

# What Makes Fiat–Shamir zkSNARKs (Updatable SRS) Simulation Extractable?

Chaya Ganesh<sup>1</sup>, Hamidreza Khoshakhlagh<sup>2</sup>, Markulf Kohlweiss<sup>3</sup>,  
Anca Nitulescu<sup>4</sup>, and Michał Zając<sup>5</sup>

<sup>1</sup> Indian Institute of Science [chaya@iisc.ac.in](mailto:chaya@iisc.ac.in)

<sup>2</sup> Aarhus University [hamidreza@cs.au.dk](mailto:hamidreza@cs.au.dk)

<sup>3</sup> University of Edinburgh and IOHK [mkohlwei@inf.ed.ac.uk](mailto:mkohlwei@inf.ed.ac.uk)

<sup>4</sup> Protocol Labs [anca@protocol.ai](mailto:anca@protocol.ai)

<sup>5</sup> Nethermind [m.p.zajac@gmail.com](mailto:m.p.zajac@gmail.com)

**Abstract.** We show that three popular universal zero-knowledge SNARKs (Plonk, Sonic, and Marlin) are updatable SRS simulation extractable NIZKs and signatures of knowledge (SoK) out-of-the-box avoiding any compilation overhead.

Towards this we generalize results for the Fiat–Shamir (FS) transformation, which turns interactive protocols into signature schemes, non-interactive proof systems, or SoK in the random oracle model (ROM). The security of the transformation relies on rewinding to extract the secret key or the witness, even in the presence of signing queries for signatures and simulation queries for proof systems and SoK, respectively. We build on this line of work and analyze multi-round FS for arguments with a structured reference string (SRS). The combination of ROM and SRS, while redundant in theory, is the model of choice for the most efficient practical systems to date. We also consider the case where the SRS is updatable and define a strong simulation extractability notion that allows for simulated proofs with respect to an SRS to which the adversary can contribute updates.

We define three properties (trapdoor-less zero-knowledge, rewinding-based knowledge soundness, and a unique response property) that are sufficient for argument systems based on multi-round FS to be also simulation extractable in this strong sense. We show that Plonk, Sonic, and Marlin satisfy these properties, and conjecture that many other argument systems such as Lunar, Basilisk, and transparent variants of Plonk fall within the reach of our main theorem.

## 1 Introduction

Zero-knowledge proof systems, which allow a prover to convince a verifier of an NP statement  $\mathbf{R}(x, w)$  without revealing anything else about the witness  $w$  have broad application in cryptography and theory of computation [9, 28, 34]. When restricted to computationally sound proof systems, also called *argument systems*<sup>6</sup>, proof size can be shorter than the size of the witness [18]. Zero-knowledge Succinct Non-interactive ARGuments of Knowledge (zkSNARKs) are zero-knowledge argument systems that additionally have two succinctness properties: small proof sizes and fast verification. Since their introduction in [47], zk-SNARKs have been a versatile design tool for secure cryptographic protocols. They became particularly relevant for blockchain applications that demand short proofs and fast verification for on-chain storage and processing. Starting with their deployment by Zcash [11], they have seen broad adoption, e.g., for privacy-preserving cryptocurrencies and scalable and private smart contracts in Ethereum.

While research on zkSNARKs has seen rapid progress [12, 14, 15, 32, 37, 38, 43, 44, 49] with many works proposing significant improvements in proof size, verifier and prover efficiency, and complexity of the public setup, less attention has been paid to non-malleable zkSNARKs and succinct signatures of knowledge [20, 22] (sometimes abbreviated SoK or referred to as SNARKY signatures [6, 40]).

*Relevance of simulation extractability.* Most zkSNARKs are shown only to satisfy a standard knowledge soundness property. Intuitively, this guarantees that a prover that creates a valid proof in isolation knows a valid witness. However, deployments of zkSNARKs in real-world applications, unless they are carefully designed to have application-specific malleability protection, e.g. [11], require a stronger property – *simulation-extractability* (SE) – that corresponds much more closely to existential unforgeability of signatures.

This correspondence is made precise by SoK, which use an NP-language instance as the public verification key. Instead of signing with the secret key, SoK signing requires knowledge of the NP-witness. Intuitively, an SoK is thus a proof of knowledge (PoK) of a witness that is tied to a message. In fact, many signatures schemes, e.g.,

<sup>6</sup> We use both terms interchangeably.

Schnorr, can be read as SoK for a specific hard relation, e.g., DL [25]. To model strong existential unforgeability of SoK signatures, even when given an oracle for obtaining signatures on different instances, an attacker must not be able to produce new signatures. Chase and Lysyanskaya [22] model this via the notion of simulation extractability which guarantees extraction of a witness even in the presence of simulated signatures.

In practice, an adversary against a zkSNARK system also has access to proofs computed by honest parties that should be modeled as simulated proofs. The definition of knowledge soundness (KS) ignores the ability of an adversary to see other valid proofs that may occur in real-world applications. For instance, in applications of zkSNARKs in privacy-preserving blockchains, proofs are posted on-chain for all blockchain participants to see. We thus argue that SE is a much more suitable notion for robust protocol design. We also claim that SE has primarily an intellectual cost, as it is harder to prove SE than KS—another analogy here is IND-CCA vs IND-CPA security for encryption. However, we will show that the proof systems we consider are SE out-of-the-box.

*Fiat–Shamir-based zkSNARKs.* Most modern zkSNARK constructions follow a modular blueprint that involves the design of an information-theoretic interactive protocol, e.g. an Interactive Oracle Proof (IOP) [13], that is then compiled via cryptographic tools to obtain an interactive argument system. This is then turned into a zkSNARK using the Fiat-Shamir transform. By additionally hashing the message, the Fiat-Shamir transform is also a popular technique for constructing signatures. While well-understood for 3-message sigma protocols and justifiable in the ROM [8], Fiat–Shamir should be used with care because there are both counterexamples in theory [35] and real-world attacks in practice when implemented incorrectly [48].

In particular, several schemes such as Sonic [46], Plonk [30], Marlin [23] follow this approach where the information-theoretic object is a multi-message algebraic variant of IOP, and the cryptographic primitive in the compiler is a polynomial commitment scheme (PC) that requires a trusted setup. To date, this blueprint lacks an analysis in the ROM in terms of simulation extractability.

*Updatable SRS zkSNARKs.* One of the downsides of many efficient zkSNARKs [24, 32, 37, 38, 43, 44, 49] is that they rely on a *trusted setup*, where there is a structured reference string (SRS) that is assumed to be generated by a trusted party. In practice, however, this assumption is not well-founded; if the party that generates the SRS is not honest, they can produce proofs for false statements. If the trusted setup assumption does not hold, knowledge soundness breaks down. Groth et al. [39] propose a setting to tackle this challenge which allows parties – provers and verifiers – to *update* the SRS.<sup>7</sup> The update protocol takes an existing SRS and contributes to its randomness in a verifiable way to obtain a new SRS. The guarantee in this *updatable setting* is that knowledge soundness holds as long as one of the parties updating the SRS is honest. The SRS is also *universal*, in that it does not depend on the relation to be proved but only on an upper bound on the size of the statement’s circuit. Although inefficient, as the SRS size is quadratic in the size of the circuit, [39] set a new paradigm for designing zkSNARKs.

The first universal zkSNARK with updatable and linear size SRS was Sonic proposed by Maller et al. in [46]. Subsequently, Gabizon, Williamson, and Ciobotaru designed Plonk [30] which currently is the most efficient updatable universal zkSNARK. Independently, Chiesa et al. [23] proposed Marlin with comparable efficiency to Plonk.

*The challenge of SE in the updatable setting.* The notion of simulation-extractability for zkSNARKs which is well motivated in practice, has not been studied in the updatable setting. Consider the following scenario: We assume a “rushing” adversary that starts off with a sequence of updates by malicious parties resulting in a subverted reference string  $srs$ . By combining their trapdoor contributions and employing the simulation algorithm, these parties can easily compute a proof to obtain a triple  $(srs, x, \pi)$  that convinces the verifier of a statement  $x$  without knowing a witness. Now, assume that at a later stage, a party produces a triple  $(srs', x, \pi')$  for the same statement with respect to an updated  $srs'$  that has an honest update contribution. We want the guarantee that this party must know a witness corresponding to  $x$ . The ability to “maul” the proof  $\pi$  from the old SRS to a proof  $\pi'$  for the new SRS without knowing a witness would clearly violate security. The natural idea is to require that honestly *updated* reference strings are indistinguishable from honestly *generated* reference strings even for parties that previously contributed updates. However, this is not sufficient as the adversary can also rush toward the end of the SRS generation ceremony to perform the last update.

A definition of SE in the updatable setting should take these additional powers of the adversary, which are not captured by existing definitions of SE, into consideration. While generic compilers [2, 42] can be applied to updatable SRS SNARKs to obtain SE, not only do they inevitably incur overheads and lead to efficiency loss, we contend that the standard definition of SE does not suffice in the updatable setting.

---

<sup>7</sup> This can be seen as an efficient player-replaceable [33] multi-party computation.

## 1.1 Our Contributions

We investigate the non-malleability properties of zkSNARK protocols obtained by FS-compiling multi-message protocols in the updatable SRS setting and give a modular approach to analyze their simulation-extractability. We make the following contributions:

- *Updatable simulation extractability (USE)*. We propose a definition of simulation extractability in the updatable SRS setting called USE, that captures the additional power the adversary gets by being able to update the SRS.
- *Theorem for USE of FS-compiled proof systems*. We define three notions in the updatable SRS and ROM model, *trapdoor-less zero-knowledge*, a *unique response* property, and *rewinding-based knowledge soundness*. Our main theorem shows that multi-message FS-compiled proof systems that satisfy these notions *are USE out-of-the box*.
- *USE for concrete zkSNARKs*. We prove that the most efficient updatable SRS SNARKS – Plonk/Sonic/Marlin – satisfy the premises of our theorem. We thus show that these zkSNARKs are updatable simulation extractable.
- *SNARKY signatures in the updatable setting*. Our results validate the folklore that the Fiat–Shamir transform is a natural means for constructing signatures of knowledge. This gives rise to the first SoK in the updatable setting and confirms that a much larger class of zkSNARKs, besides [40], can be lifted to SoK.
- *Broad applicability*. The updatable SRS plus ROM model includes both the trusted SRS and the ROM model as special cases. This implies the relevance of our theorem for transparent zkSNARKs such as Halo2 and Plonky2 that replace the polynomial commitments of Kate et al. [41] with commitments from Bulletproof [19] and STARKs [10], respectively.

## 1.2 Technical Overview

At a high level, the proof of our main theorem for updatable simulation extractability is along the lines of the simulation extractability proof for FS-compiled sigma protocol from [26]. However, our theorem introduces new notions that are more general to allow us to consider proof systems that are richer than sigma protocols and support an updatable setup. We discuss some of the technical challenges below.

Plonk, Sonic, and Marlin were originally presented as interactive proofs of knowledge that are made non-interactive via the Fiat–Shamir transform. In the following, we denote the underlying interactive protocols by  $\mathbf{P}$  (for Plonk),  $\mathbf{S}$  (for Sonic), and  $\mathbf{M}$  (for Marlin) and the resulting non-interactive proof systems by  $\mathbf{P}_{\text{FS}}$ ,  $\mathbf{S}_{\text{FS}}$ ,  $\mathbf{M}_{\text{FS}}$  respectively.

**Rewinding-Based Knowledge Soundness (RBKS)**. Following [26], one would have to show that for the protocols we consider, a witness can be extracted from sufficiently many valid transcripts with a common prefix. The standard definition of special soundness for sigma protocols requires the extraction of a witness from any two transcripts with the same first message. However, most zkSNARK protocols do not satisfy this notion. We put forth a notion analogous to special soundness that is more general and applicable to a wider class of protocols. Namely, protocols compiled using multi-round FS that rely on an (updatable) SRS.  $\mathbf{P}$ ,  $\mathbf{S}$ , and  $\mathbf{M}$  have more than three messages, and the number of transcripts required for extraction is more than two. Concretely,  $(3n + 6)$  for Plonk,  $(n + 1)$  for Sonic, and  $(2n + 3)$  for Marlin, where  $n$  is the number of constraints in the proven circuit. Hence, we do not have a pair of transcripts but a *tree of transcripts*.

Furthermore, the protocols we consider are arguments and rely on a SRS that comes with a trapdoor. An adversary in possession of the trapdoor can produce multiple valid proof transcripts potentially for false statements without knowing any witness. This is true even in the updatable setting, where a trapdoor still exists for any updated SRS. Recall that the standard special soundness definition requires witness extraction from *any* suitably structured tree of accepting transcripts. This means that there are no such trees for false statements.

Instead, we give a rewinding-based knowledge soundness definition with an extractor that proceeds in two steps. It first uses a tree building algorithm  $\mathcal{T}$  to obtain a tree of transcripts. In the second step, it uses a tree extraction algorithm  $\text{Ext}_{\text{ks}}$  to compute a witness from this tree. Tree-based knowledge soundness guarantees that it is possible to extract a witness from all (but negligibly many) trees of accepting transcripts produced by probabilistic polynomial time (PPT) adversaries. That is, if extraction from such a tree fails, then we break an underlying computational assumption. Moreover, this should hold even against adversaries that contribute to the SRS generation.

**Unique Response Protocols (UR)**. Another property required to show simulation extractability is the unique response property which says that for 3-message sigma protocols, the response of the prover (3-rd message) is determined by the first message and the challenge [27] (intuitively, the prover can only employ fresh randomness in the first message of the protocol). We cannot use this definition since the protocols we consider have multiple

rounds of randomized prover messages. In Plonk, both the first and the third messages are randomized. Although the Sonic prover is deterministic after it picks its first message, the protocol has more than 3 messages. The same holds for Marlin. We propose a generalization of the unique response property called  $k$ -UR. It requires that the behavior of the prover be determined by the first  $k$  of its messages. For our proof, it is sufficient that Plonk is 3-UR, and Sonic and Marlin are 2-UR.

**Trapdoor-Less Zero-Knowledge (TLZK).** The premises of our main theorem include two computational properties that do not mention a simulator, RBKS and UR. The theorem states that together with a suitable property for the simulator of the zero-knowledge property, they imply USE. Our key technique is to simulate simulation queries when reducing to RBKS and UR. For this it is convenient that the zero-knowledge simulator be trapdoor-less, that is can produce proofs without relying on the knowledge of the trapdoor. Simulation is based purely on the simulator's early control over the challenge. In the ROM this corresponds to a simulator that programs the random oracle and can be understood as a generalization of honest-verifier zero-knowledge for multi-message Fiat–Shamir transformed proof systems with an SRS. We say that such a proof system is  $k$ -TLZK, if the simulator only programs the  $k$ -th challenge and we construct such simulators for  $\mathbf{P}_{\text{FS}}$ ,  $\mathbf{S}_{\text{FS}}$ , and  $\mathbf{M}_{\text{FS}}$ .

Technically we will make use of the  $k$ -UR property together with the  $k$ -TLZK property to bound the probability that the tree produced by the tree builder  $\mathcal{T}$  of RBKS contains any programmed random oracle queries.

### 1.3 Related Work

There are many results on simulation extractability for non-interactive zero-knowledge proofs (NIZKs). First, Groth [36] noticed that a (black-box) SE NIZK is universally-composable (UC) [21]. Then Dodis et al. [25] introduced a notion of (black-box) *true simulation extractability* (i.e., SE with simulation of true statements only) and showed that no NIZK can be UC-secure if it does not have this property.

In the context of zkSNARKs, the first SE zkSNARK was proposed by Groth and Maller [40] and a SE zkSNARK for QAP was designed by Lipmaa [45]. Kosba et al. [42] give a general transformation from a NIZK to a black-box SE NIZK. Although their transformation works for zkSNARKs as well, the succinctness of the proof system is not preserved by this transformation. Abdolmaleki et al. [2] showed another transformation that obtains non-black-box simulation extractability but also preserves the succinctness of the argument. The zkSNARK of [38] has been shown to be SE by introducing minor modifications to the construction and making stronger assumptions [3, 17]. Recently, [6] showed that the Groth's original proof system from [38] is weakly SE and randomizable. None of these results are for zkSNARKs in the updatable SRS setting or for zkSNARKs obtained via the Fiat–Shamir transformation. The recent work of [31] shows that Fiat–Shamir transformed Bulletproofs are simulation extractable. While they show a general theorem for multi-round protocols, they do not consider a setting with an SRS, and are therefore inapplicable to zkSNARKs in the updatable SRS setting.

## 2 Definitions and Lemmas for Multi-message SRS-based Protocols

*Simulation-extractability for multi-message protocols.* Most recent SNARK schemes follow the same blueprint of constructing an interactive information-theoretic proof system that is then compiled into a public coin computationally sound scheme using cryptographic tools such as polynomial commitments, and finally made non-interactive via the Fiat–Shamir transformation. Existing results on simulation extractability (for proof systems and signatures of knowledge) for Fiat–Shamir transformed systems work for 3-message protocols without reference string that require two transcripts for standard model extraction, e.g., [26, 50, 51].

In this section, we define properties that are necessary for our analysis of multi-message protocols with a universal updatable SRS. In order to prove simulation-extractability for such protocols, we require more than just two transcripts for extraction. Moreover, in the updatable setting we consider protocols that rely on an SRS where the adversary gets to contribute to the SRS. We first recall the updatable SRS setting and the Fiat-Shamir transform for  $(2\mu + 1)$  message protocols. Next, we define trapdoor-less zero-knowledge and simulation-extractability which we base on [26] adapted to the updatable SRS setting. Then, to support multi-message SRS-based protocols compiled using the Fiat–Shamir transform, we generalize the unique response property, and define a notion of computational special soundness called rewinding-based knowledge soundness.

Let  $P$  and  $V$  be PPT algorithms, the former called the *prover* and the latter the *verifier* of a proof system. Both algorithms take a pre-agreed structured reference string  $\text{srs}$  as input. The structured reference strings we consider are (potentially) updatable, a notion we recall shortly. We focus on proof systems made non-interactive via the multi-message Fiat–Shamir transform presented below where prover and verifier are provided with a random oracle  $\mathcal{H}$ . We denote by  $\pi$  a proof created by  $P$  on input  $(\text{srs}, x, w)$ . We say that proof is accepting if  $V(\text{srs}, x, \pi)$  accepts it.

```

UpdO(intent, srsn, {ρj}j=1n)
if srs ≠ ⊥ : return ⊥
if (intent = setup) :
  (srs', ρ') ← GenSRS(R)
  Qsrs ← Qsrs ∪ {(srs', ρ')}
  return (srs', ρ')

if (intent = update) :
  b ← VerifySRS(srsn, {ρj}j=1n)
  if (b = 0) : return ⊥
  (srs', ρ') ← UpdSRS(srsn, {ρj}j=1n)
  Qsrs ← Qsrs ∪ {(srs', ρ')}
  return (srs', ρ')

if (intent = final) :
  b ← VerifySRS(srsn, {ρj}j=1n)
  if (b = 0) ∨ Qsrs(2) ∩ {ρj}i = ∅ :
    return ⊥
    srs ← srsn, return srs
else return ⊥

```

Fig. 1: The oracle defines the notion of updatable SRS setup.

Let  $R(\mathcal{A})$  denote the set of random tapes of correct length for adversary  $\mathcal{A}$  (assuming the given value of security parameter  $\lambda$ ), and let  $r \leftarrow_{\$} R(\mathcal{A})$  denote the random choice of tape  $r$  from  $R(\mathcal{A})$ .

## 2.1 Updatable SRS Setup Ceremonies

The definition of updatable SRS ceremonies of [39] requires the following algorithms.

- $(srs, \rho) \leftarrow \text{GenSRS}(\mathbf{R})$  is a PPT algorithm that takes a relation  $\mathbf{R}$  and outputs a reference string  $srs$ , and correctness proof  $\rho$ .
- $(srs', \rho') \leftarrow \text{UpdSRS}(srs, \{\rho_j\}_{j=1}^n)$  is a PPT algorithm that takes a  $srs$ , a list of update proofs and outputs an updated  $srs'$  together with a proof of correct update  $\rho'$ .
- $b \leftarrow \text{VerifySRS}(srs, \{\rho_j\}_{j=1}^n)$  takes a reference string  $srs$ , a list of update proofs, and outputs a bit indicating acceptance or not.<sup>8</sup>

In the next section, we define security notions in the updatable setting by giving the adversary access to an SRS update oracle  $\text{UpdO}$ , defined in Fig. 1. The oracle allows the adversary to control the SRS generation. A trusted setup can be expressed by the updatable setup definition simply by restricting the adversary to only call the oracle on  $\text{intent} = \text{setup}$  and  $\text{intent} = \text{final}$ . Note that a soundness adversary now has access to both the random oracle  $\mathcal{H}$  and  $\text{UpdO}$ :  $(x, \pi) \leftarrow \mathcal{A}^{\text{UpdO}, \mathcal{H}}(1^\lambda; r)$ .

*Remark on universality of the SRS.* The proof systems we consider in this work are universal. This means that both the relation  $\mathbf{R}$  and the reference string  $srs$  allows to prove arithmetic constraints defined over a particular field up to some size bound. The public instance  $x$  must determine the constraints. If  $\mathbf{R}$  comes with any auxiliary input, the latter is benign. We elide public preprocessing of constraint specific proving and verification keys. While important for performance, this modeling is not critical for security.

## 2.2 Multi-message Fiat-Shamir Compiled Provers and Verifiers

Given interactive prover and (public coin) verifier  $P', V'$  that exchange messages resulting in transcript  $\tilde{\pi} = (a_1, c_1, \dots, a_\mu, c_\mu, a_{\mu+1})$ , where  $a_i$  comes from  $P'$  and  $c_i$  comes from  $V'$ , the  $(2\mu + 1)$ -message Fiat-Shamir heuristic defines non-interactive provers and verifiers  $P, V$  as follows:

<sup>8</sup> For instance Plonk and Marlin will use the  $\text{GenSRS}$ ,  $\text{UpdSRS}$  and  $\text{VerifySRS}$  algorithms in Fig. 4.

- P behaves as  $P'$  except after sending message  $a_i, i \in [1 \dots \mu]$ , the prover does not wait for the message from the verifier but computes it locally setting  $c_i = \mathcal{H}(\tilde{\pi}[0..i])$ , where  $\tilde{\pi}[0..j] = (x, a_1, c_1, \dots, a_{j-1}, c_{j-1}, a_j)$ .<sup>9</sup> P output the non-interactive proof  $\pi = (a_1, \dots, a_\mu, a_{\mu+1})$ , that omits challenges as they can be recomputed using  $\mathcal{H}$ .
- V takes  $x$  and  $\pi$  as input and behaves as  $V'$  would but does not provide challenges to the prover. Instead it computes the challenges locally as P would, starting from  $\tilde{\pi}[0..1] = (x, a_1)$  which can be obtained from  $x$  and  $\pi$ . Then it verifies the resulting transcript  $\tilde{\pi}$  as the verifier  $V'$  would.

We note that since the verifier can compute the challenges by querying the random oracle, they do not need to be sent by the prover. Thus the  $\pi - \tilde{\pi}$  notational distinction.  $(2\mu + 1)$ -message FS-transformed NIZK proof system with an updatable SRS setup

*Notation for  $(2\mu + 1)$ -message Fiat–Shamir transformed proof systems.* Let  $\text{SRS} = (\text{GenSRS}, \text{UpdSRS}, \text{VerifySRS})$  be the algorithm of an updatable SRS ceremony. All our definitions and theorems are about non-interactive proof systems  $\Psi = (\text{SRS}, P, V, \text{Sim})$  compiled via the  $(2\mu + 1)$ -message FS transform. That is  $\pi = (a_1, \dots, a_\mu, a_{\mu+1})$  and  $\tilde{\pi} = (a_1, c_1, \dots, a_\mu, c_\mu, a_{\mu+1})$ , with  $c_i = \mathcal{H}(\tilde{\pi}[0..i])$ . We use  $\tilde{\pi}[0]$  for instance  $x$  and  $\tilde{\pi}[i], \tilde{\pi}[i].ch$  to denote prover message  $a_i$  and challenge  $c_i$  respectively.

$\text{SimO.H}(x)$	$\text{SimO.Prog}(x, h)$	$\text{SimO.P}(x, w)$	$\text{SimO.P}'(x)$
<b>if</b> $H[x] = \perp$ <b>then</b>	<b>if</b> $H[x] = \perp$ <b>then</b>	<b>assert</b> $(x, w) \in \mathbf{R}$	
$H[x] \leftarrow \$\text{Im}(\mathcal{H})$	$H[x] \leftarrow h$	$\pi \leftarrow \text{Sim}^{\text{SimO.H}, \text{SimO.Prog}}(\text{srs}, x)$	
<b>return</b> $H[x]$	$Q_{\text{prog}} \leftarrow Q_{\text{prog}} \cup \{x\}$	$Q \leftarrow Q \cup \{(x, \pi)\}$	
	<b>return</b> $H[x]$	<b>return</b> $\pi$	

Fig. 2: Simulation oracles: srs is the finalized SRS, only  $\text{SimO.P}'$  allows for simulation of false statements

### 2.3 Trapdoor-Less Zero-Knowledge (TLZK)

We call a protocol *trapdoor-less zero-knowledge* (TLZK) if there exists a simulator does not require the trapdoor, and works by programming the random oracle. Moreover, the simulator may only be allowed to program the random oracle on point  $\tilde{\pi}[0, k]$ , that is the simulator can only program the challenges that come after the  $k$ -th prover message. We call protocols which allow for such a simulation *k-programmable trapdoor-less zero-knowledge*.

Our definition of zero-knowledge for non-interactive arguments is in programmable ROM. We model this using the oracles from Fig. 2 that provide a stateful wrapper around  $\text{Sim}$ .  $\text{SimO.H}(x)$  simulates  $\mathcal{H}$  using lazy sampling,  $\text{SimO.Prog}(x, h)$  allows for programming the simulated  $\mathcal{H}$  and is available only to  $\text{Sim}$ .  $\text{SimO.P}(x, w)$  and  $\text{SimO.P}'(x)$  call the simulator. The former is used in the zero-knowledge definition and requires the statement and witness to be in the relation, the latter is used in the simulation extraction definition and does not require a witness input.

**Definition 1 (Updatable k-Programmable Trapdoor-Less Zero-Knowledge).** Let  $\Psi_{\text{FS}} = (\text{SRS}, P, V, \text{Sim})$  be a  $(2\mu + 1)$ -message FS-transformed NIZK proof system with an updatable SRS setup. We call  $\Psi_{\text{FS}}$  trapdoor-less zero-knowledge with security  $\varepsilon_{\text{zk}}$  if for any adversary  $\mathcal{A}$ ,  $|\varepsilon_0(\lambda) - \varepsilon_1(\lambda)| \leq \varepsilon_{\text{zk}}(\lambda)$ , where

$$\varepsilon_0(\lambda) = \Pr[\mathcal{A}^{\text{UpdO}, \mathcal{H}, P}(1^\lambda)], \varepsilon_1(\lambda) = \Pr[\mathcal{A}^{\text{UpdO}, \text{SimO.H}, \text{SimO.P}}(1^\lambda)].$$

If  $\varepsilon_{\text{zk}}(\lambda)$  is negligible, we say  $\Psi_{\text{FS}}$  is trapdoor-less zero-knowledge. Additionally, we say that  $\Psi_{\text{FS}}$  is *k-programmable*, if  $\text{Sim}$  before returning a proof  $\pi$  only calls  $\text{SimO.Prog}$  on  $(\tilde{\pi}[0..k], h)$ . That is, it only programs the  $k$ -th message.

*Remark 1 (TLZK vs HVZK).* We note that TLZK notion is closely related to honest-verifier zero-knowledge in the standard model. That is, if we consider an interactive proof system  $\Psi$  that is HVZK in the standard model then  $\Psi_{\text{FS}}$  is TLZK. This comes as the simulator  $\text{Sim}$  in  $\Psi$  produces a valid simulated proof by picking verifier's challenges according to a predefined distribution and  $\Psi_{\text{FS}}$ 's simulator  $\text{Sim}_{\text{FS}}$  produces its proofs similarly by picking the

<sup>9</sup> For Fiat–Shamir based SoK the message signed  $m$  is added to  $x$  before hashing.

challenges and additionally programming the random oracle to return the picked challenges. Importantly, in both  $\Psi$  and  $\Psi_{FS}$  success of the simulator does not depend on access to an SRS trapdoor.

We note that Plonk is 3-programmable TLZK, and Sonic and Marlin are 2-programmable TLZK. This follows directly from the proofs of their standard model zero-knowledge property in Lemmas 7, 11 and 14.

## 2.4 Updatable Simulation Extractability (USE)

We note that the zero-knowledge property is only guaranteed for statements in the language. For *simulation extractability* where the simulator should be able to provide simulated proofs for false statements as well, we thus use the oracle  $\text{SimO.P}'$ <sup>10</sup>.

**Definition 2 (Updatable Simulation Extractability).** Let  $\Psi_{NI} = (\text{SRS}, P, V, \text{Sim})$  be a NIZK proof system with an updatable SRS setup. We say that  $\Psi_{NI}$  is updatable simulation-extractable with security loss  $\varepsilon_{se}(\lambda, \text{acc}, q)$  if for any PPT adversary  $\mathcal{A}$  that is given oracle access to setup oracle  $\text{UpdO}$  and simulation oracle  $\text{SimO}$  and that produces an accepting proof for  $\Psi_{NI}$  with probability  $\text{acc}$ , where

$$\text{acc} = \Pr \left[ \begin{array}{l} V(\text{srs}, x, \pi) = 1 \\ \wedge (x, \pi) \notin Q \end{array} \middle| \begin{array}{l} r \leftarrow_{\$} R(\mathcal{A}) \\ (x, \pi) \leftarrow \mathcal{A}^{\text{UpdO}, \text{SimO}, \mathcal{H}, \text{SimO.P}'}(1^\lambda; r) \end{array} \right]$$

there exists an expected PPT extractor  $\text{Ext}_{se}$  such that

$$\Pr \left[ \begin{array}{l} V(\text{srs}, x, \pi) = 1, \\ (x, \pi) \notin Q, \\ \mathbf{R}(x, w) = 0 \end{array} \middle| \begin{array}{l} r \leftarrow_{\$} R(\mathcal{A}), (x, \pi) \leftarrow \mathcal{A}^{\text{UpdO}, \text{SimO}, \mathcal{H}, \text{SimO.P}'}(1^\lambda; r) \\ w \leftarrow \text{Ext}_{se}(\text{srs}, \mathcal{A}, r, Q_{\text{srs}}, Q_{\mathcal{H}}, Q) \end{array} \right] \leq \varepsilon_{se}(\lambda, \text{acc}, q)$$

Here,  $\text{srs}$  is the finalized SRS. List  $Q_{\text{srs}}$  contains all  $(\text{srs}, \rho)$  of update SRSs and their proofs, list  $Q_{\mathcal{H}}$  contains all  $\mathcal{A}$ 's queries to  $\text{SimO.H}$  and the (simulated) random oracle's answers,  $|Q_{\mathcal{H}}| \leq q$ , and list  $Q$  contains all  $(x, \pi)$  pairs where  $x$  is an instance queried to  $\text{SimO.P}'$  by the adversary and  $\pi$  is the simulator's answer.

## 2.5 Unique Response (UR) Protocols

A technical hurdle identified by Faust et al. [26] for proving simulation extraction via the Fiat–Shamir transformation is that the transformed proof system satisfies a unique response property. The original formulation by Fischlin, although suitable for applications presented in [26, 27], does not suffice in our case. First, the property assumes that the protocol has three messages, with the second being the challenge from the verifier. That is not the case we consider here. Second, it is not entirely clear how to generalize the property. Should one require that after the first challenge from the verifier, the prover's responses are fixed? That does not work since the prover needs to answer differently on different verifier's challenges, as otherwise the protocol could have fewer messages. Another problem is that the protocol could have a message, beyond the first prover's message, which is randomized. Unique response cannot hold in this case. Finally, the protocols we consider here are not in the standard model, but use an SRS.

We work around these obstacles by providing a generalized notion of the unique response property. More precisely, we say that a  $(2\mu + 1)$ -message protocol has *unique responses from  $k$* , and call it a  $k$ -UR-protocol, if it follows the definition below:

**Definition 3 (Updatable  $k$ -Unique Response Protocol).** Let  $\Psi_{FS} = (\text{SRS}, P, V, \text{Sim})$  be a  $(2\mu + 1)$ -message FS-transformed NIZK proof system with an updatable SRS setup. Let  $\mathcal{H}$  be the random oracle. We say that  $\Psi_{FS}$  has unique responses for  $k$  with security  $\varepsilon_{ur}(\lambda)$  if for any PPT adversary  $\mathcal{A}_{ur}$ :

$$\Pr \left[ \begin{array}{l} \pi \neq \pi', \tilde{\pi}[0..k] = \tilde{\pi}'[0..k], \\ V'(\text{srs}, x, \pi, c) = V'(\text{srs}, x, \pi', c) = 1 \end{array} \middle| (x, \pi, \pi', c) \leftarrow \mathcal{A}_{ur}^{\text{UpdO}, \mathcal{H}}(1^\lambda) \right] \leq \varepsilon_{ur}(\lambda)$$

where  $\text{srs}$  is the finalized SRS and  $V'(\text{srs}, x, \pi = (a_1, \dots, a_\mu, a_{\mu+1}))$  behaves as  $V(\text{srs}, x, \pi)$  except for using  $c$  as the  $k$ -th challenge instead of calling  $\mathcal{H}(\tilde{\pi}[0..k])$ . Thus,  $\mathcal{A}$  can program the  $k$ -th challenge. We say  $\Psi_{FS}$  is  $k$ -UR, if  $\varepsilon_{ur}(\lambda)$  is negligible.

<sup>10</sup> Note, that simulation extractability property where the simulator is required to give simulated proofs for true statements only is called *true simulation extractability*.

Intuitively, a protocol is  $k$ -UR if it is infeasible for a PPT adversary to produce a pair of accepting proofs  $\pi \neq \pi'$  that are the same on the first  $k$  messages of the prover.

The definition can be easily generalized to allow for programming the oracle on more than just a single point. We opted for this simplified presentation, since all the protocols analyzed in this paper require only single-point programming,

## 2.6 Rewinding-Based Knowledge Soundness (RBKS)

Before giving the definition of rewinding-based knowledge soundness for NIZK proof systems compiled via the  $2\mu + 1$ -message FS transformation, we first recall the notion of a tree of transcripts.

**Definition 4 (Tree of accepting transcripts, cf. [16]).** A  $(n_1, \dots, n_\mu)$ -tree of accepting transcripts is a tree where each node on depth  $i$ , for  $i \in [1.. \mu + 1]$ , is an  $i$ -th prover's message in an accepting transcript; edges between the nodes are labeled with challenges, such that no two edges on the same depth have the same label; and each node on depth  $i$  has  $n_i - 1$  siblings and  $n_{i+1}$  children. The tree consists of  $N = \prod_{i=1}^{\mu} n_i$  branches, where  $N$  is the number of accepting transcripts. We require  $N = \text{poly}(\lambda)$ . We refer to a  $(1, \dots, n_k = n, 1, \dots, 1)$ -tree as a  $(k, n)$ -tree.

The existence of simulation trapdoor for **P**, **S** and **M** means that they are not special sound in the standard sense. We therefore put forth the notion of rewinding-based knowledge soundness that is a computational notion. Note that in the definition below, it is implicit that each transcript in the tree is accepting with respect to a ‘‘local programming’’ of the random oracle. However, the verification of the proof output by the adversary is with respect to a non-programmed random oracle.

**Definition 5 (Updatable Rewinding-Based Knowledge Soundness).** Let  $n_1, \dots, n_\mu \in \mathbb{N}$ . Let  $\Psi_{\text{FS}} = (\text{SRS}, \text{P}, \text{V}, \text{Sim})$  be a  $(2\mu + 1)$ -message FS-transformed NIZK proof system with an updatable SRS setup for relation **R**. Let  $\mathcal{H}$  be the random oracle. We require existence of an expected PPT tree builder  $\mathcal{T}$  that eventually outputs a  $\mathbb{T}$  which is either a  $(n_1, \dots, n_\mu)$ -tree of accepting transcript or  $\perp$  and a PPT extractor  $\text{Ext}_{\text{ks}}$ . Let adversary  $\mathcal{A}_{\text{ks}}$  be a PPT algorithm, that outputs a valid proof with probability at least  $\text{acc}$ , where

$$\text{acc} = \Pr \left[ \begin{array}{l} \text{V}(\text{srs}, x, \pi) = 1 \\ \wedge (x, \pi) \notin Q \end{array} \mid \begin{array}{l} r \leftarrow_{\$} \text{R}(\mathcal{A}_{\text{ks}}) \\ (x, \pi) \leftarrow \mathcal{A}_{\text{ks}}^{\text{UpdO}, \mathcal{H}}(1^\lambda; r) \end{array} \right].$$

We say that  $\Psi_{\text{FS}}$  is  $(n_1, \dots, n_\mu)$ -rewinding-based knowledge sound with security loss  $\varepsilon_{\text{ks}}(\lambda, \text{acc}, q)$  if

$$\Pr \left[ \begin{array}{l} \text{V}(\text{srs}, x, \pi) = 1, \\ \mathbf{R}(x, w) = 0 \end{array} \mid \begin{array}{l} r \leftarrow_{\$} \text{R}(\mathcal{A}_{\text{ks}}), \\ (\text{srs}, x, \cdot) \leftarrow \mathcal{A}_{\text{ks}}^{\text{UpdO}, \mathcal{H}}(1^\lambda; r) \\ \mathbb{T} \leftarrow \mathcal{T}(\text{srs}, \mathcal{A}_{\text{ks}}, r, Q_{\text{srs}}, Q_{\mathcal{H}}), w \leftarrow \text{Ext}_{\text{ks}}(\mathbb{T}) \end{array} \right] \leq \varepsilon_{\text{ks}}(\lambda, \text{acc}, q).$$

Here,  $\text{srs}$  is the finalized SRS. List  $Q_{\text{srs}}$  contains all  $(\text{srs}, \rho)$  of updated SRSs and their proofs, and list  $Q_{\mathcal{H}}$  contains all of the adversaries queries to  $\mathcal{H}$  and the random oracle's answers,  $|Q_{\mathcal{H}}| \leq q$ .

## 3 Simulation Extractability—The General Result

Equipped with the definitional framework of Section 2, we now present the main result of this paper: a proof of simulation extractability for multi-message Fiat–Shamir-transformed NIZK proof systems.

Without loss of generality, we assume that whenever the accepting proof contains a response to a challenge from a random oracle, then the adversary queried the oracle to get it. It is straightforward to transform any adversary that violates this condition into an adversary that makes these additional queries to the random oracle and wins with the same probability.

The core conceptual insight of the proof is that the  $k$ -unique response and  $k$ -programmable trapdoor-less zero-knowledge properties together ensures that the  $k$ -th move challenges in the trees of rewinding-based knowledge soundness are fresh and do not come from the simulator. This allows us to eliminate the simulation oracle in our rewinding argument and enables us to use the existing results of [4] in later sections.

**Theorem 1 (Simulation-extractable multi-message protocols).** Let  $\Psi_{\text{FS}} = (\text{SRS}, \text{P}, \text{V}, \text{Sim})$  be a  $(2\mu + 1)$ -message FS-transformed NIZK proof system with an updatable SRS setup. If  $\Psi_{\text{FS}}$  is an updatable  $k$ -unique response protocol with security loss  $\varepsilon_{\text{ur}}$ , is updatable  $k$ -programmable trapdoor-less zero-knowledge, and updatable rewinding-based knowledge sound with security loss  $\varepsilon_{\text{ks}}$ ; Then  $\Psi_{\text{FS}}$  is updatable simulation-extractable with security loss

$$\varepsilon_{\text{se}}(\lambda, \text{acc}, q) \leq \varepsilon_{\text{ks}}(\lambda, \text{acc} - \varepsilon_{\text{ur}}(\lambda), q)$$

against any PPT adversary  $\mathcal{A}$  that makes up to  $q$  random oracle queries and returns an proof with probability at least  $\text{acc}$ .

*Proof.* Let  $(x, \pi) \leftarrow \mathcal{A}^{\text{UpdO}, \text{SimO}, \mathcal{H}, \text{SimO}, P'}(r_{\mathcal{A}})$  be the USE adversary. We show how to build an extractor  $\text{Ext}_{\text{se}}(\text{srs}, \mathcal{A}, r_{\mathcal{A}}, Q, Q_{\mathcal{H}}, Q_{\text{srs}})$  that outputs a witness  $w$ , such that  $\mathbf{R}(x, w)$  holds with high probability. To that end we define an algorithm  $\mathcal{A}_{\text{ks}}^{\text{UpdO}, \mathcal{H}}(r)$  against rewinding-based knowledge soundness of  $\Psi_{\text{FS}}$  that runs internally  $\mathcal{A}^{\text{UpdO}, \text{SimO}, \mathcal{H}, \text{SimO}, P'}(r_{\mathcal{A}})$ . Here  $r = (r_{\text{Sim}}, r_{\mathcal{A}})$  with  $r_{\text{Sim}}$  the randomness that will be used to simulate  $\text{SimO}, P'$ .

The code of  $\mathcal{A}_{\text{ks}}^{\text{UpdO}, \mathcal{H}}(r)$  hardcodes  $Q$  such that it does not use any randomness for proofs in  $Q$  as long as statements are queried in order. In this case it simply returns a proof  $\pi_{\text{Sim}}$  from  $Q$  but nevertheless queries  $\text{SimO}, \text{Prog}$  on  $(\tilde{\pi}_{\text{Sim}}[0..k], \tilde{\pi}_{\text{Sim}}[k].\text{ch})$ , i.e. it programs the  $k$ -th challenge. While it is hard to construct such an adversary without knowing  $Q$ , it clearly exists and  $\text{Ext}_{\text{se}}$  has the necessary inputs to construct  $\mathcal{A}_{\text{ks}}$ . This hardcoding guarantees that  $\mathcal{A}_{\text{ks}}$  returns the same  $(x, \pi)$  as  $\mathcal{A}$  in the experiment. Eventually,  $\text{Ext}_{\text{se}}$  uses the tree builder  $\mathcal{T}$  and extractor  $\text{Ext}_{\text{ks}}$  for  $\mathcal{A}_{\text{ks}}$  to extract the witness for  $x$ . Both guaranteed to exist (and be successful with high probability) by rewinding-based knowledge soundness. This high-level argument shows that  $\text{Ext}_{\text{se}}$  exists as well.

We now give the details of the simulation that guarantees that  $\mathcal{A}_{\text{ks}}$  is successful whenever  $\mathcal{A}$  is—except with a small security loss that we will bound late: Since  $\mathcal{A}_{\text{ks}}$  runs  $\mathcal{A}$  internally, it needs to take care of  $\mathcal{A}$ 's oracle queries.  $\mathcal{A}_{\text{ks}}$  passes on queries of  $\mathcal{A}$  to the update oracle  $\text{UpdO}$  to its own  $\text{UpdO}$  oracle and returns the result to  $\mathcal{A}$ .  $\mathcal{A}_{\text{ks}}$  internally simulates (non-hardcoded) queries to the simulator  $\text{SimO}, P'$  by running the  $\text{Sim}$  algorithm on randomness  $r_{\text{Sim}}$  of its tape.  $\text{Sim}$  requires access to oracles  $\text{SimO}, \mathcal{H}$  to compute a challenge honestly and  $\text{SimO}, \text{Prog}$  to program a challenge. Again  $\mathcal{A}_{\text{ks}}$  simulates both of these oracles internally, cf. Fig. 3, this time using the  $\mathcal{H}$  oracle of  $\mathcal{A}_{\text{ks}}$ . Note that queries of  $\mathcal{A}$  to  $\text{SimO}, \mathcal{H}$  are not programmed, but passed on to  $\mathcal{H}$ .

Importantly, all challenges in simulated proofs, up to round  $k$  are also computed honestly, i.e.  $\tilde{\pi}[i].\text{ch} = \mathcal{H}(\tilde{\pi}[0..i])$ , for  $i < k$ .

$\text{SimO}, \mathcal{H}(x)$	$\text{SimO}, \text{Prog}(x, h)$
<b>if</b> $H[x] = \perp$ <b>then</b>	<b>if</b> $H[x] = \perp$ <b>then</b>
$H[x] \leftarrow \mathcal{H}(x)$	$H[x] \leftarrow h$
<b>return</b> $H[x]$	$Q_{\text{prog}} \leftarrow Q_{\text{prog}} \cup \{x\}$
	<b>return</b> $H[x]$

Fig. 3: Simulating random oracle calls.

Eventually,  $\mathcal{A}$  outputs an instance and proof  $(x, \pi)$ .  $\mathcal{A}_{\text{ks}}$  returns the same values as long as  $\tilde{\pi}[0..i] \notin Q_{\text{prog}}$ ,  $i \in [1, \mu]$ . This models that the proof output by  $\mathcal{A}_{\text{ks}}$  must not contain any programmed queries as such a proof would not w.r.t  $\mathcal{H}$  in the RBKS experiment. If  $\mathcal{A}$  outputs a proof that does contain programmed challenges, then  $\mathcal{A}_{\text{ks}}$  aborts. We denote this event by  $E$ .

**Lemma 1.** *Probability that  $E$  happens is upper-bounded by  $\varepsilon_{\text{ur}}(\lambda)$ .*

*Proof.* We build an adversary  $\mathcal{A}_{\text{ur}}^{\text{UpdO}, \mathcal{H}}(\lambda; r)$  that has access to the random oracle  $\mathcal{H}$  and update oracle  $\text{UpdO}$ .  $\mathcal{A}_{\text{ur}}$  uses  $\mathcal{A}_{\text{ks}}$  to break the  $k$ -UR property of  $\Psi_{\text{FS}}$ .

When  $\mathcal{A}_{\text{ks}}$  outputs a proof  $\pi$  for  $x$  such that  $E$  holds,  $\mathcal{A}_{\text{ur}}$  looks through lists  $Q$  and  $Q_{\mathcal{H}}$  until it finds  $\tilde{\pi}_{\text{Sim}}[0..k]$  such that  $\tilde{\pi}[0..k] = \tilde{\pi}_{\text{Sim}}[0..k]$  and a programmed random oracle query  $\tilde{\pi}_{\text{Sim}}[k].\text{ch}$  on  $\tilde{\pi}_{\text{Sim}}[0..k]$ .  $\mathcal{A}_{\text{ur}}$  returns two proofs  $\pi$  and  $\pi_{\text{Sim}}$  for  $x$ : and the challenge  $\tilde{\pi}_{\text{Sim}}[k].\text{ch} = \tilde{\pi}[k].\text{ch}$

Importantly, both proofs are w.r.t the unique response verifier. The first, since it is a correctly computed simulated proof for which the unique response property definition allows any challenges at  $k$ . The latter, since it is an proof produced by the adversary. We have that  $\pi \neq \pi_{\text{Sim}}$  as otherwise  $\mathcal{A}$  does not win the simulation extractability game as  $\pi \in Q$ . On the other hand, if the proofs are different, then  $\mathcal{A}_{\text{ur}}$  breaks  $k$ -UR-ness of  $\Psi_{\text{FS}}$ . This happens only with probability  $\varepsilon_{\text{ur}}(\lambda)$ .  $\square$

We denote by  $\widetilde{\text{acc}}$  the probability that  $\mathcal{A}_{\text{ks}}$  outputs an proof. We note that by up-to-bad reasoning  $\widetilde{\text{acc}}$  is at most  $\varepsilon_{\text{ur}}(\lambda)$  far from the probability that  $\mathcal{A}$  outputs an proof. Thus, the probability that  $\mathcal{A}_{\text{ks}}$  outputs an proof is at least  $\widetilde{\text{acc}} \geq \text{acc} - \varepsilon_{\text{ur}}(\lambda)$ . Since  $\Psi_{\text{FS}}$  is  $\varepsilon_{\text{ks}}(\lambda, \widetilde{\text{acc}}, q)$  rewinding-based knowledge sound, there is a tree builder  $\mathcal{T}$  and

extractor  $\text{Ext}_{ks}$  that rewinds  $\mathcal{A}_{ks}$  to obtain a tree of accepting transcripts  $T$  and fails to extract the witness with probability at most  $\varepsilon_{ks}(\lambda, \widetilde{\text{acc}}, q)$ . The extractor  $\text{Ext}_{se}$  outputs the witness with the same probability.

Thus  $\varepsilon_{se}(\lambda, \text{acc}, q) = \varepsilon_{ks}(\lambda, \widetilde{\text{acc}}, q) \leq \varepsilon_{ks}(\lambda, \text{acc} - \varepsilon_{ur}, q)$ .  $\square$

*Remark 2.* Observe that our theorem does not depend on  $\varepsilon_{zk}(\lambda)$ . There is no real prover algorithm  $P$  in the experiment. Only the  $k$ -programmability of TLZK matters.

*Remark 3.* Observe that the theorem does not prescribe a tree shape for the tree builder  $\mathcal{T}$ . Interestingly, in our concrete results  $\mathcal{T}$  outputs a  $(k, *)$ -tree of accepting transcripts.

## 4 Polynomial Commitment Schemes

A polynomial commitment scheme  $\mathbf{PC} = (\text{GenSRS}, \text{Com}, \text{Op}, \text{Verify})$  consists of four algorithms and allows to commit to a polynomial  $f$  and later open the evaluation in a point  $z$  to some value  $s = f(z)$ . More formally:

$\text{GenSRS}(1^\lambda, \text{max})$ : The key generation algorithm takes in a security parameter  $\lambda$  and a parameter  $\text{max}$  which determines the maximal degree of the committed polynomial. It outputs a structured reference string  $\text{srs}$  (the commitment key). Note that  $\text{srs}$  implicitly determines  $\lambda$  and  $\text{max}$ .

$\text{Com}(\text{srs}, f)$ : The commitment algorithm  $\text{Com}(\text{srs}, f)$  takes in  $\text{srs}$  and a polynomial  $f$  with maximum degree  $\text{max}$ , and outputs a commitment  $c$ .

$\text{Op}(\text{srs}, z, s, f)$ : The opening algorithm takes as input  $\text{srs}$ , an evaluation point  $z$ , a value  $s$  and the polynomial  $f$ . It outputs an opening  $o$ .

$\text{Verify}(\text{srs}, c, z, s, o)$ : The verification algorithm takes in  $\text{srs}$ , a commitment  $c$ , an evaluation point  $z$ , a value  $s$  and an opening  $o$ . It outputs 1 if  $o$  is a valid opening for  $(c, z, s)$  and 0 otherwise.

A secure polynomial commitment  $\mathbf{PC}$  should satisfy correctness, evaluation binding, opening uniqueness, hiding and knowledge-binding. Note that since we are in the updatable setting,  $\text{srs}$  in these security definitions is the SRS that  $\mathcal{A}$  finalizes using the update oracle  $\text{UpdO}$  (See Fig. 1).

**Evaluation binding:** A PPT adversary  $\mathcal{A}$  which outputs a commitment  $c$  and evaluation points  $z$  has at most negligible chances to open the commitment to two different evaluations  $s, s'$ . That is, let  $k \in \mathbb{N}$  be the number of committed polynomials,  $l \in \mathbb{N}$  number of evaluation points,  $c \in \mathbb{G}^k$  be the commitments,  $z \in \mathbb{F}_p^l$  be the arguments the polynomials are evaluated at,  $s, s' \in \mathbb{F}_p^k$  the evaluations, and  $o, o' \in \mathbb{F}_p^l$  be the commitment openings. Then for every PPT adversary  $\mathcal{A}$

$$\Pr \left[ \begin{array}{l} \text{Verify}(\text{srs}, c, z, s, o) = 1, \\ \text{Verify}(\text{srs}, c, z, s', o') = 1, \\ s \neq s' \end{array} \middle| (c, z, s, s', o, o') \leftarrow \mathcal{A}^{\text{UpdO}}(\text{max}) \right] \leq \text{negl}(\lambda).$$

**Hiding:** We also formalize notion of  $k$ -hiding property of a polynomial commitment scheme. Let  $H$  be a set of size  $\text{max} + 1$  and  $Z_H$  its vanishing polynomial. We say that a polynomial scheme is *hiding* with security  $\varepsilon_{\text{hid}}(\lambda)$  if for every PPT adversary  $\mathcal{A}$ ,  $k \in \mathbb{N}$ , probability

$$\Pr[b' = b \mid (f_0, f_1, c, k, b') \leftarrow \mathcal{A}^{\text{UpdO}, \text{Oc}}(\text{max}, c), f_0, f_1 \in \mathbb{F}^{\text{max}}[X]] \leq \frac{1}{2} + \varepsilon(\lambda).$$

Where  $c = f'_b(\chi)$ , for a random bit  $b$  and the polynomial  $f'_b(X) = f_b + Z_H(X)(a_0 + a_1X + \dots + a_{k-1}X^{k-1})$ , and the oracle  $\text{Oc}$  on adversary's evaluation query  $x$  it adds  $x$  to initially empty set  $Q_x$  and if  $|Q_x| \leq k$ , it provides  $f'_b(x)$ .

**Commitment of knowledge:** Intuitively, when a commitment scheme is “of knowledge” then if an adversary produces a (valid) commitment  $c$ , which it can open correctly in an evaluation point, then it must know the underlying polynomial  $f$  which commits to that value. For every PPT adversary  $\mathcal{A}$  who produces commitment  $c$ , evaluation  $s$  and opening  $o$  there exists a PPT extractor  $\text{Ext}$  such that

$$\Pr \left[ \begin{array}{l} \deg f \leq \text{max}, c = \text{Com}(\text{srs}, f), \\ \text{Verify}(\text{srs}, c, z, s, o) = 1 \end{array} \middle| \begin{array}{l} c \leftarrow \mathcal{A}^{\text{UpdO}}(\text{max}), z \leftarrow \mathbb{F}_p \\ (s, o) \leftarrow \mathcal{A}(c, z), \\ f = \text{Ext}_{\mathcal{A}}(\text{srs}, c) \end{array} \right] \geq 1 - \varepsilon_k(\lambda).$$

In that case we say that  $\mathbf{PC}$  is  $\varepsilon_k(\lambda)$ -knowledge.

$\text{GenSRS}(1^\lambda, \max)$ <hr style="border: 0.5px solid black;"/> $\chi \leftarrow \mathbb{F}_p$ $\text{srs} := \left( \left[ \{\chi^i\}_{i=0}^{\max} \right]_1, [\chi]_2 \right);$ $\rho = ([\chi, \chi]_1, [\chi]_2)$ $\mathbf{return} (\text{srs}, \rho)$	$\text{UpdSRS}(\text{srs}, \{\rho_j\}_{j=1}^n)$ <hr style="border: 0.5px solid black;"/> $\text{Parse srs as } ([\{A_i\}_{i=0}^{\max}]_1, [B]_2)$ $\chi' \leftarrow \mathbb{F}_p$ $\text{srs}' := \left( \left[ \{\chi'^i A_i\}_{i=0}^{\max} \right]_1, [\chi' B]_2 \right);$ $\rho' = ([\chi' A_1, \chi']_1, [\chi']_2)$ $\mathbf{return} (\text{srs}', \rho')$
$\text{VerifySRS}(\text{srs}, \{\rho_j\}_{j=1}^n)$ <hr style="border: 0.5px solid black;"/> $\text{Parse srs as } ([\{A_i\}_{i=0}^{\max}]_1, [B]_2)$ $\text{and } \{\rho_j\}_{j=1}^n \text{ as } \left\{ (P_j, \bar{P}_j, \hat{P}_j) \right\}_{j=1}^n$ $\text{Verify Update proofs:}$ $\bar{P}_1 = P_1$ $P_j \bullet [1]_2 = P_{j-1} \bullet \hat{P}_j \quad \forall j \geq 2$ $\bar{P}_n \bullet [1]_2 = [1]_1 \bullet \hat{P}_n$ $\text{Verify SRS structure:}$ $[A_i]_1 \bullet [1]_2 = [A_{i-1}]_1 \bullet [B]_2 \text{ for all } 0 < i \leq \max$	

Fig. 4: Updatable SRS scheme SRS for  $\mathbf{PC}_p$

#### 4.1 Unique Opening Property of $\mathbf{PC}_p$

Now, we show that the batched variant of the KZG polynomial commitment scheme that is used in Plonk and Marlin, has the unique opening property.

**Lemma 2.** *Let  $\mathbf{PC}_p = (\text{GenSRS}, \text{Com}, \text{Op}, \text{VerifyBatched})$  be a batched version of a KZG polynomial commitment, cf. Fig. 5, that commits to  $m$  polynomials of degree up to  $\max$ . Let  $\mathbf{z} = (z_0, \dots, z_{l-1}) \in \mathbb{F}_p^l$  be the points the polynomials are evaluated at,  $k_i \in \mathbb{N}$  be the number of the committed polynomials to be evaluated at  $z_i$ , and  $K_i$  be the set of indices of these polynomials. Let  $\mathbf{s}_{K_i} \in \mathbb{F}_p^{k_i}$  the evaluations of polynomials at  $z_i$ , and  $\mathbf{o} = (o_0, \dots, o_{l-1}) \in \mathbb{F}_p^l$  be the commitment openings. We show that the probability an algebraic adversary  $\mathcal{A}$ , who can made up to  $q$  random oracle queries, opens the same vector of commitments in two different ways is at most  $\varepsilon_{\text{op}}(\lambda)$ , for  $\varepsilon_{\text{op}}(\lambda) \leq l \cdot \varepsilon_{\text{udlog}}(\lambda) + 1/p$ , where  $\varepsilon_{\text{udlog}}(\lambda)$  is security of the  $(\max, 1)$ -udlog assumption and  $p$  is the field used in  $\mathbf{PC}_p$ .*

*Proof.* The proof goes by game hops. In the first game the adversary wins if it presents two accepting openings of a vector of polynomials. Then, we restrict the winning condition and require that the adversary also makes the idealized batched verifier to accept the proof. In the next game, we abort if the idealized verifier rejects a proof for one of the evaluation point.

**Game 0:** In this game the adversary wins if it provides two different openings for a vector of polynomial commitments and their evaluations that are accepting by `VerifyBatched`.

**Game 1:** This game is identical to Game 0 except it is additionally aborted if the commitment opening are not accepting by `VerifyBatched'`.

**Game 0 to Game 1:** Any discrepancy between the idealized verifier rejection and real verifier acceptance allows one to break the (updatable) discrete logarithm problem. The reduction  $\mathcal{R}_{\text{udlog}}$  proceeds as follows. It answers  $\mathcal{A}$ 's queries for SRS updates according to the answers it receives from its udlog update oracle. When  $\mathcal{A}$  finalizes an SRS,  $\mathcal{R}_{\text{udlog}}$  finalizes the corresponding udlog challenge  $([1, \dots, \chi'^{\max}]_1, [1]_2)$ . We consider verification equation (\*\*\*) as a polynomial in  $X$  and the verification equation (\*) as it's evaluation at  $\chi'$ . Consider an opening such that verification equation (\*\*\*), cf. Fig. 5, does not hold, i.e. (\*\*\*) is not a zero polynomial, but (\*) does, i.e. (\*\*\*) zeroes at  $\chi'$ . Since  $\mathcal{A}$  is algebraic, all proof elements are extended by their representation as a combination of the input  $\mathbb{G}_1$ -elements. Therefore, all coefficients of the verification equation polynomial related to (\*\*\*) are known. Now,  $\mathcal{R}_{\text{udlog}}$  computes the roots of (\*\*\*), finds  $\chi'$  among them, and returns  $\chi'$ . Hence the probability that the adversary wins in Game 1, but does not win in Game 0 is upper-bounded by  $\varepsilon_{\text{udlog}}(\lambda)$ .

<p>SRS(<math>1^\lambda, \max</math>)</p> <p>cf. Fig. 4</p>	<p>Com(srs, <math>\mathbf{f}(X)</math>)</p> <hr/> <p><b>return</b> <math>[c]_1 = [\mathbf{f}(X)]_1</math></p> <div style="border: 1px solid black; padding: 2px; display: inline-block;"> <p><b>return</b> <math>\mathbf{f}(X)</math></p> </div>	<p>Op(srs, <math>\mathbf{z}, \mathbf{s}, \mathbf{f}(X), \text{aux}_0</math>)</p> <hr/> <p><math>\gamma \leftarrow \mathcal{H}(g_0(\mathbf{z}, \mathbf{s}, [c]_1, \text{aux}_0))</math></p> <p><b>for</b> <math>i \in [1 ..  \mathbf{z} ]</math> <b>do</b></p> $o_j(X) \leftarrow \sum_{i \in K_j} \gamma_j^{i-1} \frac{f_i(X) - f_i(z_j)}{X - z_j}$ <p><b>return</b> <math>\mathbf{o} = [\mathbf{o}(X)]_1</math></p> <div style="border: 1px solid black; padding: 2px; display: inline-block;"> <p><b>return</b> <math>\mathbf{o}(X)</math></p> </div>
<p>VerifyBatched(srs, <math>[c]_1, \mathbf{z}, \mathbf{s}, [\mathbf{o}(X)]_1, (\text{aux}_0, \text{aux}_1)</math>)</p> <hr/> <p><math>\gamma \leftarrow \mathcal{H}(g_0(\mathbf{z}, \mathbf{s}, [c]_1, \text{aux}_0))</math></p> <p><math>r \leftarrow \mathcal{H}(g_1([c]_1, \mathbf{z}, \mathbf{s}, [\mathbf{o}(X)]_1, \text{aux}_1))</math></p> <p>(*) <b>if</b> <math>\sum_{j=1}^{ \mathbf{z} } r^j \cdot \left[ \sum_{i \in K_j} \gamma_j^{i-1} c_i - \sum_{i \in K_j} \gamma_j^{i-1} s_{i,j} \right]_1 \cdot [1]_2 +</math></p> $\sum_{j=1}^{ \mathbf{z} } r^j z_j o_j \cdot [1]_2 \neq \left[ \sum_{j=1}^{ \mathbf{z} } r^j o_j \right]_1 \cdot [X]_2$ <b>then</b> <p><b>return</b> 0</p> <div style="border: 1px solid black; padding: 2px;"> <p>(**) <b>if</b> <math>\sum_{j=1}^{ \mathbf{z} } r^j \cdot (\sum_{i \in K_j} \gamma_j^{i-1} f_i(X) -</math></p> <math display="block">\sum_{i \in K_j} \gamma_j^{i-1} s_{i,j}) +</math> <math display="block">\sum_{j=1}^{ \mathbf{z} } r^j z_j o_j(X) \neq \sum_{j=1}^{ \mathbf{z} } r^j o_j(X) \cdot X</math> <b>then</b> <p><b>return</b> 0</p> </div> <p><b>return</b> 1.</p>	<p>Verify(srs, <math>[c]_1, \mathbf{z}, \mathbf{s}, [\mathbf{o}(X)]_1, \text{aux}_0</math>)</p> <hr/> <p><math>\gamma \leftarrow \mathcal{H}(g_0(\mathbf{z}, \mathbf{s}, [c]_1, \text{aux}_0))</math></p> <p><b>for</b> <math>j \in [1 ..  \mathbf{z} ]</math> <b>do</b></p> <p><b>if</b> <math>\left[ \sum_{i \in K_j} \gamma_j^{i-1} c_i - \sum_{i \in K_j} \gamma_j^{i-1} s_{i,j} \right]_1 \cdot [1]_2 +</math></p> $z_j o_j \cdot [1]_2 \neq [o_j]_1 \cdot [X]_2$ <b>then</b> <p><b>return</b> 0</p> <div style="border: 1px solid black; padding: 2px;"> <p><b>if</b> <math>\sum_{i \in K_j} \gamma_j^{i-1} f_i(X) - \sum_{i \in K_j} \gamma_j^{i-1} s_{i,j} +</math></p> <math display="block">z_j o_j(X) \neq o_j(X) X</math> <b>then return</b> 0 </div> <p><b>return</b> 1.</p>	

Fig. 5:  $\text{PC}_{\mathcal{P}}$  polynomial commitment scheme. Here  $|\mathbf{z}| = l$  is the number of evaluation points, the number of committed polynomials is  $m$ ,  $K_j$  is the set of polynomials that was evaluated at point  $z_j$ . Functions  $g_0$  and  $g_1$  are injective and specific to the context in which the polynomial commitment is used. (In our case, functions  $g_0$  and  $g_1$  are produce partial transcripts of the proof that utilizes the commitment scheme, aux contains all additional information that is needed by the functions.) In the boxes we describe values returned or equality computed in the ideal protocol where the verifier checks equalities on the polynomials instead of their evaluations. For algorithm Alg we denote its ideal variant by  $\text{Alg}'$ .

**Game 2:** This game is identical to Game 1 except it is additionally aborted if one of the opening is not accepting by an idealized verifier  $\text{Verify}'$ .

**Game 1 to Game 2:** The ideal verifier checks whether the following equality, for  $\{\gamma_j\}_{j=1}^l, r$  picked at random, holds:

$$\sum_{j=0}^{l-1} r^j \left( \sum_{i \in K_j} \gamma_j^{i-1} \cdot f_i(X) - \sum_{i \in K_j} \gamma_j^{i-1} \cdot s_{ij} \right) \equiv \sum_{j=0}^{l-1} r^j \mathfrak{o}_j(X)(X - z_j). \quad (1)$$

Since  $r$  has been picked as a random oracle output, probability that Eq. (1) holds while for some  $j \in [0..l-1]$

$$r^j \left( \sum_{i \in K_j} \gamma_j^{i-1} \cdot f_i(X) - \sum_{i \in K_j} \gamma_j^{i-1} \cdot s_{ij} \right) \neq r^j \mathfrak{o}_j(X)(X - z_j)$$

is  $q/p$  cf. [30]. When  $r^j \left( \sum_{i \in K_j} \gamma_j^{i-1} \cdot f_i(X) - \sum_{i \in K_j} \gamma_j^{i-1} \cdot s_{ij} \right) = r^j \mathfrak{o}_j(X)(X - z_j)$  holds, polynomial  $\mathfrak{o}_j(X)$  is uniquely determined from the uniqueness of polynomial composition.

**Conclusion:** We note that the idealized verifier  $\text{ve}(X)$  does not accept two different openings of a correct evaluation. Hence the probability that the adversary wins Game 2 is 0 and the probability that the adversary wins in Game 0 is upper-bounded by  $\varepsilon_{\text{udlog}}(\lambda) + \frac{q}{p}$ .  $\square$

## 5 Concrete SNARKs Preliminaries

*Bilinear groups.* A bilinear group generator  $\text{Pgen}(1^\lambda)$  returns public parameters  $\mathfrak{p} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \hat{e}, [1]_1, [1]_2)$ , where  $\mathbb{G}_1, \mathbb{G}_2$ , and  $\mathbb{G}_T$  are additive cyclic groups of prime order  $p = 2^{\Omega(\lambda)}$ ,  $[1]_1, [1]_2$  are generators of  $\mathbb{G}_1, \mathbb{G}_2$ , resp., and  $\hat{e} : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$  is a non-degenerate PPT-computable bilinear pairing. We assume the bilinear pairing to be Type-3, i.e., that there is no efficient isomorphism from  $\mathbb{G}_1$  to  $\mathbb{G}_2$  or from  $\mathbb{G}_2$  to  $\mathbb{G}_1$ . We use the by now standard bracket notation, i.e., we write  $[a]_i$  to denote  $a[1]_i$ . We denote  $\hat{e}([a]_1, [b]_2)$  as  $[a]_1 \bullet [b]_2$ . Thus,  $[a]_1 \bullet [b]_2 = [ab]_T$ . Since every algorithm  $\mathcal{A}$  takes as input the public parameters we skip them when describing  $\mathcal{A}$ 's input. Similarly, we do not explicitly state that each protocol starts by running  $\text{Pgen}$ .

### 5.1 Algebraic Group Model

The algebraic group model (AGM) of Fuchsbauer, Kiltz, and Loss [29] lies somewhat between the standard and generic bilinear group model. In the AGM it is assumed that an adversary  $\mathcal{A}$  can output a group element  $[y] \in \mathbb{G}$  if  $[y]$  has been computed by applying group operations to group elements given to  $\mathcal{A}$  as input. It is further assumed, that  $\mathcal{A}$  knows how to “build”  $[y]$  from those elements. More precisely, the AGM requires that whenever  $\mathcal{A}([x])$  outputs a group element  $[y]$  then it also outputs  $c$  such that  $[y] = c^\top \cdot [x]$ . Plonk, Sonic and Marlin have been shown secure using the AGM. An adversary that works in the AGM is called *algebraic*.

**Ideal Verifier and Verification Equations.** Let  $(\text{SRS}, \text{P}, \text{V}, \text{Sim})$  be a proof system. Observe that the SRS algorithms provide an SRS which can be interpreted as a set of group representation of polynomials evaluated at trapdoor elements. That is, for a trapdoor  $\chi$  the SRS contains  $[\mathfrak{p}_1(\chi), \dots, \mathfrak{p}_k(\chi)]_1$ , for some polynomials  $\mathfrak{p}_1(X), \dots, \mathfrak{p}_k(X) \in \mathbb{F}_p[X]$ . The verifier  $\text{V}$  accepts a proof  $\pi$  for instance  $x$  if (a set of) verification equation  $\text{ve}_{x,\pi}$  (which can also be interpreted as a polynomial in  $\mathbb{F}_p[X]$  whose coefficients depend on messages sent by the prover) zeroes at  $\chi$ . Following [30] we call verifiers who check that  $\text{ve}_{x,\pi}(\chi) = 0$  *real verifiers* as opposed to *ideal verifiers* who accept only when  $\text{ve}_{x,\pi}(X) = 0$ . That is, while a real verifier accepts when a polynomial *evaluates* to zero, an ideal verifier accepts only when the polynomial *is* zero.

Although ideal verifiers are impractical, they are very useful in our proofs. More precisely, we show that the idealized verifier accepts an incorrect proof (what “incorrect” means depends on the situation) with at most negligible probability (and in many cases—never); when the real verifier accepts, but not the idealized one, then a malicious prover can be used to break the underlying security assumption (in our case—a variant of dlog.)

Analogously, idealized verifier can be defined for polynomial commitment schemes.

### 5.2 Dlog Assumptions in Standard and Updatable Setting

**Definition 6** ( $((q_1, q_2)$ -dlog **assumption**). *Let  $\mathcal{A}$  be a PPT adversary that gets as input  $[1, \chi, \dots, \chi^{q_1}]_1, [1, \chi, \dots, \chi^{q_2}]_2$ , for some randomly picked  $\chi \in \mathbb{F}_p$ , the assumption requires that  $\mathcal{A}$  cannot compute  $\chi$ . That is*

$$\Pr[\chi = \mathcal{A}([1, \chi, \dots, \chi^{q_1}]_1, [1, \chi, \dots, \chi^{q_2}]_2) \mid \chi \leftarrow_{\$} \mathbb{F}_p] \leq \text{negl}(\lambda).$$

Since all our protocols and security notions are in the updatable setting, it is natural to define the dlog assumptions also in the updatable setting. That is, instead of being given a dlog challenge the adversary  $\mathcal{A}$  is given access to an update oracle as defined in Fig. 1. The honestly generated SRS is set to be a dlog challenge and the update algorithm UpdSRS re-randomizing the challenge. We define this assumptions and show a reduction between the assumptions in the updatable and standard setting.

Note that for clarity we here refer to the SRS by Ch. Further, to avoid cluttering notation, we do not make the update proofs explicit. They are generated in the same manner as the proofs in Fig. 4.

**Definition 7** ( $(q_1, q_2)$ -udlog assumption). *Let  $\mathcal{A}$  be a PPT adversary that gets oracle access to UpdO with internal algorithms (GenSRS, UpdSRS, VerifySRS), where GenSRS and UpdSRS are defined as follows:*

- GenSRS( $\lambda$ ) samples  $\chi \leftarrow_{\$} \mathbb{F}_p$  and defines  $\text{Ch} := ([1, \chi, \dots, \chi^{q_1}]_1, [1, \chi, \dots, \chi^{q_2}]_2)$ .
- UpdSRS( $\text{Ch}, \{\rho_j\}_{j=1}^n$ ) parses Ch as  $([\{A_i\}_{i=0}^{q_1}]_1, [\{B_i\}_{i=0}^{q_2}]_2)$ , samples  $\tilde{\chi} \leftarrow_{\$} \mathbb{F}_p$ , and defines  $\widetilde{\text{Ch}} := ([\{\tilde{\chi}^i A_i\}_{i=0}^{q_1}]_1, [\{\tilde{\chi}^i B_i\}_{i=0}^{q_2}]_2)$ .

Then  $\Pr[\tilde{\chi} \leftarrow \mathcal{A}^{\text{UpdO}}(\lambda)] \leq \text{negl}(\lambda)$ , where  $([\{\tilde{\chi}^i\}_{i=0}^{q_1}]_1, [\{\tilde{\chi}^i\}_{i=0}^{q_2}]_2)$  is the final Ch.

*Remark 4 (Single adversarial updates after an honest setup).* As an alternative to the updatable setting defined in Fig. 1, one can consider a slightly different model of setup, where the adversary is given an initial honestly-generated SRS and is then allowed to perform a malicious update in one-shot fashion. Groth et al. show in [39] that the two definitions are equivalent for polynomial commitment based SNARKs. We use this simpler definition in our reductions.

We show a reduction from  $(q_1, q_2)$ -dlog assumption to its variant in the updatable setting (with single adversarial update).

**Lemma 3.**  $(q_1, q_2)$ -dlog  $\Rightarrow$   $(q_1, q_2)$ -udlog.

*Proof.* We show a reduction  $\mathcal{R}$  that uses an adversary  $\mathcal{A}$  that breaks  $(q_1, q_2)$ -udlog to break  $(q_1, q_2)$ -dlog. Specifically,  $\mathcal{R}$  proceeds as follows: given a dlog instance Ch as input, it sets Ch to be the initial (honestly generated) challenge and runs  $\mathcal{A}$ . After  $\mathcal{A}$  performs its update and finalizes the dlog challenge it returns the answer  $\chi'$ . Let  $\chi_{\mathcal{A}}$  be the trapdoor contribution of adversary  $\mathcal{A}$  in its update. Reduction  $\mathcal{R}$  can extract this value from the update proof of  $\mathcal{A}$ .  $\mathcal{R}$  now returns  $\chi = \chi' \chi_{\mathcal{A}}^{-1}$  as the discrete logarithm of the original challenge Ch.  $\square$

**Generalized Forking Lemma** Although dubbed “general”, the forking lemma of [7] is not general enough for our purpose as it is useful only for protocols where a witness can be extracted from just two transcripts. To be able to extract a witness from, say, an execution of  $\mathbf{P}$  we need at least  $(3n + 6)$  valid proofs (where  $n$  is the number of constrains),  $(n + 1)$  for  $\mathbf{S}$ , and  $2n + 3$  for  $\mathbf{M}$ . Here we use a result by Attema et al. [4]<sup>11</sup> which lower-bounds the probability of generating a tree of accepting transcripts  $T$ . We restate their Proposition 2 in our notation:

**Lemma 4 (Run Time and Success Probability).** *Let  $N = n_1 \cdot \dots \cdot n_\mu$  and  $p = 2^{\Omega(\lambda)}$ . Let  $\varepsilon_{\text{err}}(\lambda) = 1 - \prod_{i=1}^{\mu} \left(1 - \frac{n_i - 1}{p}\right)$ . Assume adversary  $\mathcal{A}$  that makes up to  $q$  random oracle queries and outputs an accepting proof with probability at least  $\text{acc}$ . There exists a tree building algorithm  $\mathcal{T}$  for  $(n_1, \dots, n_\mu)$ -trees that succeeds in building a tree of accepting transcripts in expected running time  $N + q(N - 1)$  with probability at least*

$$\frac{\text{acc} - (q + 1)\varepsilon_{\text{err}}(\lambda)}{1 - \varepsilon_{\text{err}}(\lambda)}.$$

**Opening Uniqueness of Batched Polynomial Commitment Openings** To show the unique response property required by our main theorem we show that the polynomial commitment schemes employed by concrete proof systems have unique openings, which, intuitively, assures that there is only one valid opening for a given committed polynomial and evaluation point:

**Definition 8 (Unique opening property).** *Let  $m \in \mathbb{N}$  be the number of committed polynomials,  $l \in \mathbb{N}$  number of evaluation points,  $c \in \mathbb{G}^m$  be the commitments,  $z \in \mathbb{F}_p^l$  be the arguments the polynomials are evaluated at,  $K_j$*

<sup>11</sup> An earlier versions had its own forking lemma generalization. Attema et al. has a better bound.

set of indices of polynomials which are evaluated at  $z_j$ ,  $\mathbf{s}_i$  vector of evaluations of  $f_i$ , and  $\mathbf{o}_j, \mathbf{o}'_j \in \mathbb{F}_p^{K_j}$  be the commitment openings. Then for every PPT adversary  $\mathcal{A}$

$$\Pr \left[ \begin{array}{l} \text{Verify}(\text{srs}, \mathbf{c}, \mathbf{z}, \mathbf{s}, \mathbf{o}) = 1, \\ \text{Verify}(\text{srs}, \mathbf{c}, \mathbf{z}, \mathbf{s}, \mathbf{o}') = 1, \\ \mathbf{o} \neq \mathbf{o}' \end{array} \middle| (\mathbf{c}, \mathbf{z}, \mathbf{s}, \mathbf{o}, \mathbf{o}') \leftarrow \mathcal{A}^{\text{UpdO}}(\text{max}) \right] \leq \text{negl}(\lambda).$$

We show that the polynomial commitment schemes of Plonk, Sonic, and Marlin satisfy this requirement in Section 4.1.

*Remark 5.* Fig. 5, cf. Section 4.1, presents efficient variants of KZG [41] polynomial commitment schemes used in Plonk, Sonic and Marlin that support batched verification. Algorithms Com, Op, Verify take vectors as input and receive an additional arbitrary auxiliary string. This adversarially chosen string only provides additional context for the computation of challenges and allows reconstruction of proof transcripts  $\tilde{\pi}[0..i]$  for batch challenge computations. We treat auxiliary input implicitly in the definition above.

## 6 Non-malleability of Plonk

In this section, we show that  $\mathbf{P}_{\text{FS}}$  is simulation-extractable. To this end, we first use the unique opening property to show that  $\mathbf{P}_{\text{FS}}$  has the 3-UR property, cf. Lemma 5. Next, we show that  $\mathbf{P}_{\text{FS}}$  is rewinding-based knowledge sound. That is, given a number of accepting transcripts whose first 3 messages match, we can either extract a witness for the proven statement or use one of the transcripts to break the udlog assumption. This result is shown in the AGM, cf. Lemma 6. We then show that  $\mathbf{P}_{\text{FS}}$  is 3-programmable trapdoor-less ZK in the AGM, cf. Lemma 7.

Given rewinding-based knowledge soundness, 3-UR and trapdoor-less zero-knowledge of  $\mathbf{P}_{\text{FS}}$ , we invoke Theorem 1 and conclude that  $\mathbf{P}_{\text{FS}}$  is simulation-extractable.

### 6.1 Plonk Protocol Description

**The constraint system.** Assume  $C$  is a fan-in two arithmetic circuit, whose fan-out is unlimited and has  $n$  gates and  $m$  wires ( $n \leq m \leq 2n$ ). The constraint system of Plonk is defined as follows:

- Let  $\mathbf{V} = (\mathbf{a}, \mathbf{b}, \mathbf{c})$ , where  $\mathbf{a}, \mathbf{b}, \mathbf{c} \in [1..m]^n$ . Entries  $a_i, b_i, c_i$  represent indices of left, right and output wires of the circuit's  $i$ -th gate.
- Vectors  $\mathbf{Q} = (\mathbf{q}_L, \mathbf{q}_R, \mathbf{q}_O, \mathbf{q}_M, \mathbf{q}_C) \in (\mathbb{F}^n)^5$  are called *selector vectors*: (a) If the  $i$ -th gate is a multiplication gate then  $q_{L_i} = q_{R_i} = 0, q_{M_i} = 1$ , and  $q_{O_i} = -1$ . (b) If the  $i$ -th gate is an addition gate then  $q_{L_i} = q_{R_i} = 1, q_{M_i} = 0$ , and  $q_{O_i} = -1$ . (c)  $q_{C_i} = 0$  for multiplication and addition gates.<sup>12</sup>

We say that vector  $\mathbf{x} \in \mathbb{F}^m$  satisfies constraint system if for all  $i \in [1..n]$

$$\mathbf{q}_{L_i} \cdot \mathbf{x}_{a_i} + \mathbf{q}_{R_i} \cdot \mathbf{x}_{b_i} + \mathbf{q}_O \cdot \mathbf{x}_{c_i} + \mathbf{q}_{M_i} \cdot (\mathbf{x}_{a_i} \mathbf{x}_{b_i}) + \mathbf{q}_{C_i} = 0.$$

Public inputs  $(x_j)_{j=1}^\ell$  are enforced by adding the constrains

$$a_i = j, q_{L_i} = 1, q_{M_i} = q_{R_i} = q_{O_i} = 0, q_{C_i} = -x_j,$$

for some  $i \in [1..n]$ .

**Algorithms rolled out** Plonk argument system is universal. That is, it allows to verify computation of any arithmetic circuit which has up to  $n$  gates using a single SRS. However, to make computation efficient, for each circuit there is allowed a preprocessing phase which extends the SRS with circuit-related polynomial evaluations.

For the sake of simplicity of the security reductions presented in this paper, we include in the SRS only these elements that cannot be computed without knowing the secret trapdoor  $\chi$ . The rest of the preprocessed input can be computed using these SRS elements. We thus let them to be computed by the prover, verifier, and simulator separately.

Plonk *SRS generating algorithm*  $\text{GenSRS}(\mathbf{R})$ : The SRS generating algorithm picks at random  $\chi \leftarrow \mathbb{F}_p$ , computes and outputs  $\text{srs} = ([\{\chi^i\}_{i=0}^{n+5}]_1, [\chi]_2)$ .

*Preprocessing:* Let  $H = \{\omega^i\}_{i=1}^n$  be a (multiplicative)  $n$ -element subgroup of a field  $\mathbb{F}$  compound of  $n$ -th roots of unity in  $\mathbb{F}$ . Let  $L_i(X)$  be the  $i$ -th element of an  $n$ -elements Lagrange basis. During the preprocessing phase polynomials  $S_{\text{id}_j}, S_{\sigma_j}$ , for  $j \in [1..3]$ , are computed:

$$\begin{aligned} S_{\text{id}_1}(X) &= X, & S_{\sigma_1}(X) &= \sum_{i=1}^n \sigma(i) L_i(X), \\ S_{\text{id}_2}(X) &= k_1 \cdot X, & S_{\sigma_2}(X) &= \sum_{i=1}^n \sigma(n+i) L_i(X), \\ S_{\text{id}_3}(X) &= k_2 \cdot X, & S_{\sigma_3}(X) &= \sum_{i=1}^n \sigma(2n+i) L_i(X). \end{aligned}$$

<sup>12</sup> The  $q_{C_i}$  selector vector is meant to encode (input independent) constants.

Coefficients  $k_1, k_2$  are such that  $H, k_1 \cdot H, k_2 \cdot H$  are different cosets of  $\mathbb{F}^*$ , thus they define  $3 \cdot n$  different elements. Gabizon et al. [30] notes that it is enough to set  $k_1$  to a quadratic residue and  $k_2$  to a quadratic non-residue.

Furthermore, we define polynomials  $q_L, q_R, q_O, q_M, q_C$  such that

$$\begin{aligned} q_L(X) &= \sum_{i=1}^n q_{L_i} L_i(X), & q_O(X) &= \sum_{i=1}^n q_{O_i} L_i(X), \\ q_R(X) &= \sum_{i=1}^n q_{R_i} L_i(X), & q_C(X) &= \sum_{i=1}^n q_{C_i} L_i(X). \\ q_M(X) &= \sum_{i=1}^n q_{M_i} L_i(X), \end{aligned}$$

*Proving statements in  $\mathbf{P}_{FS}$*  We show how prover's algorithm  $P(\text{srs}, x = (w'_i)_{i=1}^\ell, w = (w_i)_{i=1}^{3 \cdot n})$  operates for the Fiat–Shamir transformed version of Plonk. Note that for notational convenience  $w$  also contains the public input wires  $w'_i = w_i, i \in [1 \dots \ell]$ .

**Message 1** Sample  $b_1, \dots, b_9 \leftarrow \mathbb{F}_p$ ; compute  $a(X), b(X), c(X)$  as

$$\begin{aligned} a(X) &= (b_1 X + b_2) Z_H(X) + \sum_{i=1}^n w_i L_i(X) \\ b(X) &= (b_3 X + b_4) Z_H(X) + \sum_{i=1}^n w_{n+i} L_i(X) \\ c(X) &= (b_5 X + b_6) Z_H(X) + \sum_{i=1}^n w_{2 \cdot n+i} L_i(X) \end{aligned}$$

Output polynomial commitments  $[a(\chi), b(\chi), c(\chi)]_1$ .

**Message 2** Compute challenges  $\beta, \gamma \in \mathbb{F}_p$  by querying random oracle on partial proof, that is,  $\beta = \mathcal{H}(\tilde{\pi}[0..1], 0), \gamma = \mathcal{H}(\tilde{\pi}[0..1], 1)$ .

Compute permutation polynomial  $z(X)$

$$\begin{aligned} z(X) &= (b_7 X^2 + b_8 X + b_9) Z_H(X) + L_1(X) + \\ &+ \sum_{i=1}^{n-1} \left( L_{i+1}(X) \prod_{j=1}^i \frac{(w_j + \beta \omega^{j-1} + \gamma)(w_{n+j} + \beta k_1 \omega^{j-1} + \gamma)(w_{2n+j} + \beta k_2 \omega^{j-1} + \gamma)}{(w_j + \sigma(j)\beta + \gamma)(w_{n+j} + \sigma(n+j)\beta + \gamma)(w_{2n+j} + \sigma(2n+j)\beta + \gamma)} \right) \end{aligned}$$

Output polynomial commitment  $[z(\chi)]_1$

**Message 3** Compute the challenge  $\alpha = \mathcal{H}(\tilde{\pi}[0..2])$ , compute the quotient polynomial

$$\begin{aligned} t(X) &= \\ &= (a(X)b(X)q_M(X) + a(X)q_L(X) + b(X)q_R(X) + c(X)q_O(X) + \text{Pl}(X) + q_C(X))/Z_H(X) + \\ &+ ((a(X) + \beta X + \gamma)(b(X) + \beta k_1 X + \gamma)(c(X) + \beta k_2 X + \gamma)z(X))\alpha/Z_H(X) \\ &- (a(X) + \beta S_{\sigma_1}(X) + \gamma)(b(X) + \beta S_{\sigma_2}(X) + \gamma)(c(X) + \beta S_{\sigma_3}(X) + \gamma)z(X)\alpha/Z_H(X) \\ &+ (z(X) - 1)L_1(X)\alpha^2/Z_H(X) \end{aligned}$$

Split  $t(X)$  into degree less than  $n$  polynomials  $t_{lo}(X), t_{mid}(X), t_{hi}(X)$ , such that  $t(X) = t_{lo}(X) + X^n t_{mid}(X) + X^{2n} t_{hi}(X)$ . Output  $[t_{lo}(\chi), t_{mid}(\chi), t_{hi}(\chi)]_1$ .

**Message 4** Get the challenge  $\mathfrak{z} \in \mathbb{F}_p, \mathfrak{z} = \mathcal{H}(\tilde{\pi}[0..3])$ . Compute opening evaluations  $a(\mathfrak{z}), b(\mathfrak{z}), c(\mathfrak{z}), S_{\sigma_1}(\mathfrak{z}), S_{\sigma_2}(\mathfrak{z}), t(\mathfrak{z}), z(\mathfrak{z}\omega)$ , Compute the linearization polynomial

$$\begin{aligned} r(X) &= \\ &= a(\mathfrak{z})b(\mathfrak{z})q_M(X) + a(\mathfrak{z})q_L(X) + b(\mathfrak{z})q_R(X) + c(\mathfrak{z})q_O(X) + q_C(X) \\ &+ \alpha \cdot ((a(\mathfrak{z}) + \beta \mathfrak{z} + \gamma)(b(\mathfrak{z}) + \beta k_1 \mathfrak{z} + \gamma)(c(\mathfrak{z}) + \beta k_2 \mathfrak{z} + \gamma) \cdot z(X)) \\ &- \alpha \cdot ((a(\mathfrak{z}) + \beta S_{\sigma_1}(\mathfrak{z}) + \gamma)(b(\mathfrak{z}) + \beta S_{\sigma_2}(\mathfrak{z}) + \gamma)\beta z(\mathfrak{z}\omega) \cdot S_{\sigma_3}(X)) \\ &+ \alpha^2 \cdot L_1(\mathfrak{z}) \cdot z(X) \end{aligned}$$

Output  $a(\mathfrak{z}), b(\mathfrak{z}), c(\mathfrak{z}), S_{\sigma_1}(\mathfrak{z}), S_{\sigma_2}(\mathfrak{z}), t(\mathfrak{z}), z(\mathfrak{z}\omega), r(\mathfrak{z})$ .

**Message 5** Compute the opening challenge  $v \in \mathbb{F}_p, v = \mathcal{H}(\tilde{\pi}[0..4])$ . Compute the openings for the polynomial commitment scheme

$$W_{\mathfrak{z}}(X) = \frac{1}{X - \mathfrak{z}} \left( \begin{aligned} &t_{lo}(X) + \mathfrak{z}^n t_{mid}(X) + \mathfrak{z}^{2n} t_{hi}(X) - t(\mathfrak{z}) + v(r(X) - r(\mathfrak{z})) + v^2(a(X) - a(\mathfrak{z})) \\ &+ v^3(b(X) - b(\mathfrak{z})) + v^4(c(X) - c(\mathfrak{z})) + v^5(S_{\sigma_1}(X) - S_{\sigma_1}(\mathfrak{z})) \\ &+ v^6(S_{\sigma_2}(X) - S_{\sigma_2}(\mathfrak{z})) \end{aligned} \right)$$

$$W_{\mathfrak{z}\omega}(X) = (z(X) - z(\mathfrak{z}\omega))/(X - \mathfrak{z}\omega)$$

Output  $[W_{\mathfrak{z}}(\chi), W_{\mathfrak{z}\omega}(\chi)]_1$ .

**Plonk verifier**  $V(\text{srs}, x, \pi)$ :

The Plonk verifier works as follows

1. Validate all obtained group elements.
2. Validate all obtained field elements.
3. Parse the instance as  $\{w_i\}_{i=1}^{\ell} \leftarrow x$ .
4. Compute challenges  $\beta, \gamma, \alpha, \mathfrak{z}, v, u$  from the transcript.
5. Compute zero polynomial evaluation  $Z_H(\mathfrak{z}) = \mathfrak{z}^n - 1$ .
6. Compute Lagrange polynomial evaluation  $L_1(\mathfrak{z}) = \frac{\mathfrak{z}^n - 1}{n(\mathfrak{z} - 1)}$ .
7. Compute public input polynomial evaluation  $PI(\mathfrak{z}) = \sum_{i \in [1.. \ell]} w_i L_i(\mathfrak{z})$ .
8. Compute quotient polynomials evaluations

$$t(\mathfrak{z}) = \left( r(\mathfrak{z}) + PI(\mathfrak{z}) - (a(\mathfrak{z}) + \beta S_{\sigma_1}(\mathfrak{z}) + \gamma)(b(\mathfrak{z}) + \beta S_{\sigma_2}(\mathfrak{z}) + \gamma)(c(\mathfrak{z}) + \gamma)z(\mathfrak{z}\omega)\alpha - L_1(\mathfrak{z})\alpha^2 \right) / Z_H(\mathfrak{z}).$$

9. Compute batched polynomial commitment  $[D]_1 = v[r]_1 + u[z]_1$  that is

$$[D]_1 = v \left( \begin{aligned} & a(\mathfrak{z})b(\mathfrak{z}) \cdot [q_M]_1 + a(\mathfrak{z})[q_L]_1 + b(\mathfrak{z})[q_R]_1 + c(\mathfrak{z})[q_O]_1 + \\ & + ((a(\mathfrak{z}) + \beta\mathfrak{z} + \gamma)(b(\mathfrak{z}) + \beta k_1\mathfrak{z} + \gamma)(c(\mathfrak{z}) + \beta k_2\mathfrak{z} + \gamma)\alpha + L_1(\mathfrak{z})\alpha^2) + \\ & - (a(\mathfrak{z}) + \beta S_{\sigma_1}(\mathfrak{z}) + \gamma)(b(\mathfrak{z}) + \beta S_{\sigma_2}(\mathfrak{z}) + \gamma)\alpha\beta z(\mathfrak{z}\omega) [S_{\sigma_3}(\chi)]_1 \end{aligned} \right) + u[z(\chi)]_1.$$

10. Computes full batched polynomial commitment  $[F]_1$ :

$$[F]_1 = ([t_o(\chi)]_1 + \mathfrak{z}^n [t_{mid}(\chi)]_1 + \mathfrak{z}^{2n} [t_{hi}(\chi)]_1) + u[z(\chi)]_1 + v \left( \begin{aligned} & a(\mathfrak{z})b(\mathfrak{z}) \cdot [q_M]_1 + a(\mathfrak{z})[q_L]_1 + b(\mathfrak{z})[q_R]_1 + c(\mathfrak{z})[q_O]_1 + \\ & + ((a(\mathfrak{z}) + \beta\mathfrak{z} + \gamma)(b(\mathfrak{z}) + \beta k_1\mathfrak{z} + \gamma)(c(\mathfrak{z}) + \beta k_2\mathfrak{z} + \gamma)\alpha + L_1(\mathfrak{z})\alpha^2) + \\ & - (a(\mathfrak{z}) + \beta S_{\sigma_1}(\mathfrak{z}) + \gamma)(b(\mathfrak{z}) + \beta S_{\sigma_2}(\mathfrak{z}) + \gamma)\alpha\beta z(\mathfrak{z}\omega) [S_{\sigma_3}(\chi)]_1 \end{aligned} \right) + v^2 [a(\chi)]_1 + v^3 [b(\chi)]_1 + v^4 [c(\chi)]_1 + v^5 [S_{\sigma_1}(\chi)]_1 + v^6 [S_{\sigma_2}(\chi)]_1.$$

11. Compute group-encoded batch evaluation  $[E]_1$

$$[E]_1 = \frac{1}{Z_H(\mathfrak{z})} \left[ \begin{aligned} & r(\mathfrak{z}) + PI(\mathfrak{z}) + \alpha^2 L_1(\mathfrak{z}) + \\ & - \alpha ((a(\mathfrak{z}) + \beta S_{\sigma_1}(\mathfrak{z}) + \gamma)(b(\mathfrak{z}) + \beta S_{\sigma_2}(\mathfrak{z}) + \gamma)(c(\mathfrak{z}) + \gamma)z(\mathfrak{z}\omega)) \end{aligned} \right]_1 + [vr(\mathfrak{z}) + v^2 a(\mathfrak{z}) + v^3 b(\mathfrak{z}) + v^4 c(\mathfrak{z}) + v^5 S_{\sigma_1}(\mathfrak{z}) + v^6 S_{\sigma_2}(\mathfrak{z}) + uz(\mathfrak{z}\omega)]_1.$$

12. Check whether the verification equation holds

$$([W_{\mathfrak{z}}(\chi)]_1 + u \cdot [W_{\mathfrak{z}\omega}(\chi)]_1) \bullet [\chi]_2 - (\mathfrak{z} \cdot [W_{\mathfrak{z}}(\chi)]_1 + u\mathfrak{z}\omega \cdot [W_{\mathfrak{z}\omega}(\chi)]_1 + [F]_1 - [E]_1) \bullet [1]_2 = 0. \quad (2)$$

The verification equation is a batched version of the verification equation from [41] which allows the verifier to check openings of multiple polynomials in two points (instead of checking an opening of a single polynomial at one point).

**Plonk simulator**  $\text{Sim}_{\chi}(\text{srs}, \text{td} = \chi, x)$ : We describe the simulator in Lemma 7.

## 6.2 Simulation extractability of Plonk

Due to lack of space, we provide here only theorem statements and intuition for why they hold. Full proofs are given in ??.

### Unique Response Property

**Lemma 5.** *Let  $\text{PC}_{\mathbf{P}}$  be a polynomial commitment that is  $\varepsilon_{\text{bind}}(\lambda)$ -binding and has unique opening property with loss  $\varepsilon_{\text{op}}(\lambda)$ . Then  $\text{P}_{\text{FS}}$  is 3-UR against algebraic adversaries, who makes up to  $q$  random oracle queries, with security loss  $\varepsilon_{\text{bind}}(\lambda) + \varepsilon_{\text{op}}(\lambda)$ .*

*Intuition.* We show that an adversary who can break the 3-unique response property of  $\mathbf{P}_{\text{FS}}$  can be either used to break the commitment scheme's evaluation binding or unique opening property. The former happens with the probability upper-bounded by  $\varepsilon_{\text{bind}}(\lambda)$ , the latter with the probability upper bounded by  $\varepsilon_{\text{op}}(\lambda)$ .

*Proof.* Let  $\mathcal{A}$  be an algebraic adversary tasked to break the 3-UR-ness of  $\mathbf{P}_{\text{FS}}$ . We show that the first three prover's messages determine, along with the verifiers challenges, the rest of it. We denote by  $\pi^0$  and  $\pi^1$  the two proofs that the adversary outputs. To distinguish polynomials and commitments which an honest prover would send in the proof from the polynomials and commitments computed by the adversary we write the latter using indices 0 and 1 (two indices as we have two transcripts), e.g. to describe the quotient polynomial provided by the adversary we write  $t^0$  and  $t^1$  instead of  $t$  as in the description of the protocol.

We note that since the unique response property requires from  $\pi^0$  and  $\pi^1$  that the first place they possibly differ is the 4-th prover's message, then the challenge  $\mathfrak{z}$ , that is picked by the adversary after the 3-rd message is the same in both transcripts. This challenge determines the evaluation point of polynomials  $a(X), b(X), c(X), t(X), z(X)$  which commitments are already sent.

In its fourth message, the prover provides evaluations of the aforementioned polynomials, along with evaluations of publicly known polynomials  $S_{\sigma_1}(\mathfrak{z}), S_{\sigma_2}(\mathfrak{z})$ , and evaluation of a linearization polynomial  $r(\mathfrak{z})$ .

Note that the adversary can output two accepting proofs that differ on their fourth message only if it either manages to break evaluation binding of one of the opening, or provides an incorrect opening which is accepted due to a batching error. Since the commitment scheme is evaluation binding with security loss  $\varepsilon_{\text{bind}}(\lambda)$ , the adversary can make  $\pi^0$  and  $\pi^1$  differ on the fourth message with the same probability.

Next, assume that the transcripts are the same up to the fourth message, but differ at the fifth. In that message, the adversary provides openings of the evaluations. Since the unique opening property, the adversary can open the valid evaluation of a polynomial to two different values with probability at most  $\varepsilon_{\text{op}}(\lambda)$ . (We note that for the KZG polynomial commitment scheme, as used in [30],  $\varepsilon_{\text{op}}(\lambda) \leq \varepsilon_{\text{udlog}}(\lambda) + q/p$ , cf. Lemma 2.)

By the union bound, the adversary is able to break the unique response property with probability upper bounded by  $\varepsilon_{\text{bind}}(\lambda) + \varepsilon_{\text{op}}(\lambda)$ .  $\square$

### Rewinding-Based Knowledge Soundness

**Lemma 6.**  $\mathbf{P}_{\text{FS}}$  is  $(3, 3n + 6)$ -rewinding-based knowledge sound against algebraic adversaries who make up to  $q$  random oracle queries with security loss

$$\varepsilon_{\text{ks}}(\lambda, \text{acc}, q) \leq \left( 1 - \frac{\text{acc} - (q + 1) \binom{3n+5}{p}}{1 - \frac{3n+5}{p}} \right) + (3n + 6) \cdot \varepsilon_{\text{udlog}}(\lambda),$$

Here  $\text{acc}$  is a probability that the adversary outputs an accepting proof, and  $\varepsilon_{\text{udlog}}(\lambda)$  is security of  $(n + 5, 1)$ -udlog assumption.

*Intuition.* We use Attema et al. [5, Proposition 2] to bound the probability that an algorithm  $\mathcal{T}$  does not obtain a tree of accepting transcripts in an expected number of runs. This happens with probability at most

$$1 - \frac{\text{acc} - (q + 1) \binom{3n+5}{p}}{1 - \frac{3n+5}{p}}$$

Then we analyze the case that one of the proofs in the tree  $T$  outputted by  $\mathcal{T}$  is not accepting by the ideal verifier. This discrepancy can be used to break an instance of an updatable dlog assumption which happens with probability at most  $(3n + 6) \cdot \varepsilon_{\text{udlog}}(\lambda)$ .

*Proof.* Let  $\mathcal{A}^{\mathcal{H}, \text{UpdO}}(1^\lambda; r)$  be the adversary who outputs  $(x, \pi)$  such that  $\mathbf{P}_{\text{FS}}.V$  accepts the proof. Let  $\mathcal{T}$  be a tree-building algorithm of Lemma 4 that outputs a tree  $T$ , and let  $\text{Ext}_{\text{ks}}$  be an extractor that given the tree output by  $\mathcal{T}$  reveals the witness for  $x$ . The main idea of the proof is to show that an adversary who breaks rewinding-based knowledge soundness can be used to break a udlog-problem instance. The proof goes by game hops. Note that since the tree branches after  $\mathcal{A}$ 's 3-rd message, the instance  $x$ , commitments  $[a(x), b(x), c(x), z(x), t_{\text{lo}}(x), t_{\text{mid}}(x), t_{\text{hi}}(x))]_1$ , and challenges  $\alpha, \beta, \gamma$  are the same in all the transcripts. Also, the tree branches after the third adversary's message where the challenge  $\mathfrak{z}$  is presented, thus tree  $T$  is built using different values of  $\mathfrak{z}$ . We consider the following games.

**Game 0:** In this game the adversary wins if it outputs a valid instance-proof pair  $(x, \pi)$ , and the extractor  $\text{Ext}_{\text{ks}}$  does not manage to output a witness  $w$  such that  $\mathbf{R}(x, w)$  holds.

**Game 1:** In this game the environment aborts the game if the tree building algorithm  $\mathcal{T}$  fails in building a tree of accepting transcripts  $T$ .

**Game 0 to Game 1:** By Lemma 4 probability that Game 1 is aborted, while Game 0 is not, is at most

$$1 - \frac{\text{acc} - (q + 1) \left( \frac{3n+5}{p} \right)}{1 - \frac{3n+5}{p}}.$$

**Game 2:** In this game the environment additionally aborts if at least one of its proofs in  $T$  is not accepting by an ideal verifier.

**Game 1 to Game 2:** As usual, we show a reduction that breaks an instance of a udlog assumption when Game 2 is aborted, while Game 1 is not.

Let  $\mathcal{R}_{\text{udlog}}$  be a reduction that gets as input an  $(n + 5, 1)$ -udlog instance  $[1, \dots, \chi^{n+5}]_1, [\chi]_2$ . Then it can update the instance to another one  $[1, \dots, \chi'^{n+5}]_1, [\chi']_2$ . Eventually, the reduction outputs  $\chi'$ . The reduction  $\mathcal{R}_{\text{udlog}}$  proceeds as follows. First, it builds  $\mathcal{A}$ 's SRS srs using the input udlog instance. Then it processes the adversary's update query by adding it to the list  $Q_{\text{srs}}$  and passing it to its own update oracle getting instance  $[1, \dots, \chi'^{n+5}]_1, [\chi']_2$ . The updated SRS srs' is then computed and given to  $\mathcal{A}$ .  $\mathcal{R}_{\text{udlog}}$  also takes care of the random oracle queries made by  $\mathcal{A}$ . It picks their answers honestly and write them in  $Q_{\mathcal{H}}$ . The reduction then starts  $\mathcal{T}(\text{srs}, \mathcal{A}, r, Q_{\mathcal{H}}, Q_{\text{srs}})$ .

Let  $(1, T)$  be the output returned by  $\mathcal{T}$ . Let  $x$  be a relation proven in  $T$ . Consider a transcript  $\pi \in T$  such that  $\text{ve}_{x,\pi}(X) \neq 0$ , but  $\text{ve}_{x,\pi}(\chi') = 0$ . Since  $\mathcal{A}$  is algebraic, all group elements included in  $T$  are extended by their representation as a combination of the input  $\mathbb{G}_1$ -elements. Hence, all coefficients of the verification equation polynomial  $\text{ve}_{x,\pi}(X)$  are known. Eventually, the reduction finds  $\text{ve}_{x,\pi}(X)$  zero points and returns  $\chi'$  which is one of them.

Hence, the probability that the adversary wins in Game 2 but does not win in Game 1 is upper-bounded by  $(3n + 6) \cdot \varepsilon_{\text{udlog}}(\lambda)$ .

### Conclusion:

Note that the adversary can win in Game 2 only if  $\mathcal{T}$  manages to produce a tree of accepting transcripts  $T$ , such that each of the transcripts in  $T$  is accepting by an ideal verifier. Note that since  $\mathcal{T}$  produces  $(3n + 6)$  accepting transcripts for different challenges  $\mathfrak{z}$ , it obtains the same number of different evaluations of polynomials  $a(X), b(X), c(X), z(X), t(X)$ . Since all the transcripts are accepting by an idealised verifier, the equality between polynomial  $t(X)$  and combination of polynomials  $a(X), b(X), c(X), z(X)$  defined in prover's 3-rd message description holds. Hence,  $a(X), b(X), c(X)$  encodes the valid witness for the proven statement.  $\text{Ext}_{\text{ks}}$  can recreate polynomials' coefficients by interpolation and reveal the witness given  $(3n + 6)$  evaluations.

Hence, the probability that the adversary wins in Game 0 is upper-bounded by

$$\varepsilon_{\text{ks}}(\lambda, \text{acc}, q) \leq \left( 1 - \frac{\text{acc} - (q + 1) \left( \frac{3n+5}{p} \right)}{1 - \frac{3n+5}{p}} \right) + (3n + 6) \cdot \varepsilon_{\text{udlog}}(\lambda).$$

□

## Trapdoor-Less Zero-Knowledge of Plonk

**Lemma 7.**  $\text{P}_{\text{FS}}$  is 3-programmable trapdoor-less zero-knowledge.

*Intuition.* The simulator, that does not know the SRS trapdoor can make a simulated proof by programming the random oracle. It proceeds as follows. It picks a random witness and behaves as an honest prover up to the point when a commitment to the polynomial  $t(X)$  is sent. Since the simulator picked a random witness and  $t(X)$  is a polynomial only (modulo some negligible function) when the witness is correct, it cannot compute commitment to  $t(X)$  as it is a rational function. However, the simulator can pick a random challenge  $\mathfrak{z}$  and a polynomial  $\tilde{t}(X)$  such that  $t(\mathfrak{z}) = \tilde{t}(\mathfrak{z})$ . Then the simulator continues behaving as an honest prover. We argue that such a simulated proof is indistinguishable from a real one.

*Proof.* As noted in Section 2.1, subvertible zero-knowledge implies updatable zero-knowledge. Hence, here we show that Plonk is TLZK even against adversaries who picks the SRS on its own.

The adversary  $\mathcal{A}(1^\lambda)$  picks an SRS  $\text{srs}$  and instance–witness pair  $(x, w)$  and gets a proof  $\pi$  simulated by the simulator  $\text{Sim}$  which proceeds as follows.

For its 1-st message the simulator picks randomly both the randomizers  $b_1, \dots, b_6$  and sets  $w_i = 0$  for  $i \in [1 \dots 3n]$ . Then  $\text{Sim}$  outputs  $[a(\chi), b(\chi), c(\chi)]_1$ . For the first challenge, the simulator picks permutation argument challenges  $\beta, \gamma$  randomly.

For its 2-nd message, the simulator computes  $z(X)$  from the newly picked randomizers  $b_7, b_8, b_9$  and coefficients of polynomials  $a(X), b(X), c(X)$ . Then it evaluates  $z(X)$  honestly and outputs  $[z(\chi)]_1$ . Challenge  $\alpha$  that should be sent by the verifier after the simulator’s 2 message is picked by the simulator at random.

In its 3-rd message the simulator starts by picking at random a challenge  $\mathfrak{z}$ , which in the real proof comes as a challenge from the verifier sent *after* the prover sends its 3-rd message. Then  $\text{Sim}$  computes evaluations  $a(\mathfrak{z}), b(\mathfrak{z}), c(\mathfrak{z}), S_{\sigma_1}(\mathfrak{z}), S_{\sigma_2}(\mathfrak{z}), \text{Pl}(\mathfrak{z}), L_1(\mathfrak{z}), Z_H(\mathfrak{z}), z(\mathfrak{z}\omega)$  and computes  $t(X)$  honestly. Since for a random  $a(X), b(X), c(X), z(X)$  the constraint system is (with overwhelming probability) not satisfied and the constraints-related polynomials are not divisible by  $Z_H(X)$ , hence  $t(X)$  is a rational function rather than a polynomial. Then, the simulator evaluates  $t(X)$  at  $\mathfrak{z}$  and picks randomly a degree- $(3n + 15)$  polynomial  $\tilde{t}(X)$  such that  $t(\mathfrak{z}) = \tilde{t}(\mathfrak{z})$  and publishes a commitment  $[\tilde{t}_{\text{lo}}(\chi), \tilde{t}_{\text{mid}}(\chi), \tilde{t}_{\text{hi}}(\chi)]_1$ . After that the simulator outputs  $\mathfrak{z}$  as a challenge.

For the next message, the simulator computes polynomial  $r(X)$  as an honest prover would, cf. Section 6.1 and evaluates  $r(X)$  at  $\mathfrak{z}$ .

The rest of the evaluations are already computed, thus  $\text{Sim}$  simply outputs  $a(\mathfrak{z}), b(\mathfrak{z}), c(\mathfrak{z}), S_{\sigma_1}(\mathfrak{z}), S_{\sigma_2}(\mathfrak{z}), t(\mathfrak{z}), z(\mathfrak{z}\omega)$ . After that it picks randomly the challenge  $v$ , and prepares the last message as an honest prover would. Eventually,  $\text{Sim}$  and outputs the final challenge,  $u$ , by picking it at random as well.

We argue about zero-knowledge as usual. The property holds since the polynomials that has witness elements at their coefficients are randomized by at least two randomizers and are evaluated at at most two points; and the simulator computes all polynomials as an honest prover would.  $\square$

## Simulation Extractability of $\mathbf{P}_{\text{FS}}$

Since Lemmas 5 to 7 hold,  $\mathbf{P}$  is 3-UR, rewinding-based knowledge sound and trapdoor-less zero-knowledge. We now make use of Theorem 1 and show that  $\mathbf{P}_{\text{FS}}$  is simulation-extractable as defined in Definition 2.

**Corollary 1 (Simulation extractability of  $\mathbf{P}_{\text{FS}}$ ).**  $\mathbf{P}_{\text{FS}}$  is updatable simulation-extractable against any PPT adversary  $\mathcal{A}$  who makes up to  $q$  random oracle queries and returns an accepting proof with probability at least  $\text{acc}$  with extraction failure probability

$$\varepsilon_{\text{se}}(\lambda, \text{acc}, q) \leq \left(1 - \frac{\text{acc} - \varepsilon_{\text{ur}}(\lambda) - (q + 1)\varepsilon_{\text{err}}(\lambda)}{1 - \varepsilon_{\text{err}}(\lambda)}\right) + (3n + 6) \cdot \varepsilon_{\text{udlog}}(\lambda),$$

where  $\varepsilon_{\text{err}}(\lambda) = \frac{3n+5}{p}$ ,  $\varepsilon_{\text{ur}}(\lambda) \leq \varepsilon_{\text{bind}}(\lambda) + \varepsilon_{\text{op}}(\lambda)$ ,  $p$  is the size of the field, and  $n$  is the number of constrains in the circuit.

## 7 Non-malleability of $\mathbf{S}_{\text{FS}}$

### 7.1 Preliminaries

**Definition 9 ( $(q_1, q_2)$ -ldlog assumption).** Let  $\mathcal{A}$  be a PPT adversary that gets as input  $[\chi^{-q_1}, \dots, 1, \dots, \chi^{q_1}]_1, [\chi^{-q_2}, \dots, 1, \dots, \chi^{q_2}]_2$ , for some randomly picked  $\chi \in \mathbb{F}_p$ , the assumption requires that  $\mathcal{A}$  cannot compute  $\chi$ . That is

$$\Pr[\chi = \mathcal{A}([\chi^{-q_1}, \dots, 1, \dots, \chi^{q_1}]_1, [\chi^{-q_2}, \dots, 1, \dots, \chi^{q_2}]_2) \mid \chi \leftarrow_{\$} \mathbb{F}_p] \leq \text{negl}(\lambda).$$

**Definition 10 ( $(q_1, q_2)$ -uldlog assumption).** Let  $\mathcal{A}$  be a PPT adversary that gets oracle access to  $\text{UpdO}$  with internal algorithms  $(\text{GenSRS}_{\text{ldlog}}, \text{UpdSRS}_{\text{ldlog}}, \text{VerifySRS})$ , where  $\text{GenSRS}_{\text{ldlog}}$  and  $\text{UpdSRS}_{\text{ldlog}}$  are defined as follows:

- $\text{GenSRS}_{\text{ldlog}}(\lambda)$  samples  $\chi \leftarrow_{\$} \mathbb{F}_p$  and defines  $\text{Ch} := ([\chi^{-q_1}, \dots, 1, \chi, \dots, \chi^{q_1}]_1, [\chi^{-q_2}, \dots, 1, \chi, \dots, \chi^{q_2}]_2)$ .
- $\text{UpdSRS}_{\text{ldlog}}(\text{Ch}, \{\rho_j\}_{j=1}^n)$  parses  $\text{Ch}$  as  $([\{A_i\}_{i=-q_1}^{q_1}]_1, [\{B_i\}_{i=-q_2}^{q_2}]_2)$ , samples  $\tilde{\chi} \leftarrow_{\$} \mathbb{F}_p$ , and defines  $\tilde{\text{Ch}} := ([\{\tilde{\chi}^i A_i\}_{i=-q_1}^{q_1}]_1, [\{\tilde{\chi}^i B_i\}_{i=-q_2}^{q_2}]_2)$ .

Then

$$\Pr[\bar{\chi} \leftarrow \mathcal{A}^{\text{UpdO}}(\lambda)] \leq \text{negl}(\lambda),$$

where  $\left( [\{\bar{\chi}^i\}_{i=-q_1}^{q_1}]_1, [\{\bar{\chi}^i\}_{i=-q_2}^{q_2}]_2 \right)$  is the finalized challenge.

## 7.2 Sonic Protocol Rolled-out

In this section we present Sonic's constraint system and algorithms. Reader familiar with them may jump directly to the next section.

<b>GenSRS</b> ( $1^\lambda, \text{max}$ ) <hr/> $\alpha, \chi \leftarrow \mathbb{F}_p^2$ <b>return</b> $\left[ \{\chi^i\}_{i=-Q_{\text{mult}}}^{Q_{\text{mult}}}, \{\alpha\chi^i\}_{i=-Q_{\text{mult}}, i \neq 0}^{Q_{\text{mult}}} \right]_1,$ $\left[ \{\chi^i, \alpha\chi^i\}_{i=-Q_{\text{mult}}}^{Q_{\text{mult}}} \right]_2, [\alpha]_T$	<b>Com</b> (srs, max, $f(X)$ ) <hr/> $c(X) \leftarrow \alpha \cdot X^{d-\text{max}} f(X)$ <b>return</b> $[c]_1 = [c(\chi)]_1$
<b>Op</b> (srs, $z, s, f(X)$ ) <hr/> $o(X) \leftarrow \frac{f(X) - f(z)}{X - z}$ <b>return</b> $[o(\chi)]_1$	<b>Verify</b> (srs, max, $[c]_1, z, s, [o(\chi)]_1$ ) <hr/> <b>if</b> $[o(\chi)]_1 \bullet [\alpha\chi]_2 + [s - zo(\chi)]_1 \bullet [\alpha]_2 =$ $[c]_1 \bullet [\chi^{-d+\text{max}}]_2$ <b>then return 1</b> <b>else return 0.</b>

Fig. 6:  $\text{PC}_S$  polynomial commitment scheme.

**The Constraint System** Fig. 6 presents a variant of KZG [41] polynomial commitment schemes used in Sonic. Sonic's system of constraints composes of three  $Q_{\text{mult}}$ -long vectors  $\mathbf{a}, \mathbf{b}, \mathbf{c}$  which corresponds to left and right inputs to multiplication gates and their outputs. It hence holds  $\mathbf{a} \cdot \mathbf{b} = \mathbf{c}$ .

There is also  $Q_{\text{lin}}$  linear constraints of the form

$$\mathbf{a}\mathbf{u}_q + \mathbf{b}\mathbf{v}_q + \mathbf{c}\mathbf{w}_q = k_q,$$

where  $\mathbf{u}_q, \mathbf{v}_q, \mathbf{w}_q$  are vectors for the  $q$ -th linear constraint with instance value  $k_q \in \mathbb{F}_p$ . Furthermore define polynomials

$$\begin{aligned} \mathbf{u}_i(Y) &= \sum_{q=1}^{Q_{\text{lin}}} Y^{q+Q_{\text{mult}}} u_{q,i}, & \mathbf{w}_i(Y) &= -Y^i - Y^{-i} + \sum_{q=1}^{Q_{\text{lin}}} Y^{q+Q_{\text{mult}}} w_{q,i}, \\ \mathbf{v}_i(Y) &= \sum_{q=1}^{Q_{\text{lin}}} Y^{q+Q_{\text{mult}}} v_{q,i}, & \mathbf{k}(Y) &= \sum_{q=1}^{Q_{\text{lin}}} Y^{q+Q_{\text{mult}}} k_q. \end{aligned} \quad (3)$$

Sonic constraint system requires that

$$\mathbf{a}^\top \cdot \mathbf{u}(Y) + \mathbf{b}^\top \cdot \mathbf{v}(Y) + \mathbf{c}^\top \cdot \mathbf{w}(Y) + \sum_{i=1}^{Q_{\text{mult}}} a_i b_i (Y^i + Y^{-i}) - \mathbf{k}(Y) = 0. \quad (4)$$

In Sonic we will use commitments to the following polynomials.

$$\begin{aligned} r(X, Y) &= \sum_{i=1}^{Q_{\text{mult}}} (a_i X^i Y^i + b_i X^{-i} Y^{-i} + c_i X^{-i-Q_{\text{mult}}} Y^{-i-Q_{\text{mult}}}) \\ s(X, Y) &= \sum_{i=1}^{Q_{\text{mult}}} (u_i(Y) X^{-i} + v_i(Y) X^i + w_i(Y) X^{i+Q_{\text{mult}}}) \\ t(X, Y) &= r(X, 1)(r(X, Y) + s(X, Y)) - \mathbf{k}(Y). \end{aligned}$$

Polynomials  $r(X, Y)$ ,  $s(X, Y)$ ,  $t(X, Y)$  are designed such that  $t(0, Y) = \mathbf{a}^\top \cdot \mathbf{u}(Y) + \mathbf{b}^\top \cdot \mathbf{v}(Y) + \mathbf{c}^\top \cdot \mathbf{w}(Y) + \sum_{i=1}^{\mathcal{Q}_{\text{mult}}} a_i b_i (Y^i + Y^{-i}) - k(Y)$ . That is, the prover is asked to show that  $t(0, Y) = 0$ , cf. Eq. (4).

Furthermore, the commitment system in Sonic is designed such that it is infeasible for a PPT algorithm to commit to a polynomial with non-zero constant term.

**Algorithms Rolled out** Sonic *SRS generation*  $\text{GenSRS}(\mathbf{R})$ . The SRS generating algorithm picks randomly  $\alpha, \chi \leftarrow_{\$} \mathbb{F}_p$  and outputs

$$\text{srs} = \left( [\{\chi^i\}_{i=-d}^d, \{\alpha\chi^i\}_{i=-d, i \neq 0}]_1, [\{\chi^i, \alpha\chi^i\}_{i=-d}^d]_2, [\alpha]_T \right)$$

Sonic *prover*  $\text{P}(\text{srs}, x, w = \mathbf{a}, \mathbf{b}, \mathbf{c})$ .

**Message 1** The prover picks randomly randomizers  $c_{\mathcal{Q}_{\text{mult}}+1}, c_{\mathcal{Q}_{\text{mult}}+2}, c_{\mathcal{Q}_{\text{mult}}+3}, c_{\mathcal{Q}_{\text{mult}}+4} \leftarrow_{\$} \mathbb{F}_p$ . Sets  $r(X, Y) \leftarrow r(X, Y) + \sum_{i=1}^4 c_{\mathcal{Q}_{\text{mult}}+i} X^{-2\mathcal{Q}_{\text{mult}}-i}$ . Commits to  $r(X, 1)$  and outputs  $[r]_1 \leftarrow \text{Com}(\text{srs}, \mathcal{Q}_{\text{mult}}, r(X, 1))$ . Then it computes challenge  $y = \mathcal{H}(\tilde{\pi}[0..1])$ .

**Message 2** P commits to  $t(X, y)$  and outputs  $[t]_1 \leftarrow \text{Com}(\text{srs}, d, t(X, y))$ . Then it gets a challenge  $z = \mathcal{H}(\tilde{\pi}[0..2])$ .

**Message 3** The prover computes commitment openings. That is, it outputs

$$\begin{aligned} [o_a]_1 &= \text{Op}(\text{srs}, z, r(z, 1), r(X, 1)) \\ [o_b]_1 &= \text{Op}(\text{srs}, yz, r(yz, 1), r(X, 1)) \\ [o_t]_1 &= \text{Op}(\text{srs}, z, t(z, y), t(X, y)) \end{aligned}$$

along with evaluations  $a' = r(z, 1)$ ,  $b' = r(yz, 1)$ ,  $t' = t(z, y)$ . Then it engages in the signature of correct computation playing the role of the helper, i.e. it commits to  $s(X, y)$  and sends the commitment  $[s]_1$ , commitment opening

$$[o_s]_1 = \text{Op}(\text{srs}, z, s(z, y), s(X, y)),$$

and  $s' = s(z, y)$ . Then it obtains a challenge  $u = \mathcal{H}(\tilde{\pi}[0..3])$ .

**Message 4** For the next message the prover computes  $[c]_1 \leftarrow \text{Com}(\text{srs}, d, s(u, Y))$  and computes commitments' openings

$$\begin{aligned} [w]_1 &= \text{Op}(\text{srs}, u, s(u, y), s(X, y)), \\ [q_y]_1 &= \text{Op}(\text{srs}, y, s(u, y), s(u, Y)), \end{aligned}$$

and returns  $[w]_1, [q_y]_1, s = s(u, y)$ . Eventually the prover gets the last challenge  $z' = \mathcal{H}(\tilde{\pi}[0..4])$ .

**Message 5** For the final message, P computes opening  $[q_{z'}]_1 = \text{Op}(\text{srs}, z', s(u, z'), s(u, X))$  and outputs  $[q_{z'}]_1$ .

Sonic *verifier*  $\text{V}(\text{srs}, x, \pi)$ . The verifier in Sonic runs as subroutines the verifier for the polynomial commitment. That is it sets  $t' = a'(b' + s') - k(y)$  and checks the following:

$$\begin{array}{ll} \text{PC}_{\mathcal{S}}.V(\text{srs}, \mathcal{Q}_{\text{mult}}, [r]_1, z, a', [o_a]_1), & \text{PC}_{\mathcal{S}}.V(\text{srs}, d, [s]_1, u, s, [w]_1), \\ \text{PC}_{\mathcal{S}}.V(\text{srs}, \mathcal{Q}_{\text{mult}}, [r]_1, yz, b', [o_b]_1), & \text{PC}_{\mathcal{S}}.V(\text{srs}, d, [c]_1, y, s, [q_y]_1), \\ \text{PC}_{\mathcal{S}}.V(\text{srs}, d, [t]_1, z, t', [o_t]_1), & \text{PC}_{\mathcal{S}}.V(\text{srs}, d, [c]_1, z', s(u, z'), [q_{z'}]_1), \\ \text{PC}_{\mathcal{S}}.V(\text{srs}, d, [s]_1, z, s', [o_s]_1), & \end{array}$$

and accepts the proof iff all the checks holds. Note that the value  $s(u, z')$  that is recomputed by the verifier uses separate challenges  $u$  and  $z'$ . This enables the batching of many proof and outsourcing of this part of the proof to an untrusted helper.

### 7.3 Unique Opening Property of $\text{PC}_{\mathcal{S}}$

**Lemma 8.**  $\text{PC}_{\mathcal{S}}$  has the unique opening property in the AGM.

*Proof.* Let  $z \in \mathbb{F}_p$  be the attribute the polynomial is evaluated at,  $[c]_1 \in \mathbb{G}$  be the commitment,  $s \in \mathbb{F}_p$  the evaluation value, and  $o \in \mathbb{G}$  be the commitment opening. We need to show that for every PPT adversary  $\mathcal{A}$  probability

$$\Pr \left[ \begin{array}{l} \text{Verify}(\text{srs}, [c]_1, z, s, [o]_1) = 1, \\ \text{Verify}(\text{srs}, [c]_1, z, \tilde{s}, [\tilde{o}]_1) = 1 \end{array} \middle| ([c]_1, z, s, [o]_1, [\tilde{o}]_1) \leftarrow \mathcal{A}^{\text{UpdO}}(1^\lambda, \text{max}) \right]$$

is at most negligible.

As noted in [30, Lemma 2.2] it is enough to upper bound the probability of the adversary succeeding against the ideal verifier, who verifies equality between polynomials.

For a polynomial  $f$ , its degree upper bound  $\text{max}$ , evaluation point  $z$ , evaluation result  $s$ , and opening  $[o(X)]_1$  the ideal verifier checks that

$$\alpha(X^{\text{d-max}} f(X) \cdot X^{-\text{d+max}} - s) \equiv \alpha \cdot o(X)(X - z), \quad (5)$$

what is equivalent to

$$f(X) - s \equiv o(X)(X - z). \quad (6)$$

Since  $o(X)(X - z) \in \mathbb{F}_p[X]$  then from the uniqueness of polynomial composition, there is only one  $o(X)$  that fulfills the equation above.  $\square$

#### 7.4 Unique Response Property

The unique response property of  $\mathbf{S}_{\text{FS}}$  follows from the unique opening property of the polynomial commitment scheme  $\mathbf{PC}_S$ .

**Lemma 9.** *If a polynomial commitment scheme  $\mathbf{PC}_S$  is evaluation binding with security loss  $\varepsilon_{\text{bind}}(\lambda)$  and has unique openings property with security loss  $\varepsilon_{\text{op}}(\lambda)$ , then  $\mathbf{S}_{\text{FS}}$  is 2-UR against algebraic adversaries with security loss*

$$2 \cdot \varepsilon_{\text{bind}}(\lambda) + \varepsilon_{\text{op}}(\lambda).$$

*Proof.* Let  $\mathcal{A}$  be an algebraic adversary tasked to break the 2-UR-ness of  $\mathbf{S}_{\text{FS}}$ . We note that to show 2-UR-ness is enough to show that the first prover's message determines, along with the verifiers challenges, the rest of it. We denote by  $\pi^0$  and  $\pi^1$  the two proofs for the same statement the adversary outputs. To distinguish polynomials and commitments which an honest prover sends in the proof from the polynomials and commitments computed by the adversary we write the latter using indices 0 and 1 (two indices as we have two transcripts), e.g. to describe the quotient polynomial provided by the adversary we write  $t^0$  and  $t^1$  instead of  $t$  as in the description of the protocol.

We note that since the unique response property requires from  $\pi^0$  and  $\pi^1$  that the first place they possibly differ is the 3-th prover's message, then the challenge  $z$ , that is picked by the adversary after the 2-rd message is the same in both transcripts. This challenge determines the evaluation point of polynomials  $r(X, 1)$ ,  $t(X, y)$  which commitments are already sent.

In its third message, the prover provides evaluations of these polynomials along with their openings at  $z$  or  $yz$ . Note that the adversary can output two accepting proofs that differ on their third message only if it manages to break evaluation binding of one of the opening. Since the commitment scheme is evaluation binding with security loss  $\varepsilon_{\text{bind}}(\lambda)$ , the adversary can make  $\pi^0$  and  $\pi^1$  differ on the third message with probability at most  $\varepsilon_{\text{bind}}(\lambda)$ .

Similarly, in its fourth message, the prover provides an evaluation at  $u$  of polynomial  $s(X, y)$ , an evaluation at  $y$  of  $s(u, Y)$ , and the corresponding openings. Note that the adversary can output two accepting proofs that differ on their fourth message only if it manages to break evaluation binding of one of the opening. Since the commitment scheme is evaluation binding with security loss  $\varepsilon_{\text{bind}}(\lambda)$ , the adversary can make  $\pi^0$  and  $\pi^1$  differ on the fourth message with probability at most  $\varepsilon_{\text{bind}}(\lambda)$ .

Next, assume that the transcripts are the same up to the fourth message, but differ at the fifth. In that message, the adversary provides openings of the evaluations. Since the unique opening property, the adversary can open the valid evaluation of a polynomial to two different values with probability at most  $\varepsilon_{\text{op}}(\lambda)$ .

By the union bound, the adversary is able to break the unique response property with probability upper bounded by  $2\varepsilon_{\text{bind}}(\lambda) + \varepsilon_{\text{op}}(\lambda)$ .  $\square$

## 7.5 Rewinding-Based Knowledge Soundness

**Lemma 10.**  $\mathbf{S}_{\text{FS}}$  is  $(2, n + 1)$ -rewinding-based knowledge sound against algebraic adversaries who make up to  $q$  random oracle queries with security loss

$$\varepsilon_{\text{ks}}(\lambda, \text{acc}, q) \leq \left( 1 - \frac{\text{acc} - (q + 1) \binom{n}{p}}{1 - \frac{n}{p}} \right) + (n + 1) \cdot \varepsilon_{\text{uldlog}}(\lambda),$$

Here  $\text{acc}$  is a probability that the adversary outputs an accepting proof, and  $\varepsilon_{\text{uldlog}}(\lambda)$  is the security of  $(\max, \max)$ -uldlog assumption.

Let  $\mathcal{A}^{\mathcal{H}, \text{UpdO}}(1^\lambda; r)$  be the adversary who outputs  $(x, \pi)$  such that  $\mathbf{S}_{\text{FS}}.V$  accepts the proof. Let  $\mathcal{T}$  be the tree-building algorithm of Lemma 4 that outputs a tree  $T$ , and let  $\text{Ext}_{\text{ks}}$  be an extractor that given the tree output by  $\mathcal{T}$  reveals the witness for  $x$ . The main idea of the proof is to show that an adversary who breaks rewinding-based knowledge soundness can be used to break a uldlog-problem instance. The proof goes by game hops. Note that since the tree branches after  $\mathcal{A}$ 's 2-nd message, the instance  $x$ , commitments  $[r(\chi, 1), t(\chi, y)]_1$ , and challenge  $y$  are the same in all the transcripts. Also, the tree branches after the second adversary's message where the challenge  $z$  is presented, thus tree  $T$  is built using different values of  $z$ . We consider the following games.

**Game 0:** In this game, the adversary wins if it outputs a valid instance-proof pair  $(x, \pi)$ , and the extractor  $\text{Ext}_{\text{ks}}$  does not manage to output a witness  $w$  such that  $\mathbf{R}(x, w)$  holds.

**Game 1:** In this game, the environment aborts the game if the tree building algorithm  $\mathcal{T}$  fails in building a tree of accepting transcripts  $T$ .

**Game 0 to Game 1:** By Lemma 4 probability that Game 1 is aborted, while Game 0 is not, is, at most

$$1 - \frac{\text{acc} - (q + 1) \binom{n}{p}}{1 - \frac{n}{p}}.$$

**Game 2:** In this game the environment additionally aborts if at least one of its proofs in  $T$  is not accepting by an ideal verifier.

**Game 1 to Game 2:** As usual, we show a reduction that breaks an instance of a uldlog assumption when Game 2 is aborted, while Game 1 is not.

Let  $\mathcal{R}_{\text{uldlog}}$  be a reduction that gets as input an  $(\max, \max)$ -uldlog instance  $[\chi^{-\max}, \dots, 1, \dots, \chi^{\max}]_1, [\chi^{-\max}, \dots, 1, \dots, \chi^{\max}]_2$ . Then it can update the instance to another one  $[\chi'^{-\max}, \dots, 1, \dots, \chi'^{\max}]_1, [\chi'^{-\max}, \dots, 1, \dots, \chi'^{\max}]_2$ . Eventually, the reduction outputs  $\chi'$ . The reduction  $\mathcal{R}_{\text{uldlog}}$  proceeds as follows. First, it builds  $\mathcal{A}$ 's SRS  $\text{srs}$  using the input uldlog instance. Then it processes the adversary's update query by adding it to the list  $Q_{\text{srs}}$  and passing it to its own update oracle getting instance  $[\chi'^{-\max}, \dots, 1, \dots, \chi'^{\max}]_1, [\chi'^{-\max}, \dots, 1, \dots, \chi'^{\max}]_2$ . The updated SRS  $\text{srs}'$  is then computed and given to  $\mathcal{A}$ .  $\mathcal{R}_{\text{uldlog}}$  also takes care of the random oracle queries made by  $\mathcal{A}$ . It picks their answers honestly and writes them in  $Q_{\mathcal{H}}$ . The reduction then starts  $\mathcal{T}(\text{srs}, \mathcal{A}, r, Q_{\mathcal{H}}, Q_{\text{srs}})$ .

Let  $(1, T)$  be the output returned by  $\mathcal{T}$ . Let  $x$  be a relation proven in  $T$ . Consider a transcript  $\pi \in T$  such that  $\text{ve}_{x, \pi}(X) \neq 0$ , but  $\text{ve}_{x, \pi}(\chi') = 0$ . Since  $\mathcal{A}$  is algebraic, all group elements included in  $T$  are extended by their representation as a combination of the input  $\mathbb{G}_1$ -elements. Hence, all coefficients of the verification equation polynomial  $\text{ve}_{x, \pi}(X)$  are known. Eventually, the reduction finds  $\text{ve}_{x, \pi}(X)$  zero points and returns  $\chi'$  which is one of them.

Hence, the probability that the adversary wins in Game 2 but does not win in Game 1 is upper-bounded by  $(n + 1) \cdot \varepsilon_{\text{uldlog}}(\lambda)$ .

**Conclusion:** Note that the adversary can win in Game 2 only if  $\mathcal{T}$  manages to produce a tree of accepting transcript  $T$ , such that each of the transcripts in  $T$  is accepting by an ideal verifier. Note that since  $\mathcal{T}$  produces  $(n + 1)$  accepting transcripts for different challenges  $z$ , it obtains the same number of different evaluations of polynomial  $t(z, y)$  what allows to extract the witness, cf. [46].

Hence, the probability that the adversary wins in Game 0 is upper-bounded by

$$\varepsilon_{\text{ks}}(\lambda, \text{acc}, q) \leq \left( 1 - \frac{\text{acc} - (q + 1) \binom{n}{p}}{1 - \frac{n}{p}} \right) + (n + 1) \cdot \varepsilon_{\text{uldlog}}(\lambda).$$

## 7.6 Trapdoor-Less Zero-Knowledge of Sonic

**Lemma 11.**  $\mathbf{S}_{\text{FS}}$  is 2-programmable trapdoor-less zero-knowledge.

*Proof.* The simulator proceeds as follows.

1. Pick randomly vectors  $\mathbf{a}$ ,  $\mathbf{b}$  and set

$$\mathbf{c} = \mathbf{a} \cdot \mathbf{b}. \quad (7)$$

2. Pick randomizers  $c_{Q_{\text{mult}}+1}, \dots, c_{Q_{\text{mult}}+4}$ , honestly compute polynomials  $r(X, Y)$ ,  $r'(X, Y)$ ,  $s(X, Y)$  and pick randomly challenges  $y, z$ .
3. Output commitment  $[r]_1 \leftarrow \text{Com}(\text{srs}, Q_{\text{mult}}, r(X, 1))$  and challenge  $y$ .
4. Compute

$$\begin{aligned} a' &= r(z, 1), \\ b' &= r(z, y), \\ s' &= s(z, y). \end{aligned}$$

5. Pick polynomial  $t(X, Y)$  such that

$$\begin{aligned} t(X, y) &= r(X, 1)(r(X, y) + s(X, y)) - k(Y) \\ t(0, y) &= 0 \end{aligned}$$

6. Output commitment  $[t]_1 = \text{Com}(\text{srs}, d, t(X, y))$  and challenge  $z$ .
7. Continue following the protocol.

We note that the simulation is perfect. This comes since, except polynomial  $t(X, Y)$  all polynomials are computed following the protocol. For polynomial  $t(X, Y)$  we observe that in a case of both real and simulated proof the verifier only learns commitment  $[t]_1 = t(\chi, y)$  and evaluation  $t' = t(z, y)$ . Since the simulator picks  $t(X, Y)$  such that

$$t(X, y) = r(X, 1)(r(X, y) + s(X, y)) - k(Y)$$

Values of  $[t]_1$  are equal in both proofs. Furthermore, the simulator picks its polynomial such that  $t(0, y) = 0$ , hence it does not need the trapdoor to commit to it. (Note that the proof system's SRS does not allow to commit to polynomials which have non-zero constant term).  $\square$

*Remark 6.* As noted in [46], Sonic is statistically subversion zero-knowledge (Sub-ZK). As noted in [1], one way to achieve subversion zero-knowledge is to utilize an extractor that extracts a SRS trapdoor from a SRS-generator. Unfortunately, a NIZK made subversion zero-knowledge by this approach cannot achieve perfect Sub-ZK as one has to count in the probability of extraction failure. However, with the simulation presented in Lemma 11, the trapdoor is not required for the simulator as it is able to simulate the execution of the protocol just by picking appropriate (honest) verifier's challenges. This result transfers to  $\mathbf{S}_{\text{FS}}$ , where the simulator can program the random oracle to provide challenges that fits it.

## 7.7 Simulation Extractability of $\mathbf{S}_{\text{FS}}$

Since Lemmas 9 to 11 hold,  $\mathbf{S}_{\text{FS}}$  is 2-UR, rewinding-based knowledge sound and trapdoor-less zero-knowledge. We now make use of Theorem 1 and show that  $\mathbf{S}_{\text{FS}}$  is simulation-extractable as defined in Definition 2.

**Corollary 2 (Simulation extractability of  $\mathbf{S}_{\text{FS}}$ ).**  $\mathbf{S}_{\text{FS}}$  is updatable simulation-extractable against any PPT adversary  $\mathcal{A}$  who makes up to  $q$  random oracle queries and returns an accepting proof with probability at least  $\text{acc}$  with extraction failure probability

$$\varepsilon_{\text{se}}(\lambda, \text{acc}, q) \leq \left( 1 - \frac{\text{acc} - \varepsilon_{\text{ur}}(\lambda) - (q+1)\varepsilon_{\text{err}}(\lambda)}{1 - \varepsilon_{\text{err}}(\lambda)} \right) + (n+1) \cdot \varepsilon_{\text{uldlog}}(\lambda),$$

where  $\varepsilon_{\text{err}}(\lambda) = \frac{n}{p}$ ,  $p$  is the size of the field, and  $n$  is the number of constrains in the circuit.

## 8 Non-malleability of Marlin

We show that Marlin is simulation-extractable. To that end, we show that Marlin has all the required properties: has unique response property, is rewinding-based knowledge sound, and its simulator can provide indistinguishable proofs without a trapdoor, just by programming the random oracle.

### 8.1 Marlin Protocol Rolled-out

Marlin uses RICS as its arithmetization method. Given instance  $x$ , witness  $w$  and  $|H| \times |H|$  matrices  $A, B, C$ , the prover shows that  $A(x^\top, w^\top)^\top \circ B(x^\top, w^\top)^\top = C(x^\top, w^\top)^\top$ , where  $\circ$  denotes entry-wise product.

We assume that the matrices have at most  $|K|$  non-zero entries. Obviously,  $|K| \leq |H|^2$ . Let  $b = 3$ , the upper-bound of polynomial evaluations the prover has to provide for each of the sent polynomials. Denote by  $d$  an upper-bound for  $\{|H| + 2b - 1, 2|H| + b - 1, 6|K| - 6\}$ .

The idea of showing that the constraint system is fulfilled is as follows. Denote by  $z = (x, w)$ . The prover computes polynomials  $z_A(X), z_B(X), z_C(X)$  which encode vectors  $Az, Bz, Cz$  and have degree  $< |H|$ . Importantly, when constraints are fulfilled,  $z_A(X)z_B(X) - z_C(X) = h_0(X)Z_H(X)$ , for some  $h_0(X)$  and vanishing polynomial  $Z_H(X)$ . The prover sends commitments to these polynomials and shows that they have been computed correctly. More precisely, it shows that

$$\forall M \in \{A, B, C\}, \forall \kappa \in H, z_M(\kappa) = \sum_{\iota \in H} M[\kappa, \iota]z(\iota). \quad (8)$$

The ideal verifier checks the following equalities

$$\begin{aligned} h_3(\beta_3)Z_K(\beta_3) &= a(\beta_3) - b(\beta_3)(\beta_3g_3(\beta_3) + \sigma_3/|K|) \\ r(\alpha, \beta_2)\sigma_3 &= h_2(\beta_2)Z_H(\beta_2) + \beta_2g_2(\beta_2) + \sigma_2/|H| \\ s(\beta_1) + r(\alpha, \beta_1)\left(\sum_M \eta_M z_M(\beta_1)\right) - \sigma_2z(\beta_1) &= h_1(\beta_1)Z_H(\beta_1) + \beta_1g_1(\beta_1) + \sigma_1/|H| \\ z_A(\beta_1)z_B(\beta_1) - z_C(\beta_1) &= h_0(\beta_1)Z_H(\beta_1) \end{aligned} \quad (9)$$

where  $g_i(X), h_i(X), i \in [1..3], a(X), b(X), \sigma_1, \sigma_2, \sigma_3$  are polynomials and variables required by the sumcheck protocol which allows the verifier to efficiently verify that Eq. (8) holds.

### 8.2 Unique Response Property

**Lemma 12.** *Let PC be a commitment of knowledge that is evaluation binding with security loss  $\varepsilon_{\text{bind}}(\lambda)$  and has unique opening property with security loss  $\varepsilon_{\text{op}}(\lambda)$ . Then  $\mathbf{M}_{\text{FS}}$  is 2-UR against algebraic adversaries with security loss  $2 \cdot \varepsilon_{\text{bind}}(\lambda) + \varepsilon_{\text{op}}(\lambda)$ .*

*Proof.* The proof is similar to the proof of Lemma 5 and Lemma 9. An adversary who can break the 2-unique response property of  $\mathbf{M}_{\text{FS}}$  can be either used to break the commitment scheme's evaluation binding or unique opening property. The former happens with the probability upper-bounded by  $2 \cdot \varepsilon_{\text{bind}}(\lambda)$ , the latter with probability at most  $\varepsilon_{\text{op}}(\lambda)$ . By the union bound, the adversary is able to break the unique response property with probability upper bounded by  $2 \cdot \varepsilon_{\text{bind}}(\lambda) + \varepsilon_{\text{op}}(\lambda)$ .  $\square$

### 8.3 Rewinding-Based Knowledge Soundness

**Lemma 13.**  *$\mathbf{M}_{\text{FS}}$  is  $(2, 2n + 3)$ -rewinding-based knowledge sound against algebraic adversaries who make up to  $q$  random oracle queries with security loss*

$$\varepsilon_{\text{ks}}(\lambda, \text{acc}, q) \leq \left(1 - \frac{\text{acc} - (q + 1) \binom{2n+2}{p}}{1 - \frac{2n+2}{p}}\right) + (2n + 3) \cdot \varepsilon_{\text{udlog}}(\lambda),$$

Here  $\text{acc}$  is a probability that the adversary outputs an acceptable proof, and  $\varepsilon_{\text{udlog}}(\lambda)$  is the security of  $(2n + 2, 1)$ -udlog assumption.

*Proof.* The proof is similar to the proof of Lemma 6 and Lemma 10. We use Attema et al. [5, Proposition 2] to bound the probability that the tree-building algorithm  $\mathcal{T}$  does not obtain a tree of acceptable transcript in an expected number of runs. This happens with probability at most

$$1 - \frac{\text{acc} - (q + 1) \binom{2n+2}{p}}{1 - \frac{2n+2}{p}}$$

Let  $T$  be the tree output by  $\mathcal{T}$ . If one of the proofs in  $T$  is not accepting by the ideal verifier, one can break an instance of an updatable dlog assumption which happens with probability at most  $(2n + 3) \cdot \varepsilon_{\text{udlog}}(\lambda)$ . In the case that all the transcripts are accepting by the ideal verifier, but  $\text{Ext}_{\text{ks}}$  fails to extract a valid witness from  $T$ , one can break the soundness of the ideal verifier in one of the transcripts. Taking a union bound completes the proof.  $\square$

#### 8.4 Trapdoor-Less Zero-Knowledge of Marlin

**Lemma 14.**  $\mathbf{M}_{\text{FS}}$  is 2-programmable trapdoor-less zero-knowledge.

*Proof.* The simulator follows the protocol except that it picks the challenges  $\alpha, \eta_A, \eta_B, \eta_C, \beta_1, \beta_2, \beta_3$  before it picks the polynomials it sends.

First, it picks  $\tilde{z}_A(X), \tilde{z}_B(X)$  at random and  $\tilde{z}_C(X)$  such that  $\tilde{z}_A(\beta_1)\tilde{z}_B(\beta_1) = \tilde{z}_C(\beta_1)$ . Given the challenges and polynomials  $\tilde{z}_A(X), \tilde{z}_B(X), \tilde{z}_C(X)$  the simulator computes  $\sigma_1 \leftarrow \sum_{\kappa \in H} s(\kappa) + r(\alpha, X)(\sum_{M \in \{A, B, C\}} \eta_M \tilde{z}_M(X)) - \sum_{M \in \{A, B, C\}} \eta_M r_M(\alpha, X) \tilde{z}(X)$ .

Then the simulator starts the protocol and follows it, except it programs the random oracle such that on partial transcripts, it returns the challenges already picked by Sim.

#### 8.5 Simulation Extractability of $\mathbf{M}_{\text{FS}}$

Since Lemmas 12 to 14 hold,  $\mathbf{M}_{\text{FS}}$  is 2-UR, rewinding-based knowledge sound and trapdoor-less zero-knowledge. By making use of Theorem 1, we conclude that  $\mathbf{M}_{\text{FS}}$  is simulation-extractable as defined in Definition 2.

**Corollary 3 (Simulation extractability of  $\mathbf{M}_{\text{FS}}$ ).**  $\mathbf{M}_{\text{FS}}$  is updatable simulation-extractable against any PPT adversary  $\mathcal{A}$  who makes up to  $q$  random oracle queries and returns an acceptable proof with probability at least  $\text{acc}$  with extraction failure probability

$$\varepsilon_{\text{se}}(\lambda, \text{acc}, q) \leq \left(1 - \frac{\text{acc} - \varepsilon_{\text{ur}}(\lambda) - (q + 1)\varepsilon_{\text{err}}(\lambda)}{1 - \varepsilon_{\text{err}}(\lambda)}\right) + (2n + 3) \cdot \varepsilon_{\text{udlog}}(\lambda),$$

where  $\varepsilon_{\text{err}}(\lambda) = \frac{2n+2}{p}$ ,  $p$  is the size of the field, and  $n$  is the number of constrains in the circuit.

## References

1. B. Abdolmaleki, K. Bagheri, H. Lipmaa, and M. Zajac. A subversion-resistant SNARK. In T. Takagi and T. Peyrin, editors, *ASIACRYPT 2017, Part III*.
2. B. Abdolmaleki, S. Ramacher, and D. Slamanig. Lift-and-shift: Obtaining simulation extractable subversion and updatable SNARKs generically. In J. Ligatti, X. Ou, J. Katz, and G. Vigna, editors, *ACM CCS 2020*.
3. S. Atapoor and K. Bagheri. Simulation extractability in groth's zk-SNARK. Cryptology ePrint Archive, Report 2019/641, 2019. <https://eprint.iacr.org/2019/641>.
4. T. Attema, S. Fehr, and M. Klooß. Fiat-shamir transformation of multi-round interactive proofs. Cryptology ePrint Archive, Report 2021/1377.
5. T. Attema, S. Fehr, and M. Klooß. Fiat-shamir transformation of multi-round interactive proofs. Cryptology ePrint Archive, Report 2021/1377, 2021. <https://eprint.iacr.org/2021/1377>.
6. K. Bagheri, M. Kohlweiss, J. Siim, and M. Volkhov. Another look at extraction and randomization of groth's zk-SNARK. Cryptology ePrint Archive, Report 2020/811, 2020. <https://eprint.iacr.org/2020/811>.
7. M. Bellare and G. Neven. Multi-signatures in the plain public-key model and a general forking lemma. In A. Juels, R. N. Wright, and S. De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 390–399. ACM Press, Oct. / Nov. 2006.
8. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In D. E. Denning, R. Pyle, R. Ganesan, R. S. Sandhu, and V. Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, Nov. 1993.
9. M. Ben-Or, O. Goldreich, S. Goldwasser, J. Håstad, J. Kilian, S. Micali, and P. Rogaway. Everything provable is provable in zero-knowledge. In S. Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 37–56. Springer, Heidelberg, Aug. 1990.
10. E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev. Scalable, transparent, and post-quantum secure computational integrity. Cryptology ePrint Archive, Report 2018/046, 2018. <https://eprint.iacr.org/2018/046>.
11. E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014.
12. E. Ben-Sasson, A. Chiesa, D. Genkin, E. Tromer, and M. Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108. Springer, Heidelberg, Aug. 2013.
13. E. Ben-Sasson, A. Chiesa, and N. Spooner. Interactive oracle proofs. In M. Hirt and A. D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 31–60. Springer, Heidelberg, Oct. / Nov. 2016.

14. E. Ben-Sasson, A. Chiesa, E. Tromer, and M. Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In K. Fu and J. Jung, editors, *USENIX Security 2014*, pages 781–796. USENIX Association, Aug. 2014.
15. N. Bitansky, A. Chiesa, Y. Ishai, R. Ostrovsky, and O. Paneth. Succinct non-interactive arguments via linear interactive proofs. In A. Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 315–333. Springer, Heidelberg, Mar. 2013.
16. J. Bootle, A. Cerulli, P. Chaidos, J. Groth, and C. Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In M. Fischlin and J.-S. Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 327–357. Springer, Heidelberg, May 2016.
17. S. Bowe and A. Gabizon. Making groth’s zk-SNARK simulation extractable in the random oracle model. *Cryptology ePrint Archive*, Report 2018/187, 2018. <https://eprint.iacr.org/2018/187>.
18. G. Brassard, D. Chaum, and C. Crépeau. Minimum disclosure proofs of knowledge. *Journal of computer and system sciences*, 37(2):156–189, 1988.
19. B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. 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.
20. J. Camenisch and M. Stadler. Efficient group signature schemes for large groups (extended abstract). In B. S. Kaliski Jr., editor, *CRYPTO ’97*, volume 1294 of *LNCS*, pages 410–424. Springer, Heidelberg, Aug. 1997.
21. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. *Cryptology ePrint Archive*, Report 2000/067, 2000. <https://eprint.iacr.org/2000/067>.
22. M. Chase and A. Lysyanskaya. On signatures of knowledge. In C. Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 78–96. Springer, Heidelberg, Aug. 2006.
23. A. Chiesa, Y. Hu, M. Maller, P. Mishra, N. Vesely, and N. P. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 738–768. Springer, Heidelberg, May 2020.
24. G. Danezis, C. Fournet, J. Groth, and M. Kohlweiss. Square span programs with applications to succinct NIZK arguments. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 532–550. Springer, Heidelberg, Dec. 2014.
25. Y. Dodis, K. Haralambiev, A. López-Alt, and D. Wichs. Efficient public-key cryptography in the presence of key leakage. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 613–631. Springer, Heidelberg, Dec. 2010.
26. S. Faust, M. Kohlweiss, G. A. Marson, and D. Venturi. On the non-malleability of the Fiat-Shamir transform. In S. D. Galbraith and M. Nandi, editors, *INDOCRYPT 2012*, volume 7668 of *LNCS*, pages 60–79. Springer, Heidelberg, Dec. 2012.
27. M. Fischlin. Communication-efficient non-interactive proofs of knowledge with online extractors. In V. Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 152–168. Springer, Heidelberg, Aug. 2005.
28. L. Fortnow. The complexity of perfect zero-knowledge (extended abstract). In A. Aho, editor, *19th ACM STOC*, pages 204–209. ACM Press, May 1987.
29. G. Fuchsbauer, E. Kiltz, and J. Loss. The algebraic group model and its applications. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Heidelberg, Aug. 2018.
30. A. Gabizon, Z. J. Williamson, and O. Ciobotaru. PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. *Cryptology ePrint Archive*, Report 2019/953, 2019. <https://eprint.iacr.org/2019/953>.
31. C. Ganesh, C. Orlandi, M. Pancholi, A. Takahashi, and D. Tschudi. Fiat-shamir bulletproofs are non-malleable (in the algebraic group model). *Cryptology ePrint Archive*, Report 2021/1393, 2021. <https://ia.cr/2021/1393>.
32. R. Gennaro, C. Gentry, B. Parno, and M. Raykova. Quadratic span programs and succinct NIZKs without PCPs. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.
33. Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. *Cryptology ePrint Archive*, Report 2017/454, 2017. <https://eprint.iacr.org/2017/454>.
34. O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). In *27th FOCS*, pages 174–187. IEEE Computer Society Press, Oct. 1986.
35. S. Goldwasser and Y. T. Kalai. On the (in)security of the Fiat-Shamir paradigm. In *44th FOCS*, pages 102–115. IEEE Computer Society Press, Oct. 2003.
36. J. Groth. Fully anonymous group signatures without random oracles. In K. Kurosawa, editor, *ASIACRYPT 2007*, volume 4833 of *LNCS*, pages 164–180. Springer, Heidelberg, Dec. 2007.
37. J. Groth. Short pairing-based non-interactive zero-knowledge arguments. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 321–340. Springer, Heidelberg, Dec. 2010.
38. J. Groth. On the size of pairing-based non-interactive arguments. In M. Fischlin and J.-S. Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.
39. J. Groth, M. Kohlweiss, M. Maller, S. Meiklejohn, and I. Miers. Updatable and universal common reference strings with applications to zk-SNARKs. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 698–728. Springer, Heidelberg, Aug. 2018.
40. J. Groth and M. Maller. Snarky signatures: Minimal signatures of knowledge from simulation-extractable SNARKs. In J. Katz and H. Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 581–612. Springer, Heidelberg, Aug. 2017.

41. A. Kate, G. M. Zaverucha, and I. Goldberg. Constant-size commitments to polynomials and their applications. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Heidelberg, Dec. 2010.
42. A. Kosba, Z. Zhao, A. Miller, Y. Qian, H. Chan, C. Papamanthou, R. Pass, a. shelat, and E. Shi. How to use SNARKs in universally composable protocols. Cryptology ePrint Archive, Report 2015/1093, 2015. <https://eprint.iacr.org/2015/1093>.
43. H. Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In R. Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 169–189. Springer, Heidelberg, Mar. 2012.
44. H. Lipmaa. Succinct non-interactive zero knowledge arguments from span programs and linear error-correcting codes. In K. Sako and P. Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 41–60. Springer, Heidelberg, Dec. 2013.
45. H. Lipmaa. Key-and-argument-updatable QA-NIZKs. Cryptology ePrint Archive, Report 2019/333, 2019. <https://eprint.iacr.org/2019/333>.
46. M. Maller, S. Bowe, M. Kohlweiss, and S. Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings. In L. Cavallaro, J. Kinder, X. Wang, and J. Katz, editors, *ACM CCS 2019*, pages 2111–2128. ACM Press, Nov. 2019.
47. S. Micali. CS proofs (extended abstracts). In *35th FOCS*, pages 436–453. IEEE Computer Society Press, Nov. 1994.
48. J. Miller. Coordinated disclosure of vulnerabilities affecting girault, bulletproofs, and plonk. <https://blog.trailofbits.com/2022/04/13/part-1-coordinated-disclosure-of-vulnerabilities-affecting-girault-bulletproofs-2022>.
49. B. Parno, J. Howell, C. Gentry, and M. Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013.
50. D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, June 2000.
51. L. Rotem and G. Segev. Tighter security for schnorr identification and signatures: A high-moment forking lemma for  $\Sigma$ -protocols. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 222–250, Virtual Event, Aug. 2021. Springer, Heidelberg.