^{3n} (w_i + Y\mathbb{H}'[i] + Z), \quad g(Y, Z) = \prod_{i=1}^{3n} (w_i + Y\mathbb{H}'[\sigma(i)] + Z).$$

Just as in Lemma 4, each factor $w_i + Y\mathbb{H}'[\sigma(i)] + Z$ of f is equal to exactly one factor $w_{i'} + Y\mathbb{H}'[i'] + Z$. Thus, $\mathbb{H}'[\sigma(i)] = \mathbb{H}'[i']$ and $w_i = w_{i'}$. The first identity implies that $\sigma(i) = i'$ and the second implies that $w_i = w_{i'} = w_{\sigma(i)}$. Since this holds for all $i \in [1, 3n]$, we have proven the lemma. \square

We prove one more small result before proving our main result.

Lemma 6. *For all s, i, j ,*

$$(w_s + \beta_i S_{\sigma 1}(\omega^s) + \gamma_{ij}) \cdot (w_{n+s} + \beta_i S_{\sigma 2}(\omega^s) + \gamma_{ij}) \cdot (w_{2n+s} + \beta_i S_{\sigma 3}(\omega^s) + \gamma_{ij}) \neq 0.$$

Proof. Consider $w_s + \beta_i S_{\sigma 1}(\omega^s) + \gamma_{ij}$ as an example. We already noted that $w_s + S_{\sigma 1}(\omega^s)Y + Z$ is an irreducible polynomial, which means it has no roots in \mathbb{F} . Thus, (β_i, γ_{ij}) is not a root and $w_s + \beta_i S_{\sigma 1}(\omega^s) + \gamma_{ij} \neq 0$. Similarly, the other two factors are non-zero. Hence, their product is non-zero. \square

We will inductively show that

$$z_{ij}(\omega^{s+1}) = \prod_{t=1}^s \frac{(w_t + \beta_i \omega^t + \gamma_{ij}) \cdot (w_{n+t} + \beta_i k_1 \omega^t + \gamma_{ij}) \cdot (w_{2n+t} + \beta_i k_2 \omega^t + \gamma_{ij})}{(w_t + \beta_i S_{\sigma 1}(\omega^t) + \gamma_{ij}) \cdot (w_{n+t} + \beta_i S_{\sigma 2}(\omega^t) + \gamma_{ij}) \cdot (w_{2n+t} + \beta_i S_{\sigma 3}(\omega^t) + \gamma_{ij})}. \quad (23)$$

Since we already showed that $z_{ij}(\omega) = 1$, the claim holds for $s = 0$. Suppose the statement holds for $z_{ij}(\omega^s)$. From $F_{1ij}(\omega^i) = 0$ (see Eq. (7)), we conclude

$$z_{ij}(\omega^{s+1}) = \frac{z_{ij}(\omega^s) \cdot (w_s + \beta_i \omega^s + \gamma_{ij}) \cdot (w_{n+s} + \beta_i k_1 \omega^s + \gamma_{ij}) \cdot (w_{2n+s} + \beta_i k_2 \omega^s + \gamma_{ij})}{(w_s + \beta_i S_{\sigma 1}(\omega^s) + \gamma_{ij}) \cdot (w_{n+s} + \beta_i S_{\sigma 2}(\omega^s) + \gamma_{ij}) \cdot (w_{2n+s} + \beta_i S_{\sigma 3}(\omega^s) + \gamma_{ij})}.$$

Note that the division is well-defined according to Lemma 6 and Eq. (23) follows by expanding $z_{ij}(\omega^s)$.

Since $\omega^{n+1} = \omega$, we have $z_{ij}(\omega^{n+1}) = 1$, implying that $\prod_{t=1}^n (w_t + \beta_i \omega^t + \gamma_{ij}) \cdot (w_{n+t} + \beta_i k_1 \omega^t + \gamma_{ij}) \cdot (w_{2n+t} + \beta_i k_2 \omega^t + \gamma_{ij}) = \prod_{t=1}^n (w_t + \beta_i S_{\sigma 1}(\omega^t) + \gamma_{ij}) \cdot (w_{n+t} + \beta_i S_{\sigma 2}(\omega^t) + \gamma_{ij}) \cdot (w_{2n+t} + \beta_i S_{\sigma 3}(\omega^t) + \gamma_{ij})$. We can conclude from Lemma 5 that $w_i = w_{\sigma(i)}$ for all $i \in [1, 3n]$. Hence, w also satisfies the final constraint in Eq. (6). \square

D.6 Fiat-Shamir Transform

Recall the following theorem from [AFK22].

Theorem 10 ([AFK22]). *Let Π be a $(\kappa_1, \dots, \kappa_\mu)$ -out-of- (N_1, \dots, N_μ) -special-sound interactive proof. Then, the Fiat-Shamir transformation $\text{FS}[\Pi]$ of Π is knowledge-sound with knowledge error*

$$\text{Er}_{\text{fs}}(Q) = (Q + 1) \cdot \text{Er},$$

where Q is the number of random oracle queries the adversary makes and $\text{Er} = 1 - \prod_{i=1}^{\mu} \left(1 - \frac{\kappa_i - 1}{N_i}\right)$.

Importantly, in all our security proofs, k_i/N_i is negligible for all i , resulting in a negligible $\text{Er}_{\text{fs}}(Q)$. This holds since all verifier challenges (including the evaluation point \mathfrak{z}), used as branches in the transcript tree, in **SanPlonk** and **Plonk** are chosen randomly from \mathbb{F} (or, a large subset of \mathbb{F}); thus, $N_i \approx |\mathbb{F}|$ and $k_i/N_i = \text{negl}(\lambda)$.

While Theorem 10 applies to proof systems, it also extends to argument systems as explained in [AFKR23, Remark 1]. This is because the interactive argument system Π can be seen as a proof of knowledge for a slightly different relation: the knowledge of a witness for the underlying relation OR a solution to some computationally hard problem (in some cases more than one hard problem). Note that all of the computational special-soundness proofs in this paper showed that there exists a DPT extractor Ext_{ss} and a PPT \mathcal{A} such that given an accepting transcript tree \mathcal{T} , either Ext_{ss} outputs a witness or \mathcal{A} outputs a solution to some hard problem. The proof system's special soundness extractor for the OR relation runs internally both Ext_{ss} and \mathcal{A} on \mathcal{T} . It returns the witness if Ext_{ss} returns the witness and otherwise returns the output of \mathcal{A} . Applying now the Fiat-Shamir transform, we obtain a non-interactive knowledge-sound proof system $\text{FS}[\Pi]$ for the OR-relation with the knowledge error $\text{Er}_{\text{fs}}(Q)$ as described in Theorem 10. Furthermore, the obtained $\text{FS}[\Pi]$ is also an argument system for the original relation since we can reduce the security to the underlying assumption with a factor $\text{Er}_{\text{fs}}(Q)$ loss.

E Zero-Knowledge of **SanPlonk**

We recall the zero-knowledge property of a non-interactive argument.⁸ Below $\mathcal{UR}_{\mathbf{p},n}$ is a family of binary relations parameterized by the system parameters \mathbf{p} and an integer n . For $\mathcal{R} \in \mathcal{UR}_{\mathbf{p},n}$ and $\mathbf{srs} \in \text{range}(\text{KGen}(\mathbf{p}, n))$.

A non-interactive argument is (statistical) *zero-knowledge* if there exists a PPT simulator Sim , s.t. for all unbound $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, all $\mathbf{p} \in \text{range}(\text{Pgen})$, all $n \in \text{poly}(\lambda)$,

$$\Pr \left[\begin{array}{l} \mathcal{A}_2(st, \pi) = 1 \wedge \\ \mathcal{R}(\mathbf{x}, \mathbf{w}) \wedge \mathcal{R} \in \mathcal{UR}_{\mathbf{p},n}; \end{array} \middle| \begin{array}{l} (\mathbf{srs}, \text{td}_{\mathbf{srs}}) \leftarrow \text{KGen}(\mathbf{p}, n); \\ (\mathcal{R}, \mathbf{x}, \mathbf{w}, st) \leftarrow \mathcal{A}_1(\mathbf{srs}); \\ \pi \leftarrow \text{P}(\mathcal{R}, \mathbf{srs}, \mathbf{x}, \mathbf{w}) \end{array} \right] \approx_s \Pr \left[\begin{array}{l} \mathcal{A}_2(st, \pi) = 1 \wedge \\ \mathcal{R}(\mathbf{x}, \mathbf{w}) \wedge \mathcal{R} \in \mathcal{UR}_{\mathbf{p},n} \end{array} \middle| \begin{array}{l} (\mathbf{srs}, \text{td}_{\mathbf{srs}}) \leftarrow \text{KGen}(\mathbf{p}, n); \\ (\mathcal{R}, \mathbf{x}, \mathbf{w}, st) \leftarrow \mathcal{A}_1(\mathbf{srs}); \\ \pi \leftarrow \text{Sim}(\mathcal{R}, \mathbf{srs}, \text{td}_{\mathbf{srs}}, \mathbf{x}) \end{array} \right].$$

Here, \approx_s denotes the statistical distance as a function of λ . Π is *perfect zero-knowledge* if the above probabilities are equal.

⁸ We proved the special soundness for interactive arguments, and since then, we could use Theorem 10 (which has a pretty complicated proof) to obtain knowledge-soundness for non-interactive arguments. For zero-knowledge, it is easy to prove a direct result for non-interactive arguments, so we do that.

Recently, Sefranek [Sef24] showed that an earlier version of Plonk did not satisfy statistical zero-knowledge since the polynomials $\mathbf{t}_{\text{lo}}(X)$, $\mathbf{t}_{\text{mid}}(X)$, $\mathbf{t}_{\text{hi}}(X)$ were not randomized. He corrects this mistake and proves statistical zero-knowledge of the corrected Plonk. The correction is a part of the Plonk now, [GWC19]. Sefranek also mentions that **SmallPlonk**, which does not split $\mathbf{t}(X)$, is simulatable. Our paper contains the current (corrected) version of Plonk, so we will not repeat the zero-knowledge proofs of Plonk and **SmallPlonk** and instead refer the reader to [Sef24]. However, we provide a similar proof of zero-knowledge for **SanPlonk**.

We recall the following well-known lemma; see, e.g., [Sef24].

Lemma 7 ([Sef24]). *Let $f(X) \in \mathbb{F}[X]$ and x_1, \dots, x_k be distinct values in $\mathbb{F} \setminus \mathbb{H}$. Assume $f(X) = \mathbf{f}(X) + \varrho(X)Z_{\mathbb{H}}(X)$ for $\varrho(X) \leftarrow_{\mathbb{S}} \mathbb{F}_{\leq k-1}[X]$. Then, $(\tilde{f}(x_1), \dots, \tilde{f}(x_k))$ is distributed uniformly over \mathbb{F}^k ,*

The same claim also holds for $\tilde{f}(X) = \mathbf{f}(X) + \varrho(X)$; moreover, then it is true even when $x_1, \dots, x_k \in \mathbb{F}$, not just in $\mathbb{F} \setminus \mathbb{H}$ (in Lemma 7, we need that $Z_{\mathbb{H}}(x_i) \neq 0$). We use both results to prove the statistical zero-knowledge of **SanPlonk**.

Theorem 11. *SanPlonk has statistical zero-knowledge.*

Proof. Consider the following simulation strategy. Recall that **SanPlonk**'s transcript is $([a, b, c]_1; \beta, \gamma; [z]_1; \alpha; [t_{\text{lo}}, t_{\text{mid}}, t_{\text{hi}}]_1; \mathfrak{z}; \bar{a}, \bar{b}, \bar{c}, \bar{s}_{\sigma 1}, \bar{s}_{\sigma 2}, \bar{z}_{\omega}; \delta; \bar{t}_3; v; [W_3, W_{3\omega}]_1)$. Similarly to Plonk's simulator in [Sef24], given the verifier's random challenges, **SanPlonk**'s simulator works like an honest prover with four crucial differences. First, it creates the commitments a, b, c, z and their openings by sampling them randomly. Second, since $\mathbf{t}(X)$, $W_3(X)$, $W_{3\omega}(X)$ might not be polynomials, it computes $[\mathbf{t}(x), W_3(x), W_{3\omega}(x)]_1$ by using the trapdoor. Third, it computes $[t_{\text{lo}}, t_{\text{mid}}, t_{\text{hi}}]_1$ by sampling two of the values at random and setting the third one so that $[t_{\text{lo}}, t_{\text{mid}}, t_{\text{hi}}]_1$ agrees with the previously computed value of $[\mathbf{t}(x)]_1$. Fourth, the sanitization value \bar{t}_3 is sampled randomly. In Fig. 18, we present the full simulator for **SanPlonk** as a zk-SNARK, including the definition of the verifier's challenges as the outputs of the random oracle. (For brevity, we refer to Eq. (7) for the formulas of $F_i(X)$ and $F(X)$.)

From Lemma 7 it follows that in the honest proof $\mathbf{a}(x)$, $\mathbf{b}(x)$, $\mathbf{c}(x)$, $\mathbf{z}(x)$, $\mathbf{a}(\mathfrak{z})$, $\mathbf{b}(\mathfrak{z})$, $\mathbf{c}(\mathfrak{z})$, $\mathbf{z}(\mathfrak{z}\omega)$, and $\mathbf{z}(x\omega)$ are distributed uniformly randomly and independently. Here, the last element $\bar{z}_{x\omega} = \mathbf{z}(x\omega)$ is not part of the proof transcript, but it is necessary for determining a unique $[\mathbf{t}(x)]_1$. Furthermore, in the honest protocol $\mathbf{t}_{\text{hi}}(X) := \mathbf{t}_{\text{hi}}'(X) - b_{11}$, $\mathbf{t}_{\text{mid}}(X) := \mathbf{t}_{\text{mid}}'(X) - b_{10} - b_{12}X + b_{11}X^n$, and $\mathbf{t}_{\text{lo}}(X) := \mathbf{t}_{\text{lo}}'(X) + b_{10}X^n + b_{12}X^{n+1}$ (as always, we highlight the changes compared to Plonk). Thus, $\mathbf{t}_{\text{hi}}(x)$, $\mathbf{t}_{\text{mid}}(x)$, and $\mathbf{t}_{\text{lo}}(\mathfrak{z})$ are distributed uniformly at random and independently since (resp.) b_{11}, b_{10} , and b_{12} are sampled uniformly at random. In the case of $\mathbf{t}_{\text{lo}}(\mathfrak{z})$, it holds under the assumption that $\mathfrak{z}^{n+1} \neq 0$ (otherwise $b_{12}\mathfrak{z}^{n+1} = 0$). Note that the proof does not reveal $t_{\text{lo}}(\mathfrak{z})$. However, $t_{\text{lo}}(\mathfrak{z})$ being uniformly random and independent of other elements guarantees that $\bar{t}_3 = \mathbf{t}_{\text{lo}}(\mathfrak{z}) + \delta \mathbf{t}_{\text{mid}}(\mathfrak{z}) + \delta^2 \mathbf{t}_{\text{hi}}(\mathfrak{z})$ is uniformly random and independent. In the simulator, we also pick all the mentioned random elements uniformly at random and independently.

1. $a, b, c \leftarrow_{\mathbb{S}} \mathbb{F}$.
2. $\beta \leftarrow H(\mathbf{srs}, \mathbb{X}, [a, b, c]_1, 0)$; $\gamma \leftarrow H(\mathbf{srs}, \mathbb{X}, [a, b, c]_1, 1)$.
3. $z \leftarrow_{\mathbb{S}} \mathbb{F}$.
4. $\alpha \leftarrow H(\mathbf{srs}, \mathbb{X}, [a, b, c, z]_1)$.
5. Sample $\bar{z}_{x\omega} \leftarrow_{\mathbb{S}} \mathbb{F}$ (a candidate value for $z(x\omega)$ used to compute $F_1(x)$ in the next step).
6. For $F_i(X)$ defined as in Eq. (7), set $[t(x)]_1 \leftarrow \frac{1}{Z_{\mathbb{H}}(x)} [F_0(x) + \alpha F_1(x) + \alpha^2 F_2(x)]_1$.
7. $t_{hi}, t_{mid} \leftarrow_{\mathbb{S}} \mathbb{F}$; $[t_{lo}]_1 \leftarrow [t(x) - t_{mid}x^n - t_{hi}x^{2n}]_1$. Abort if $Z_{\mathbb{H}}(x) = 0$.
8. $\mathfrak{z} \leftarrow H(\mathbf{srs}, \mathbb{X}, [a, b, c, z, t_{lo}, t_{mid}, t_{hi}]_1)$.
9. Abort if $\mathfrak{z} = x$ or $\mathfrak{z}\omega = x$.
10. $\bar{a}, \bar{b}, \bar{c}, \bar{z}_{\omega} \leftarrow_{\mathbb{S}} \mathbb{F}$; $\bar{s}_{\sigma_1} \leftarrow S_{\sigma_1}(\mathfrak{z})$; $\bar{s}_{\sigma_2} \leftarrow S_{\sigma_2}(\mathfrak{z})$.
11. $\delta \leftarrow H(\mathbf{srs}, \mathbb{X}, [a, b, c, z, t_{lo}, t_{mid}, t_{hi}]_1, \bar{a}, \bar{b}, \bar{c}, \bar{s}_{\sigma_1}, \bar{s}_{\sigma_2}, \bar{z}_{\omega})$.
12. $\bar{t}_3 \leftarrow_{\mathbb{S}} \mathbb{F}$.
13. $v \leftarrow H(\mathbf{srs}, \mathbb{X}, [a, b, c, z, t_{lo}, t_{mid}, t_{hi}]_1, \bar{a}, \bar{b}, \bar{c}, \bar{s}_{\sigma_1}, \bar{s}_{\sigma_2}, \bar{z}_{\omega}, \bar{t}_3)$.
14. $W \leftarrow r + v(a - \bar{a}) + v^2(b - \bar{b}) + v^3(c - \bar{c}) + v^4(s_{\sigma_1} - \bar{s}_{\sigma_1}) + v^5(s_{\sigma_2} - \bar{s}_{\sigma_2})$.
15. $[W_3]_1 \leftarrow \frac{1}{x - \bar{z}_3} [W + v^6(t_{lo} + \delta t_{mid} + \delta^2 t_{hi} - \bar{t}_3)]_1$.
16. $[W_{3\omega}]_1 \leftarrow \frac{1}{x - \bar{z}_{3\omega}} [z - \bar{z}_{\omega}]_1$.
17. Return $([a, b, c]_1; \beta, \gamma; [z]_1; \alpha; [t_{lo}, t_{mid}, t_{hi}]_1; \mathfrak{z}; \bar{a}, \bar{b}, \bar{c}, \bar{s}_{\sigma_1}, \bar{s}_{\sigma_2}, \bar{z}_{\omega}; \delta; \bar{t}_3; v; [W_3, W_{3\omega}]_1)$.

Fig. 18. SanPlonk's simulator.

The remaining proof elements have only one possible satisfiable value. Namely, there is precisely one possible value of $t_{lo}(x)$ such that $t_{lo}(x) + x^n t_{mid}(x) + x^{2n} t_{hi}(x) = t(x)$ and only one possible value for opening proofs $[W_3, W_{3\omega}]_1$ such that the verification equation is satisfied. The simulator computes these elements accordingly. The public polynomials are evaluated honestly as $\bar{s}_{\sigma_1} \leftarrow S_{\sigma_1}(\mathfrak{z})$ and $\bar{s}_{\sigma_2} \leftarrow S_{\sigma_2}(\mathfrak{z})$. Also, the challenges β, γ, \dots are computed from the correct distribution. The simulator fails when either $x^{n+1} = 0$, $\mathfrak{z} = x$, $\mathfrak{z}\omega = x$, or $Z_{\mathbb{H}}(x) = 0$ (the last three conditions are needed to avoid division by 0 in the simulator); this happens only with a negligible probability. \square