

On the (In)security of SNARKs in the Presence of Oracles

Dario Fiore, Anca Nitulescu

► **To cite this version:**

Dario Fiore, Anca Nitulescu. On the (In)security of SNARKs in the Presence of Oracles. TCC 2016-B - Fourteenth IACR Theory of Cryptography Conference , Oct 2016, Beijing, China. 2016, Theory of Cryptography 14th International Conference, TCC 2016-B, Beijing, China, November 1-3, 2016, Proceedings. <hal-01378013>

HAL Id: hal-01378013

<https://hal.inria.fr/hal-01378013>

Submitted on 8 Oct 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the (In)security of SNARKs in the Presence of Oracles^{*}

Dario Fiore¹ and Anca Nitulescu²

¹ IMDEA Software Institute, Madrid, Spain
`dario.fiore@imdea.org`

² CNRS, ENS, INRIA, and PSL, Paris, France
`anca.nitulescu@ens.fr`

Abstract. In this work we study the feasibility of knowledge extraction for succinct non-interactive arguments of knowledge (SNARKs) in a scenario that, to the best of our knowledge, has not been analyzed before. While prior work focuses on the case of adversarial provers that may receive (statically generated) *auxiliary information*, here we consider the scenario where adversarial provers are given *access to an oracle*. For this setting we study if and under what assumptions such provers can admit an extractor. Our contribution is mainly threefold.

First, we formalize the question of extraction in the presence of oracles by proposing a suitable proof of knowledge definition for this setting. We call SNARKs satisfying this definition O-SNARKs. Second, we show how to use O-SNARKs to obtain formal and intuitive security proofs for three applications (homomorphic signatures, succinct functional signatures, and SNARKs on authenticated data) where we recognize an issue while doing the proof under the standard proof of knowledge definition of SNARKs. Third, we study whether O-SNARKs exist, providing both negative and positive results. On the negative side, we show that, assuming one way functions, there do not exist O-SNARKs in the standard model for every signing oracle family (and thus for general oracle families as well). On the positive side, we show that when considering signature schemes with appropriate restrictions on the message length O-SNARKs for the corresponding signing oracles exist, based on classical SNARKs and assuming extraction with respect to specific distributions of auxiliary input.

^{*} This article is based on an earlier article which appears in the proceedings of TCC 2016-B, © IACR 2016.

Table of Contents

On the (In)security of SNARKs in the Presence of Oracles	1
<i>Dario Fiore and Anca Nitulescu</i>	
1 Introduction	3
1.1 Extraction in the Presence of Oracles	4
1.2 An Overview of Our Results	5
1.3 Organization	8
2 Preliminaries	8
2.1 Succinct Non-Interactive Arguments	8
3 SNARKs in the presence of oracles	11
3.1 O-SNARKs: SNARKs in the presence of oracles	11
3.2 Non-Adaptive O-SNARKs	12
4 On the existence of O-SNARKs	13
4.1 O-SNARKs in the random oracle model from Micali’s CS proofs	14
4.2 Impossibility of O-SNARKs for every family of oracles, in the standard model	14
4.3 O-SNARKs for signing oracles from SNARKs in the random oracle model	16
4.4 O-SNARKs for signing oracles from SNARKs	19
4.5 O-SNARKs for (pseudo)random oracles from SNARKs	21
5 Applications of O-SNARKs	23
5.1 Homomorphic Signatures	23
5.2 Succinct Functional Signatures	33
5.3 SNARKs on authenticated data	42
A Additional Preliminaries	45
A.1 Digital Signatures	45
A.2 Universal One-Way Hash Functions	45
B The SNARK definition used in [BGI14]	46

1 Introduction

Succinct Arguments. Proof systems [GMR89] are fundamental in theoretical computer science and cryptography. Extensively studied aspects of proof systems are the expressivity of provable statements and the efficiency. Related to efficiency, it has been shown that statistically-sound proof systems are unlikely to allow for significant improvements in communication [BHZ87, GH98, GVW02, Wee05]. When considering proof systems for NP this means that, unless some complexity-theoretic collapses occur, in a statistically sound proof system any prover has to communicate, roughly, as much information as the size of the NP witness. The search of ways to beat this bound motivated the study of *computationally-sound* proof systems, also called *argument systems* [BCC88]. Assuming existence of collision-resistant hash functions, Kilian [Kil92] showed a four-message interactive argument for NP. In this protocol, membership of an instance x in an NP language with NP machine M can be proven with communication and verifier’s running time bounded by $p(\lambda, |M|, |x|, \log t)$, where λ is a security parameter, t is the NP verification time of machine M for the instance x , and p is a *universal* polynomial. Argument systems of this kind are called *succinct*.

Succinct Non-Interactive Arguments. Starting from Kilian’s protocol, Micali [Mic94] constructed a *one-message* succinct argument for NP whose soundness is set in the random oracle model. The fact that one-message succinct arguments are unlikely to exist for hard-enough languages in the plain model motivated the consideration of *two-message* non-interactive arguments, in which the verifier generates its message (a common reference string, if this can be made publicly available) ahead of time and independently of the statement to be proved. Such systems are called *succinct non-interactive arguments* (SNARGs) [GW11]. Several SNARGs constructions have been proposed [CL08, Mie08, Gro10, BCCT12, Lip12, BCC⁺14, GGPR13, BCI⁺13, PHGR13, BSCG⁺13, BCTV14] and the area of SNARGs has become popular in the last years with the proposal of constructions which gained significant improvements in efficiency. Noteworthy is that all such constructions are based on non-falsifiable assumptions [Nao03], a class of assumptions that is likely to be inherent in proving the security of SNARGs (without random oracles), as shown by Gentry and Wichs [GW11].

Almost all SNARGs are also *arguments of knowledge*—so called SNARKs [BCCT12, BCC⁺14]. Intuitively speaking, this property (which replaces soundness) says that every prover producing a convincing proof must “know” a witness. On the one hand, proof of knowledge turns out to be useful in many applications, such as delegation of computation where the untrusted worker contributes its own input to the computation, or recursive proof composition [Val08, BCCT13]. On the other hand, the formalization of proof of knowledge in SNARKs is a delicate point. Typically, the concept that the prover “must know” a witness is expressed by assuming that such knowledge can be efficiently extracted from the prover by means of a so-called *knowledge extractor*. In SNARKs, extractors are inherently non-black-box and proof of knowledge requires that for every adversarial prover \mathcal{A} generating an accepting proof π there must be an extractor $\mathcal{E}_{\mathcal{A}}$ that, given the same input of \mathcal{A} , outputs a valid witness.

Extraction with Auxiliary Input. Unfortunately, stated as above, proof of knowledge is insufficient for being used in many applications. The problem is that, when using SNARKs in larger cryptographic protocols, adversarial provers may get additional information which can contribute to the generation of adversarial proofs. To address this problem, a stronger, and more useful, definition of proof of knowledge requires that for any adversary \mathcal{A} there is an extractor $\mathcal{E}_{\mathcal{A}}$ such that, for any honestly generated crs and any polynomial-size auxiliary input aux , whenever $\mathcal{A}(\text{crs}, aux)$

returns an accepting proof, $\mathcal{E}_A(\text{crs}, \text{aux})$ outputs a valid witness. This type of definition is certainly more adequate when using SNARKs in larger cryptographic protocols, but it also introduces other subtleties. As first discussed in [HT98], extraction in the presence of arbitrary auxiliary input can be problematic, if not implausible. Formal evidence of this issue has been recently given in [BCPR14, BP15]. Bitansky et al. [BCPR14] show that, assuming indistinguishability obfuscation, there do not exist extractable one-way functions (and thus SNARKs) with respect to arbitrary auxiliary input of unbounded polynomial length. Boyle and Pass [BP15] generalize this result showing that assuming collision-resistant hash functions and differing-input obfuscation, there is a fixed auxiliary input distribution for which extractable one-way functions do not exist.

1.1 Extraction in the Presence of Oracles

In this work we continue the study on the feasibility of extraction by looking at a scenario that, to the best of our knowledge, has not been explicitly analyzed before. We consider the case in which adversarial provers run in interactive security experiments where they are given *access to an oracle*. For this setting we study if and under what assumptions such provers can admit an extractor.

Before giving more detail on our results, let us discuss a motivation for analyzing this scenario. To keep the presentation simple, here we give a motivation via a hypothetical example; more concrete applications are discussed later.

A CASE STUDY APPLICATION. Consider an application where Alice gets a collection of signatures generated by Bob, and she has to prove to a third party that she owns a valid signature of Bob on some message m such that $P(m) = 1$. Let us say that this application is secure if Alice, after asking for signatures on several messages, cannot cheat letting the third party accept for a false statement (i.e., $P(m) = 0$, or $P(m) = 1$ but Alice did not receive a signature on m). If messages are large and one wants to optimize bandwidth, SNARKs can be a perfect candidate solution for doing such proofs,³ i.e., Alice can generate a proof of knowledge of (m, σ) such that “ (m, σ) verifies with Bob’s public key and $P(m) = 1$ ”.

AN ATTEMPT OF SECURITY PROOF. Intuitively, the security of this protocol should follow easily from the proof of knowledge of the SNARK and the unforgeability of the signature scheme. However, somewhat surprisingly, the proof becomes quite subtle. Let us consider a cheating Alice that always outputs a proof for a statement in the language.⁴ If Alice is still cheating, then it must be that she is using a signature on a message that she did not query – in other words a forgery. Then one would like to reduce such a cheating Alice to a forger for the signature scheme. To do this, one would proceed as follows. For any Alice one defines a forger that, on input the verification key vk , generates the SNARK crs , gives (crs, vk) to Alice, and simulate’s Alice’s queries using its own signing oracle. When Alice comes with the cheating proof, the forger would need an extractor for Alice in order to obtain the forgery from her. However, even if we see Alice as a SNARK prover with auxiliary input vk , Alice does not quite fit the proof of knowledge definition in which adversaries have no oracles. To handle similar cases, one typically shows that for every, interactive, Alice there is a non-interactive algorithm \mathcal{B} that runs Alice simulating her oracles (i.e., \mathcal{B} samples the signing key) and returns the same output. The good news is that for such \mathcal{B} one can claim the existence of an extractor \mathcal{E}_B as it fits the proof of knowledge definition. The issue is though that \mathcal{E}_B expects

³ Further motivation can be to keep the privacy of m by relying on zero-knowledge SNARKs.

⁴ The other case of statements not in the language can be easily reduced to the soundness of the SNARK.

the same input of \mathcal{B} , which includes the secret signing key. This means that our candidate forger mentioned above (which does not have the secret key) cannot run $\mathcal{E}_{\mathcal{B}}$.

APPLICATIONS THAT NEED EXTRACTION WITH ORACLES. Besides the above example, this issue can show up essentially in every application of SNARKs in which adversaries have access to oracles with a secret state, and one needs to run an extractor during an experiment (e.g., a reduction) where the secret state of the oracle is not available. For instance, we recognize this issue while trying to formally prove the security of a “folklore” construction of homomorphic signatures based on SNARKs and digital signatures that is mentioned in several papers (e.g., [BF11, GW13, CF13, GVW15]). The same issue appears in a generic construction of SNARKs on authenticated data in [BBFR15] (also informally discussed in [BCCT12]), where the security proof uses the existence of an extractor for the oracle-aided prover, but without giving particular justification. A similar issue also appears in the construction of succinct functional signatures of [BGI14]. To be precise, in [BGI14] the authors provide a (valid) proof but under a stronger definition of SNARKs in which the adversarial prover and the extractor are independent PPT machines without *common* auxiliary input: a notion for which we are not aware of standard model constructions. In contrast, if one attempts to prove the succinct functional signatures of [BGI14] using the standard definition of SNARKs, one incurs the same issues illustrated above, i.e., the proof would not go through.

In this work we address this problem by providing both negative and positive results to the feasibility of extraction in the presence of oracles. On one hand, our negative results provide an explanation of why the above proofs do not go through so easily. On the other hand, our positive results eventually provide some guidelines to formally state and prove the security of the cryptographic constructions mentioned above (albeit with various restrictions).

1.2 An Overview of Our Results

Defining SNARKs in the presence of oracles. As a first step, we formalize the definition of non-black-box extraction in the presence of oracles by proposing a notion of *SNARKs in the presence of oracles* (O-SNARKs, for short). In a nutshell, an O-SNARK is like a SNARK except that adaptive proof of knowledge must hold with respect to adversaries that have access to an oracle \mathcal{O} sampled from some oracle family \mathbb{O} .⁵ Slightly more in detail, we require that for any adversary $\mathcal{A}^{\mathcal{O}}$ with access to \mathcal{O} there is an extractor $\mathcal{E}_{\mathcal{A}}$ such that, whenever $\mathcal{A}^{\mathcal{O}}$ outputs a valid proof, $\mathcal{E}_{\mathcal{A}}$ outputs a valid witness, by running on the same input of \mathcal{A} , plus the transcript of oracle queries-answers of \mathcal{A} .

Existence of O-SNARKs. Once having defined their notion, we study whether O-SNARKs exist and under what assumptions. Below we summarize our results.

O-SNARKS IN THE RANDOM ORACLE MODEL. As a first positive result, we show that the construction of Computationally Sounds (CS) proofs of Micali [Mic00] yields an O-SNARK for *every* oracle family, in the random oracle model. This result follows from the work of Valiant [Val08] which shows that Micali’s construction already allows for extraction. More precisely, using the power of the random oracle model, Valiant shows a *black-box* extractor. This powerful extractor can then be used to build an O-SNARK extractor that works for any oracle family.

INSECURITY OF O-SNARKS FOR EVERY ORACLE FAMILY, IN THE STANDARD MODEL. Although the above result gives a candidate O-SNARK, it only works in the random oracle model, and it is

⁵ The notion is parametrized by the family \mathbb{O} , i.e., we say Π is an O-SNARK for \mathbb{O} .

tailored to one construction [Mic00]. It is therefore interesting to understand whether extraction with oracles is feasible in the *standard model*. And it would also be interesting to see if this is possible based on the classical SNARK notion. Besides its theoretical interest, the latter question has also a practical motivation since there are several efficient SNARK constructions proposed in the last years that one might like to use in place of CS proofs. Our first result in this direction is that assuming existence of one way functions (OWFs) there do not exist O-SNARKs for NP with respect to *every* oracle family. More precisely, we show the following:

Theorem 1 (Informal). *Assume OWFs exist. Then for any polynomial $p(\cdot)$ there is an unforgeable signature scheme Σ_p such that any candidate O-SNARK, that is correct and succinct with proofs of length bounded by $p(\cdot)$, cannot satisfy adaptive proof of knowledge with respect to signing oracles corresponding to Σ_p .*

The above result shows the existence of an oracle family for which O-SNARKs do not exist. A basic intuition behind it is that oracles provide additional auxiliary input to adversaries and, as formerly shown in [BCPR14, BP15], this can create issues for extraction. In fact, to obtain our result we might also have designed an oracle that simply outputs a binary string following a distribution with respect to which extraction is impossible due to [BCPR14, BP15]. However, in this case the result should additionally assume the existence of indistinguishability (or differing-input) obfuscation. In contrast, our result shows that such impossibility holds by only assuming existence of OWFs, which is a much weaker assumption.

In addition to ruling out existence of O-SNARKs for general oracles, our theorem also rules out their existence for a more specific class of oracle families – *signing oracles* – that is motivated by the three applications mentioned earlier.⁶ Its main message is thus that one cannot assume existence of O-SNARKs that work with *any* signature scheme. This explains why the security proofs of the primitives considered earlier do not go through, if one wants to base it on an arbitrary signature scheme.

EXISTENCE OF O-SNARKS FOR SPECIFIC FAMILIES OF SIGNING ORACLES. We study ways to circumvent our impossibility result for signing oracles of Theorem 1. Indeed, the above result can be interpreted as saying that there exist (perhaps degenerate) signature schemes such that there are no O-SNARKs with respect to the corresponding signing oracle family. This is not ruling out that O-SNARKs may exist for specific signature schemes, or – even better – for specific classes of signature schemes. We provide the following results:

1. Hash-and-sign signatures, where the hash is a random oracle, yield “safe oracles”, i.e., oracles for which any SNARK is an O-SNARK for that oracle, in the ROM.
2. Turning to the standard model setting, we show that any classical SNARK is an O-SNARK for signing oracles if the message space of the signature scheme is properly bounded, and O-SNARK adversaries query “almost” the entire message space. This positive result is useful in applications that use SNARKs with signing oracles, under the condition that adversaries make signing queries on almost all messages.

NON-ADAPTIVE O-SNARKS. Finally, we consider a relaxed notion of O-SNARKs in which adversaries are required to declare in advance (i.e., before seeing the common reference string) all

⁶ We do believe that many more applications along the same line – proving knowledge of valid signatures – are conceivable. Two recent examples which considered our work in such a setting are [DLFKP16, NT16].

the oracle queries. For this weaker notion we show that, in the standard model, every SNARK (for arbitrary auxiliary inputs) is a non-adaptive O-SNARK.

Applications of O-SNARKs. A nice feature of the O-SNARK notion is that it lends itself to easy and intuitive security proofs in all those applications where one needs to execute extractors in interactive security games with oracles. We show that by replacing SNARKs with O-SNARKs (for appropriate oracle families) we can formally prove the security of the constructions of homomorphic signatures, succinct functional signatures and SNARKs on authenticated data that we mentioned in the previous section. By combining these O-SNARK-based constructions with our existence results mentioned earlier we eventually reach conclusions about the possible secure instantiations of these constructions. The first option is to instantiate them by using Micali’s CS proofs as an O-SNARK: this solution essentially yields secure instantiations in the random oracle model that work with a specific proof system [Mic00] (perhaps not the most efficient one in practice). The second option is to instantiate them using hash-and-sign signatures, apply our result on hash-and-sign signatures mentioned above, and then conjecture that replacing the random oracle with a suitable hash function preserves the overall security.⁷ Third, one can instantiate the constructions using a classical SNARK scheme Π and signature scheme Σ , and then conjecture that Π is also an O-SNARK with respect to the family of signing oracles corresponding to Σ . Compared to the first solution, the last two ones have the advantage that one could use some of the recently proposed efficient SNARKs (e.g., [PHGR13, BSCG⁺13]); on the other hand, these solutions have the drawback that security is based only on a heuristic argument. Finally, as a fourth option we provide security proofs of these primitives under a weak, non-adaptive, notion where adversaries declare all their queries in advance. Security in this weaker model can be proven assuming *non-adaptive O-SNARKs*, and thus classical SNARKs. The advantage of this fourth option is that one obtains a security proof for these instantiations based on clear – not newly crafted – assumptions, although under a much weaker security notion. Finally, worth noting is that we cannot apply the positive result on O-SNARK for signing oracles to the O-SNARK-based constructions of homomorphic signatures, functional signatures and SNARKs on authenticated data that we provide, and thus conclude their security under classical SNARKs. The inapplicability is due to the aforementioned restriction of our result, for which adversaries have to query almost the entire message space.⁸

Interpretation of our results. In line with recent work [BCPR14, BP15] on the feasibility of extraction in the presence of auxiliary input, our results indicate that additional care must be taken when considering extraction *in the presence of oracles*. While for auxiliary input impossibility of extraction is known under obfuscation-related assumptions, in the case of oracles we show that extraction becomes impossible even by only assuming one-way functions. Our counterexamples are of artificial nature and do not rule out the feasibility of extraction in the presence of “natural, benign” oracles. Nevertheless, our impossibility results provide formal evidence of why certain security proofs do not go through, and bring out important subtle aspects of security proofs. Given the importance of provable security and considered the increasing popularity of SNARKs in more practical scenarios, we believe these results give a message that is useful to protocol designers and of interest to the community at large.

⁷ The need of this final heuristic step is that hash-and-sign signatures use a random oracle in verification and in our applications the SNARK is used to prove knowledge of valid signatures, i.e., one would need a SNARK for $\text{NP}^{\mathcal{O}}$.

⁸ The exact reason is rather technical and requires to see the precise definitions and constructions of these primitives first. For the familiar reader, the intuition is that in these primitives/constructions an adversary that queries almost the entire message space of the underlying signature scheme becomes able to trivially break their security.

1.3 Organization

The paper is organized as follows. In Section 2 we recall notation and definitions used in the rest of our work. Section 3 introduces the notion of O-SNARKs, Section 4 includes positive and negative results about the existence of O-SNARKs, and in Section 5 we give three applications where our new notion turns out to be useful.

2 Preliminaries

Notation. We denote with $\lambda \in \mathbb{N}$ the security parameter. We say that a function $\epsilon(\lambda)$ is *negligible* if it vanishes faster than the inverse of any polynomial in λ . If not explicitly specified otherwise, negligible functions are negligible with respect to λ . If S is a set, $x \xleftarrow{\$} S$ denotes the process of selecting x uniformly at random in S . If \mathcal{A} is a probabilistic algorithm, $x \xleftarrow{\$} \mathcal{A}(\cdot)$ denotes the process of running \mathcal{A} on some appropriate input and assigning its output to x . For binary strings x and y , we denote by $x|y$ their concatenation and by x_i the i -th bit of x . For a positive integer n , we denote by $[n]$ the set $\{1, \dots, n\}$. For a random-access machine M we denote by $\#M(x, w)$ the number of execution steps needed by M to accept on input (x, w) .

The Universal Relation and NP Relations. We recall the notion of universal relation from [BG08], here adapted to the case of non-deterministic computations.

Definition 1. *The universal relation is the set $\mathcal{R}_{\mathcal{U}}$ of instance-witness pairs $(y, w) = ((M, x, t), w)$, where $|y|, |w| \leq t$ and M is a random-access machine such that $M(x, w)$ accepts after running at most t steps. The universal language $\mathcal{L}_{\mathcal{U}}$ is the language corresponding to $\mathcal{R}_{\mathcal{U}}$.*

For any constant $c \in \mathbb{N}$, \mathcal{R}_c denotes the subset of $\mathcal{R}_{\mathcal{U}}$ of pairs $(y, w) = ((M, x, t), w)$ such that $t \leq |x|^c$. \mathcal{R}_c is a “generalized” NP relation that is decidable in some fixed time polynomial in the size of the instance.

2.1 Succinct Non-Interactive Arguments

In this section we provide formal definitions for the notion of succinct non-interactive arguments of knowledge (SNARKs).

Definition 2 (SNARGs). *A succinct non-interactive argument (SNARG) for a relation $\mathcal{R} \subseteq \mathcal{R}_{\mathcal{U}}$ is a triple of algorithms $\Pi = (\text{Gen}, \text{Prove}, \text{Ver})$ working as follows*

$\text{Gen}(1^\lambda, T) \rightarrow \text{crs}$: *on input a security parameter $\lambda \in \mathbb{N}$ and a time bound $T \in \mathbb{N}$, the generation algorithm outputs a common reference string $\text{crs} = (\text{prs}, \text{vst})$ consisting of a public prover reference string prs and a verification state vst .*

$\text{Prove}(\text{prs}, y, w) \rightarrow \pi$: *given a prover reference string prs , an instance $y = (M, x, t)$ with $t \leq T$ and a witness w s.t. $(y, w) \in \mathcal{R}$, this algorithm produces a proof π .*

$\text{Ver}(\text{vst}, y, \pi) \rightarrow b$: *on input a verification state vst , an instance y , and a proof π , the verifier algorithm outputs $b = 0$ (reject) or $b = 1$ (accept).*

and satisfying completeness, succinctness, and (adaptive) soundness as described below:

- **Completeness.** For every time bound $T \in \mathbb{N}$, every valid $(y, w) \in \mathcal{R}$ with $y = (M, x, t)$ and $t \leq T$, there exists a negligible function negl such that

$$\Pr \left[\text{Ver}(\text{vst}, y, \pi) = 0 \mid \begin{array}{l} (\text{prs}, \text{vst}) \leftarrow \text{Gen}(1^\lambda, T) \\ \pi \leftarrow \text{Prove}(\text{prs}, y, w) \end{array} \right] \leq \text{negl}(\lambda)$$

- **Succinctness.** There exists a fixed polynomial $p(\cdot)$ independent of \mathcal{R} such that for every large enough security parameter $\lambda \in \mathbb{N}$, every time bound $T \in \mathbb{N}$, and every instance $y = (M, x, t)$ such that $t \leq T$, we have

- Gen runs in time $\begin{cases} p(\lambda + \log T) & \text{for a fully-succinct SNARG} \\ p(\lambda + T) & \text{for a pre-processing SNARG} \end{cases}$
- Prove runs in time $\begin{cases} p(\lambda + |M| + |x| + t + \log T) & \text{for fully-succinct SNARG} \\ p(\lambda + |M| + |x| + T) & \text{for pre-processing SNARG} \end{cases}$
- Ver runs in time $p(\lambda + |M| + |x| + \log T)$
- a honestly generated proof has size $|\pi| = p(\lambda + \log T)$.

- **Adaptive Soundness.** For every non-uniform \mathcal{A} of size $s(\lambda) = \text{poly}(\lambda)$ there is a negligible function $\epsilon(\lambda)$ such that for every time bound $T \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} \text{Ver}(\text{vst}, y, \pi) = 1 \\ \wedge y \notin \mathcal{L}_{\mathcal{R}} \end{array} \mid \begin{array}{l} (\text{prs}, \text{vst}) \leftarrow \text{Gen}(1^\lambda, T) \\ (y, \pi) \leftarrow \mathcal{A}(\text{prs}) \end{array} \right] \leq \epsilon(\lambda)$$

Furthermore, we say that Π has (s, ϵ) -adaptive soundness if the above condition holds for concrete values s and ϵ .

The notion of SNARG can be extended to be an argument of knowledge (a SNARK) by replacing soundness by an appropriate proof of knowledge property.

Definition 3 (SNARKs [BCC⁺14]). A succinct non-interactive argument of knowledge (SNARK) for a relation $\mathcal{R} \subseteq \mathcal{R}_{\mathcal{U}}$ is a triple of algorithms $\Pi = (\text{Gen}, \text{Prove}, \text{Ver})$ that constitutes a SNARG (as per Definition 2) except that soundness is replaced by the following property:

- **Adaptive Proof of Knowledge.** For every non-uniform prover \mathcal{A} of size $s(\lambda) = \text{poly}(\lambda)$ there exists a non-uniform extractor $\mathcal{E}_{\mathcal{A}}$ of size $t(\lambda) = \text{poly}(\lambda)$ and a negligible function $\epsilon(\lambda)$ such that for every auxiliary input $\text{aux} \in \{0, 1\}^{\text{poly}(\lambda)}$, and every time bound $T \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} \text{Ver}(\text{vst}, y, \pi) = 1 \\ \wedge \\ (y, w) \notin \mathcal{R} \end{array} \mid \begin{array}{l} (\text{prs}, \text{vst}) \leftarrow \text{Gen}(1^\lambda, T) \\ (y, \pi) \leftarrow \mathcal{A}(\text{prs}, \text{aux}) \\ w \leftarrow \mathcal{E}_{\mathcal{A}}(\text{prs}, \text{aux}) \end{array} \right] \leq \epsilon(\lambda)$$

Furthermore, we say that Π satisfies (s, t, ϵ) -adaptive proof of knowledge if the above condition holds for concrete values (s, t, ϵ) .

Remark 1 (Publicly verifiable vs. designated verifier). If security (adaptive PoK) holds against adversaries that have also access to the verification state vst (i.e., \mathcal{A} receives the whole crs) then the SNARK is called *publicly verifiable*, otherwise it is *designated verifier*. For simplicity, in the remainder of this work all definitions are given for the publicly verifiable setting; the corresponding designated-verifier variants are easily obtained by giving to the adversary only the prover state prs .

Remark 2 (About extraction and auxiliary input). First, we stress that in the PoK property the extractor $\mathcal{E}_{\mathcal{A}}$ takes exactly the same input of \mathcal{A} , including its random tape. Second, the PoK definition can also be relaxed to hold with respect to auxiliary inputs from specific distributions (instead of arbitrary ones). Namely, let \mathcal{Z} be a probabilistic algorithm (called the auxiliary input generator) that outputs a string aux , and let compactly denote this process as $aux \leftarrow \mathcal{Z}$. Then we say that adaptive proof of knowledge holds for \mathcal{Z} if the above definition holds for auxiliary inputs sampled according to $\mathcal{Z} - aux \leftarrow \mathcal{Z}$ – where \mathcal{Z} is also a non-uniform polynomial-size algorithm. More formally, we have the following definition.

Definition 4 (\mathcal{Z} -auxiliary input SNARKs). *Π is called a \mathcal{Z} -auxiliary input SNARK if Π is a SNARK as in Definition 3 except that adaptive proof of knowledge holds for auxiliary input $aux \leftarrow \mathcal{Z}$.*

For ease of exposition, in our proofs we compactly denote by $\text{AdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, \mathcal{Z})$ the adaptive proof of knowledge experiment executed with adversary \mathcal{A} , extractor $\mathcal{E}_{\mathcal{A}}$ and auxiliary input generator \mathcal{Z} , and that outputs 1 if $\text{Ver}(\text{vst}, y, \pi) = 1$ and $(y, w) \notin \mathcal{R}$. See below its description:

```

AdPoK( $\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, \mathcal{Z}$ )
   $aux \leftarrow \mathcal{Z}(1^\lambda); \text{crs} \leftarrow \text{Gen}(1^\lambda, T)$ 
   $(y, \pi) \leftarrow \mathcal{A}(\text{crs}, aux)$ 
   $w \leftarrow \mathcal{E}_{\mathcal{A}}(\text{crs}, aux)$ 
  if  $\text{Ver}(\text{crs}, y, \pi) = 1 \wedge (y, w) \notin \mathcal{R}$  return 1
  else return 0

```

We say that Π satisfies adaptive proof of knowledge for \mathcal{Z} -auxiliary input if for every non-uniform \mathcal{A} of size $s(\lambda) = \text{poly}(\lambda)$ there is a non-uniform extractor of size $t(\lambda) = \text{poly}(\lambda)$ and a negligible function $\epsilon(\lambda)$ such that for every time bound T we have

$$\Pr[\text{AdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, \mathcal{Z}) \Rightarrow 1] \leq \epsilon.$$

Furthermore, Π has (s, t, ϵ) -adaptive proof of knowledge for \mathcal{Z} -auxiliary input if the above condition holds for concrete (s, t, ϵ) .

SNARKs FOR NP. A SNARK for the universal relation $\mathcal{R}_{\mathcal{U}}$ is called a universal SNARK. SNARKs for NP are instead SNARKs in which the verification algorithm Ver takes as additional input a constant $c > 0$, and adaptive proof of knowledge is restricted to hold only for relations $\mathcal{R}_c \subset \mathcal{R}_{\mathcal{U}}$. More formally,

Definition 5 (SNARKs for NP). *A SNARK for NP is a tuple of algorithms $\Pi = (\text{Gen}, \text{Prove}, \text{Ver})$ satisfying Definition 3 except that the adaptive proof of knowledge property is replaced by the following one:*

- **Adaptive Proof of Knowledge for NP.** *For every non-uniform polynomial-size prover \mathcal{A} there exists a non-uniform polynomial-size extractor $\mathcal{E}_{\mathcal{A}}$ such that for every large enough $\lambda \in \mathbb{N}$, every auxiliary input $aux \in \{0, 1\}^{\text{poly}(\lambda)}$, and every time bound $T \in \mathbb{N}$, and every constant $c > 0$,*

$$\Pr \left[\begin{array}{l} \text{Ver}_c(\text{vst}, y, \pi) = 1 \\ \wedge \\ (y, w) \notin \mathcal{R}_c \end{array} \middle| \begin{array}{l} \text{crs} \leftarrow \text{Gen}(1^\lambda, T) \\ (y, \pi) \leftarrow \mathcal{A}(\text{crs}, aux) \\ w \leftarrow \mathcal{E}_{\mathcal{A}}(\text{crs}, aux) \end{array} \right] \leq \text{negl}(\lambda)$$

In the case of fully-succinct SNARKs for NP, it is not necessary to provide a time bound as one can set $T = \lambda^{\log \lambda}$. In this case we can write $\text{Gen}(1^\lambda)$ as a shorthand for $\text{Gen}(1^\lambda, \lambda^{\log \lambda})$.

Zero-Knowledge SNARKs. Here we recall the zero-knowledge definition for SNARKs.

Definition 6 (zk-SNARKs). $\Pi = (\text{Gen}, \text{Prove}, \text{Ver})$ is a (statistical) zero-knowledge SNARK for a relation $\mathcal{R} \subseteq \mathcal{R}_{\mathcal{U}}$ if Π is a SNARK and, moreover, satisfies the following property:

ZERO-KNOWLEDGE. There exists a stateful interactive polynomial-size simulator $S = (S^{\text{crs}}, S^{\text{Prove}})$ such that for all stateful interactive distinguishers \mathcal{D} , for every large enough security parameter $\lambda \in \mathbb{N}$, every auxiliary input $\text{aux} \in \{0, 1\}^{\text{poly}(\lambda)}$, and every time bound $T \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} (y, w) \in \mathcal{R} \\ \wedge \\ \mathcal{D}(\pi) = 1 \end{array} \middle| \begin{array}{l} (\text{prs}, \text{vst}) \leftarrow \text{Gen}(1^\lambda, T) \\ (y, w) \leftarrow \mathcal{D}(\text{prs}, \text{vst}, \text{aux}) \\ \pi \leftarrow \text{Prove}(\text{prs}, y, w) \end{array} \right] - \Pr \left[\begin{array}{l} (y, w) \in \mathcal{R} \\ \wedge \\ \mathcal{D}(\pi) = 1 \end{array} \middle| \begin{array}{l} (\text{prs}, \text{vst}, \text{tr}) \leftarrow S^{\text{crs}}(1^\lambda, T) \\ (y, w) \leftarrow \mathcal{D}(\text{prs}, \text{vst}, \text{aux}) \\ \pi \leftarrow S^{\text{Prove}}(\text{prs}, \text{tr}, y, \text{aux}) \end{array} \right] \leq \text{negl}(\lambda)$$

3 SNARKs in the presence of oracles

In this section we formalize the notion of extraction in the presence of oracles for SNARKs. We do this by proposing a suitable adaptive proof of knowledge definition, and we call a SNARK satisfying this definition a *SNARK in the presence of oracles* (O-SNARK, for short). As we shall see, the advantage of O-SNARKs is that this notion lends itself to easy and intuitive security proofs in all those applications where one needs to execute extractors in interactive security games with oracles (with a secret state). Below we provide the definition while the existence of O-SNARKs is discussed in Section 4.

3.1 O-SNARKs: SNARKs in the presence of oracles

Let $\mathbb{O} = \{\mathcal{O}\}$ be a family of oracles. We denote by $\mathcal{O} \leftarrow \mathbb{O}$ the process of sampling an oracle \mathcal{O} from the family \mathbb{O} according to some (possibly probabilistic) process. For example, \mathbb{O} can be a random oracle family, i.e., $\mathbb{O} = \{\mathcal{O} : \{0, 1\}^\ell \rightarrow \{0, 1\}^L\}$ for all possible functions from ℓ -bits strings to L -bits strings, in which case $\mathcal{O} \leftarrow \mathbb{O}$ consists of choosing a function \mathcal{O} uniformly at random in \mathbb{O} . As another example, \mathbb{O} might be the signing oracle corresponding to a signature scheme, in which case the process $\mathcal{O} \leftarrow \mathbb{O}$ consists of sampling a secret key of the signature scheme according to the key generation algorithm (and possibly a random tape for signature generation in case the signing algorithm is randomized).

For any oracle family \mathbb{O} , we define an O-SNARK Π for \mathbb{O} as follows.

Definition 7 (Z-auxiliary input O-SNARKs for \mathbb{O}). We say that Π is a Z-auxiliary input O-SNARK for the oracle family \mathbb{O} , if Π satisfies the properties of completeness and succinctness as in Definition 3, and the following property of adaptive proof of knowledge for \mathbb{O} :

- **Adaptive Proof of Knowledge for \mathbb{O} .** Consider the following experiment for security parameter $\lambda \in \mathbb{N}$, time bound $T \in \mathbb{N}$, adversary \mathcal{A} , extractor $\mathcal{E}_{\mathcal{A}}$, auxiliary input generator \mathcal{Z} and oracle family \mathbb{O} :

```

O-AdPoK( $\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, \mathcal{Z}, \mathbb{O}$ )
   $\text{aux} \leftarrow \mathcal{Z}(1^\lambda)$ ;  $\mathcal{O} \leftarrow \mathbb{O}$ ;  $\text{crs} \leftarrow \text{Gen}(1^\lambda, T)$ 
   $(y, \pi) \leftarrow \mathcal{A}^{\mathcal{O}}(\text{crs}, \text{aux})$ 
   $w \leftarrow \mathcal{E}_{\mathcal{A}}(\text{crs}, \text{aux}, \text{qt})$ 
  if  $\text{Ver}(\text{crs}, y, \pi) = 1 \wedge (y, w) \notin \mathcal{R}$  return 1
  else return 0

```

where $\text{qt} = \{q_i, \mathcal{O}(q_i)\}$ is the transcript of all oracle queries and answers made and received by \mathcal{A} during its execution.

Π satisfies adaptive proof of knowledge with respect to oracle family \mathbb{O} and auxiliary input from \mathcal{Z} if for every non-uniform oracle prover $\mathcal{A}^{\mathcal{O}}$ of size $s(\lambda) = \text{poly}(\lambda)$ making at most $Q(\lambda) = \text{poly}(\lambda)$ queries there exists a non-uniform extractor $\mathcal{E}_{\mathcal{A}}$ of size $t(\lambda) = \text{poly}(\lambda)$ and a negligible function $\epsilon(\lambda)$ such that for every time bound T ,

$$\Pr[\text{O-AdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, \mathcal{Z}, \mathbb{O}) \Rightarrow 1] \leq \epsilon(\lambda)$$

Furthermore, we say that Π satisfies (s, t, Q, ϵ) -adaptive proof of knowledge with respect to oracle family \mathbb{O} and auxiliary input from \mathcal{Z} if the above condition holds for concrete values (s, t, Q, ϵ) .

3.2 Non-Adaptive O-SNARKs

In this section we define a relaxation of O-SNARKs in which the adversary is non-adaptive in making its queries to the oracle. Namely, we consider adversaries that first declare all their oracle queries q_1, \dots, q_Q and then run on input the common reference string as well as the queries' outputs $\mathcal{O}(q_1), \dots, \mathcal{O}(q_Q)$. More formally,

Definition 8 (\mathcal{Z} -auxiliary input non-adaptive O-SNARKs for \mathbb{O}). We say that Π is a \mathcal{Z} -auxiliary input non-adaptive O-SNARK for the oracle family \mathbb{O} , if Π satisfies the properties of completeness and succinctness as in Definition 3, and the following property of non-adaptive queries proof of knowledge for \mathbb{O} :

- **Non-Adaptive Proof of Knowledge for \mathbb{O} .** Consider the following experiment for security parameter $\lambda \in \mathbb{N}$, time bound $T \in \mathbb{N}$, adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, extractor $\mathcal{E}_{\mathcal{A}}$, auxiliary input generator \mathcal{Z} and oracle family \mathbb{O} :

```

O-NonAdPoK( $\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, \mathcal{Z}, \mathbb{O}$ )
  ( $q_1, \dots, q_Q, st$ )  $\leftarrow \mathcal{A}_1(1^\lambda)$ 
   $aux \leftarrow \mathcal{Z}(1^\lambda)$ ;  $\mathcal{O} \leftarrow \mathbb{O}$ ;  $crs \leftarrow \text{Gen}(1^\lambda, T)$ 
   $qt = (q_1, \mathcal{O}(q_1), \dots, q_Q, \mathcal{O}(q_Q))$ 
  ( $y, \pi$ )  $\leftarrow \mathcal{A}_2(st, crs, aux, qt)$ 
   $w \leftarrow \mathcal{E}_{\mathcal{A}}(crs, aux, qt)$ 
  if  $\text{Ver}(crs, y, \pi) = 1 \wedge (y, w) \notin \mathcal{R}$  return 1
  else return 0

```

where st is simply a state information shared between \mathcal{A}_1 and \mathcal{A}_2 .

Π satisfies non-adaptive proof of knowledge with respect to oracle family \mathbb{O} and auxiliary input from \mathcal{Z} if for every non-uniform prover $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ of size $s(\lambda) = \text{poly}(\lambda)$ making at most $Q(\lambda) = \text{poly}(\lambda)$ non-adaptive queries there exists a non-uniform extractor $\mathcal{E}_{\mathcal{A}}$ of size $t(\lambda) = \text{poly}(\lambda)$ and a negligible function $\epsilon(\lambda)$ such that for every time bound T ,

$$\Pr[\text{O-NonAdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, \mathcal{Z}, \mathbb{O}) \Rightarrow 1] \leq \epsilon(\lambda)$$

Furthermore, we say that Π satisfies (s, t, Q, ϵ) -non-adaptive proof of knowledge with respect to oracle family \mathbb{O} and auxiliary input from \mathcal{Z} if the above condition holds for concrete values (s, t, Q, ϵ) .

It is also possible to define a stronger variant of the above definition in which \mathcal{A}_1 is given (adaptive) oracle access to \mathcal{O} , whereas \mathcal{A}_2 has no access to \mathcal{O} , except for the query transcript obtained by \mathcal{A}_1 . It is not hard to see that the result given in the following paragraph works under this intermediate definition as well.

Existence of Non-Adaptive O-SNARKs from SNARKs. Below we prove a simple result showing that non-adaptive O-SNARKs follow directly from classical SNARKs for which the proof of knowledge property holds for arbitrary auxiliary input distributions.

The idea of the proof is that the second stage adversary \mathcal{A}_2 of non-adaptive O-SNARKs is very much like a classical SNARK adversary that makes no queries and receives a certain auxiliary input which contains the set of oracle queries chosen by \mathcal{A}_1 with corresponding answers. The fact that the auxiliary input includes the set of queries chosen by \mathcal{A}_1 , which is an arbitrary adversary, implies that the SNARK must support arbitrary, *not necessarily benign*, auxiliary inputs (i.e., it is not sufficient to fix an auxiliary input distribution that depends only on the oracle family \mathbb{O}).

Theorem 2. *Let \mathbb{O} be any oracle family. If Π is a SNARK satisfying (s, t, ϵ) -adaptive PoK (for arbitrary auxiliary input), then Π is a non-adaptive O-SNARK for \mathbb{O} satisfying (s, t, Q, ϵ) -non-adaptive PoK.*

Proof. Given the first stage adversary \mathcal{A}_1 , and the oracle family \mathbb{O} we define the following auxiliary input distribution:

```

 $\mathcal{Z}_{\mathcal{A}_1, \mathbb{O}}(1^\lambda)$ 
   $(\{q_1, \dots, q_Q\}, st) \leftarrow \mathcal{A}_1(1^\lambda)$ 
   $\mathcal{O} \stackrel{\$}{\leftarrow} \mathbb{O}$ 
  return  $\langle st, \{q_i, \mathcal{O}(q_i)\}_{i=1}^Q \rangle$ 

```

Then, for any $(\mathcal{A}_1, \mathcal{A}_2)$ we can build the following SNARK adversary \mathcal{B} taking $z \leftarrow \mathcal{Z}_{\mathcal{A}_1, \mathbb{O}}$:

```

 $\mathcal{B}(\text{crs}, z)$ 
  Parse  $z = \langle st, q_1, y_1, \dots, q_Q, y_Q \rangle$ 
  Run  $\mathcal{A}_2(st, \text{crs}, \text{qt} = (q_1, y_1, \dots, q_Q, y_Q)) \rightarrow (y, \pi)$ 
  Return the same  $(y, \pi)$  returned by  $\mathcal{A}_2$ .

```

Since Π is by assumption a SNARK, for \mathcal{B} there exists an extractor $\mathcal{E}_{\mathcal{B}}$ such that, for *any* auxiliary input (and in particular for auxiliary input from $\mathcal{Z}_{\mathcal{A}_1, \mathbb{O}}$), it holds

$$\Pr[\text{O-AdPoK}(\lambda, T, \mathcal{B}, \mathcal{E}_{\mathcal{B}}, \mathcal{Z}_{\mathcal{A}_1, \mathbb{O}}) \Rightarrow 1] \leq \epsilon(\lambda)$$

Finally, we simply define $\mathcal{E}_{\mathcal{A}} = \mathcal{E}_{\mathcal{B}}$. Since \mathcal{B} 's simulation of \mathcal{A}_2 is perfect, it is easy to see that for any $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ this extractor $\mathcal{E}_{\mathcal{A}}$ is such that

$$\Pr[\text{O-NonAdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, \mathbb{O}) \Rightarrow 1] \leq \epsilon(\lambda)$$

□

4 On the existence of O-SNARKs

In this section we study whether O-SNARKs exist and under what assumptions. In the following sections we give both positive and negative answers to this question.

4.1 O-SNARKs in the random oracle model from Micali’s CS proofs

In this section we briefly discuss how the construction of CS proofs of Micali [Mic00] can be seen as an O-SNARK for *any* oracle family, albeit in the random oracle model. To see this, we rely on the result of Valiant [Val08] who shows that Micali’s construction is a “CS proof of knowledge” in the random oracle model. The main observation is in fact that Valiant’s proof works by showing a *black-box* extractor working for any prover.

Proposition 1. *Let \mathbb{O} be any oracle family and RO be a family of random oracles. Let Π_{Mic} be the CS proof construction from [Mic00]. Then Π_{Mic} is an O-SNARK for (RO, \mathbb{O}) , in the random oracle model.*

Proof (Sketch). Let \mathcal{E}^{RO} be Valiant’s black-box extractor⁹ which takes as input the code of the prover and outputs a witness w . For any adversary $\mathcal{A}^{\text{RO}, \mathbb{O}}$ we can define its extractor $\mathcal{E}_{\mathcal{A}}$ as the one that, on input the query transcript qt of \mathcal{A} , executes $w \leftarrow \mathcal{E}^{\text{RO}}(\mathcal{A})$ by simulating all the random oracle queries of \mathcal{E}^{RO} using qt , and finally outputs the same w . The reason why qt suffices to $\mathcal{E}_{\mathcal{A}}$ for simulating random oracle queries to \mathcal{E}^{RO} is that Valiant’s extractor \mathcal{E}^{RO} makes exactly the same queries of the prover.

4.2 Impossibility of O-SNARKs for every family of oracles, in the standard model

In this section we show that, *in the standard model*, there do not exist O-SNARKs with respect to *every* family of oracles. We show this under the assumption that universal one-way hash functions (and thus one-way functions [Rom90]) exist. To show the impossibility, we describe an oracle family in the presence of which any candidate O-SNARK that is correct and succinct cannot satisfy adaptive proof of knowledge with respect to that oracle family. Our impossibility result is shown for designated-verifier O-SNARKs, and thus implies impossibility for publicly verifiable ones as well (since every publicly verifiable O-SNARK is also designated-verifier secure). More specifically, we show the impossibility by means of a signing oracle family. Namely, we show a secure signature scheme Σ_p such that every correct and succinct O-SNARK Π cannot satisfy adaptive proof of knowledge in the presence of the signing oracle corresponding to Σ_p . Interestingly, such a result not only shows that extraction cannot work for general families of oracles, but also for families of signing oracles, a class which is relevant to several applications.

For every signature scheme $\Sigma = (\text{kg}, \text{sign}, \text{vfy})$ we let \mathbb{O}_{Σ} be the family of oracles $\mathcal{O}(m) = \text{sign}(\text{sk}, m)$, where every family member \mathcal{O} is described by a secret key sk of the signature scheme, i.e., the process $\mathcal{O} \leftarrow \mathbb{O}_{\Sigma}$ corresponds to obtaining sk through a run of $(\text{sk}, \text{vk}) \xleftarrow{\$} \text{kg}(1^\lambda)$. For the sake of simplicity, we also assume that the oracle allows for a special query, say $\mathcal{O}('vk')$,¹⁰ whose answer is the verification key vk .

Theorem 3. *Assume that one-way functions exist. Then for every polynomial $p(\cdot)$ there exists a UF-CMA-secure signature scheme Σ_p such that every candidate designated-verifier O-SNARK Π for NP, that is correct and succinct with proofs of length bounded by $p(\cdot)$, does not satisfy adaptive proof of knowledge with respect to \mathbb{O}_{Σ_p} .*

⁹ The CS proofs of knowledge definition used by Valiant considers adversaries that are non-adaptive in choosing the statement. However it is easy to see that the construction and the proof work also for the adaptive case.

¹⁰ Here vk is an arbitrary choice; any symbol not in \mathcal{M} would do so. Introducing the extra query simplifies the presentation, otherwise vk should be treated as an auxiliary input from a distribution generated together with the oracle sampling.

AN INTUITION OF THE RESULT. Before delving into the details of the proof, we provide the main intuition of this result. This intuition does not use signature schemes but includes the main ideas that will be used in the signature counterexample. Given a UOWHF function family \mathcal{H} , consider the NP binary relation $\bar{R}_{\mathcal{H}} = \{((h, x), w) : h \in \mathcal{H}, h(w) = x\}$, let Π be a SNARK for NP and consider $p(\cdot)$ the polynomial for which Π is succinct. The idea is to show an oracle family $\tilde{\mathcal{O}}$ and an adversary $\bar{\mathcal{A}}$ for which there is no extractor unless \mathcal{H} is not a universal one-way family. For every polynomial $p(\cdot)$, the oracle family contains oracles \mathcal{O}_p that given a query q , interpret q as the description of a program $\mathcal{P}(\cdot, \cdot)$, samples a random member of the hash family $h \xleftarrow{\$} \mathcal{H}$, a random w , computes $x = h(w)$, and outputs (h, x) along with $\pi \leftarrow \mathcal{P}((h, x), w)$. If $\mathcal{P}(\cdot, \cdot) = \text{Prove}(\text{prs}, \cdot, \cdot)$, then the oracle is simply returning an hash image with a proof of knowledge of its (random) preimage. The adversary $\bar{\mathcal{A}}^{\mathcal{O}_p}$ is the one that on input prs , simply asks one query $q = \mathcal{P}(\cdot, \cdot) = \text{Prove}(\text{prs}, \cdot, \cdot)$, gets $((h, x), \pi) \leftarrow \mathcal{O}_p(q)$ and outputs $((h, x), \pi)$. Now, the crucial point that entails the non-existence of an extractor is that, provided that the input w is sufficiently longer than π , every valid extractor for such $\bar{\mathcal{A}}$ that outputs a valid w' immediately implies a collision (w, w') for h .¹¹ Finally, to prevent adversarially chosen \mathcal{P} from revealing too much information, we require the oracle to check the length of π , and the latter is returned only if $|\pi| \leq p(\lambda)$.

Proof (Proof of Theorem 3). The proof consists of two main steps. First, we describe the construction of the signature scheme Σ_p based on any other UF-CMA-secure signature scheme $\hat{\Sigma}$ with message space $\mathcal{M} = \{0, 1\}^*$ (that exists assuming OWFs [Lam79, Rom90]), and show that Σ_p is UF-CMA-secure. Σ_p uses also an UOWHF family \mathcal{H} . Second, we show that, when considering the oracle family \mathbb{O}_{Σ_p} corresponding to the signature scheme Σ_p , a correct Π with succinctness $p(\cdot)$ cannot be an O-SNARK for \mathbb{O}_{Σ_p} , i.e., we show an efficient O-SNARK adversary $\mathcal{A}_p^{\mathcal{O}}$ (with access to a Σ_p signing oracle $\mathcal{O}(\cdot) = \text{sign}(\text{sk}, \cdot)$), for which there is no extractor unless \mathcal{H} is not one-way.

The counterexample signature scheme Σ_p . Let $\hat{\Sigma}$ be any UF-CMA-secure scheme with message space $\mathcal{M} = \{0, 1\}^*$. Let $\mathbb{H} = \{\mathcal{H}\}_{\lambda}$ be a collection of function families $\mathcal{H} = \{h : \{0, 1\}^{L(\lambda)} \rightarrow \{0, 1\}^{\ell(\lambda)}\}$ where each \mathcal{H} is an universal one-way hash family with $L(\lambda) \geq p(\lambda) + \ell(\lambda) + \lambda$. Let $M_{\mathcal{H}}((h, x), w)$ be the machine that on input $((h, x), w)$ accepts iff $h(w) = x$, and $\mathcal{R}_{\mathcal{H}}$ be the NP relation consisting of all pairs (y, w) such that, for $y = (M_{\mathcal{H}}, (h, x), t)$, $M_{\mathcal{H}}((h, x), w)$ accepts in at most t steps.

The scheme Σ_p has message space $\mathcal{M} = \{0, 1\}^*$; its algorithms work as follows:

$\text{kg}(1^\lambda)$: Run $(\hat{\text{vk}}, \hat{\text{sk}}) \leftarrow \hat{\Sigma}.\text{kg}(1^\lambda)$, set $\text{vk} = \hat{\text{vk}}$, $\text{sk} = \hat{\text{sk}}$.

$\text{sign}(\text{sk}, m)$: Signing works as follows

- generate $\hat{\sigma} \leftarrow \hat{\Sigma}.\text{sign}(\hat{\text{sk}}, m)$;
- sample $h \xleftarrow{\$} \mathcal{H}$ and $w \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$;
- compute $x = h(w)$, $t = \#M_{\mathcal{H}}((h, x), w)$, and set $y = (M_{\mathcal{H}}, (h, x), t)$;
- interpret m as the description of program $\mathcal{P}(\cdot, \cdot)$ and thus run $\pi \leftarrow \mathcal{P}(y, w)$;
- if $|\pi| \leq p(\lambda)$, set $\pi' = \pi$, else set $\pi' = 0$;
- output $\sigma = (\hat{\sigma}, h, x, \pi')$.

$\text{vfy}(\text{vk}, m, \sigma)$: Parse $\sigma = (\hat{\sigma}, h, x, \pi')$ and return the output of $\hat{\Sigma}.\text{vfy}(\hat{\text{vk}}, m, \hat{\sigma})$.

It is trivial to check that, as long as $\hat{\Sigma}$ is a UF-CMA-secure scheme, Σ_p is also UF-CMA-secure. Moreover, remark that the scheme Σ_p does not depend on the specific O-SNARK construction Π but only on the universal polynomial $p(\cdot)$ bounding its succinctness.

¹¹ This relies on the fact that sufficiently many bits of w remain unpredictable, even given π .

Impossibility of O-SNARKs for \mathbb{O}_{Σ_p} . To show that Π is not an O-SNARK for \mathbb{O}_{Σ_p} (under the assumption that \mathcal{H} is universally one-way), we prove that there is an adversary $\mathcal{A}_p^{\mathcal{O}}$ such that every candidate extractor \mathcal{E} fails in the adaptive proof of knowledge game.

Lemma 1. *If \mathcal{H} is universally one way then every Π for NP that is correct and succinct with proofs of length $p(\cdot)$ is not a designated-verifier O-SNARK for \mathbb{O}_{Σ_p} .*

Proof. Let $\mathcal{A}_p^{\mathcal{O}}$ be the following adversary: on input prs , encode the Prove algorithm of Π with hardcoded prs as a program $\mathcal{P}(\cdot, \cdot) := \text{Prove}(\text{prs}, \cdot, \cdot)$; let q be \mathcal{P} 's description, and make a single query $\sigma = (\hat{\sigma}, h, x, \pi') \leftarrow \mathcal{O}(q)$; return (y, π') where $y = (M_{\mathcal{H}}, (h, x), t)$ is appropriately reconstructed. We show that for every polynomial-size extractor \mathcal{E} it holds

$$\Pr[\text{O-AdPoK}(\lambda, \mathcal{A}_p, \mathcal{E}, \mathbb{O}_{\Sigma_p}) \Rightarrow 0] \leq \nu_{\mathcal{H}}(\lambda) + 2^{-\lambda}$$

where $\nu_{\mathcal{H}}(\lambda) = \mathbf{Adv}_{\mathcal{B}, \mathcal{H}}^{\text{UOWHF}}(\lambda)$ is the advantage of any adversary \mathcal{B} against \mathcal{H} 's universal one-wayness. This means that there is no extractor unless \mathcal{H} is not an universal one-way family.

We proceed by contradiction assuming the existence of a polynomial-size extractor \mathcal{E} such that the above probability is greater than some non-negligible ϵ . We show how to build an adversary \mathcal{B} that breaks universal one-wayness of \mathcal{H} with non-negligible probability.

\mathcal{B} first chooses an hash input $w \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$, and then receives an instance h of \mathcal{H} . Next, \mathcal{B} generates $(\text{prs}, \text{vst}) \leftarrow \text{Gen}(1^\lambda)$ and $(\widehat{\text{vk}}, \widehat{\text{sk}}) \leftarrow \widehat{\Sigma}.\text{kg}(1^\lambda)$, and runs $\mathcal{A}_p^{\mathcal{O}}(\text{prs})$ simulating the oracle \mathcal{O} on the single query $q := \mathcal{P}(\cdot, \cdot) = \text{Prove}(\text{crs}, \cdot, \cdot)$ asked by \mathcal{A}_p . In particular, to answer the query \mathcal{B} uses the secret key $\widehat{\text{sk}}$ to generate $\hat{\sigma}$, and computes $x = h(w)$ using the function h received from its challenger, and the input w chosen earlier. Notice that such a simulation can be done perfectly in a straightforward way, and that \mathcal{A}_p 's output is the pair (y, π) created by \mathcal{B} . Next, \mathcal{B} runs the extractor $w' \leftarrow \mathcal{E}(\text{prs}, \text{qt} = (\mathcal{P}(\cdot, \cdot), (\hat{\sigma}, h, x, \pi)))$, and outputs w' .

By correctness of Π it holds that the pair (y, π) returned by \mathcal{A}_p satisfies $\text{Ver}(\text{vst}, y, \pi) = 1$. Thus, by our contradiction assumption, with probability $\geq \epsilon(\lambda)$, \mathcal{E} outputs w' such that $(y, w') \in \mathcal{R}_{\mathcal{H}}$. Namely, $h(w') = x = h(w)$. To show that this is a collision, we argue that, information-theoretically, $w' \neq w$ with probability $\geq 1 - 1/2^\lambda$. This follows from the fact that w is randomly chosen of length $L(\lambda) \geq p(\lambda) + \ell(\lambda) + \lambda$ and the only information about w which is leaked to \mathcal{E} is through π and $x = h(w)$, an information of length at most $p(\lambda) + \ell(\lambda)$. Therefore there are at least λ bits of entropy in w , from which $\Pr[w' = w] \leq 2^{-\lambda}$ over the random choice of w . Hence, \mathcal{B} can break the universal one-wayness of \mathcal{H} with probability $\geq \epsilon(\lambda) - 2^{-\lambda}$. \square

4.3 O-SNARKs for signing oracles from SNARKs in the random oracle model

In this section we show that it is possible to “immunize” any signature scheme in such a way that any classical SNARK is also an O-SNARK for the signing oracle corresponding to the transformed scheme. The idea is very simple and consists into applying the hash-then-sign approach using a hash function that will be modeled as a random oracle. A limitation of this result is that, since the verification algorithm uses a random oracle, in all those applications where the SNARK is used to prove knowledge of valid signatures, one would need a SNARK for $\text{NP}^{\mathcal{O}}$. Hence, the best one can do is to conjecture that this still works when replacing the random oracle with a suitable hash function.

Let us now state formally our result. To this end, for any signature scheme Σ and polynomial $Q(\cdot)$ we define $\mathcal{Z}_{Q, \Sigma}$ as the distribution on tuples $\langle \text{vk}, m_1, \sigma_1, \dots, m_Q, \sigma_Q \rangle$ obtained by running the following probabilistic algorithm:

```

 $\mathcal{Z}_{Q,\Sigma}(1^\lambda)$ 
  let  $Q = Q(\lambda)$ 
   $(\text{sk}, \text{vk}) \leftarrow \text{kg}(1^\lambda)$ 
   $\tilde{\mathcal{M}} \xleftarrow{\$} \text{MsgSample}(\mathcal{M}, Q)$ 
  let  $\tilde{\mathcal{M}} = \{m_1, \dots, m_Q\}$ 
  for  $i = 1$  to  $Q$  do :
     $\sigma_i \leftarrow \text{sign}(\text{sk}, m_i)$ 
  return  $\langle \text{vk}, \{m_i, \sigma_i\}_{i=1}^Q \rangle$ 

```

where $\text{MsgSample}(\mathcal{M}, Q)$ is an algorithm that returns Q distinct messages, each randomly chosen from \mathcal{M} .

Theorem 4. *Let Σ be a UF-CMA-secure signature scheme, and \mathcal{H} be a family of hash functions modeled as a random oracle. Let \mathcal{U}_n be the uniform distribution over strings of length n , and $\mathcal{Z}_{Q,\Sigma}$ be the distribution defined above, where Q is any polynomial in the security parameter. Then there exists a signature scheme $\Sigma_{\mathcal{H}}$ such that every $(\mathcal{Z}, \mathcal{U}, \mathcal{Z}_{\Sigma, Q})$ -auxiliary input SNARK Π is a \mathcal{Z} -auxiliary input O-SNARK for $(\mathbb{O}_{\mathcal{H}}, \mathbb{O}_{\Sigma_{\mathcal{H}}})$ where $\mathbb{O}_{\mathcal{H}}$ is a random oracle.*

Proof. The proof consists of two steps. First we define the signature scheme $\Sigma_{\mathcal{H}}$ and prove its security from Σ . Next, we show that Π is an O-SNARK for $(\mathbb{O}_{\mathcal{H}}, \mathbb{O}_{\Sigma_{\mathcal{H}}})$.

THE SIGNATURE SCHEME $\Sigma_{\mathcal{H}}$. The signature scheme $\Sigma_{\mathcal{H}}$ is essentially an application of the hash-then-sign paradigm to scheme Σ . The only difference is that, instead of hashing only messages, we also hash a short random string of λ bits.

Let $\Sigma = (\Sigma.\text{kg}, \Sigma.\text{sign}, \Sigma.\text{vfy})$ be a signature scheme with message space $\mathcal{M} = \{0, 1\}^L$, and let $\mathcal{H}\{H : \{0, 1\}^* \rightarrow \{0, 1\}^L\}$ be a family of hash functions modeled as a random oracle. The signature scheme $\Sigma_{\mathcal{H}} = (\Sigma_{\mathcal{H}}.\text{kg}, \Sigma_{\mathcal{H}}.\text{sign}, \Sigma_{\mathcal{H}}.\text{vfy})$ works as follows:

$\Sigma_{\mathcal{H}}.\text{kg}(1^\lambda)$: Run $(\text{sk}, \text{vk}) \leftarrow \Sigma.\text{kg}(1^\lambda)$. Output (sk, vk) .

$\Sigma_{\mathcal{H}}.\text{sign}(\text{sk}, m)$: Sample $s \xleftarrow{\$} \{0, 1\}^\lambda$, compute $h \leftarrow H(s|m)$, $\hat{\sigma} \leftarrow \Sigma.\text{sign}(\text{sk}, h)$, and output $\sigma = (s, \hat{\sigma})$.

$\Sigma_{\mathcal{H}}.\text{vfy}(\text{vk}, m, \sigma)$: Parse $\sigma = (s, \hat{\sigma})$, compute $h \leftarrow H(s|m)$, return the same output of $\Sigma.\text{vfy}(\text{vk}, h, \hat{\sigma})$.

Lemma 2. *If Σ is an UF-CMA-secure signature scheme, so is $\Sigma_{\mathcal{H}}$ in the random oracle model.*

The security proof of the lemma is straightforward and is omitted.

Π IS AN O-SNARK FOR $(\mathbb{O}_{\mathcal{H}}, \mathbb{O}_{\Sigma_{\mathcal{H}}})$. We are going to prove¹² that for every non-uniform polynomial-size oracle adversary \mathcal{A} there exists a non-uniform polynomial-size extractor $\mathcal{E}_{\mathcal{A}}$ and a negligible function $\epsilon(\lambda)$ such that for every time bound T ,

$$\Pr[\text{O-AdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, (\mathbb{O}_{\mathcal{H}}, \mathbb{O}_{\Sigma_{\mathcal{H}}})) \Rightarrow 1] \leq \epsilon(\lambda)$$

As a first step, we show that for every non-uniform polynomial-size algorithm \mathcal{A} , making q queries to $\mathbb{O}_{\mathcal{H}}$ and Q queries to $\mathbb{O}_{\Sigma_{\mathcal{H}}}$ (where both $q, Q = \text{poly}(\lambda)$), there is a non-uniform, polynomial-size, non-interactive algorithm \mathcal{B} .

\mathcal{B} takes as input (crs, o, z) , where $o \leftarrow \mathcal{U}_{q \cdot L + Q \cdot \lambda}$, $z \leftarrow \mathcal{Z}_{Q,\Sigma}$, and these are parsed as follows: $o = \langle r'_1, \dots, r'_q, s_1, \dots, s_Q \rangle$ with $r'_j \in \{0, 1\}^L$ and $s_i \in \{0, 1\}^\lambda$, $z = \langle \text{vk}, r_1, \hat{\sigma}_1, \dots, r_Q, \hat{\sigma}_Q \rangle$.

¹² For simplicity we are ignoring the additional auxiliary input distribution \mathcal{Z} , as it is easy to see that it essentially “carries over”.

$\mathcal{B}(\text{crs}, o, z)$

Run $\mathcal{A}^{\mathcal{O}_{\mathcal{H}}, \mathcal{O}_{\Sigma_{\mathcal{H}}}}(\text{crs}) \rightarrow (y, \pi)$ and simulate queries to $\mathcal{O}_{\mathcal{H}}$ and $\mathcal{O}_{\Sigma_{\mathcal{H}}}$ as follows:

Query $\mathcal{O}_{\mathcal{H}}(x_j), j \in \{1, \dots, q\}$:

Answer with $\mathcal{O}_{\mathcal{H}}(x_j) = r'_j$ using r'_j from o

Query $\mathcal{O}_{\Sigma_{\mathcal{H}}}(m_i), i \in \{1, \dots, Q\}$:

If $\mathcal{O}_{\mathcal{H}}(s_i|m_i)$ has been already answered before, abort.

Else, set $\mathcal{O}_{\mathcal{H}}(s_i|m_i) = r_i$, and answer $\mathcal{O}_{\Sigma_{\mathcal{H}}}(m_i) = (s_i, \hat{\sigma}_i)$,
where $r_i, \hat{\sigma}_i$ are taken from z and s_i from o .

Return the same (y, π) returned by \mathcal{A} .

The simulation provided by \mathcal{B} to \mathcal{A} is perfect unless \mathcal{B} aborts during a signing query. This event – let us formally call it **Abort** – happens if there exist $i \in [Q]$ and $j \in [q]$ such that $x_j = s_i|m_i$. However, over the random choice of $s_i \xleftarrow{\$} \{0, 1\}^\lambda$ in the i -th query (indeed s_i was never used before and is thus hidden to \mathcal{A}), we have that

$$\Pr[\text{Abort}] \leq \sum_{i=1}^Q \frac{q}{2^\lambda} \leq \frac{q \cdot Q}{2^\lambda}$$

Hence, if \mathcal{A} succeeds in producing an instance-proof pair (y, π) , so does \mathcal{B} with overwhelming probability.

Then, by the adaptive PoK property we have that for every such \mathcal{B} there is an extractor $\mathcal{E}_{\mathcal{B}}$ such that for every polynomial n, Q

$$\Pr[\text{O-AdPoK}(\lambda, T, \mathcal{B}, \mathcal{E}_{\mathcal{B}}, (\mathcal{U}_n, \mathcal{Z}_{Q, \Sigma})) \Rightarrow 1] \leq \epsilon(\lambda)$$

for a negligible ϵ .

So far, we have that for every \mathcal{A} there is an extractor $\mathcal{E}_{\mathcal{B}}$. In what follows, we use this $\mathcal{E}_{\mathcal{B}}$ to define the extractor $\mathcal{E}_{\mathcal{A}}$. The extractor $\mathcal{E}_{\mathcal{A}}$ takes as input crs and the transcript

$$\text{qt} = \langle x_1, r'_1, \dots, x_q, r'_q, r_1, \dots, r_Q, m_1, s_1, \hat{\sigma}_1, \dots, m_Q, s_Q, \hat{\sigma}_Q \rangle$$

of queries made by \mathcal{A} , and proceeds as follows:

$\mathcal{E}_{\mathcal{A}}(\text{crs}, \text{qt})$

Run $w \leftarrow \mathcal{E}_{\mathcal{B}}(\text{crs}, o, z)$

Return the same w returned by $\mathcal{E}_{\mathcal{B}}$

Basically, it rearranges the data in qt to fulfill the format of (o, z) from distributions $\mathcal{U}_{qL+Q\lambda}, \mathcal{Z}_{Q, \Sigma}$. It is not hard to see that from the above construction we have

$$\begin{aligned} \Pr[\text{O-AdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_{\mathcal{A}}, (\mathcal{O}_{\mathcal{H}}, \mathcal{O}_{\Sigma_{\mathcal{H}}})) \Rightarrow 1] &\leq \Pr[\text{O-AdPoK}(\lambda, T, \mathcal{B}, \mathcal{E}_{\mathcal{B}}, (\mathcal{U}_n, \mathcal{Z}_{Q, \Sigma})) \Rightarrow 1] + \frac{qQ}{2^\lambda} \\ &\leq \epsilon(\lambda) + \frac{qQ}{2^\lambda} \end{aligned}$$

□

4.4 O-SNARKs for signing oracles from SNARKs

In this section we give a positive result showing that any SNARK Π is an O-SNARK for the signing oracle of signature scheme Σ if: (i) the message space of Σ is appropriately bounded (to be polynomially or at most superpolynomially large); (ii) Π tolerates auxiliary input consisting of the public key of Σ plus a collection of signatures on randomly chosen messages; (iii) one considers O-SNARK adversaries that query the signing oracle on almost the entire message space. Furthermore, in case of superpolynomially large message spaces, one needs to assume sub-exponential hardness for Π .

The intuition behind this result is to simulate the O-SNARK adversary by using a (non-interactive) SNARK adversary that receives the public key and a set of signatures on (suitably chosen) messages as its auxiliary input. If these messages exactly match¹³ those queried by the O-SNARK adversary, the simulation is perfect. However, since the probability of matching exactly all the $Q = \text{poly}(\lambda)$ queries may decrease exponentially in Q (making the simulation meaningless), we show how to put proper bounds so that the simulation can succeed with probability depending only on the message space size.

More formally, our result is stated as follows. Let Σ be a signature scheme with message space \mathcal{M} , and let $Q := Q(\cdot)$ be a function of the security parameter. Let $\mathcal{Z}_{Q,\Sigma}$ be the following auxiliary input distribution

```

 $\mathcal{Z}_{Q,\Sigma}(1^\lambda)$ 
  let  $Q = Q(\lambda)$ 
   $(\text{sk}, \text{vk}) \leftarrow \text{kg}(1^\lambda)$ 
   $\tilde{\mathcal{M}} \xleftarrow{\$} \text{MsgSample}(\mathcal{M}, Q)$ 
  let  $\tilde{\mathcal{M}} = \{m_1, \dots, m_Q\}$ 
  for  $i = 1$  to  $Q$  do :
     $\sigma_i \leftarrow \text{sign}(\text{sk}, m_i)$ 
  return  $\langle \text{vk}, \{m_i, \sigma_i\}_{i=1}^Q \rangle$ 

```

where $\text{MsgSample}(\mathcal{M}, Q)$ is a probabilistic algorithm that returns a subset $\tilde{\mathcal{M}} \subseteq \mathcal{M}$ of cardinality Q chosen according to some strategy that we discuss later. At this point we only assume a generic strategy such that $\delta(|\mathcal{M}|, Q) = \Pr[\text{MsgSample}(\mathcal{M}, Q) = \mathcal{M}^*]$ for any $\mathcal{M}^* \subseteq \mathcal{M}$ of cardinality Q .

Theorem 5. *Let Σ be a signature scheme with message space \mathcal{M} , let \mathbb{O}_Σ be the associated family of signing oracles, and let $\mathcal{Z}_{Q,\Sigma}$ be as defined above. If Π is a $\mathcal{Z}_{Q,\Sigma}$ -auxiliary input SNARK satisfying (s, t, ϵ) -adaptive PoK, then Π is an O-SNARK for \mathbb{O}_Σ satisfying (s', t', Q, ϵ') -adaptive PoK, where $\epsilon' = \epsilon/\delta(|\mathcal{M}|, Q)$, $s' = s - O(Q \cdot \log |\mathcal{M}|)$, and $t' = t$.*

Proof. The basic idea for the proof is to show that for any O-SNARK adversary \mathcal{A} against Π there is a non-interactive SNARK adversary $\mathcal{B}_\mathcal{A}$ against Π that can perfectly simulate \mathcal{A} , provided that the auxiliary input from $\mathcal{Z}_{Q,\Sigma}$ “matches” the queries made by \mathcal{A} . More precisely, for any adversary $\mathcal{A}^\mathcal{O}$ making Q oracle queries we define the following adversary $\mathcal{B}_\mathcal{A}$:

```

 $\mathcal{B}_\mathcal{A}(\text{crs}, \text{aux} = \langle \text{vk}, \{m_i, \sigma_i\}_{i=1}^Q \rangle)$ 
  Run  $(y, \pi) \leftarrow \mathcal{A}^\mathcal{O}(\text{crs})$  and simulate every query  $\mathcal{O}(m)$  as follows:
  let  $\tilde{\mathcal{M}} = \{m_1, \dots, m_Q\}$ 

```

¹³ We note that the proof requires an exact match and it is not sufficient that the O-SNARK adversary’s queries are a subset of the sampled messages. A more precise explanation of this fact is given at the end of the proof.

```

    if  $m = 0$  return  $vk$ 
    else if  $m = m_i \in \tilde{\mathcal{M}}$  return  $\sigma_i$ 
    else return  $\perp$ 
return  $(y, \pi)$ 

```

As one can see, the adversary \mathcal{B}_A fits the syntax of a SNARK adversary in the adaptive proof of knowledge property, by which (if Π is a SNARK) there exists an extractor $\mathcal{E}_{\mathcal{B}_A}$ such that for every time bound T and every Q

$$\Pr[\text{AdPoK}(\lambda, T, \mathcal{B}_A, \mathcal{E}_{\mathcal{B}_A}, \mathcal{Z}_{Q, \Sigma}) \Rightarrow 1] \leq \epsilon$$

For every \mathcal{A} we define the extractor $\mathcal{E}_A = \mathcal{E}_{\mathcal{B}_A}$, and our goal is to bound

$$\Pr[\text{O-AdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_A, \mathbb{O}_{\text{sign}}) \Rightarrow 1] \leq \epsilon'$$

To prove this we proceed by contradiction. Namely, let us assume that there is \mathcal{A} (and thus corresponding $\mathcal{E}_A = \mathcal{E}_{\mathcal{B}_A}$) such that:

$$\Pr[\text{O-AdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_A, \mathbb{O}_{\text{sign}}) \Rightarrow 1] \geq \epsilon'$$

We define G_1 to be the following variation of experiment $\text{O-AdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_A, \mathbb{O}_{\text{sign}})$: it generates $aux' \leftarrow \mathcal{Z}_{Q, \Sigma}$ (using the same keypair generated when sampling oracle $\mathcal{O} \leftarrow \mathbb{O}_{\text{sign}}$), and at the end G_1 outputs 1 if $qt = aux'$ holds in addition to $\text{Ver}(vst, y, \pi) = 1 \wedge (y, w) \notin \mathcal{R}$. Let QueryMatch be the event that in G_1 it occurs $qt = aux'$. Clearly, $\Pr[\text{QueryMatch}] = \delta(Q, |\mathcal{M}|)$, and thus

$$\begin{aligned} \Pr[G_1 \Rightarrow 1] &= \Pr[\text{QueryMatch}] \cdot \Pr[\text{O-AdPoK}(\lambda, T, \mathcal{A}, \mathcal{E}_A, \mathbb{O}_{\text{sign}}) \Rightarrow 1] \\ &\geq \delta(Q, |\mathcal{M}|) \cdot \epsilon' = \epsilon \end{aligned}$$

To conclude the proof, and reach the contradiction, we show below that

$$\Pr[\text{AdPoK}(\lambda, T, \mathcal{B}_A, \mathcal{E}_{\mathcal{B}_A}, \mathcal{Z}_{Q, \Sigma}) \Rightarrow 1] \geq \Pr[G_1 \Rightarrow 1]$$

To see this, consider experiment $\text{AdPoK}(\lambda, T, \mathcal{B}_A, \mathcal{E}_{\mathcal{B}_A}, \mathcal{Z}_{Q, \Sigma})$ and let $\text{QueryMatch}'$ be the event that, inside the execution of \mathcal{B}_A , all the messages queried by \mathcal{A} coincide with those in aux . Since $\mathcal{E}_A = \mathcal{E}_{\mathcal{B}_A}$ one can see that experiment G_1 is basically identical to $\text{AdPoK}(\lambda, T, \mathcal{B}_A, \mathcal{E}_{\mathcal{B}_A}, \mathcal{Z}_{Q, \Sigma}) \Rightarrow 1 \wedge \text{QueryMatch}'$. Namely,

$$\begin{aligned} \Pr[\text{AdPoK}(\lambda, T, \mathcal{B}_A, \mathcal{E}_{\mathcal{B}_A}, \mathcal{Z}_{Q, \Sigma}) \Rightarrow 1] &\geq \Pr[\text{AdPoK}(\lambda, T, \mathcal{B}_A, \mathcal{E}_{\mathcal{B}_A}, \mathcal{Z}_{Q, \Sigma}) \Rightarrow 1 \wedge \text{QueryMatch}'] \\ &= \Pr[G_1 \Rightarrow 1] \geq \epsilon \end{aligned}$$

which concludes the proof. \square

As a final remark, the reason why we need messages in $\tilde{\mathcal{M}}$ to exactly match the queries of \mathcal{A} is that we construct \mathcal{A} 's extractor from \mathcal{B}_A 's extractor – $\mathcal{E}_A = \mathcal{E}_{\mathcal{B}_A}$. Hence, in order for \mathcal{E}_A to call $\mathcal{E}_{\mathcal{B}_A}$, it needs to know all its inputs, which is not the case if \mathcal{B} has a superset of the signed messages known to \mathcal{A} .

Implications of Theorem 5. The statement of Theorem 5 is parametrized by values $|\mathcal{M}|, Q$ and the function $\delta(|\mathcal{M}|, Q)$, which in turn depends on the query guessing strategy. As for the $\text{MsgSample}(\mathcal{M}, Q)$ algorithm, let us consider the one that *samples a random subset* $\tilde{\mathcal{M}} \subseteq \mathcal{M}$ of cardinality Q . For this algorithm we have $\delta(|\mathcal{M}|, Q) = \frac{1}{\binom{|\mathcal{M}|}{Q}}$.

Notice that $\delta(|\mathcal{M}|, Q)$ is governing the success probability of our reduction, and thus we would like this function not to become negligible. However, since $Q = \text{poly}(\lambda)$ is a parameter under the choice of the adversary, it might indeed be the case that $\delta(|\mathcal{M}|, Q) \approx 2^{-Q} \approx 2^{-\lambda}$, which would make our reduction meaningless. To avoid this bad case, we restrict our attention to adversaries for which $Q = |\mathcal{M}| - c$ for some constant $c \geq 1$, i.e., adversaries that ask for signatures on the entire message but a constant number of messages. For this choice of Q we indeed have that $\delta(|\mathcal{M}|, Q) = \frac{1}{|\mathcal{M}|^c}$ depends only on the cardinality of $|\mathcal{M}|$. This gives us the following corollary:

Corollary 1. *Let Σ be a signature scheme with message space \mathcal{M} where $|\mathcal{M}| = \text{poly}(\lambda)$ (resp. $|\mathcal{M}| = \lambda^{\omega(1)}$), and let $Q = |\mathcal{M}| - c$ for constant $c \in \mathbb{N}$. If Π is a polynomially (resp. sub-exponentially) secure $\mathcal{Z}_{Q, \Sigma}$ -auxiliary input SNARK, then Π is an O-SNARK for \mathbb{O}_{Σ} (for adversaries making Q queries).*

4.5 O-SNARKs for (pseudo)random oracles from SNARKs

In this section we show a positive result on the existence of O-SNARKs for (pseudo)random oracles, based on classical SNARKs that are assumed to satisfy proof of knowledge with respect to randomly-distributed auxiliary input. While we do not have a direct application of this result, we think that it can be useful and of independent interest.

Let $\mathbb{O} = \{\mathcal{O} : \{0, 1\}^{\ell} \rightarrow \{0, 1\}^L\}$ be a random oracle family, i.e., a family where sampling a member $\mathcal{O} \leftarrow \mathbb{O}$ consists of sampling a suitably long random tape (of size $2^{\ell} \cdot L$). Let us stress that here, when we refer to a random oracle family, we do not necessarily consider the random oracle model. We simply consider the case of an interactive game in which \mathcal{A} has oracle access to a random function.

Theorem 6. *Let \mathbb{O} be a random oracle family. Let \mathcal{Z} be some distribution, \mathcal{U}_n be the uniform distribution over strings of length n , and denote by $(\mathcal{Z}, \mathcal{U})$ the distribution over pairs (aux, o) such that $aux \leftarrow \mathcal{Z}$ and $o \leftarrow \mathcal{U}$. If Π is a $(\mathcal{Z}_q, \mathcal{U})$ -auxiliary input SNARK for every $q = \text{poly}(\lambda)$, then Π is a \mathcal{Z} -auxiliary input O-SNARK for \mathbb{O} .*

Proof. Completeness and succinctness follow immediately. So, we are left to prove that (adaptive) proof of knowledge implies the corresponding property in the presence of an oracle $\mathcal{O} \leftarrow \mathbb{O}$.

Formally, we have to show that for any efficient oracle prover $\mathcal{A}^{\mathcal{O}}$ there exists an efficient extractor algorithm $\mathcal{E}_{\mathcal{A}}$ such that the joint probability that $\mathcal{A}^{\mathcal{O}}$ outputs a valid proof but $\mathcal{E}_{\mathcal{A}}$ returns a wrong witness is negligible. At a high level, we do this proof by showing that $\mathcal{E}_{\mathcal{A}}$'s existence is equivalent to the existence of a SNARK extractor $\bar{\mathcal{E}}$, which is assured under the adaptive proof of knowledge whenever $\mathcal{A}^{\mathcal{O}}$ exists.

First, for any $\mathcal{A}^{\mathcal{O}}$ we construct another adversary $\bar{\mathcal{A}}$. Precisely, let \mathcal{A} be a non-uniform algorithm that takes as input (crs, aux) and (without loss of generality) makes exactly Q queries to \mathcal{O} for some $Q = \text{poly}(\lambda)$. Then, we can define an adversary $\bar{\mathcal{A}}$ that on input $(\text{crs}, \overline{aux})$ – where $\overline{aux} = (aux, o) \leftarrow (\mathcal{Z}, \mathcal{U}_{Q \cdot L})$ – works as follows:

$\bar{\mathcal{A}}(\text{crs}, \overline{aux})$

Run $\mathcal{A}^{\mathcal{O}}(\text{crs}, aux) \rightarrow (y, \pi)$ simulating \mathcal{O} queries as follows:
 given the i -th query $\mathcal{O}(q_i)$, answer with the substring
 $\mathcal{O}(q_i) = a_i = o_{L \cdot (i-1) + 1} | \cdots | o_{L \cdot i}$
 Return the same (y, π) returned by \mathcal{A}

The simulation of \mathcal{O} provided by $\bar{\mathcal{A}}$ to \mathcal{A} is clearly perfect. Hence, if $\mathcal{A}^{\mathcal{O}}$ succeeds in producing an accepting instance-proof pair (y, π) , so does $\bar{\mathcal{A}}$. Then, by the adaptive PoK property we have that for every such $\bar{\mathcal{A}}$ there is an extractor $\bar{\mathcal{E}}$ that takes the same input of $\bar{\mathcal{A}}$ and outputs a value w such that the probability that $\bar{\mathcal{A}}$ is successful in outputting an accepting pair (y, π) and the value w returned by $\bar{\mathcal{E}}$ is a wrong witness (i.e., $(y, w) \notin \mathcal{R}$) is bounded by a negligible function ϵ .

Basically, above we showed that for every $\mathcal{A}^{\mathcal{O}}$ there exists an extractor $\bar{\mathcal{E}}$. As a last step, we use this $\bar{\mathcal{E}}$ to show the existence of the extractor $\mathcal{E}_{\mathcal{A}}$.

For simplicity, parse the transcript of queries qt made by $\mathcal{A}^{\mathcal{O}}$ as a pair (\mathbf{q}, \mathbf{a}) , where \mathbf{q} are all the queries and \mathbf{a} all the corresponding answers, in order. Then we define the extractor $\mathcal{E}_{\mathcal{A}}$ in the following way.

$\mathcal{E}_{\mathcal{A}}(\text{crs}, aux, qt)$

Run $w \leftarrow \bar{\mathcal{E}}(\text{crs}, \overline{aux})$ where $\overline{aux} = (aux, \mathbf{a})$
 Return the same w returned by $\bar{\mathcal{E}}$

It is easy to see that, except for some syntactic changes, $\mathcal{E}_{\mathcal{A}}$ is the same as $\bar{\mathcal{E}}$. Combining all the above arguments, we have that for every \mathcal{A} there is an extractor $\mathcal{E}_{\mathcal{A}}$ such that adaptive proof of knowledge for \mathbb{O} holds with probability ϵ . \square

The case of pseudorandom functions As an interesting corollary of Theorem 6 we show that a similar result holds even for the case in which adversaries get access to a pseudorandom function as an oracle.

Corollary 2. *Let $F : \{0, 1\}^{\kappa} \times \{0, 1\}^{\ell} \rightarrow \{0, 1\}^L$ be a family of pseudorandom functions. Let \mathbb{O}_F be the family of oracles $\mathcal{O}_F(x) = F_K(x)$, where every family member \mathcal{O} is described by a seed K of the pseudorandom function, and thus the process $\mathcal{O}_F \leftarrow \mathbb{O}_F$ corresponds to sampling a random seed $K \xleftarrow{\$} \{0, 1\}^{\kappa}$. Let \mathcal{Z} be some distribution, \mathcal{U}_n be the uniform distribution over strings of length n , and denote by $(\mathcal{Z}, \mathcal{U})$ the distribution over pairs (aux, o) such that aux follows \mathcal{Z} and o follows \mathcal{U} .*

If Π is a $(\mathcal{Z}_q, \mathcal{U})$ -auxiliary input adaptive SNARK for every $q = \text{poly}(\lambda)$, and F is pseudorandom, then Π is a \mathcal{Z} -auxiliary input O-SNARK for \mathbb{O}_F .

Proof. The proof of the corollary follows in two steps. The first step relies directly on Theorem 6 to see that Π is an O-SNARK for the family \mathbb{O} of random oracles. This means that for every $\mathcal{A}^{\mathcal{O}}$ there is an extractor $\mathcal{E}_{\mathcal{A}}$ such that adaptive proof of knowledge for \mathbb{O} holds with probability δ .

The second step consists, informally, in replacing a random oracle with a PRF oracle and then showing that the success probability cannot change too much, unless the function is not pseudorandom.

For every $\tilde{\mathcal{A}}^{\mathcal{O}_F}$, we can construct another adversary $\mathcal{A}^{\mathcal{O}}$ that simply runs $\tilde{\mathcal{A}}^{\mathcal{O}_F}$, simulates all queries using its oracle, and returns the same output of $\tilde{\mathcal{A}}$. By adaptive proof of knowledge for \mathbb{O} there is an extractor $\mathcal{E}_{\mathcal{A}}$. We then define the extractor $\mathcal{E}_{\tilde{\mathcal{A}}} = \mathcal{E}_{\mathcal{A}}$.

Let us now consider the adaptive PoK experiment with $\tilde{\mathcal{A}}^{\mathcal{O}_F}$ and $\mathcal{E}_{\tilde{\mathcal{A}}}$ and let $\tilde{\delta}$ be the probability that the final condition holds. It is easy to prove that by the pseudorandomness of F there is a negligible function ν such that $\tilde{\delta} \leq \delta + \nu$. \square

5 Applications of O-SNARKs

In this section we show three applications of O-SNARKs for building homomorphic signatures [BF11], succinct functional signatures [BGI14], and SNARKs on authenticated data [BBFR15].

Generally speaking, our results show constructions of these primitives based on a signature scheme Σ and a succinct non-interactive argument Π , and show their security by assuming that Π is an O-SNARK for signing oracles corresponding to Σ . Once these results are established, we can essentially reach the following conclusions about the possible secure instantiations of these constructions. First, one can instantiate them by using Micali’s CS proofs as O-SNARK (cf. Section 4.1): this solution essentially yields secure instantiations in the random oracle model that work with a specific proof system (perhaps not the most efficient one in practice). Second, one can instantiate them with a classical SNARK and a hash-and-sign signature scheme (cf. Section 4.3), and conjecture that replacing the random oracle with a suitable hash function preserves the overall security. Third, one can instantiate the constructions using a classical SNARK construction Π and signature scheme Σ , and then conjecture that Π is an O-SNARK with respect to the family of signing oracles corresponding to Σ . Compared to the first solution, the last two ones have the advantage that one could use some of the recently proposed efficient SNARK schemes (e.g., [PHGR13, BSCG⁺13]); on the other hand these solutions have the drawback that the security of the instantiations would be heavily based on a heuristic argument. Finally, a fourth option that we provide are security proofs of these primitives which consider only non-adaptive adversaries (i.e., adversaries that declare all their queries in advance). In this case we can prove security based on non-adaptive O-SNARKs, and thus based on classical SNARKs (applying our Theorem 2). The advantage of this fourth option is that one obtains a security proof for these instantiations based on classical, not new, assumptions, although the proof holds only for a much weaker security notion.

5.1 Homomorphic Signatures

As first application of O-SNARKs we revisit a “folklore” construction of homomorphic signatures from SNARKs. This construction has been mentioned several times in the literature (e.g., [BF11, GW13, CF13, CFW14, GVW15]) and is considered as the ‘straightforward’ approach for constructing this primitive. In this section we formalize this construction, and notice that its security proof is quite subtle as one actually incurs the extraction issues that we mentioned in the introduction. Namely, one needs to run an extractor in an interactive security game in the presence of a signing oracle. Here we solve this issue by giving a simple proof based on our notion of O-SNARKs (for families of signing oracles).

Definition of Homomorphic Signatures. We begin by recalling the definition of homomorphic signatures. The definition below can be seen as the public key version of the notion of homomorphic message authenticators for labeled programs of Gennaro and Wichs [GW13].

LABELED PROGRAMS [GW13]. A labeled program consists of a tuple $\mathcal{P} = (F, \tau_1, \dots, \tau_n)$ such that $F : \mathcal{M}^n \rightarrow \mathcal{M}$ is a function on n variables (e.g., a circuit), and $\tau_i \in \{0, 1\}^\ell$ is the label of the i -th variable input of F . Let $F_{id} : \mathcal{M} \rightarrow \mathcal{M}$ be the canonical identity function and $\tau \in \{0, 1\}^\ell$

be a label. We consider $\mathcal{I}_\tau = (F_{id}, \tau)$ as the identity program for input label τ . Given t labeled programs $\mathcal{P}_1, \dots, \mathcal{P}_t$ and a function $G : \mathcal{M}^t \rightarrow \mathcal{M}$, the composed program \mathcal{P}^* is the one obtained by evaluating G on the outputs of $\mathcal{P}_1, \dots, \mathcal{P}_t$, and is compactly denoted as $\mathcal{P}^* = G(\mathcal{P}_1, \dots, \mathcal{P}_t)$. The labeled inputs of \mathcal{P}^* are all distinct labeled inputs of $\mathcal{P}_1, \dots, \mathcal{P}_t$, i.e., all inputs with the same label are grouped together in a single input of the new program.

Definition 9 (Homomorphic Signatures for Labeled Programs). *A homomorphic signature scheme HomSig is a tuple of probabilistic, polynomial-time algorithms ($\text{HomKG}, \text{HomSign}, \text{HomVer}, \text{HomEval}$) that work as follows*

$\text{HomKG}(1^\lambda)$ takes a security parameter λ and outputs a public key VK and a secret key SK . The public key VK defines implicitly a message space \mathcal{M} , the label space \mathcal{L} , and a set \mathcal{F} of admissible functions.

$\text{HomSign}(\text{SK}, \tau, m)$ takes a secret key SK , a (unique) label $\tau \in \mathcal{L}$ and a message $m \in \mathcal{M}$, and it outputs a signature σ .

$\text{HomEval}(\text{VK}, F, (\sigma_1, \dots, \sigma_n))$ takes a public key VK , a function $F \in \mathcal{F}$ and a tuple of signatures $(\sigma_1, \dots, \sigma_n)$. It outputs a new signature σ .

$\text{HomVer}(\text{VK}, \mathcal{P}, m, \sigma)$ takes a public key VK , a labeled program $\mathcal{P} = (F, (\tau_1 \dots \tau_n))$ with $F \in \mathcal{F}$, a message $m \in \mathcal{M}$, and a signature σ . It outputs either 0 (reject) or 1 (accept).

and satisfy authentication correctness, evaluation correctness, succinctness, and security, as described below.

- **Authentication Correctness.** Informally, we require that signatures generated by $\text{HomSign}(\text{SK}, \tau, m)$ verify correctly for m as the output of the identity program $\mathcal{I} = (F_{id}, \tau)$. Formally, HomSig has authentication correctness if for all key pairs $(\text{SK}, \text{VK}) \leftarrow \text{HomKG}(1^\lambda)$, any label $\tau \in \mathcal{L}$, message $m \in \mathcal{M}$, and any signature $\sigma \leftarrow \text{HomSign}(\text{SK}, \tau, m)$, $\text{HomVer}(\text{VK}, \mathcal{I} = (F_{id}, \tau), m, \sigma)$ outputs 1 with all but negligible probability.
- **Evaluation Correctness.** Intuitively, we require that running the evaluation algorithm on signatures $(\sigma_1, \dots, \sigma_n)$, where σ_i is a signature for m_i on label τ_i , produces a signature σ which verifies for $F(m_1, \dots, m_n)$. Formally, fix a key pair $(\text{SK}, \text{VK}) \leftarrow \text{HomKG}(1^\lambda, \mathcal{L})$, a function $G : \mathcal{M}^t \rightarrow \mathcal{M}$ and any set of program/message/signature triples $\{(\mathcal{P}_i, m_i, \sigma_i)\}_{i=1 \dots t}$ such that $\text{HomVer}(\text{VK}, \mathcal{P}_i, m_i, \sigma_i) = 1$. If $m^* = G(m_1 \dots m_t)$, $\mathcal{P}^* = G(\mathcal{P}_1, \dots, \mathcal{P}_t)$ and $\sigma^* = \text{HomEval}(\text{VK}, G, (\sigma_1, \dots, \sigma_t))$, then $\text{HomVer}(\text{VK}, \mathcal{P}^*, m^*, \sigma^*) = 1$ holds with all but negligible probability.
- **Succinctness.** For every large enough security parameter $\lambda \in \mathbb{N}$, there is a polynomial $p(\cdot)$ such that for every $(\text{SK}, \text{VK}) \leftarrow \text{HomKG}(1^\lambda)$ the output size of HomSign and HomEval is bounded by $p(\lambda)$ for any choice of their inputs.
- **Security.** A homomorphic signature scheme HomSig is secure if for every PPT adversary \mathcal{A} there is a negligible function ϵ such that $\Pr[\text{Exp}_{\mathcal{A}, \text{HomSig}}^{\text{HomSig-UF}}(\lambda) = 1] \leq \epsilon(\lambda)$ where the experiment $\text{Exp}_{\mathcal{A}, \text{HomSig}}^{\text{HomSig-UF}}(\lambda)$ is described in the following:

Key generation: Run $(\text{VK}, \text{SK}) \leftarrow \text{HomKG}(1^\lambda)$ and give VK to \mathcal{A} .

Signing queries: \mathcal{A} can adaptively submit queries of the form (τ, m) , where $\tau \in \mathcal{L}$ and $m \in \mathcal{M}$.

The challenger initializes an empty list T and proceeds as follows:

- * If (τ, m) is the first query with label τ , then the challenger computes $\sigma \leftarrow \text{HomSign}(\text{SK}, \tau, m)$, returns σ to \mathcal{A} and updates the list of queries $T \leftarrow T \cup \{(\tau, m)\}$.
- * If $(\tau, m) \in T$ (i.e., the adversary had already queried the tuple (τ, m)), then the challenger replies with the same signature generated before.

- * If T contains a tuple (τ, m_0) for some different message $m_0 \neq m$, then the challenger ignores the query.

Note that each label τ can be queried only once.

Forgery: After the adversary is done with the queries of the previous stage, it outputs a tuple $(\mathcal{P}^*, m^*, \sigma^*)$. Finally, the experiment outputs 1 iff the tuple returned by the adversary is a forgery (as defined below).

Forgeries are tuples $(\mathcal{P}^* = (F^*, (\tau_1^*, \dots, \tau_n^*)), m^*, \sigma^*)$ such that $\text{HomVer}(\text{VK}, \mathcal{P}^*, m^*, \sigma^*) = 1$ and they satisfy one the following conditions:

- * *Type 1 Forgery:* There is $i \in [n]$ such that $(\tau_i^*, \cdot) \notin T$ (i.e., no message m has ever been signed w.r.t. label τ_i^* during the experiment).
- * *Type 2 Forgery:* All labels τ_i^* have been queried— $\forall i \in [n], (\tau_i^*, m_i) \in T$ —but $m^* \neq F^*(m_1, \dots, m_n)$ (i.e., m^* is not the correct output of the labeled program \mathcal{P}^* when executed on the previously signed messages (m_1, \dots, m_n)).

Context-Hiding. Another property which is useful for homomorphic signatures is context-hiding. Intuitively, this property says that a signature on the output of a function does not reveal anything about its inputs, beyond what can be trivially learned by the verifier. Here we recall a (statistical) version of the definition proposed in [BF11] (also adapted to our syntax).

Definition 10 (Weak Context-Hiding). A homomorphic signature scheme HomSig is weakly context hiding if there exists a PPT simulator $S = (S^{\text{KG}}, S^{\text{Eval}})$ such that, for any fixed choice of function F , tuple of messages $m_1, \dots, m_n \in \mathcal{M}$, set of labels $\tau_1, \dots, \tau_n \in \mathcal{L}$, it holds that for any distinguisher \mathcal{D} and large enough security parameter $\lambda \in \mathbb{N}$:

$$\Pr \left[\mathcal{D}(\text{VK}, \text{SK}, \{\sigma_i\}_{i \in [n]}, \bar{\sigma}) = 1 \mid \begin{array}{l} (\text{VK}, \text{SK}) \leftarrow \text{HomKG}(1^\lambda) \\ \sigma_i \leftarrow \text{HomSign}(\text{SK}, \tau_i, m_i) \forall i \in [n] \\ \bar{\sigma} \leftarrow \text{HomEval}(\text{VK}, F, (\sigma_1, \dots, \sigma_n)) \end{array} \right] - \\ \Pr \left[\mathcal{D}(\text{VK}, \text{SK}, \{\sigma_i\}_{i \in [n]}, \bar{\sigma}) = 1 \mid \begin{array}{l} (\text{VK}, \text{SK}) \leftarrow S^{\text{KG}}(1^\lambda) \\ \sigma_i \leftarrow \text{HomSign}(\text{SK}, \tau_i, m_i) \forall i \in [n] \\ \bar{\sigma} \leftarrow S^{\text{Eval}}(\mathcal{P}, F(m_1, \dots, m_n)) \end{array} \right] \leq \text{negl}(\lambda)$$

Note that since the definition holds for fixed $F, \{m_i\}_i, \{\tau_i\}_i$, these values can also be given to \mathcal{D} . The notion is called *weak* context-hiding in contrast to a *strong* notion where one can also hide the fact that a homomorphic evaluation took place.

Homomorphic Signatures from O-SNARKs. To build the homomorphic signature we use a regular signature scheme Σ (cf. Appendix A.1 for the definition) and a fully-succinct O-SNARK Π for NP. The resulting scheme is homomorphic for all functions F whose running time is upper bounded by some fixed polynomial $t_{\mathcal{F}}(\cdot)$, and the scheme is 1-hop, i.e., it is not possible to apply HomEval on signatures obtained from other executions of HomEval .¹⁴

DEFINING THE MACHINE $M_{\Sigma, F}$. Let Σ be a signature scheme, and F be the description of a function $F : \mathcal{X}^n \rightarrow \mathcal{X}$ where \mathcal{X} is some appropriate domain (e.g., $\mathcal{X} = \{0, 1\}^\mu$). Then $M_{\Sigma, F}(x, w)$

¹⁴ Previous work hinted the possibility of achieving multi-hop homomorphic signatures by using SNARKs with recursive composition. However, given the issues we already notice in using classical SNARKs, it is unclear to us whether such a multi-hop construction would allow for a proof.

is the random-access machine that works as follows. It takes inputs (x, w) where values x are of the form $x = (\text{vk}, m, \tau_1, \dots, \tau_n)$ where vk is a public key of the scheme Σ , $m \in \mathcal{X}$ is a message and $\tau_i \in \{0, 1\}^\ell$ are labels, for $1 \leq i \leq n$. The values w are instead tuples $w = (m_1, \sigma_1, \dots, m_n, \sigma_n)$ where for every $i \in [n]$, $m_i \in \mathcal{X}$ is a message and σ_i is a signature of the scheme Σ . On input such a pair (x, w) , $M_{\Sigma, F}(x, w)$ accepts iff

$$m = F(m_1, \dots, m_n) \wedge \text{vfy}(\text{vk}, \tau_i | m_i, \sigma_i) = 1, \forall i = 1, \dots, n$$

Associated to such machine there is also a polynomial time bound $t_{\Sigma, \mathcal{F}}(k) = k^{e_{\Sigma, \mathcal{F}}}$, such that $M_{\Sigma, F}$ rejects if it does more than $t_{\Sigma, F}(|x|)$ steps. Finally, we note that given a polynomial bound $t_{\mathcal{F}}(k) = k^{e_{\mathcal{F}}}$ on the running time of every F supported by the scheme, a polynomial bound $t_{\Sigma}(k) = k^{e_{\Sigma}}$ on the running time of Σ 's verification algorithm, and values n, μ, ℓ , one can efficiently deduce the constant exponent $e_{\Sigma, \mathcal{F}}$ for the time bound $t_{\Sigma, \mathcal{F}}(|x|) = |x|^{e_{\Sigma, \mathcal{F}}}$.

We call \mathcal{R}_{Σ} the NP binary relation consisting of all pairs (y, w) such that, parsing $y = (M_{\Sigma, F}, x, t)$, $M_{\Sigma, F}(x, w)$ accepts in at most t steps and $t \leq t_{\Sigma, \mathcal{F}}(|x|)$.

THE CONSTRUCTION. Let $\Sigma = (\text{kg}, \text{sign}, \text{vfy})$ be a signature scheme and $\Pi = (\text{Gen}, \text{Prove}, \text{Ver})$ be a fully-succinct O-SNARK for NP. The homomorphic signature scheme $\text{HomSig}[\Sigma, \Pi]$ is defined as follows.

HomKG (1^λ) : Run $(\text{sk}, \text{vk}) \leftarrow \text{kg}(1^\lambda)$ and $\text{crs} \leftarrow \text{Gen}(1^\lambda)$. Define $\text{SK} = \text{sk}$ and $\text{VK} = (\text{vk}, \text{crs})$. Let the message be $\mathcal{M} = \{0, 1\}^\mu$ and the label space be $\mathcal{L} = \{0, 1\}^\ell$. Output (SK, VK) .

HomSign (SK, τ, m) : Run $\sigma \leftarrow \text{sign}(\text{sk}, \tau | m)$. Output $\bar{\sigma} = (\text{signature}, (\tau, m, \sigma))$.

HomEval $(\text{VK}, m, F, (\bar{\sigma}_1, \dots, \bar{\sigma}_n))$: Parse every $\bar{\sigma}_i = (\text{signature}, (\tau_i, m_i, \sigma_i))$, compute $m = F(m_1, \dots, m_n)$, reconstruct an instance $y = (M_{\Sigma, F}, x, t)$ where $x = (\text{vk}, m, \tau_1, \dots, \tau_n)$ and $t = |x|^{e_{\Sigma, \mathcal{F}}}$, and the witness $w = (m_1, \sigma_1, \dots, m_n, \sigma_n)$. Finally, run $\pi \leftarrow \text{Prove}(\text{crs}, y, w)$ and output $\bar{\sigma} = (\text{proof}, \pi)$.

HomVer $(\text{VK}, \mathcal{P} = (F, (\tau_1, \dots, \tau_n)), m, \bar{\sigma})$: Parse the signature $\bar{\sigma} = (\text{flag}, \cdot)$ and output the bit b computed as follows:

If $\bar{\sigma} = (\text{signature}, (\tau, m, \sigma))$ and $\mathcal{P} = \mathcal{I} = (F_{\text{id}}, \tau)$ run $\text{vfy}(\text{vk}, \tau | m, \sigma) \rightarrow b$.

If $\bar{\sigma} = (\text{proof}, \pi)$ run $\text{Ver}_{e_{\Sigma, \mathcal{F}}}(\text{crs}, y, \pi) \rightarrow b$ where $y = (M_{\Sigma, F}, x = (\text{vk}, m, \tau_1, \dots, \tau_n), |x|^{e_{\Sigma, \mathcal{F}}})$.

Recall that in a SNARK for NP, Ver_c is given a constant $c > 0$ and only works for relation \mathcal{R}_c .

In what follows we show that the scheme above is a homomorphic signature.

Authentication Correctness. For the present scheme the definition asks that signatures $\bar{\sigma} = (\text{signature}, (\tau, m, \sigma))$ as generated by $\text{HomSign}(\text{SK}, \tau, m)$ verify correctly for m as the output of the identity program $\mathcal{I} = (F_{\text{id}}, \tau)$. Observe that for all key pairs $(\text{SK}, \text{VK}) \leftarrow \text{HomKG}(1^\lambda)$, any label $\tau \in \mathcal{L}$ and any $\bar{\sigma} \leftarrow \text{HomSign}(\text{SK}, \tau, m)$, the algorithm $\text{HomVer}(\text{VK}, \mathcal{I} = (F_{\text{id}}, \tau), \bar{\sigma})$ performs the verification by running $\text{vfy}(\text{vk}, \tau | m, \sigma)$. By the correctness of the signature scheme Σ , for any signature $\sigma \leftarrow \text{sign}(\text{sk}, \tau | m)$, $\text{vfy}(\text{vk}, \tau | m, \sigma)$ outputs 1 with all but negligible probability.

Therefore authentication correctness of $\text{HomSig}(\Sigma, \Pi)$ follows.

Evaluation Correctness. To see evaluation correctness, fix a key pair $(\text{SK}, \text{VK}) \leftarrow \text{HomKG}(1^\lambda)$, and a set of triples of the form $\{(\tau_i, m_i, \bar{\sigma}_i)\}_{i=1..n}$ where $\bar{\sigma}_i = (\text{signature}, (\tau_i, m_i, \sigma_i))$ and for which $\text{HomVer}(\text{VK}, \mathcal{I}_i = (F_{\text{id}}, \tau_i), m_i, \bar{\sigma}_i) = 1$.

Let $\bar{\sigma}^* \leftarrow \text{HomEval}(\text{VK}, F^*, (\bar{\sigma}_1, \dots, \bar{\sigma}_n))$, where $\bar{\sigma}^* = (\text{proof}, \pi)$ with π a proof generated by $\text{Prove}(\text{crs}, y, w)$ by setting $y = (M_{\Sigma, F}, x = (\text{vk}, m^*, \tau_1, \dots, \tau_n), |x|^{e_{\Sigma, \mathcal{F}}})$ and $w = (m_1, \sigma_1, \dots, m_n, \sigma_n)$.

If $m^* = F^*(m_1 \dots m_n)$ (for an F^* that is in the class of supported functions) and since $\text{HomVer}(\text{VK}, \mathcal{I}_i, m_i, \bar{\sigma}_i)$ means $\text{vfy}(\text{vk}, \tau_i | m_i, \sigma_i) = 1$, we have that $(y, w) \in \mathcal{R}_\Sigma$. Thus $\text{HomVer}(\text{VK}, \mathcal{P}^*, m^*, \bar{\sigma}^*) = 1$ holds with all but negligible probability by the correctness of the SNARK $\text{Ver}_{e_{\Sigma, \mathcal{F}}}(\text{crs}, y, \pi) = 1$.

Succinctness. As the output of the HomEval algorithm is a proof π obtained by running the Prove algorithm, the succinctness of $\text{HomSig}(\Sigma, \Pi)$ follows from the succinctness of Π .

Security. As in Section 4.2, for every signature scheme $\Sigma = (\text{kg}, \text{sign}, \text{vfy})$ we denote by \mathbb{O}_Σ the family of oracles $\mathcal{O}(m) = \text{sign}(\text{sk}, m)$ (where the verification key is returned as output of a special query $\mathcal{O}('vk')$). We show the security of the scheme $\text{HomSig}[\Sigma, \Pi]$ via the following theorem.

Theorem 7. *Let Σ be a signature scheme. If Π is an O-SNARK for \mathbb{O}_Σ , and Σ is UF-CMA-secure, then $\text{HomSig}[\Sigma, \Pi]$ is a secure homomorphic signature scheme.*

Proof. To prove that this scheme is a secure homomorphic signature, we assume by contradiction the existence of an efficient adversary \mathcal{A} that is able to output a forgery of one of the two types with non-negligible probability δ . Starting from this algorithm \mathcal{A} , we further construct a successful forger \mathcal{B} against the UF-CMA-security of Σ , which leads to a contradiction. Along the way of this reduction, we also rely on the fact that Π is an O-SNARK satisfying adaptive proof of knowledge for \mathbb{O}_Σ .

By looking at the definition of the security experiment $\text{Exp}_{\mathcal{A}, \text{HomSig}}^{\text{HomSig-UF}}(\lambda)$ for the HomSig scheme, adversary \mathcal{A} is almost equivalent to an O-SNARK adversary $\tilde{\mathcal{A}}^\mathcal{O}$ for oracle $\mathcal{O} \leftarrow \mathbb{O}_{\text{sign}}$. Stated more formally, for every adversary \mathcal{A} against HomSig that outputs a forgery $\mathcal{P}^* = (F^*, (\tau_1^*, \dots, \tau_n^*), m^*, \bar{\sigma}^* = (\text{proof}, \pi))$, it is possible to construct another adversary $\tilde{\mathcal{A}}^\mathcal{O}(\text{crs})$ working as follows: it queries $\text{vk} \leftarrow \mathcal{O}('vk')$; it runs $\mathcal{A}^{\text{HomSig}(\text{SK}, \cdot)}(\text{VK} = (\text{vk}, \text{crs}))$ simulating \mathcal{A} 's oracle queries using its own oracle \mathcal{O} ; finally it returns the value (y, π) , where $y = (M_{\Sigma, F^*}, x, |x|^{e_{\Sigma, \mathcal{F}}})$, with $x = (\text{vk}, m^*, \tau_1^*, \dots, \tau_n^*)$, is obtained from \mathcal{A} 's output. The adversary $\tilde{\mathcal{A}}^\mathcal{O}$ perfectly fits the O-SNARK definition by which we know that there exists an extractor $\mathcal{E}_{\tilde{\mathcal{A}}}$ that, given the same input of $\tilde{\mathcal{A}}^\mathcal{O}$ and the transcript of oracle queries/answers made and received by $\tilde{\mathcal{A}}^\mathcal{O}$, outputs a correct witness w (i.e., such that $(y, w) \in \mathcal{R}_\Sigma$) with all but negligible probability.

Hence, we have that for every successful adversary \mathcal{A} against HomSig there exists extractor $\mathcal{E}_{\tilde{\mathcal{A}}}$, that takes the very same input of \mathcal{A} (plus, the list of oracle answers). Starting from this adversary \mathcal{A} , we construct the forger $\mathcal{B}^\mathcal{O}$ that breaks the UF-CMA security of Σ . We build \mathcal{B} (which gets the public key vk and can make queries to $\mathcal{O} = \text{sign}(\text{sk}, \cdot)$ oracle) as follows:

$\mathcal{B}^\mathcal{O}(\text{vk}) :$

Initialize $\text{qt} \leftarrow ('vk', \text{vk})$

Generate $\text{crs} \leftarrow \text{Gen}(1^\lambda)$ and run $\mathcal{A}(\text{VK} = (\text{crs}, \text{vk}))$

Simulate queries (τ, m) to \mathcal{O} as follows:

 query $\sigma \leftarrow \mathcal{O}(\tau | m)$ and add $(\tau | m, \sigma)$ to qt

 output σ

When \mathcal{A} outputs $(\mathcal{P}^* = (F^*, (\tau_1^*, \dots, \tau_n^*), m^*, \bar{\sigma}^*)$ parse $\bar{\sigma}^* = (\text{proof}, \pi^*)$

Run $\mathcal{E}_{\tilde{\mathcal{A}}}(\text{crs}, \text{qt})$ and obtain the witness $w = (m_1^*, \sigma_1^*, \dots, m_n^*, \sigma_n^*)$

Check that $(y, w) \in \mathcal{R}_\Sigma$, i.e., that $M_\Sigma(x, w) = 1$ in at most $|x|^{e_{\Sigma, \mathcal{F}}}$ steps

where $M_\Sigma(x, w) := (m^* = F(m_1^*, \dots, m_n^*) \wedge \text{vfy}(\text{vk}, \tau_i^* | m_i^*, \sigma_i^*) = 1, \forall i \in [n])$

 [Fail] Abort if $(y, w) \notin \mathcal{R}_\Sigma$.

Else proceed:

- [Type 1] If $\exists j \in [n]$ such that $(\tau_j^*|\cdot, \cdot) \notin \text{qt}$ return $(\tau_j^*|m_j^*, \sigma_j^*)$
- [Type 2] If $\forall i \in [n]$ there is a tuple $(\tau_i^*|m_i, \sigma_i) \in \text{qt}$
and there $\exists j \in [n]$ such that $m_j^* \neq m_j$ return $(\tau_j^*|m_j^*, \sigma_j^*)$

Let us now show that whenever \mathcal{A} succeeds in the simulation described above, \mathcal{B} is also successful in breaking the UF-CMA security of Σ , unless the “Fail” event happens.

First, let us condition on the event that \mathcal{B} does not abort, i.e., “Fail” does not occur, that is the extractor $\mathcal{E}_{\tilde{\mathcal{A}}}$ is correct in returning a valid witness. If \mathcal{A} outputs the first type of forgery, since the witness w is valid— $\text{vfy}(\text{vk}, \tau_i^*|m_i^*, \sigma_i^*) = 1, \forall i = 1, \dots, n$ —it follows that the signature forgery $(\tau_j^*|m_j^*, \sigma_j^*)$ is a valid one.

If the forgery returned by \mathcal{A} is of the second type, recall this means that $m^* \neq F(m_1, \dots, m_n)$ where all inputs of F are the ones in the transcript qt , i.e., qt contains the tuples $(\tau_1^*|m_1, \sigma_1), \dots, (\tau_n^*|m_n, \sigma_n)$. Combining this with the validity of w we have that $F(m_1^*, \dots, m_n^*) = m^* \neq F(m_1, \dots, m_n)$, and thus $(m_1^*, \dots, m_n^*) \neq (m_1, \dots, m_n)$. This means that there exists at least one j such that $(\tau_j^*|m_j, \sigma) \in \text{qt}$ but $m_j^* \neq m_j$. Moreover, by witness validity it also holds $\text{vfy}(\text{vk}, \tau_j^*|m_j^*, \sigma_j^*) = 1$. Thus we can conclude that, even in this case, $(\tau_j^*|m_j^*, \sigma_j^*)$ is a valid forgery for Σ .

So far, we have proved that whenever the adversary \mathcal{A} is able to output a valid forgery and \mathcal{B} does not abort, then \mathcal{B} is a successful adversary against the UF-CMA security of Σ . However, whenever \mathcal{A} is successful (with non-negligible probability), by the O-SNARK definition we have that “Fail” occurs with negligible probability at most ϵ . Therefore, if \mathcal{A} is successful with probability at least δ , then \mathcal{B} is successful with non-negligible probability $\geq \delta - \epsilon$. \square

Non-adaptive security. Alternatively, one can modify the previous proof to show that the scheme has security against homomorphic signature adversaries that make non-adaptive signing queries, assuming the weaker assumption that Π is a *non-adaptive* O-SNARK (see Definition 8). In particular, combining this change with the result of Theorem 2 one obtains the following:

Theorem 8. *If Π is a SNARK, and Σ is a UF-CMA-secure signature scheme, then $\text{HomSig}[\Sigma, \Pi]$ is secure against adversaries that make non-adaptive signing queries.*

Proof. The proof is very similar to that of Theorem 7 and thus we only sketch the main differences from that proof. To work with non-adaptive adversaries the only main change is that for every non-adaptive adversary \mathcal{A} one can define a corresponding non-adaptive O-SNARK adversary $\tilde{\mathcal{A}}$. Then the only difference is that the non-adaptive queries of \mathcal{A} can be used to define the non-adaptive queries of $\tilde{\mathcal{A}}$. The rest of the proof proceeds analogously.

Remark 3 (On the applicability of Corollary 1). We note that we cannot combine the positive result of Corollary 1 with Theorem 7 to conclude that the security of the homomorphic signature scheme holds under classical SNARKs. The inapplicability of Corollary 1 is due to its restriction for which adversaries have to query almost the entire message space. By looking at the HomSig construction (and the definition of homomorphic signatures too) one can note that an adversary who queries almost the entire message space of the underlying signature scheme can trivially break the security (for example he could obtain signatures on two distinct messages under the same label).

Insecurity of $\text{HomSig}[\Sigma^*, \Pi]$. Here we show the existence of a signature scheme Σ^* for which $\text{HomSig}[\Sigma^*, \Pi]$ is insecure. Note that this insecurity result does not contradict our Theorem 7 as it is indeed possible to show that Σ^* is in the class of schemes for which the existence of an

O-SNARK is ruled out. Rather, this counterexample shows that the issue with proving the security of the homomorphic signature construction is not a mere difficulty in doing the proof, but that the construction can actually become insecure, if one simply relies on an *arbitrary* signature scheme.

Construction of Σ^* . Consider the **HomSig** construction in which messages are bits (i.e., $\mu = 1$) and labels are strings of length $\ell = \text{poly}(\lambda)$, such that the $\text{Prove}(\text{crs}, \cdot, \cdot)$ algorithm of Π is at most $(\ell + 1)$ -bits long. Let $\widehat{\Sigma}$ be any UF-CMA-secure scheme with message space $\mathcal{M} = \{0, 1\}^{\ell+1}$. We construct the signature scheme Σ^* from $\widehat{\Sigma}$ as follows:

$\text{kg}(1^\lambda)$: Run $(\widehat{\text{vk}}, \widehat{\text{sk}}) \xleftarrow{\$} \widehat{\Sigma}.\text{kg}(1^\lambda)$, set $\text{vk} = \widehat{\text{vk}}$, $\text{sk} = \widehat{\text{sk}}$, and also set $\delta = p(\lambda) + \lambda$.

$\text{sign}(\text{sk}, m)$: Signing works as follows

- sample $r \xleftarrow{\$} \{0, 1\}^\delta$, and compute $\hat{\sigma} \leftarrow \widehat{\Sigma}.\text{sign}(\widehat{\text{sk}}, r|m)$;
- sample $m' \xleftarrow{\$} \{0, 1\}^{\ell+1}$, $r' \xleftarrow{\$} \{0, 1\}^\delta$, compute $\hat{\sigma}' \leftarrow \widehat{\Sigma}.\text{sign}(\widehat{\text{sk}}, r'|m')$, and set $\sigma' = (r', \hat{\sigma}', 0, 0)$;
- parse $m' \in \{0, 1\}^{\ell+1}$ bit by bit as $(m'_1 | \dots | m'_{\ell+1})$, set $\tau = m'_1 | \dots | m'_\ell$, $x = (\text{vk}, m'_{\ell+1}, \tau)$, $w = (m'_{\ell+1}, \sigma')$, and let $\mathcal{I} : \{0, 1\} \rightarrow \{0, 1\}$ be the identity function;
- let $t = \#M_{\Sigma^*, \mathcal{I}}(x, w)$, and set $y = (M_{\Sigma^*, \mathcal{I}}(x, t))$;
- Interpret m as the description of program $\mathcal{P}(\cdot, \cdot)$ and thus run $\pi \leftarrow \mathcal{P}(y, w)$;
- if $|\pi| > p(\lambda)$ set $\pi' = 0$ and $y = 0$, else $\pi' = \pi$.
- output $\sigma = (r, \hat{\sigma}, m', \pi')$.

$\text{vfy}(\text{vk}, m, \sigma)$: Parse $\sigma = (r, \hat{\sigma}, m', \pi')$ where r is δ -bits long, and return the output of $\widehat{\Sigma}.\text{vfy}(\widehat{\text{vk}}, r|m, \hat{\sigma})$.

Before proving the security of Σ^* , we provide some intuitions about the rationale of the above construction:

- The signing algorithm consists of two main steps. First, given the message m , we sign $r|m$ (i.e., we prepend the random string r to m) using $\widehat{\Sigma}.\text{sign}$. Second, we generate another signature on a randomly chosen message $m' \in \{0, 1\}^{\ell+1}$. This is done following the same process as for m , i.e., by sampling a random $r' \xleftarrow{\$} \{0, 1\}^\delta$ and signing $r'|m'$ using $\widehat{\Sigma}.\text{sign}$.
- Then, we construct a theorem (y, w) for the relation \mathcal{R}_{Σ^*} with respect to the identity function \mathcal{I} . In particular, recall the definition of $M_{\Sigma^*, \mathcal{I}}(x, w)$ from Section 5.1: this is the machine that on inputs $x = (\text{vk}, m, \tau)$, where vk is a public key of the scheme Σ^* , m is a bit and $\tau \in \{0, 1\}^\ell$ and $w = (m^*, \sigma^*)$ accepts iff

$$m = \mathcal{I}(m^*) \wedge \text{vfy}(\text{vk}, \tau|m, \sigma^*) = 1$$

- Finally, one interprets the input message m as a program \mathcal{P} description and runs $\pi \leftarrow \mathcal{P}(y, w)$. Its output, together with m' , is included in the signature as an extra information. As one can see, if $\mathcal{P}(\cdot, \cdot) = \text{Prove}(\text{crs}, \cdot, \cdot)$ then π is a valid proof. Yet, before returning π we make sure that (regardless of its validity) π is short enough and we set the new value π' .
- The meaning of δ : The reason of adding the random string r to the signature is to increase the entropy of the signatures generated by Σ^* . This is crucial for the signatures σ' that are used to generate π . Basically, we want to make sure that π cannot leak enough information about σ' . Since the output of the program is of length p , taking r' to be sufficiently long – of $p(\lambda) + \lambda$ bits – guarantees that some information about σ' is inevitably lost.

Proof of UF-CMA-security for Σ^* . We prove that Σ^* is secure as long as $\widehat{\Sigma}$ is secure:

Lemma 3. *If $\widehat{\Sigma}$ is UF-CMA-secure scheme, then Σ^* is UF-CMA-secure. More precisely, for any PPT adversary \mathcal{A} that has advantage $\epsilon(\lambda)$ in breaking the security of Σ^* by making Q signing queries there is a PPT adversary $\hat{\mathcal{A}}$ that breaks the security of $\widehat{\Sigma}$ with advantage $> \epsilon(\lambda) - Q/2^\lambda$ by making $2Q$ queries.*

Below we mention the steps for the proof of the lemma. The idea of the proof is that $\hat{\mathcal{A}}$ runs $(m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{sign}(\widehat{\text{sk}}, \cdot)}(\widehat{\text{vk}})$ and simulates every signing query m by executing $\text{sign}(\text{sk}, m)$ except that the executions of $\widehat{\Sigma}.\text{sign}(\widehat{\text{sk}}, \cdot)$ are replaced by queries to $\hat{\mathcal{A}}$'s signing oracle (this is why every signing query produces two queries on $\hat{\mathcal{A}}$'s side). The only tricky part of the proof is to show that the message m^* used in the forgery leads to a *new* message $\hat{m}^* = r^*|m^*$ in the game played by $\hat{\mathcal{A}}$, i.e., that $r^*|m^* \neq r_i|m_i$ for all the $r_i|m_i$ queried by $\hat{\mathcal{A}}$. We argue that this is the case with overwhelming probability, based on the random choices of all values r' in the signing query simulation. In fact, one should note that \mathcal{A} never gets to see the value r' used to generate σ' ; moreover since r' is δ -bits long and \mathcal{A}_p sees at most $p(\lambda)$ bits of information of it, then λ bits of r' are always lost. Therefore the probability that $r^* = r'$ is bounded by the probability that \mathcal{A} predicts correctly r' .

Proof. Assume by contradiction that there exists an adversary \mathcal{A} that has non-negligible advantage $\epsilon(\lambda)$ against the UF-CMA security of Σ_p while running in polynomial time and making Q queries to the $\text{sign}(\text{sk}, \cdot)$ oracle. Starting from this adversary \mathcal{A} , we construct a PPT adversary $\hat{\mathcal{A}}$ that is able to break the UF-CMA security of $\widehat{\Sigma}$ with non-negligible advantage and by making at most $2Q$ queries to its signing oracle $\widehat{\Sigma}.\text{sign}(\widehat{\text{sk}}, \cdot)$.

We define $\hat{\mathcal{A}}$ (which gets the public key $\widehat{\text{vk}}$ and makes queries to $\widehat{\Sigma}.\text{sign}(\widehat{\text{sk}}, \cdot)$ oracle) as follows:

$\hat{\mathcal{A}}^{\widehat{\Sigma}.\text{sign}(\widehat{\text{sk}}, \cdot)}(\widehat{\text{vk}})$

Run $(m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{sign}(\widehat{\text{sk}}, \cdot)}(\widehat{\text{vk}})$ and simulate queries m to $\text{sign}(\widehat{\text{sk}}, \cdot)$ as follows:

sample $r \xleftarrow{\$} \{0, 1\}^\delta$, and query $\hat{\sigma} \leftarrow \widehat{\Sigma}.\text{sign}(\widehat{\text{sk}}, r|m)$;
sample $m' \xleftarrow{\$} \{0, 1\}^{\ell+1}$, $r' \xleftarrow{\$} \{0, 1\}^\delta$;
query $\hat{\sigma}' \leftarrow \widehat{\Sigma}.\text{sign}(\widehat{\text{sk}}, r'|m')$, and set $\sigma' = (r', \hat{\sigma}', 0, 0)$;
parse $m' \in \{0, 1\}^{\ell+1}$ as $(m'_1 | \dots | m'_{\ell+1})$;
set $\tau = m'_1 | \dots | m'_\ell$, $x = (\text{vk}, m'_{\ell+1}, \tau)$, $w = (m'_{\ell+1}, \sigma')$;
let $t = \#M_{\Sigma^*, \mathcal{I}}(x, w)$, and set $y = (M_{\Sigma^*, \mathcal{I}}, x, t)$, where I is the identity;
interpret m as the description of program $\mathcal{P}(\cdot, \cdot)$ and thus run $\pi \leftarrow \mathcal{P}(y, w)$;
if $|\pi| > p(\lambda)$ set $\pi' = 0$ and $m' = 0$, else $\pi' = \pi$;
output $\sigma = (r, \hat{\sigma}, m', \pi')$.

Parse $\sigma^* = (r^*, \hat{\sigma}^*, \cdot, \cdot)$ and return $(r^*|m^*, \hat{\sigma}^*)$

Let us now show that whenever \mathcal{A} succeeds in the simulation described above, $\hat{\mathcal{A}}$ succeeds in breaking the UF-CMA security of the scheme $\widehat{\Sigma}$, with all but negligible probability. To this end we have to first show that the simulation provided by $\hat{\mathcal{A}}$ works correctly, and then show that $\hat{\mathcal{A}}$ outputs a valid forgery as long as \mathcal{A}_p outputs a forgery.

To ease the analysis, consider the set of all queries (and corresponding responses) made by \mathcal{A} :

$$\mathcal{Q}^* = \{m_i, \sigma_i = (r_i, \hat{\sigma}_i, m'_i, \pi_i) \mid i = 1 \dots Q\}$$

Then the set of $\hat{\mathcal{A}}$'s queries is $\widehat{\mathcal{Q}} = \widehat{\mathcal{Q}}^* \cup \widehat{\mathcal{Q}}'$ with

$$\widehat{\mathcal{Q}}^* = \{(\hat{m}_i = r_i|m_i, \hat{\sigma}_i) \mid i = 1 \dots Q\}$$

$$\hat{\mathcal{Q}}' = \{(\hat{m}'_j = r'_j | m'_j, \hat{\sigma}'_j) \mid j = 1 \dots Q\}$$

Precisely, the first set $\hat{\mathcal{Q}}^*$ consists of all signing queries asked by $\hat{\mathcal{A}}$ to its oracle for signing the messages m_i queried by \mathcal{A} . The second set $\hat{\mathcal{Q}}'$ instead, comprises the extra queries asked by $\hat{\mathcal{A}}$ in the simulation for signing the sampled messages m'_j .

It is easy to see that $\hat{\mathcal{A}}$ provides a perfect simulation to \mathcal{A} as $\hat{\mathcal{A}}$ can correctly answer every query of \mathcal{A} using its own signing oracle.

So the main fact to show is that the message m^* used in the forgery leads to a *new* message \hat{m}^* in the game played by $\hat{\mathcal{A}}$.

Let $(m^*, \sigma^* = (r^*, \hat{\sigma}^*, \cdot, \cdot))$ be a valid forgery for Σ^* (i.e., $m^* \neq m_i$, for all $m_i \in \mathcal{Q}^*$), and let us consider the following undesired cases:

1. $(\hat{m}^* = r^* | m^*, \cdot) \in \hat{\mathcal{Q}}^*$: Since $m^* \neq m_i \forall m_i \in \mathcal{Q}^*$ this case cannot occur even if the corresponding strings r_i and r^* match.
2. $(\hat{m}^* = r^* | m^*, \cdot) \in \hat{\mathcal{Q}}'$: It must be that $\hat{m}^* = r^* | m^* = r'_j | m'_j = \hat{m}'_j$ for some $j \in \{1, \dots, Q\}$. In what follows we bound the probability that such equality happens and show that it is negligible.

Both r^* and r'_j are parsed as strings of the *same* length $\delta = p(\lambda) + \lambda$. Hence, $r^* | m^* = r'_j | m'_j$ immediately implies $m^* = m'_j$, which may be possible since m^* is of adversarial choice. To bound the probability of match we thus only look at the event that $r^* = r'_j$.

Now, the crucial observation is that the adversary \mathcal{A} never sees the strings r'_j explicitly, and thus the probability of the match can be upper bounded by the probability that the adversary \mathcal{A} guesses correctly the string $r'_j \in \{0, 1\}^\delta$ where $\delta = p(\lambda) + \lambda$.

Below we argue that this happens with negligible probability $\leq \frac{Q}{2^\lambda}$. For $j \in \{1, \dots, Q\}$ let Bad_j be the event that $r^* = r'_j$ for $(\hat{m}'_j, \hat{\sigma}'_j) \in \hat{\mathcal{Q}}'$ and let $\text{Bad} = \bigvee_{j=1}^Q \text{Bad}_j$. Using the union bound we have:

$$\Pr[\text{Bad}] = \Pr \left[\bigvee_{j=1}^Q \text{Bad}_j \right] \leq \sum_{j=1}^Q \Pr[\text{Bad}_j]$$

Now, we will bound the probability of Bad_j for *any* fixed j . The value $\pi_j \leftarrow \mathcal{P}(y, w)$ the adversary \mathcal{A} gets from the j -th query is the only one that can reveal some information about $\sigma'_j = (r'_j, \hat{\sigma}'_j, 0, 0)$ and implicitly about r'_j . We show that the information gained from π_j does not give any advantage to the adversary.

Let us represent the process of running $\pi \leftarrow \mathcal{P}(y, (m'_\lambda, \sigma'_j = (R, \hat{\sigma}'_j, 0, 0)))$ and returning π only if $|\pi| \leq p(\lambda)$ as a function $f(R)$ such that $f : \{0, 1\}^\delta \rightarrow \{0, 1\}^p$. Namely, we fix \mathcal{P} and all its inputs but R . Observe that, for random chosen inputs, any such f is essentially performing a lossy data compression of $\delta - p(\lambda) = \lambda$ bits.

We have that $\forall f : \{0, 1\}^\delta \rightarrow \{0, 1\}^p$, the probability that any algorithm $\tilde{\mathcal{A}}$ guesses the random string r'_j on input $f(r'_j)$ is less than $\frac{2^p}{2^\delta} = \frac{1}{2^\lambda}$. The same holds for \mathcal{A} and f defined as before. Hence, summing up these probabilities for all j and observing that $Q = \text{poly}(\lambda)$, we obtain:

$$\Pr[\text{Bad}] \leq \sum_{j=1}^Q \Pr[\text{Bad}_j] \leq \frac{Q}{2^\lambda}$$

that is negligible.

The proof is concluded by observing that

$$\mathbf{Adv}_{\hat{\mathcal{A}}, \hat{\Sigma}}^{\text{UF-CMA}}(\lambda) \geq \epsilon(\lambda)(1 - \Pr[\text{Bad}]) \geq \epsilon(\lambda) - \frac{Q}{2^\lambda}$$

□

Proof of insecurity of $\text{HomSig}[\Sigma^*, \Pi]$. Below we show an adversary \mathcal{A}^* that has non-negligible probability of winning in $\mathbf{Exp}_{\mathcal{A}^*, \text{HomSig}[\Sigma^*, \Pi]}^{\text{HomSig-UF}}(\lambda)$. Our adversary wins by producing a Type 1 forgery with respect to an identity function. For ease of exposition, we give the following security experiment which is a specialization of $\mathbf{Exp}^{\text{HomSig-UF}}$ to adversaries that only output Type 1 forgeries for identity functions:

Experiment $\mathbf{Exp}_{\mathcal{A}, \text{HomSig}[\Sigma, \Pi]}^{\text{Type-1, Id}}(\lambda)$

$(\text{SK}, \text{VK}) \xleftarrow{\$} \text{HomKG}(1^\lambda)$

$(\mathcal{I}^* = (F_{id}, \tau^*), m^*, \bar{\sigma}^*) \xleftarrow{\$} \mathcal{A}^{\text{HomSign}(\text{SK}, \cdot)}(\text{VK})$

If $\text{HomVer}(\text{VK}, \mathcal{I}^*, m^*, \bar{\sigma}^*) = 1$ and τ^* is “new”, output 1

Else output 0

Lemma 4. *Let Π be a SNARK, and Σ^* be the UF-CMA-secure signature scheme defined above. Then there exists an efficient adversary \mathcal{A}^* such that*

$$\Pr[\mathbf{Exp}_{\mathcal{A}^*, \text{HomSig}[\Sigma^*, \Pi]}^{\text{Type-1, Id}}(\lambda) = 1] = 1 - \text{negl}(\lambda)$$

Proof. Below is the description of our adversary \mathcal{A}^* :

Adversary $\mathcal{A}^{\text{HomSign}(\text{SK}, \cdot)}(\text{VK})$

- 1 Query the signing oracle on $(\tau, m) := \text{P}$
where P is the description of $\text{Prove}(\text{crs}, \cdot, \cdot)$
- 2 Parse the answer $\sigma = (r, \hat{\sigma}, m', \pi')$, and $m' = (m'_1 | \dots | m'_{\ell+1})$
- 3 Set $\tau^* = (m'_1 | \dots | m'_\ell)$ and $m^* = m'_{\ell+1}$
- 4 Return $(\mathcal{I}^* = (F_{id}, \tau^*), m^*, \bar{\sigma}^* = (\text{proof}, \pi'))$

Note that, except with probability $2^{-\ell}$, it holds $\tau^* \neq \tau$. Moreover, by the correctness of Π (and its succinctness), the answer to \mathcal{A}^* 's oracle query contains a valid proof π' that verifies for the identity function with message m^* and label τ^* . In other words the output of \mathcal{A}^* constitutes a Type-1 forgery with probability $1 - 2^{-\ell}$. □

Context-Hiding of HomSig.

Theorem 9. *If Π is a zero-knowledge O-SNARK then HomSig is weakly context-hiding.*

Proof. To show weakly context-hiding for $\text{HomSig}[\Sigma, \Pi]$ we define a simulator $S^{\text{Hide}} = (S^{\text{KG}}, S^{\text{Eval}})$. For this purpose we use the PPT simulator $S^{\Pi} = (S^{\text{crs}}, S^{\text{Prove}})$ from the zero-knowledge of Π .

$S^{\text{KG}}(1^\lambda)$:

Run $(\text{vk}, \text{sk}) \leftarrow \text{kg}(1^\lambda)$

Run $(\text{prs}, \text{vst}, \text{tr}) \leftarrow S^{\text{crs}}(1^\lambda, T)$

Set $\text{SK} = (\text{sk}, \text{prs})$ and $\text{VK} = (\text{vk}, \text{vst})$

Output (VK, SK)

$S^{\text{Eval}}(\mathcal{P} = (F, \tau_1, \dots, \tau_n), m^*) :$

Set $y = (M_{\Sigma, F}, x = (\text{vk}, m^*, \tau_1 \dots \tau_n), |x|^{e_{\Sigma, \mathcal{F}}})$

Run $\pi \leftarrow S^{\text{Prove}}(\text{prs}, \text{tr}, y)$

Output $\bar{\sigma} = (\text{proof}, \pi)$

For any distinguisher $\mathcal{D}^{\text{Hide}}$ against the weakly context-hiding of $\text{HomSign}[\Sigma, \Pi]$, we can easily construct a distinguisher \mathcal{D}^{Π} against zero knowledge O-SNARK property:

$\mathcal{D}^{\Pi}(\text{crs}) :$

Generate a pair $(\text{sk}, \text{vk}) \leftarrow \text{kg}(1^\lambda)$

Set $\text{SK} = (\text{sk}, \text{prs})$ and $\text{VK} = (\text{vk}, \text{vst})$

Let $F, m_1, \dots, m_n, \tau_1 \dots \tau_n$ be the fixed tuple:

Run $\sigma_i \leftarrow \text{HomSign}(\text{sk}, \tau_i, m_i) \forall i \in [n]$

Output (y, w) to its challenger, and get back π

(where $y = (M_{\Sigma, F}, x = (\text{vk}, m^* = F(m_1, \dots, m_n), \tau_1 \dots \tau_n), |x|^{e_{\Sigma, \mathcal{F}}})$

and $w = (m_1, \sigma_1, \dots, m_n, \sigma_n)$)

run $b \leftarrow \mathcal{D}^{\text{Hide}}(\text{VK}, \text{SK}, \sigma_1, \dots, \sigma_n, \tilde{\sigma} = (\text{proof}, \pi))$

Output b .

Clearly, if \mathcal{D}^{Π} receives π and crs generated using the real algorithms, it simulates the real game to $\mathcal{D}^{\text{Hide}}$. Otherwise, if these values are generated through the zero-knowledge simulator, then the view of $\mathcal{D}^{\text{Hide}}$ is identical to the case where it receives values generated using our context hiding simulator described above. Therefore, the distinguishing advantage of \mathcal{D}^{Π} in distinguishing between a real or a simulated proof is the same as that of the algorithm $\mathcal{D}^{\text{Hide}}$ in distinguish the two distributions in the answers from the context-hiding definition.

5.2 Succinct Functional Signatures

As second application of O-SNARKs we revisit the construction of succinct functional signatures of Boyle, Goldwasser, and Ivan [BGI14]. In [BGI14] this construction is proven secure using a notion of SNARKs which significantly differs from the standard one [BCC⁺14] (for completeness we recall this notion in Appendix B). To the best of our knowledge, there are no known instantiations of SNARKs under this definition, in the standard model (and is not clear whether it is possible to find some). On the other hand, if one wants to prove the security of this construction using the classical SNARK definition, the security proof incurs the same subtleties related to running an extractor in the presence of a signing oracle.

In this section, we revisit the construction of [BGI14], and we prove its security using O-SNARKs. Interestingly, this proof differs a little from the one of homomorphic signature as here we have to consider O-SNARKs for *multiple* signing oracles.

Definition 11 (Functional Signatures [BGI14]). *A functional signature scheme FS for a message space \mathcal{M} and function family $\mathcal{F} = \{f : \mathcal{D}_f \rightarrow \mathcal{M}\}$ is a tuple of probabilistic, polynomial-time algorithms (FS.Setup, FS.KeyGen, FS.Sign, FS.Ver) that work as follows*

$\text{FS.Setup}(1^\lambda)$ takes a security parameter λ and outputs a master verification key mvk and a master secret key msk .

FS.KeyGen(msk, f) takes the master secret key msk and a function $f \in \mathcal{F}$ (represented as a circuit) and it outputs a signing key sk_f for f .

FS.Sign(mvk, f, sk_f, m) takes as input a function $f \in \mathcal{F}$, a signing key sk_f , and a message $m \in \mathcal{D}_f$, and it outputs $(f(m), \sigma)$ where σ represents a signature on $f(m)$.

FS.Ver(mvk, m^*, σ) takes as input the master verification key mvk, a message $m^* \in \mathcal{M}$ and a signature σ , and outputs either 1 (accept) or 0 (reject).

and satisfy correctness, unforgeability, and function privacy as described below.

- **Correctness.** A functional signature scheme is correct if the following holds with probability 1:

$$\forall f \in \mathcal{F}, \forall m \in \mathcal{D}_f, (\text{msk}, \text{mvk}) \leftarrow \text{FS.Setup}(1^\lambda), \text{sk}_f \leftarrow \text{FS.KeyGen}(\text{msk}, f),$$

$$(m^*, \sigma) \leftarrow \text{FS.Sign}(\text{mvk}, f, \text{sk}_f, m), \text{FS.Ver}(\text{mvk}, m^*, \sigma) = 1$$

- **Unforgeability.** A functional signature scheme is unforgeable if for every PPT adversary \mathcal{A} there is a negligible function ϵ such that $\Pr[\text{Exp}_{\mathcal{A}, \text{FS}}^{\text{FS-UF}}(\lambda) = 1] \leq \epsilon(\lambda)$ where the experiment $\text{Exp}_{\mathcal{A}, \text{FS}}^{\text{FS-UF}}(\lambda)$ is described in the following:

Key generation: Generate $(\text{msk}, \text{mvk}) \leftarrow \text{FS.Setup}(1^\lambda)$, and gives mvk to \mathcal{A} .

Queries: The adversary is allowed to adaptively query a key generation oracle \mathcal{O}_{key} and a signing oracle $\mathcal{O}_{\text{sign}}$, that share a dictionary D indexed by tuples $(f, i) \in \mathcal{F} \times \mathbb{N}$, whose entries are signing keys. For answering these queries, the challenger proceeds as follows:

- $\mathcal{O}_{\text{key}}(f, i)$:
 - * If $(f, i) \in D$ (i.e., the adversary had already queried the tuple (f, i)), then the challenger replies with the same key sk_f^i generated before.
 - * Otherwise, generate a new $\text{sk}_f^i \leftarrow \text{FS.KeyGen}(\text{msk}, f)$, add the entry $(f, i) \rightarrow \text{sk}_f^i$ in D , and return sk_f^i .
- $\mathcal{O}_{\text{sign}}(f, i, m)$:
 - * If there is an entry for the key (f, i) in D , then the challenger generates a signature on $f(m)$ using this key, i.e., $\sigma \leftarrow \text{FS.Sign}(\text{mvk}, f, \text{sk}_f^i, m)$.
 - * Otherwise, generate a new key $\text{sk}_f^i \leftarrow \text{FS.KeyGen}(\text{msk}, f)$, add an entry $(f, i) \rightarrow \text{sk}_f^i$ to D , and generate a signature on $f(m)$ using this key, i.e., $\sigma \leftarrow \text{FS.Sign}(\text{mvk}, f, \text{sk}_f^i, m)$.

Forgery: After the adversary is done with its queries, it outputs a pair (m^*, σ) , and the experiment outputs 1 iff the following conditions hold

- * $\text{FS.Ver}(\text{mvk}, m^*, \sigma) = 1$.
- * there does not exist m such that $m^* = f(m)$ for any f which was sent as a query to the \mathcal{O}_{key} oracle.
- * there does not exist a pair (f, m) such that (f, m) was a query to the $\mathcal{O}_{\text{sign}}$ oracle and $m^* = f(m)$.

- **Function privacy.** Intuitively, function privacy requires that the distribution of signatures on a message m that are generated via different keys sk_f should be computationally indistinguishable, even given the secret keys and master signing key. More formally, a functional signature scheme has function privacy if for every PPT adversary \mathcal{A} there is a negligible function ν such that $\Pr[\text{Exp}_{\mathcal{A}, \text{FS}}^{\text{FS-FPri}}(\lambda) = 1] \leq \nu(\lambda)$ where experiment $\text{Exp}_{\mathcal{A}, \text{FS}}^{\text{FS-FPri}}(\lambda)$ works as follows:

- The challenger generates a key pair $(\text{mvk}, \text{msk}) \leftarrow \text{FS.Setup}(1^\lambda)$ and gives (mvk, msk) to \mathcal{A} .
- The adversary chooses a function f_0 and receives an (honestly generated) secret key $\text{sk}_{f_0} \leftarrow \text{FS.KeyGen}(\text{msk}, f_0)$.

- The adversary chooses a second function f_1 such that $|f_0| = |f_1|$ (where padding can be used if there is a known upper bound) and receives an (honestly generated) secret key $\text{sk}_{f_1} \leftarrow \text{FS.KeyGen}(\text{msk}, f_1)$.
- The adversary chooses a pair of values (m_0, m_1) such that $|m_0| = |m_1|$ and $f_0(m_0) = f_1(m_1)$.
- The challenger selects a random bit $b \leftarrow \{0, 1\}$ and computes a signature on the image message $m^* = f_0(m_0) = f_1(m_1)$ using secret key sk_{f_b} , and gives the resulting signature $\sigma \leftarrow \text{FS.Sign}(\text{sk}_{f_b}, m_b)$ to \mathcal{A} .
- The adversary outputs a bit b' , and the experiment outputs 1 iff $b' = b$.

Definition 12 (Succinct Functional Signatures). A functional signature scheme is called succinct if there exists a polynomial $s(\cdot)$ such that, for every security parameter $\lambda \in \mathbb{N}$, $f \in \mathcal{F}$, $m \in \mathcal{D}_f$, it holds with probability 1 over $(\text{mvk}, \text{msk}) \leftarrow \text{FS.Setup}(1^\lambda)$, $\text{sk}_f \leftarrow \text{FS.KeyGen}(\text{msk}, f)$, $(f(m), \sigma) \leftarrow \text{FS.Sign}(\text{sk}_f, m)$ that $|\sigma| \leq s(\lambda, |f(m)|)$. In particular, the size of σ is independent of the function's size, $|f|$, and the function's input size, $|m|$.

Succinct Functional Signatures from O-SNARKs. In the following we show a construction for message space \mathcal{M} and family of functions $\mathcal{F} = \{f : \mathcal{D}_f \rightarrow \mathcal{M}\}$ whose running time is bounded by some fixed polynomial $t_{\mathcal{F}}(|m|)$. To build the scheme, we use two UF-CMA-secure signature schemes, $\Sigma_0 = (\text{kg}_0, \text{sign}_0, \text{vfy}_0)$ for message space \mathcal{M}_0 and $\Sigma' = (\text{kg}', \text{sign}', \text{vfy}')$ for message space \mathcal{D} , together with a fully succinct zero-knowledge O-SNARK $\Pi = (\text{Gen}, \text{Prove}, \text{Ver})$ for the NP language L defined below. While in [BGI14] a single signature scheme is used, we prefer to use two different ones as this allows for a more precise statement since we will need to apply different restrictions to \mathcal{M}_0 and \mathcal{D} to obtain a precise proof.

DEFINING THE RELATION \mathcal{R}_L . Let M_L be a random-access machine as defined below, and $t_L(k) = k^{e_L}$ be a polynomial. \mathcal{R}_L is the binary relation consisting of all pairs (y, w) such that, parsing $y = (M_L, x, t)$, $M_L(x, w)$ accepts in at most t steps and $t \leq t_L(|x|)$. The values x are of the form $x = (m^*, \text{mvk}_0)$ where mvk_0 is a public key of the scheme Σ_0 , and $m^* \in \mathcal{M}$ is a message. The values w are instead tuples $w = (m, f, \text{vk}', \sigma_{\text{vk}'}, \sigma_m)$ such that $m \in \mathcal{D}_f$ with $\mathcal{D}_f \subset \mathcal{D}$, and $\sigma_{\text{vk}'}, \sigma_m$ are signatures for the schemes Σ_0 and Σ' respectively. On input such a pair (x, w) , $M_L(x, w)$ is the random-access machine that accepts iff the following conditions (1)&(2)&(3) hold:

- (1) $m^* = f(m)$
- (2) $\text{vfy}'(\text{vk}', m, \sigma_m) = 1$
- (3) $\text{vfy}_0(\text{mvk}_0, f|\text{vk}', \sigma_{\text{vk}'}) = 1$

Given polynomial bounds on the running times of verification algorithms vfy' and vfy_0 , and a (fixed) bound $t_{\mathcal{F}}(\cdot)$ on the size and running time of every $f \in \mathcal{F}$, one can deduce a polynomial time bound $t_L(|x|) = |x|^{e_L}$ for the machine M_L .

THE CONSTRUCTION. Using the signature schemes Σ_0, Σ' and a fully-succinct zero-knowledge O-SNARK Π for NP, we construct the functional signature scheme $\text{FS}[\Sigma_0, \Sigma', \Pi] = (\text{FS.Setup}, \text{FS.KeyGen}, \text{FS.Sign}, \text{FS.Ver})$ as follows:

$\text{FS.Setup}(1^\lambda)$:

Generate a pair of keys for Σ_0 : $(\text{msk}_0, \text{mvk}_0) \leftarrow \text{kg}_0(1^\lambda)$.

Generate a crs for Π : $\text{crs} \leftarrow \text{Gen}(1^\lambda)$.

Set the master secret key $\text{msk} = \text{msk}_0$, and the master verification key $\text{mvk} = (\text{mvk}_0, \text{crs})$.

FS.KeyGen(msk, f) :

Generate a new key pair $(\text{sk}', \text{vk}') \leftarrow \text{kg}'(1^\lambda)$ for the scheme Σ' .
 Compute $\sigma_{\text{vk}'} \leftarrow \text{sign}_0(\text{msk}_0, f|\text{vk}')$, and let the certificate c be $c = (f, \text{vk}', \sigma_{\text{vk}'})$.
 Output $\text{sk}_f = (\text{sk}', c)$.

FS.Sign($\text{mvk}, f, \text{sk}_f, m$) :

Parse sk_f as $(\text{sk}', c = (f, \text{vk}', \sigma_{\text{vk}'}))$.
 Sign m using sk' in Σ' : $\sigma_m \leftarrow \text{sign}'(\text{sk}', m)$.
 Set $y = (M_L, x, t)$ with $x = (\text{mvk}_0, f(m))$, $t = |x|^{e_L}$, and $w = (m, f, \text{vk}', \sigma_{\text{vk}'}, \sigma_m)$.
 Run $\pi \leftarrow \text{Prove}(\text{crs}, y, w)$ and output $(m^* = f(m), \pi)$.

FS.Ver(mvk, m^*, π) :

Parse $\text{mvk} = (\text{mvk}_0, \text{crs})$ and set $y = (M_L, x, t)$ where $x = (\text{mvk}_0, m^*)$ and $t = |x|^{e_L}$.
 Output the same bit returned by $\text{Ver}_{e_L}(\text{crs}, y, \pi)$.

Correctness. It is not hard to see that as long as Σ_0, Σ' and Π are correct, then FS is also correct.

Succinctness. This property immediately follows from the succinctness of Π .

Unforgeability. We prove the security of FS under the unforgeability of schemes Σ_0 and Σ' and using the notion of O-SNARKs for a specific family of oracles $\mathbb{O}_{\text{m}\Sigma, Q}$ that we define below.

$\mathbb{O}_{\text{m}\Sigma, Q}$ is parametrized by the algorithms of the signature schemes Σ_0, Σ' and by a polynomial $Q = Q(\lambda)$. Every member \mathcal{O} of $\mathbb{O}_{\text{m}\Sigma, Q}$ is described by a set of secret keys $\text{msk}_0, \text{sk}'_1, \dots, \text{sk}'_Q$ (i.e., the process of sampling $\mathcal{O} \leftarrow \mathbb{O}$ consists of running $(\text{mvk}_0, \text{msk}_0) \xleftarrow{\$} \text{kg}_0(1^\lambda)$ and $(\text{vk}'_i, \text{sk}'_i) \xleftarrow{\$} \text{kg}'_1(1^\lambda), \forall i \in [Q]$). The oracle \mathcal{O} works as follows:

$$\mathcal{O}(i, 'vk') = \begin{cases} \text{mvk}_0 & \text{If } i = 0, \\ \text{vk}'_i & \text{otherwise.} \end{cases} \quad \mathcal{O}(i, 'sk') = \begin{cases} \perp & \text{If } i = 0, \\ \text{sk}'_i & \text{otherwise.} \end{cases}$$

$$\mathcal{O}(i, m) = \begin{cases} (\text{Cnt}, \text{sign}_0(\text{msk}_0, m|\text{vk}'_{\text{Cnt}})), \text{Cnt} \leftarrow \text{Cnt} + 1 & \text{If } i = 0 \text{ and } \text{Cnt} \leq Q, \\ \perp & \text{If } i = 0 \text{ and } \text{Cnt} > Q, \\ \text{sign}'(\text{sk}'_i, m) & \text{otherwise.} \end{cases}$$

For the sake of simplicity we compactly denote $\mathcal{O}_0(\cdot) = \mathcal{O}(0, \cdot)$ and $\mathcal{O}'_i(\cdot) = \mathcal{O}(i, \cdot)$ for all $i > 0$. From the above description, note that oracle \mathcal{O}_0 is stateful and we assume it starts with $\text{Cnt} = 1$.

Finally, we point out that for some technical reasons that we mention in Remark 5 at the end of this section, it is not possible to use the notion of O-SNARK for a *single* signing oracle to prove the security of the functional signature scheme. This is the reason why we explicitly considered O-SNARKs for this more complex family of multiple signing oracles.

Theorem 10. *If Π is an O-SNARK for $\mathbb{O}_{\text{m}\Sigma, Q}$ for every $Q = \text{poly}(\lambda)$, and Σ_0, Σ' are UF-CMA-secure, then $\text{FS}[\Sigma_0, \Sigma', \Pi]$ is an unforgeable functional signature.*

Proof. Our proof consists of the following steps:

1. We show that for every successful \mathcal{A}_{FS} against the unforgeability of FS there exists an O-SNARK adversary $\tilde{\mathcal{A}}$ for an oracle from $\mathbb{O}_{\text{m}\Sigma, Q}$ such that $\tilde{\mathcal{A}}$ outputs a valid proof with the same (non-negligible) probability of success of \mathcal{A}_{FS} . By the adaptive proof of knowledge for $\mathbb{O}_{\text{m}\Sigma, Q}$ we then

obtain that for such $\tilde{\mathcal{A}}$ there exists a suitable extractor $\mathcal{E}_{\tilde{\mathcal{A}}}$ that outputs a valid witness with all but negligible probability.

2. From the previous point, considering adversary $\tilde{\mathcal{A}}$ and the corresponding extractor, we can partition adversary-extractor pairs in two types: (1) those that yield a witness w containing a pair (f, vk'_f) that was never signed before, and (2) those that yield w containing (f, vk'_f) that was signed before. We show that adversaries of type (1) can be used to break the security of the signature scheme Σ_0 , whereas adversaries of type (2) can be used to break the security of Σ' .

EXISTENCE OF AN EXTRACTOR FOR \mathcal{A}_{FS} . Consider any adversary \mathcal{A}_{FS} that while running in $\mathbf{Exp}_{\mathcal{A}_{\text{FS}}, \text{FS}}^{\text{FS-UF}}$ it outputs (m^*, π^*) and makes the experiment generate Q secret keys of the scheme Σ' . For every such \mathcal{A}_{FS} we show there exists another adversary $\tilde{\mathcal{A}}^{\mathcal{O}}$ that, on input crs , and given oracle $\mathcal{O} \leftarrow \mathbb{O}_{\text{m}\Sigma, Q}$, outputs a pair (y, π^*) . We describe $\tilde{\mathcal{A}}^{\mathcal{O}}$ below. During its execution it maintains a dictionary D similar to the one in the definition of $\mathbf{Exp}_{\mathcal{A}_{\text{FS}}, \text{FS}}^{\text{FS-UF}}$, except that instead of storing mappings like $(f, i) \rightarrow \text{sk}'_f$, it maps a pair (f, i) to a triple (j, sk'_j, c) where $j \in [Q]$. Intuitively, this means that a queried pair (f, i) is associated to oracle \mathcal{O}'_j .

$\tilde{\mathcal{A}}^{\mathcal{O}}(\text{crs}) :$

Query $\text{mvk}_0 \leftarrow \mathcal{O}_0('vk')$ and run $\mathcal{A}_{\text{FS}}^{\mathcal{O}_{\text{key}}, \mathcal{O}_{\text{sign}}}(\text{mvk} = (\text{crs}, \text{mvk}_0))$

Simulate queries (f, i) to \mathcal{O}_{key} as follows:

if $[(f, i) \rightarrow (j, \text{sk}'_j, c)] \in D$: output $\text{sk}'_f = (\text{sk}'_j, c)$

else if $[(f, i) \rightarrow (j, \cdot, c)] \in D$:

ask $\text{sk}'_j \leftarrow \mathcal{O}'_j('sk')$ and output $\text{sk}'_f = (\text{sk}'_j, c)$

else:

ask $(j, \sigma_{\text{vk}'_j}) \leftarrow \mathcal{O}_0(f)$, $\text{vk}'_j \leftarrow \mathcal{O}'_j('vk')$, $\text{sk}'_j \leftarrow \mathcal{O}'_j('sk')$

add $(f, i) \rightarrow (j, \text{sk}'_j, c)$ to D with $c = (f, \text{vk}'_j, \sigma_{\text{vk}'_j})$

output $\text{sk}'_f = (\text{sk}'_j, c)$

Simulate queries (f, i, m) to $\mathcal{O}_{\text{sign}}$ as follows:

if (f, i) not assigned in D :

ask $(j, \sigma_{\text{vk}'_j}) \leftarrow \mathcal{O}_0(f)$, $\text{vk}'_j \leftarrow \mathcal{O}'_j('vk')$

add $(f, i) \rightarrow (j, \cdot, c)$ to D with $c = (f, \text{vk}'_j, \sigma_{\text{vk}'_j})$

ask $\sigma_m \leftarrow \mathcal{O}'_j(m)$

set $x = (\text{mvk}_0, f(m))$, $t = |x|^{e_L}$, $w = (m, f, \text{vk}'_j, \sigma_{\text{vk}'_j}, \sigma_m)$

run $\pi \leftarrow \text{Prove}(\text{prs}, (M_L, x, t), w)$

output $(f(m), \pi)$

if $[(f, i) \rightarrow (j, \cdot, c)] \in D$: parse $c = (f, \text{vk}'_j, \sigma_{\text{vk}'_j})$

ask $\sigma_m \leftarrow \mathcal{O}'_j(m)$

set $x = (\text{mvk}_0, f(m))$, $t = |x|^{e_L}$, $w = (m, f, \text{vk}'_j, \sigma_{\text{vk}'_j}, \sigma_m)$

run $\pi \leftarrow \text{Prove}(\text{prs}, (M_L, x, t), w)$

output $(f(m), \pi)$

When \mathcal{A}_{FS} outputs (m^*, π^*)

set $y = (M_L, x = (\text{mvk}_0, m^*), t = |x|^{e_L})$

output (y, π^*)

As one can see, given the definition of oracles from $\mathbb{O}_{m\Sigma, Q}$, the simulation provided by $\tilde{\mathcal{A}}$ to \mathcal{A}_{FS} is perfect. So, whenever \mathcal{A}_{FS} outputs a valid forgery $\tilde{\mathcal{A}}$ outputs a pair (y, π^*) that verifies correctly. Moreover, defined in this way, the adversary $\tilde{\mathcal{A}}^{\mathcal{O}}$ fits the definition of adaptive proof of knowledge for $\mathbb{O}_{m\Sigma, Q}$ by which we know that there exists an extractor $\mathcal{E}_{\tilde{\mathcal{A}}}$ that, given the same input of $\tilde{\mathcal{A}}^{\mathcal{O}}$ and the transcript of oracle queries/answers made and received by $\tilde{\mathcal{A}}^{\mathcal{O}}$, outputs a witness w such that the probability that (y, π^*) verifies and $(y, w) \notin \mathcal{R}_L$ is negligible.

We define the following hybrid games that involve running $\tilde{\mathcal{A}}, \mathcal{E}_{\tilde{\mathcal{A}}}$:

G_1 is the same experiment as $\text{O-AdPoK}(\lambda, \tilde{\mathcal{A}}, \mathcal{E}_{\tilde{\mathcal{A}}}, \mathbb{O}_{m\Sigma, Q})$ except that its outcome is defined differently. G_1 outputs 1 iff $\text{Ver}(\text{crs}, y, \pi) = 1$ and the value m^* inside y constitutes a forgery according to the oracle queries made by $\tilde{\mathcal{A}}$ during the game.

By construction of $\tilde{\mathcal{A}}$ from \mathcal{A}_{FS} it holds

$$\Pr[\mathbf{Exp}_{\mathcal{A}_{\text{FS}}, \text{FS}}^{\text{FS-UF}}(\lambda) = 1] = \Pr[\text{G}_1 \Rightarrow 1] \quad (1)$$

G_2 is the same as G_1 except that in order to output 1 it additionally checks that $(y, w) \in \mathcal{R}_L$.

Essentially, the outcome of G_2 differs from that of G_1 only in G_2 $(y, w) \notin \mathcal{R}_L$. Hence,

$$\Pr[\text{G}_1 \Rightarrow 1] - \Pr[\text{G}_2 \Rightarrow 1] \leq \Pr[\text{O-AdPoK}(\lambda, \tilde{\mathcal{A}}, \mathcal{E}_{\tilde{\mathcal{A}}}, \mathbb{O}_{m\Sigma, Q}) \Rightarrow 1] \quad (2)$$

Moreover, let us define the following two events in game G_2 .

Let $w = (m, f, vk', \sigma_{vk'}, \sigma_m)$ be the witness returned by $\mathcal{E}_{\tilde{\mathcal{A}}}$:

Ev_1 occurs if $\forall j \in [Q] : vk' \neq vk'_j$, or $\exists j \in [Q] : vk' = vk'_j$ but $\tilde{\mathcal{A}}$ never made a query $\mathcal{O}_0(f)$ that returned (j, \cdot) ;

Ev_2 occurs if $vk' = vk'_j$ for some $j \in [Q]$ and $\tilde{\mathcal{A}}$ did make a query $(j, \sigma) \leftarrow \mathcal{O}_0(f)$.

Clearly it holds

$$\Pr[\text{G}_2 \Rightarrow 1] = \Pr[\text{G}_2 \Rightarrow 1 \wedge \text{Ev}_1] + \Pr[\text{G}_2 \Rightarrow 1 \wedge \text{Ev}_2] \quad (3)$$

In the remaining part of the proof we show that both $\Pr[\text{G}_2 \Rightarrow 1 \wedge \text{Ev}_1]$ and $\Pr[\text{G}_2 \Rightarrow 1 \wedge \text{Ev}_2]$ are negligible under the assumption that, respectively, Σ_0 and Σ' are unforgeable.

Claim 1 *For every efficient adversary \mathcal{A}_{FS} there is an efficient forger \mathcal{F}_0 such that $\Pr[\text{G}_2 \Rightarrow 1 \wedge \text{Ev}_1] = \mathbf{Adv}_{\mathcal{F}_0, \Sigma_0}^{\text{UF-CMA}}(\lambda)$.*

Proof. Let \mathcal{A}_{FS} be an adversary that runs in $\mathbf{Exp}_{\mathcal{A}_{\text{FS}}, \text{FS}}^{\text{FS-UF}}(\lambda)$, and let $\tilde{\mathcal{A}}, \mathcal{E}_{\tilde{\mathcal{A}}}$ be the pair of algorithms built out of \mathcal{A}_{FS} as defined before. Below we show how to build an efficient forger \mathcal{F}_0 out of $\tilde{\mathcal{A}}, \mathcal{E}_{\tilde{\mathcal{A}}}$ so that its probability of forging against Σ_0 is at least $\Pr[\text{G}_2 \Rightarrow 1 \wedge \text{Ev}_1]$. \mathcal{F}_0 gets the public key mvk_0 and has access to oracle $\mathcal{O}_{\Sigma_0} = \text{sign}_0(\text{msk}_0, \cdot)$.

$\mathcal{F}_0^{\mathcal{O}_{\Sigma_0}}(\text{mvk}_0)$:

Initialize $\text{qt} \leftarrow \emptyset, \mathcal{T} \leftarrow \emptyset, \text{Cnt} \leftarrow 1$

Generate $(\text{sk}'_i, \text{vk}'_i) \xleftarrow{\$} \text{kg}'(1^\lambda) \forall i \in [Q]$

Generate $\text{crs} = (\text{prs}, \text{vst}) \leftarrow \text{Gen}(1^\lambda)$ and run $\tilde{\mathcal{A}}^{\mathcal{O}}(\text{crs})$

Simulate all queries to \mathcal{O}'_i using $\text{sk}'_i, \text{vk}'_i$ and add all queries-answers to qt

Simulate queries $\mathcal{O}_0(m)$ as follows:

if $m = 'vk'$: output mvk_0

else:

ask $\sigma \leftarrow \mathcal{O}_{\Sigma_0}(m|\text{vk}'_{\text{Cnt}})$
 add $(m|\text{vk}'_j, \sigma)$ to qt , and add $m|\text{vk}'_j$ to \mathcal{T}
 increment $\text{Cnt} \leftarrow \text{Cnt} + 1$
 output σ

Let (y, π^*) be $\tilde{\mathcal{A}}$'s output

Run $w \leftarrow \mathcal{E}_{\tilde{\mathcal{A}}}(\text{crs}, \text{qt})$

Check that $(y, w) \in \mathcal{R}_L$:

[Fail] Abort if this does not hold.

Else parse $w = (m, f, \text{vk}', \sigma_{\text{vk}'}, \sigma_m)$ and proceed:

[A] If $(f|\text{vk}') \notin \mathcal{T}$ return $(f|\text{vk}', \sigma_{\text{vk}'})$.

[B] If $(f|\text{vk}') \in \mathcal{T}$ abort.

Algorithm \mathcal{F}_0 can perfectly simulate \mathbf{G}_2 to $\tilde{\mathcal{A}}$ and $\mathcal{E}_{\tilde{\mathcal{A}}}$. Furthermore, it is easy to see that if \mathbf{G}_2 outputs 1 and Ev_1 occurs, then \mathcal{F}_0 's simulation ends up exactly in case (A), that is \mathcal{F}_0 returns a signature $\sigma_{\text{vk}'}$ on a new message $f|\text{vk}'$. Since $(y, w) \in \mathcal{R}_L$ one has that $\sigma_{\text{vk}'}$ is valid, and thus is a forgery.

Finally, it is worth noting that for this simulation the adversary $\tilde{\mathcal{A}}$ can even ask $\mathcal{O}'_j('sk')$ for *all* j oracles without affecting our reduction. \square

Claim 2 For every efficient adversary \mathcal{A}_{FS} there is an efficient forger \mathcal{F}' such that $\Pr[\mathbf{G}_2 \Rightarrow 1 \wedge \text{Ev}_2] \leq Q \cdot \mathbf{Adv}_{\mathcal{F}', \Sigma'}^{\text{UF-CMA}}(\lambda)$.

Proof. Let \mathcal{A}_{FS} be an adversary that runs in $\mathbf{Exp}_{\mathcal{A}_{\text{FS}}, \text{FS}}^{\text{FS-UF}}(\lambda)$, and let $\tilde{\mathcal{A}}, \mathcal{E}_{\tilde{\mathcal{A}}}$ be the pair of algorithms built out of \mathcal{A}_{FS} as defined before. Below we show how to build an efficient forger \mathcal{F}' out of $\tilde{\mathcal{A}}, \mathcal{E}_{\tilde{\mathcal{A}}}$ so that its probability of forging against Σ' is at least $\Pr[\mathbf{G}_2 \Rightarrow 1 \wedge \text{Ev}_2]/Q$. \mathcal{F}' gets a public key vk' and has access to oracle $\mathcal{O}_{\Sigma'} = \text{sign}'(\text{sk}', \cdot)$.

$\mathcal{F}'^{\mathcal{O}'}$ (vk') :

Initialize $\text{qt} \leftarrow \emptyset, \mathcal{T} \leftarrow \emptyset, \text{Cnt} \leftarrow 0$

Generate $\text{crs} = (\text{prs}, \text{vst}) \leftarrow \text{Gen}(1^\lambda)$

Generate a pair $(\text{msk}_0, \text{mvk}_0) \leftarrow \text{kg}_0(1^\lambda)$

Choose a random $q \xleftarrow{\$} \{1, \dots, Q\}$

Generate $(\text{sk}'_i, \text{vk}'_i) \xleftarrow{\$} \text{kg}'(1^\lambda) \forall i \in [Q] \setminus \{q\}$

Run $\tilde{\mathcal{A}}^{\mathcal{O}}(\text{crs})$

Simulate all queries to \mathcal{O}_0 using $\text{mvk}_0, \text{msk}_0$:

add all queries-answers to qt and all signed messages $m|\text{vk}'$ to \mathcal{T}

Simulate all queries to \mathcal{O}'_i using $\text{vk}'_i, \text{sk}'_i$ for all $i \in [Q] \setminus \{q\}$

add all queries-answers to qt

Simulate queries $\mathcal{O}'_q(m)$ as follows:

if $m = 'vk'$: output vk' and add $('vk', \text{vk}')$ to qt

else if $m = 'sk'$: Abort

else: ask $\sigma \leftarrow \mathcal{O}_{\Sigma'}(m)$ and add (m, σ) to qt

output σ

Let (y, π^*) be $\tilde{\mathcal{A}}$'s output

Run $w \leftarrow \mathcal{E}_{\tilde{\mathcal{A}}}(\text{crs}, \text{qt})$

Check that $(y, w) \in \mathcal{R}_L$:

[Fail] Abort if this does not hold.

Else parse $w = (m, f, \text{vk}^*, \sigma_{\text{vk}^*}, \sigma_m)$ and proceed:

[A] If $(f|\text{vk}^*) \notin \mathcal{T}$ Abort.

[B] If $(f|\text{vk}^*) \in \mathcal{T}$ and $\text{vk}^* \neq \text{vk}'$ Abort.

[C] If $(f|\text{vk}^*) \in \mathcal{T}$ and $\text{vk}^* = \text{vk}'$ return (m, σ_m) .

As one can see, unless it aborts, algorithm \mathcal{F}' can perfectly simulate \mathbf{G}_2 to $\tilde{\mathcal{A}}$ and $\mathcal{E}_{\tilde{\mathcal{A}}}$. Furthermore, it is easy to see that if \mathbf{G}_2 outputs 1, Ev_2 occurs, and there is no abort while answering queries, then the simulation of \mathcal{F}' ends up in cases (B) or (C). However since $(f|\text{vk}^*) \in \mathcal{T}$, we have that $\text{vk}^* = \text{vk}'_j$ for some $j \in [Q]$ (where we let $\text{vk}'_q = \text{vk}'$). So, if there is no abort at all, we have that $\text{vk}^* = \text{vk}'$ and thus \mathcal{F}' returns a valid signature σ on a message m (recall that validity follows from $(y, w) \in \mathcal{R}_L$). By definition of \mathbf{G}_2 we also have that if it outputs 1, then the message m^* in y constitutes a forgery according to the definition of $\mathbf{Exp}^{\text{FS-UF}}$. In particular, it holds that for the given $f|\text{vk}'$, there was no signing query $\mathcal{O}'_q(m)$ such that $m^* = f(m)$. Therefore, if m is such that $m^* = f(m)$ (again this follows from $(y, w) \in \mathcal{R}_L$), then m cannot have been queried to \mathcal{O}'_q , i.e., \mathcal{F}' never queried m to its signing oracle. From this we have that, as long as \mathbf{G}_2 outputs 1, Ev_2 occurs and there is no abort, then \mathcal{F}' outputs a valid forgery. To conclude the proof, we observe that \mathcal{F}' does not abort with probability $1/Q$ which is the probability that the guess of q , for which $\text{vk}^* = \text{vk}'$, is correct. Therefore, we have that $\mathbf{Adv}_{\mathcal{F}', \Sigma'}^{\text{UF-CMA}}(\lambda) = \Pr[\mathbf{G}_2 \Rightarrow 1 \wedge \text{Ev}_2]/Q$.

Finally, we note that the above proof works even if the adversary $\tilde{\mathcal{A}}$ queries $\mathcal{O}'_j('sk')$ on all oracles but the q -th one. This observation will be useful when we discuss the existence of O-SNARKs for this oracle family. \square

Putting together the bounds in equations (1), (2) and (3), with the results of Claims 1 and 2, eventually we obtain:

$$\Pr[\mathbf{Exp}_{\mathcal{A}_{\text{FS}}, \text{FS}}^{\text{FS-UF}}(\lambda) = 1] \leq \Pr[\text{O-AdPoK}(\lambda, \tilde{\mathcal{A}}, \mathcal{E}_{\tilde{\mathcal{A}}}, \mathbb{O}_{\text{m}\Sigma, Q}) \Rightarrow 1] + \mathbf{Adv}_{\mathcal{F}_0, \Sigma_0}^{\text{UF-CMA}}(\lambda) + \mathbf{Adv}_{\mathcal{F}', \Sigma'}^{\text{UF-CMA}}(\lambda)$$

which shows that any efficient adversary has at most negligible probability of breaking the security of scheme FS under the assumption that Π is an O-SNARK for $\mathbb{O}_{\text{m}\Sigma, Q}$ and the schemes Σ_0, Σ' are unforgeable. \square

Non-adaptive unforgeability. Similarly to the homomorphic signature case, it is possible to show that the functional signature scheme achieves security against (functional signature) adversaries that make *non-adaptive* signing queries (i.e., all queries are declared at the beginning of the game). This weaker security can be proven assuming that Π is a *non-adaptive* O-SNARK (see Definition 8). Combining this change with the result of Theorem 2 we obtain the following:

Theorem 11. *If Π is a SNARK and Σ_0, Σ' are UF-CMA-secure signature schemes, then $\text{FS}[\Sigma_0, \Sigma', \Pi]$ is a functional signature where unforgeability holds against adversaries that make non-adaptive signing queries.*

Proof. The proof of the theorem can be obtained via straightforward modifications to the proof of Theorem 10. Having in mind the intuition provided earlier, the main idea is that to work with non-adaptive adversaries, one can define a non-adaptive O-SNARK adversary $\tilde{\mathcal{A}}$ for every non-adaptive functional signature adversary \mathcal{A} . In particular, the non-adaptive queries of \mathcal{A} can be used to define the non-adaptive queries of $\tilde{\mathcal{A}}$. The rest of the proof proceeds analogously.

Function privacy. We show that the functional signature construction satisfies function privacy provided that the O-SNARK is zero-knowledge.

Theorem 12. *If Π is a zero-knowledge O-SNARK then FS satisfies function privacy.*

Proof. We show that for every adversary $\mathcal{A}_{\text{priv}}$ against the function privacy experiment $\mathbf{Exp}_{\mathcal{A}_{\text{priv}}, \text{FS}}^{\text{FS-FPri}}(\lambda)$, we can construct a distinguisher algorithm \mathcal{D} against the zero knowledge property of Π .

Consider the following two hybrid experiments:

\mathbf{G}_0 is the same as $\mathbf{Exp}_{\mathcal{A}_{\text{priv}}, \text{FS}}^{\text{FS-FPri}}(\lambda)$. In particular, the crs for Π is generated honestly using Gen ; and the challenge functional signature $\sigma = (f_b(m_b))$ is generated as $\sigma \leftarrow \text{FS.Sign}(\text{sk}_{f_b}, m_b)$, i.e., by running $\pi \leftarrow \text{Prove}(\text{prs}, (M_L, x, t), w)$ where $x = (\text{mvk}_0, f_b(m_b))$, $t = |x|^{e_L}$, and $w = (m_b, f_b, \text{vk}', \sigma_{\text{vk}'}, \sigma)$.

\mathbf{G}_1 is the same as \mathbf{G}_0 except that one uses the zero-knowledge simulator algorithm S in order to generate both the crs and the proof in the challenge. Namely, $(\text{prs}, \text{vst}, \text{tr}) \leftarrow S^{\text{crs}}(1^\lambda)$, and the challenge signature is generated by running $\pi \leftarrow S^{\text{Prove}}(\text{aux}, \text{prs}, (M_L, x, t), \text{tr})$ for $x = (\text{mvk}_0, m')$, $t = |x|^{e_L}$, where $m' = f_0(m_0) = f_1(m_1)$.

Denote by win_0 and win_1 the advantage of the adversary $\mathcal{A}_{\text{priv}}$ in guessing the bit b in \mathbf{G}_0 , and \mathbf{G}_1 , respectively. Clearly $\text{win}_1 = 1/2$ since the bit b is not used at all, and thus the view of $\mathcal{A}_{\text{priv}}$ is independent of b . To complete the proof we show that under the assumption that Π is zero-knowledge, the following holds:

Claim 3 $\text{win}_0 - \text{win}_1 \leq \text{negl}(\lambda)$.

To prove this, we show that for any $\mathcal{A}_{\text{priv}}$ such that $\text{win}_0 - \text{win}_1 = \epsilon$ is non-negligible there is a distinguisher \mathcal{D} that succeeds against the zero-knowledge property of Π with the same advantage ϵ . \mathcal{D} is defined as follows:

$\mathcal{D}(\text{crs}) :$

Generate a pair $(\text{msk}_0, \text{mvk}_0) \leftarrow \text{kg}_0(1^\lambda)$

Run $\mathcal{A}_{\text{priv}}(\text{mvk} = (\text{crs}, \text{mvk}_0))$

$\mathcal{A}_{\text{priv}}$ adaptively chooses function queries f_0, f_1 and message pairs m_0, m_1 such that $f_0(m_0) = f_1(m_1)$:

For each f_b asked by $\mathcal{A}_{\text{priv}}$, return the secret key $\text{sk}_{f_b} \leftarrow \text{FS.KeyGen}(\text{msk}, f_b)$

To answer the challenge \mathcal{D} proceeds as follows:

pick $b \xleftarrow{\$} \{0, 1\}$

set $x = (\text{mvk}_0, f_b(m_b))$, $t = |x|^{e_L}$, and $w = (m_b, f_b, \text{vk}', \sigma_{\text{vk}'}, \sigma)$

output (y, w) (where $y = (M_L, x, t)$) to its challenger, and get back π

return $(f_b(m_b), \pi)$ to $\mathcal{A}_{\text{priv}}$

Let b' be $\mathcal{A}_{\text{priv}}$'s output

If $b' = b$ output 1, else output 0.

Note that when \mathcal{D} receives crs and π that are generated using the real algorithms, then \mathcal{D} is perfectly simulating \mathbf{G}_0 to $\mathcal{A}_{\text{priv}}$. Otherwise, if \mathcal{D} receives crs and π that are generated using the simulator, then \mathcal{D} perfectly simulates \mathbf{G}_1 . Therefore it is easy to see that \mathcal{D} 's advantage is $\text{win}_0 - \text{win}_1$.

Remark 4 (On the applicability of Corollary 1). For the same reasons discussed in Remark 3, it is not possible to apply the result of Corollary 1 to conclude that the (adaptive) security of the functional signature scheme holds under classical SNARKs.

Remark 5. [On the use of multiple signing oracles] In order to prove the security of the functional signature scheme, one might be tempted to use the notion of O-SNARK with a single signing oracle. Precisely, one might use O-SNARKs for \mathbb{O}_{Σ_0} when making a reduction to Σ_0 and O-SNARKs for $\mathbb{O}_{\Sigma'}$ when making a reduction to Σ' . Unfortunately, this approach does not work for an intricate technical reason that we explain here. Intuitively, assume that one wants to build an O-SNARK adversary $\tilde{\mathcal{A}}$ that has access to a single signing oracle, say from \mathbb{O}_{Σ_0} . Then the secret keys needed to simulate all the other oracles have to be given to $\tilde{\mathcal{A}}$ as part of its auxiliary input ($\tilde{\mathcal{A}}$ needs them to simulate \mathcal{A}_{FS}). At this point the issue is that such secret keys in fact give an efficient way to compute a witness for several y in the relation \mathcal{R}_L . Therefore, if the extractor gets these secret keys as auxiliary information, we then have no guarantee that, while doing a reduction to the unforgeability of the signature scheme, the extractor will output a witness of the form we expect.

5.3 SNARKs on authenticated data

As another application of O-SNARKs we consider the generic construction of SNARKs on authenticated data that is given in [BBFR15]. Since this construction is very similar to the homomorphic signature scheme that we present in Section 5.1, we only provide an informal discussion of this application. In [BBFR15] Backes et al. introduce the notion of SNARKs on authenticated data to capture in an explicit way the possibility of performing (zero-knowledge) proofs about statements that are authenticated by third parties, i.e., to prove that $(x, w) \in \mathcal{R}$ for some x for which there is a valid signature. While the main focus of that work is on a concrete construction based on quadratic arithmetic programs, the authors also show a generic construction based on SNARKs and digital signatures. Roughly speaking, this construction consists in letting the prover use a SNARK to prove a statement of the form “ $\exists x, w, \sigma : (x, w) \in \mathcal{R} \wedge \text{vfy}(\text{vk}, \tau | x, \sigma) = 1$ ”, for some public label τ of the statement. The formalization of their model is rather similar to that of homomorphic signatures in this paper (e.g., they also use labels). Noticeable differences are that their construction uses pre-processing SNARKs for arithmetic circuit satisfiability, and that to handle several functions they use different SNARK instantiations (one per function).

In [BBFR15] the security proof of this generic construction is only sketched, and in particular they use the existence of an extractor for an adversary that interacts with a signing oracle without providing a particular justification on its existence. With a more careful look, it is possible to see that this security proof incurs the same issue of extraction in the presence of oracles. Using the same techniques that we developed in this paper for the homomorphic signature scheme,¹⁵ it is possible to prove the security of that generic construction using O-SNARKs for signing oracles (or non-adaptive security based on classical SNARKs). In conclusion, for this construction one can either conjecture that a specific SNARK scheme (e.g., [PHGR13]) is secure in the presence of oracles, or, more conservatively, argue only the non-adaptive security of the primitive under the existence of classical SNARKs.

¹⁵ The only major difference is that one has to consider a specification of our definitions to the case of pre-processing SNARKs.

Acknowledgements. We would like to thank Manuel Barbosa and Bogdan Warinschi for valuable discussions on this work, and the anonymous reviewers of Crypto 2016 and TCC 2016-B for their useful comments and suggestions. This work was partially supported by the European Union’s Horizon 2020 Research and Innovation Programme under grant agreement 688722 (NEXTLEAP), the Spanish Ministry of Economy under project reference TIN2015-70713-R (DEDETIS) and a Juan de la Cierva fellowship to Dario Fiore, by the Madrid Regional Government under project N-Greens (ref. S2013/ICE-2731), and by the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 339563 CryptoCloud).

References

- BBFR15. Michael Backes, Manuel Barbosa, Dario Fiore, and Raphael M. Reischuk. ADSNARK: Nearly practical and privacy-preserving proofs on authenticated data. In *2015 IEEE Symposium on Security and Privacy*, pages 271–286. IEEE Computer Society Press, 2015.
- BCC88. Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.*, 37(2):156–189, October 1988.
- BCC⁺14. Nir Bitansky, Ran Canetti, Alessandro Chiesa, Shafi Goldwasser, Huijia Lin, Aviad Rubinfeld, and Eran Tromer. The hunting of the SNARK. Cryptology ePrint Archive, Report 2014/580, 2014. <http://eprint.iacr.org/2014/580>.
- BCCT12. Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In Shafi Goldwasser, editor, *ITCS 2012*, pages 326–349. ACM, January 2012.
- BCCT13. Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKS and proof-carrying data. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 111–120. ACM Press, June 2013.
- BCI⁺13. Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. Succinct non-interactive arguments via linear interactive proofs. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 315–333. Springer, March 2013.
- BCPR14. Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In David B. Shmoys, editor, *46th ACM STOC*, pages 505–514. ACM Press, May / June 2014.
- BCTV14. Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 276–294. Springer, August 2014.
- BF11. Dan Boneh and David Mandell Freeman. Homomorphic signatures for polynomial functions. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 149–168. Springer, May 2011.
- BG08. Boaz Barak and Oded Goldreich. Universal arguments and their applications. *SIAM Journal on Computing*, 38(5):1661–1694, 2008.
- BGI14. Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, March 2014.
- BHZ87. Ravi B. Boppana, Johan Hastad, and Stathis Zachos. Does co-NP have short interactive proofs? *Information Processing Letters*, 25(2):127–132, 1987.
- BP15. Elette Boyle and Rafael Pass. Limits of extractability assumptions with distributional auxiliary input. ASIACRYPT, 2015.
- BSCG⁺13. Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108. Springer, August 2013.
- CF13. Dario Catalano and Dario Fiore. Practical homomorphic MACs for arithmetic circuits. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 336–352. Springer, May 2013.
- CFW14. Dario Catalano, Dario Fiore, and Bogdan Warinschi. Homomorphic signatures with efficient verification for polynomial functions. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 371–389. Springer, August 2014.

- CL08. Giovanni Di Crescenzo and Helger Lipmaa. Succinct NP proofs from an extractability assumption. In *Logic and Theory of Algorithms, 4th Conference on Computability in Europe, CiE 2008, Athens, Greece, June 15-20, 2008, Proceedings*, pages 175–185, 2008.
- DLFKP16. Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, and Bryan Parno. Cinderella: Turning shabby x. 509 certificates into elegant anonymous credentials with the magic of verifiable computation. In *IEEE Symposium on Security & Privacy*, 2016.
- GGPR13. Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, May 2013.
- GH98. Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Information Processing Letters*, 67(4):205 – 214, 1998.
- GMR89. Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- Gro10. Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 321–340. Springer, December 2010.
- GVW02. Oded Goldreich, Salil Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *computational complexity*, 11(1-2):1–53, 2002.
- GVW15. Sergey Gorbunov, Vinod Vaikuntanathan, and Daniel Wichs. Leveled fully homomorphic signatures from standard lattices. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 469–477. ACM Press, June 2015.
- GW11. Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108. ACM Press, June 2011.
- GW13. Rosario Gennaro and Daniel Wichs. Fully homomorphic message authenticators. In Kazuo Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 301–320. Springer, December 2013.
- HT98. Satoshi Hada and Toshiaki Tanaka. On the existence of 3-round zero-knowledge protocols. In Hugo Krawczyk, editor, *CRYPTO’98*, volume 1462 of *LNCS*, pages 408–423. Springer, August 1998.
- Kil92. Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *24th ACM STOC*, pages 723–732. ACM Press, May 1992.
- Lam79. Leslie Lamport. Constructing digital signatures from a one-way function. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory, October 1979.
- Lip12. Helger Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 169–189. Springer, March 2012.
- Mic94. Silvio Micali. CS proofs (extended abstracts). In *35th FOCS*, pages 436–453. IEEE Computer Society Press, November 1994.
- Mic00. Silvio Micali. Computationally sound proofs. *SIAM Journal on Computing*, 30(4):1253–1298, 2000.
- Mie08. Thilo Mie. Polylogarithmic two-round argument systems. *J. Mathematical Cryptology*, 2(4):343–363, 2008.
- Nao03. Moni Naor. On cryptographic assumptions and challenges (invited talk). In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 96–109. Springer, August 2003.
- NT16. Assa Naveh and Eran Tromer. Photoproof: Cryptographic image authentication for any set of permissible transformations. In *IEEE Symposium on Security & Privacy*, 2016.
- NY89. Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In *21st ACM STOC*, pages 33–43. ACM Press, May 1989.
- PHGR13. Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013.
- Rom90. John Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd ACM STOC*, pages 387–394. ACM Press, May 1990.
- Val08. Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In Ran Canetti, editor, *TCC 2008*, volume 4948 of *LNCS*, pages 1–18. Springer, March 2008.
- Wee05. Hoeteck Wee. On round-efficient argument systems. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *ICALP 2005*, volume 3580 of *LNCS*, pages 140–152. Springer, July 2005.

A Additional Preliminaries

A.1 Digital Signatures

A digital signature scheme Σ consists of a triple of algorithms $\Sigma = (\text{kg}, \text{sign}, \text{vfy})$ working as follows:

$\text{kg}(1^\lambda)$ the key generation takes as input the security parameter λ and returns a pair of keys (sk, vk) .
 $\text{sign}(\text{sk}, m)$ on input a signing key sk and a message m , the signing algorithm produces a signature σ .

$\text{vfy}(\text{vk}, m, \sigma)$ given a triple vk, m, σ the verification algorithm tests if σ is a valid signature on m with respect to verification key vk .

The standard security notion for digital signatures, unforgeability against chosen-message attacks (UF-CMA, for short) is defined by the following experiment:

Experiment $\mathbf{Exp}_{\mathcal{F}, \Sigma}^{\text{UF-CMA}}(\lambda)$

$(\text{sk}, \text{vk}) \xleftarrow{\$} \text{kg}(1^\lambda)$

$(m^*, \sigma^*) \xleftarrow{\$} \mathcal{F}^{\text{sign}(\text{sk}, \cdot)}(\text{vk})$

If $\text{vfy}(\text{vk}, m^*, \sigma^*) = 1$ and m^* is “new” then output 1

Else output 0

A message $(m^*$ is said “new” if it is different from all the messages m_i that the adversary queried to the signing oracle $\text{sign}(\text{sk}, \cdot)$ during the experiment. The advantage of a forger \mathcal{F} in breaking the unforgeability against chosen-message attacks (SUF-CMA) of Σ is $\mathbf{Adv}_{\mathcal{F}, \Sigma}^{\text{SUF-CMA}}(\lambda) = \Pr[\mathbf{Exp}_{\mathcal{F}, \Sigma}^{\text{SUF-CMA}}(\lambda) = 1]$.

Definition 13 (SUF-CMA security). A digital signature scheme Σ is UF-CMA-secure if for any PPT forger \mathcal{F} , $\mathbf{Adv}_{\mathcal{F}, \Sigma}^{\text{SUF-CMA}}(\lambda) \leq \text{negl}(\lambda)$.

A stronger notion of security is *strong unforgeability* against chosen-message attacks (SUF-CMA). This notion is defined by considering a security experiment slightly different than $\mathbf{Exp}_{\mathcal{F}, \Sigma}^{\text{UF-CMA}}(\lambda)$. Instead of checking whether m^* is “new”, one checks whether the pair (m^*, σ^*) is “new”, i.e., if (m^*, σ^*) is different from all the pairs (m_i, σ_i) obtained by \mathcal{F} from the signing oracle.

A.2 Universal One-Way Hash Functions

Let us recall the notion of universal one-way hash function (UOWHF) families [NY89]. Let λ denote the security parameter.

Definition 14 (Universal one-way hash family). A collection of function families $\mathbb{H} = \{\mathcal{H}\}_\lambda$ where each \mathcal{H} is a function family $\mathcal{H} = \{h : \{0, 1\}^{q(\lambda)} \rightarrow \{0, 1\}^{\ell(\lambda)}\}$ is an universal one-way hash family if:

Efficient. The functions $q(\lambda)$ and $\ell(\lambda)$ are polynomially-bounded; furthermore, given λ and $x \in \{0, 1\}^{q(\lambda)}$ the value $h(x)$ can be computed in $\text{poly}(\lambda)$ time.

Compressing. For all λ we have that $q(\lambda) > \ell(\lambda)$.

Universal one-way. For all PPT algorithms \mathcal{A} , the following is negligible (in λ):

$$\Pr[x \leftarrow \mathcal{A}(1^\lambda); h \xleftarrow{\$} \mathcal{H}; x' \leftarrow \mathcal{A}(1^\lambda, h, x) : x, x' \in \{0, 1\}^{q(\lambda)} \wedge x \neq x' \wedge h(x) = h(x')].$$

B The SNARK definition used in [BGI14]

Here we recall the definition of SNARK considered in the work of Boyle, Goldwasser and Ivan [BGI14]. This definition is almost the same as our Definition 3 except that adaptive proof of knowledge is formulated as follows:

- **(Strong) Adaptive Proof of Knowledge.** Π is a SNARK for a language $L \in \text{NP}$ with witness relation \mathcal{R} if there exists a negligible function $\mu(\cdot)$ such that for all PPT provers \mathcal{A} there exists a PPT algorithm $\mathcal{E}_{\mathcal{A}} = (\mathcal{E}_{\mathcal{A}}^1, \mathcal{E}_{\mathcal{A}}^2)$ such that:
 - (i) for every \mathcal{D}

$$\left| \Pr \left[\mathcal{D}(\text{crs}, \text{vst}) = 1 \mid (\text{crs}, \text{vst}) \leftarrow \text{Gen}(1^\lambda) \right] - \Pr \left[\mathcal{D}(\text{crs}, \text{vst}) = 1 \mid (\text{crs}, \text{vst}, \text{td}) \leftarrow \mathcal{E}_{\mathcal{A}}^1(1^\lambda) \right] \right| = \mu(\lambda)$$

and (ii)

$$\Pr \left[\begin{array}{l} \text{Ver}(\text{crs}, y, \pi) = 1 \\ \wedge \\ (y, w) \notin \mathcal{R} \end{array} \mid \begin{array}{l} (\text{crs}, \text{td}) \leftarrow \mathcal{E}_{\mathcal{A}}^1(1^\lambda) \\ (y, \pi) \leftarrow \mathcal{A}(\text{crs}) \\ w \leftarrow \mathcal{E}_{\mathcal{A}}^2(\text{td}, \text{crs}, y, \pi) \end{array} \right] = \mu(\lambda)$$

where the probabilities are taken over $(\text{crs}, \text{td}) \leftarrow \mathcal{E}_{\mathcal{A}}^1(1^\lambda, T)$ and the random coins of $\mathcal{E}_{\mathcal{A}}^2$. Besides the fact that the extractor can take as additional input the trapdoor td , a major difference between this definition and the more standard one seems that in the above definition the extractor $\mathcal{E}_{\mathcal{A}}^2$ does not necessarily take the random tape of \mathcal{A} as part of its input. To the best of our knowledge, we are not aware of constructions under this definition, in the standard model.