

# Predictable Arguments of Knowledge

Antonio Faonio<sup>1</sup>, Jesper Buus Nielsen<sup>1</sup>, and Daniele Venturi<sup>2</sup>

<sup>1</sup> Department of Computer Science, Aarhus University, Aarhus, Denmark

<sup>2</sup> Sapienza, University of Rome, Rome, Italy

**Abstract.** We initiate a formal investigation on the power of *predictability* for argument of knowledge systems for  $NP$ . Specifically, we consider private-coin argument systems where the answer of the prover can be predicted, given the private randomness of the verifier; we call such protocols Predictable Arguments of Knowledge (PAoK).

Our study encompasses a full characterization of PAoK, showing that such arguments can be made extremely laconic, with the prover sending a single bit, and assumed to have only one round (i.e., two messages) of communication without loss of generality.

We additionally explore PAoK satisfying additional properties (including zero-knowledge and the possibility of re-using the same challenge across multiple executions with the prover), present several constructions of PAoK relying on different cryptographic tools, and discuss applications to cryptography.

Acknowledgement. The first author and second author acknowledge support by European Research Council Starting Grant 279447. The first author and second author acknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61361136003) for the Sino-Danish Center for the Theory of Interactive Computation.

## 1 Introduction

Consider the classical proof system for Graphs Non-Isomorphism where, on common input two graphs  $(G_0, G_1)$ , the verifier chooses a random bit  $b$ , and sends a uniformly random permutation of the graph  $G_b$  to the prover. If the two graphs are not isomorphic the prover replies correctly sending back the value  $b$ .

A peculiar property of the above proof system is that the verifier knows in advance the answer of the prover, i.e., the answer given by the prover is *predictable*. Another property is that it uses only one round of communication and that the prover sends a single bit. Following the work of Goldreich *et al.* [30] we call a proof system with these properties *extremely laconic*.

In this paper, we study the notion of predictability in interactive proof systems for  $NP$ . More specifically, we focus on the cryptographic setting where the prover's strategy is efficiently computable and, moreover, we aim for the notion of knowledge soundness, where any convincing polynomial-time prover must "know" the witness relative to the instance being proven.

We formalize this notion of Predictable Arguments of Knowledge (PAoK), explore their properties and applications, and provide several constructions based on various cryptographic tools and assumptions.

**Our Contributions and Techniques.** We proceed to describe our results and techniques in more details.

*Characterizing PAoK.* Syntactically a PAoK is a multi-round protocol  $(\mathcal{P}, \mathcal{V})$  where in each round: (i) The verifier  $\mathcal{V}$ , given the instance  $x$  and private coins  $r$ , generates a challenge  $c$  (that is sent to  $\mathcal{P}$ ) together with a predicted answer  $b$ ; (ii) The prover  $\mathcal{P}$ , given  $(x, w, c)$ , generates an answer  $a$ . The prover is said to convince the verifier if and only if  $a = b$  in all rounds.

Apart from being complete—meaning that an honest prover convinces the verifier with overwhelming probability—PAoK satisfy the standard property of *knowledge soundness*. Informally, this means that given any successful prover convincing the verifier on instance  $x$  with probability  $\epsilon$ , there exists an efficient extractor recovering a witness for  $x$  with probability polynomially related to  $\epsilon$ . Looking ahead, our definition of knowledge soundness is parametrized by a so-called instance sampler. Intuitively this means that only instances sampled through the sampler are extractable, and allows to consider more fine-grained flavours of extractability.<sup>1</sup>

Our first result is that PAoK can always be made extremely laconic, both in term of round complexity and of message complexity (i.e., the number of bits sent by the prover). Such a characterization is obtained as follows:

- First, we show that one can collapse any multi-round PAoK into a one-round PAoK with higher message complexity. Let  $(\mathcal{P}, \mathcal{V})$  be a  $\rho$ -round PAoK, where  $\mathcal{V}$  generates several challenges  $(c_1, \dots, c_\rho)$  with  $c_i$  used during round  $i$ .<sup>2</sup> We turn  $(\mathcal{P}, \mathcal{V})$  into a one-round predictable argument  $(\tilde{\mathcal{P}}, \tilde{\mathcal{V}})$  where the multi-round PAoK is “cut” at a random index  $i^* \in [\rho]$ ; this essentially means that  $\tilde{\mathcal{V}}$  runs  $\mathcal{V}$  and forwards  $(c_1, \dots, c_{i^*})$ , whereas  $\tilde{\mathcal{P}}$  runs  $\mathcal{P}$  and replies with  $(a_1, \dots, a_{i^*})$ . One can show that, if the initial PAoK has knowledge error  $\epsilon$ , the transformed PAoK has knowledge error  $\epsilon/\rho$ . The latter can finally be made negligible via parallel repetition. It is important to notice that parallel repetition, in general, does not amplify soundness for argument systems [5,39]. However, it is well known that for secret-coin one-round arguments (such as PAoK), parallel repetition amplifies (knowledge) soundness at an exponential rate [5].

---

<sup>1</sup> Similar fine-grained definitions have already been considered in the literature, e.g., for differing-inputs obfuscation [6].

<sup>2</sup> It is easy to see that generating all the challenges at the same time, independently of the prover’s answers, is without loss of generality.

- Second, we show how to reduce the prover’s answer length to a single bit<sup>3</sup> as follows. Let  $(\mathcal{P}, \mathcal{V})$  be a PAoK with  $\ell$ -bit answers. We define a new PAoK  $(\mathcal{P}', \mathcal{V}')$  where the verifier  $\mathcal{V}'$  runs  $\mathcal{V}$  in order to obtain a pair  $(c, b)$ , samples randomness  $r$ , and defines the new predicted answer to be the inner product between  $b$  and  $r$ . Given challenge  $(c, r)$  the prover  $\mathcal{P}'$  simply runs  $\mathcal{P}$  in order to obtain  $a$  and defines the answer to be the inner product between  $a$  and  $r$ . Knowledge soundness follows by the Goldreich-Levin hard-core bit theorem [28].

Interestingly, we can wrap up the two results together showing that any PAoK, no matter of the round or message complexity, can be made extremely laconic.

*Constructions.* Next, we turn to constructing PAoK. Our starting point is the observation that full-fledged PAoK for a relation  $R$  imply (and in fact are equivalent to) extractable witness encryption [31] (Ext-WE) for the same relation  $R$ . Briefly, a witness encryption scheme allows to encrypt an arbitrary message using a statement  $x$  belonging to an  $NP$ -language  $L$ ; decryption can be performed by anyone knowing a valid witness  $w$  for  $x$ . Extractable security means that from any adversary breaking semantic security of the encryption scheme, we can obtain an extractor computing a valid witness for  $x$ .

The equivalence between PAoK and Ext-WE can be seen as follows:

- From Ext-WE to PAoK we encrypt a random bit  $a$  using the encryption scheme and then ask the prover to return  $a$ .
- From PAoK to Ext-WE, we first make the PAoK extremely laconic, then we generate a challenge/answer pair  $(c, a)$  for the PAoK, and encrypt a single bit  $\beta$  as  $(c, a \oplus \beta)$ .<sup>4</sup>

In light of the recent work by Garg *et al.* [23], the above result can be seen as a negative result. In particular, [23] shows that, under the conjecture that a certain special-purpose obfuscator exists, it is impossible to have an Ext-WE scheme for a specific  $NP$  relation. The reason for this depends on the auxiliary information that an adversary might have on the input: The assumed special-purpose obfuscator could be used to obfuscate the auxiliary input in a way that allows to decrypt ciphertexts, without revealing any information about the witness. As stated in [23], such a negative result can be interpreted as an “implausibility result” on the existence of Ext-WE with arbitrary auxiliary input for all of  $NP$ . Given the equivalence between PAoK and Ext-WE such an implausibility result carries over to PAoK as well.<sup>5</sup>

<sup>3</sup> This further justifies our interest to arguments (as opposed to *proofs*) for  $NP$  as Goldreich *et al.* [30] showed that unless the polynomial-time hierarchy collapses there does not exist a laconic proof system for all  $NP$ .

<sup>4</sup> Domain extension for Ext-WE can be obtained by encrypting each bit of a message individually.

<sup>5</sup> Very recently, Bellare *et al.* [7] show that assuming sub-exponential one-way functions and sub-exponential indistinguishability obfuscation, differing-input obfuscation for Turing Machines [2] is impossible. While this result adds another negative evidence, it does not apply directly to Ext-WE.

Motivated by the above discussion, we propose two constructions of PAoK that circumvent the implausibility result of [23] by either restricting to specific  $NP$  relations, or by focusing on PAoK where knowledge soundness is only required to hold for a specific class of instance samplers (and thus for restricted auxiliary inputs). More in details:

- We show a simple connection between PAoK and so-called Extractable Hash-Proof Systems<sup>6</sup> [42] (Ext-HPS): Given an Ext-HPS for a relation  $R$  it is possible to construct a PAoK for a related relation  $R'$  in a natural way.
- We can construct a PAoK for a specific instance sampler by assuming a weak<sup>7</sup> form of differing-inputs obfuscation. The challenge  $c$  corresponds to an obfuscation of the circuit that hard-wires the instance  $x$  and a random value  $b$ , and upon input  $w$  returns  $b$  if and only if  $(x, w)$  is in the relation.

Interestingly, we can show that, for the special case of so-called random self-reducible relations,<sup>8</sup> a PAoK with knowledge soundness w.r.t. the instance sampler that corresponds to the algorithm for re-randomizing an instance in the language, can be generically leveraged to obtain a full-fledged PAoK (with arbitrary auxiliary input) for any  $NP$ -relation that is random-self reducible.

*Zero-Knowledge PAoK.* Notice that, as opposed to standard arguments, predictable arguments are non-trivial to construct even without requiring them to be zero-knowledge (or even witness indistinguishable).<sup>9</sup> Nevertheless, it is possible (and interesting) to consider PAoK that additionally satisfy the zero-knowledge property. It is well known that argument systems with a deterministic prover, such as PAoK, cannot be zero-knowledge in the plain model [29]. Motivated by this, given any PAoK (for some fixed relation), we propose two different transformations to obtain a zero-knowledge PAoK (for the same relation):

- The first transformation is in the non-programmable random oracle model. Here we exploit the fact that PAoK are honest-verifier zero-knowledge. Our strategy is to force the malicious verifier to act honestly; we achieve this by having the prover check that the challenge was honestly generated using randomness provided by the random oracle. In case the check fails the prover will not reveal the answer, but instead it will output a special symbol  $\perp$ . To ensure knowledge soundness we define the check to be dependent on the prover’s message, in such a way that a malicious prover cannot obtain the (private) randomness of the verifier in case it does not already know the correct answer.

---

<sup>6</sup> The connection between Hash Proof Systems and Witness Encryption was already noted by [23].

<sup>7</sup> Namely, following the terminology in [6], extractability only holds for a specific class of circuit samplers, related to the underlying instance sampler.

<sup>8</sup> Roughly speaking, a random self-reducible relation is a relation for which average-case hardness implies worst-case hardness.

<sup>9</sup> This is because the trivial protocol where the prover forwards a witness is not predictable.

- The second transformation is in the common random string (CRS) model, and works as follows. The verifier sends the challenge  $c$  together with a non-interactive zero-knowledge proof  $\pi$  that  $c$  is “well formed” (i.e., there exists random coins  $r$  such that the verifier of the underlying PAoK with coins  $r$  returns a pair  $(c, b)$ ).

We leave it as an interesting open problem to construct a witness indistinguishable PAoK in the plain model.

*Predictable ZAP.* In the basic definition of PAoK, the verifier generates the challenge  $c$  (together with the predicted answer  $b$ ) depending on the instance  $x$  being proven. We also look at the special case where the challenge is generated in an instance-independent manner, together with a trapdoor that later allows to predict the prover’s answer  $a$ . The goal here is to have the *same* challenge being used across multiple executions of a PAoK with the prover.

Protocols of this type have been already considered in the literature under the name of ZAP [17]. There are however a few crucial differences: (i) ZAP are public-coin, whereas predictable arguments are secret-coin; (ii) ZAP are witness indistinguishable, whereas predictable arguments are interesting even without requiring such a property. Hence, we formalize the notion of Predictable ZAP (PZAP) which is a kind of secret-coin ZAP in which the prover’s answer can be predicted (given the secret coins of the verifier and some trapdoor), and the same challenge can be re-used across multiple executions. We insist on PZAP satisfying knowledge soundness, but we do not require them to be witness indistinguishable; the definition of knowledge soundness features a malicious prover that can adaptively choose the target instance while keeping oracle access to the verifier algorithm. We also consider a weaker flavour, where the prover has no access to the verifier. We give a construction of PZAP relying on the recently introduced tool of Extractable Witness PRF [43]. We also show that weak PZAP can be generically leveraged to PZAP using standard cryptographic tools. This result shows that, under some standard cryptographic assumptions, for any construction of weak PZAP there exists *another* construction satisfying the definition of PZAP. It is interesting to understand if given a construction of weak PZAP the construction itself already satisfies the definition of PZAP. We give a negative evidence for this question. Namely, we show a black-box separation between weak PZAP and PZAP, ruling out a large class of black-box reductions from the former to the latter.

*Applications.* Although we find the concept of PAoK to be interesting in its own right, we also discuss applications of PAoK to proving lower bounds in two different cryptographic settings:

- Leakage-tolerant interactive protocols (as introduced by Bitanski, Canetti and Halevi [9]) are interactive protocols whose security degrades gracefully in the presence of arbitrary leakage on the state of the players. Previous work [36] showed that any leakage-tolerant interactive protocol for secure message transmission, tolerating leakage of poly-logarithmic size on

the state of the receiver, needs to have secret keys which are as long as the total number of bits transmitted using that key. Using PAoK, we can strengthen this negative result to hold already for leakage of a constant number of bits. Details are deferred in the full version of the paper [21].

- Non-malleable codes (as introduced by Dziembowski *et al.* [18]) allow to encode a message in such a way that the decoding of a tampered codeword either yields the original message or a completely unrelated value.

Previous work [22] showed an interesting application of non-malleable codes to protecting arbitrary computation (carried out by a von Neumann architecture) against tampering attacks. This result requires to assume a leakage- and tamper-free CPU which is used to carry out “simple” operations on a constant number of encodings.

A natural idea to weaken the assumption of a leakage-proof CPU, would be to design a code which remains non-malleable even given a small amount of leakage on the encoded message. Subsequent to our work [19], the concept of PAoK has been exploited to show that such non-malleable codes tolerating leakage from the encoding process cannot exist (under the assumption that collision-resistant hash functions exist).

**Giving up on Knowledge Extraction.** As already discussed above, the implausibility result of Garg *et al.* [23] has negative implications on some of our results. We were able to circumvent these implications by either constructing PAoK for restricted relations, or by considering weaker flavours of extractability. Yet another way to circumvent the implausibility result of [23] is to give up on knowledge soundness and to consider instead standard computational soundness (i.e., a computationally bounded malicious prover cannot convince the verifier into accepting a false statement).

Let us call a multi-round, predictable, computationally sound interactive protocol a *predictable argument*. It is easy to see that all our results for PAoK continue to hold for predictable arguments. In particular: (i) Predictable arguments can be assumed w.l.o.g. to be extremely laconic; (ii) There exists a predictable argument for a relation  $R$  if and only if there exists a (non-extractable) witness encryption scheme for  $R$ ; (iii) We can construct a predictable argument for a relation  $R$  given any hash-proof system for  $R$ ;<sup>10</sup> (iv) Computationally sound PZAP can be obtained based on any (non-extractable) Witness PRF.

**Additional Related Work.** A study of interactive proofs with laconic provers was done already in [27,30]. They did not investigate proofs of *knowledge*, though. As explained above our notion of PAoK is intimately related to extractable witness encryption, as first proposed by Goldwasser *et al.* [31]— where it is

<sup>10</sup> We note that, in the other direction, predictable arguments seem to imply some kind of hash-proof system where “statistical smoothness” is replaced by “computational smoothness.” We leave it as an interesting direction for future research to explore potential applications of such “computationally smooth” hash-proof systems and their connection to trapdoor hash-proof system (see Benhamouda *et al.* [8]).

argued that the construction of Garg *et al.* [24] is extractable. See [1,16] for more recent work on witness encryption.

In [25], Garg *et al.* introduce the concept of Efficiently Extractable Non-Interactive Instance-Dependent Commitment Scheme (Ext-NI-ID Commitment for short). The primitive resembles the concept of PAoK, however there is a crucial difference. Ext-NI-ID Commitments are statistical hiding, this implies that an Ext-NI-ID can be used to construct a Predictable Argument with “statistical soundness” for the same language, however, the reverse implication does not hold.

The problem we faced to amplify knowledge soundness of PAoK shares similarities with the problem of amplifying computational soundness for argument systems. Although it is well known that parallel repetition does not work in general [5,39], there are some exceptions such as 3-message arguments [5,13], public-coin arguments [38,15], and simulatable arguments [33,14] (a generalization of both 3-message and public-coin). Relevant to ours is the work of Haitner on random-terminating arguments [32].

**Roadmap.** We start by setting some basic notation, in Section 2. The definition of PAoK, together with their characterization in terms of round-complexity and amount of prover communication, can be found in Section 3. In Section 4 we explore constructions of PAoK for random self-reducible relations. The two compilers yielding zero-knowledge PAoK in the CRS model and in the non-programmable random oracle model are presented in Section 5. In Section 6 we investigate the concept of predictable ZAP. Finally, in Section 7, we discuss a few interesting open problems related to our work.

## 2 Preliminaries

For  $a, b \in \mathbb{R}$ , we let  $[a, b] = \{x \in \mathbb{R} : a \leq x \leq b\}$ ; for  $a \in \mathbb{N}$  we let  $[a] = \{1, 2, \dots, a\}$ . If  $x$  is a string, we denote its length by  $|x|$ ; if  $\mathcal{X}$  is a set,  $|\mathcal{X}|$  represents the number of elements in  $\mathcal{X}$ . When  $x$  is chosen randomly in  $\mathcal{X}$ , we write  $x \leftarrow \mathcal{X}$ . When  $\mathcal{A}$  is an algorithm, we write  $y \leftarrow \mathcal{A}(x)$  to denote a run of  $\mathcal{A}$  on input  $x$  and output  $y$ ; if  $\mathcal{A}$  is randomized, then  $y$  is a random variable and  $\mathcal{A}(x; r)$  denotes a run of  $\mathcal{A}$  on input  $x$  and randomness  $r$ . An algorithm  $\mathcal{A}$  is *probabilistic polynomial-time* (PPT) if  $\mathcal{A}$  is randomized and for any input  $x, r \in \{0, 1\}^*$  the computation of  $\mathcal{A}(x; r)$  terminates in at most  $\text{poly}(|x|)$  steps. Vectors and matrices are typeset in boldface. For a vector  $\mathbf{v} = (v_1, \dots, v_n)$  we sometimes write  $\mathbf{v}[i]$  for the  $i$ -th element of  $\mathbf{v}$ . We use  $\text{Maj}$  to denote the majority function.

Throughout the paper we let  $\kappa \in \mathbb{N}$  denote the security parameter. We say that a function  $\nu : \mathbb{N} \rightarrow [0, 1]$  is negligible in the security parameter, if  $\nu(\kappa) = \kappa^{-\omega(1)}$ . A function  $\mu : \mathbb{N} \rightarrow [0, 1]$  is noticeable in the security parameter, if there exists a positive polynomial  $p(\cdot)$  such that  $\nu(\kappa) \geq 1/p(\kappa)$  for infinitely many  $\kappa \geq \kappa_0$ .

Let  $X$  and  $Y$  be a pair of random variables. The statistical distance between  $X$  and  $Y$  is defined as  $\Delta(X, Y) := \max_{\mathcal{D}} |\Pr[\mathcal{D}(X) = 1] - \Pr[\mathcal{D}(Y) = 1]|$ , where the maximum is taken over all (possibly unbounded) distinguishers. In case the maximum is taken over all PPT distinguishers, we sometimes speak of computational distance. For two ensembles  $\mathcal{X} = \{X_\kappa\}_{\kappa \in \mathbb{N}}$  and  $\mathcal{Y} = \{Y_\kappa\}_{\kappa \in \mathbb{N}}$ , we write  $\mathcal{X} \equiv \mathcal{Y}$  to denote that  $\mathcal{X}$  and  $\mathcal{Y}$  are identically distributed,  $\mathcal{X} \stackrel{s}{\approx} \mathcal{Y}$  to denote that  $\mathcal{X}$  and  $\mathcal{Y}$  are statistically close (i.e., their statistical distance is bounded by a negligible function of the security parameter), and  $\mathcal{X} \stackrel{c}{\approx} \mathcal{Y}$  to denote that  $\mathcal{X}$  and  $\mathcal{Y}$  are computationally indistinguishable.

*Interactive Protocols.* Let  $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$  be an NP-relation, naturally defining a language  $L_R := \{x : \exists w \text{ s.t. } (x, w) \in R\}$ . We are typically interested in efficiently samplable relations, for which there exists a PPT algorithm `SamR` taking as input the security parameter (and random coins  $r$ ) and outputting a pair  $(x, w) \in R$ . An interactive protocol  $\Pi = (\mathcal{P}, \mathcal{V})$  for  $R$  features a prover  $\mathcal{P}$  (holding a value  $x \in L_R$  together with a corresponding witness  $w$ ) and a verifier  $\mathcal{V}$  (holding  $x$ ), where the goal of the prover is to convince the verifier that  $x \in L_R$ . At the end of the protocol execution, the verifier outputs either `acc` or `rej`. We write  $\langle \mathcal{P}(1^\kappa, x, w), \mathcal{V}(1^\kappa, x) \rangle$  for the random variable corresponding to the verifier's verdict, and  $\mathcal{P}(1^\kappa, x, w) \stackrel{c}{\approx} \mathcal{V}(1^\kappa, x)$  for the random variable corresponding to a transcript of protocol  $\Pi$  on input  $(x, w)$ .

Unless stated otherwise, all interactive protocols considered in this paper are *secret-coin*, meaning that the verifier's strategy depends on a secretly kept random tape. We also call  $\Pi$  a  $\rho$ -round protocol if the protocol consists of  $\rho$  rounds, where each round features a message from the verifier to the prover and viceversa.

### 3 Predictable Arguments of Knowledge

We start by defining Predictable Arguments of Knowledge (PAoK) as multi-round interactive protocols in which the verifier generates a challenge (to be sent to the prover) and can at the same time predict the prover's answer to that challenge; we insist on (computational) extractable security, meaning that from any prover convincing a verifier with some probability we can extract a witness with probability related to the prover's success probability.

The main result of this section is that PAoK can be assumed without loss of generality to be extremely laconic (i.e., the prover sends a single bit and the protocol consists of a single round of communication). More in detail, in Section 3.1, we show that any multi-round PAoK can be squeezed into a one-round PAoK. In Section 3.2 we show that, for any  $\ell \in \mathbb{N}$ , the existence of a PAoK where the prover answer is of length  $\ell$  bits implies the existence of a laconic PAoK.

*The Definition.* In a multi-round protocol the verifier produces many challenges  $\mathbf{c} = (c_1, \dots, c_\rho)$ . W.l.o.g. in a predictable argument, we can assume that all

the challenges are generated together and then forwarded one-by-one to the prover; this is because the answers are known *in advance*. Specifically, a  $\rho$ -round predictable argument is fully specified by a tuple of algorithms  $\Pi = (\text{Chall}, \text{Resp})$ , as described below:

1.  $\mathcal{V}$  samples  $(\mathbf{c}, \mathbf{b}) \leftarrow_{\$} \text{Chall}(1^\kappa, x)$ , where  $\mathbf{c} := (c_1, \dots, c_\rho)$  and  $\mathbf{b} := (b_1, \dots, b_\rho)$ .
2. For all  $i \in [\rho]$  in increasing sequence:
  - $\mathcal{V}$  forwards  $c_i$  to  $\mathcal{P}$ ;
  - $\mathcal{P}$  computes  $(a_1, \dots, a_i) := \text{Resp}(1^\kappa, x, w, c_1, \dots, c_i)$  and forwards  $a_i$  to  $\mathcal{V}$ ;
  - $\mathcal{V}$  checks that  $a_i = b_i$ , and returns **rej** if this is not the case.
3. If all challenges are answered correctly,  $\mathcal{V}$  returns **acc**.

Notice that the algorithm  $\text{Resp}$  takes as input all challenges up-to round  $i$  in order to generate the  $i$ -th answer.<sup>11</sup>

We say that prover  $\mathcal{P}$  and verifier  $\mathcal{V}$ , running the protocol above, *execute a PAoK*  $\Pi$  upon input security parameter  $1^\kappa$ , common input  $x$ , and prover's private input  $w$ ; we denote with  $\langle \mathcal{P}(1^\kappa, x, w), \mathcal{V}(1^\kappa, x) \rangle_\Pi$  (or, when  $\Pi$  is clear from the context, simply  $\langle \mathcal{P}(1^\kappa, x, w), \mathcal{V}(1^\kappa, x) \rangle$ ) the output of such interaction. We say that a prover  $\mathcal{P}$  *succeeds* on the instance  $x$  and auxiliary input  $w$  if  $\langle \mathcal{P}(1^\kappa, x, w), \mathcal{V}(1^\kappa, x) \rangle = \text{acc}$ . We give a granular definition of extractability that is parametrized by an efficient instance sampler  $\mathcal{S}$ , and that roughly says that the protocol is sound and moreover sampled instances are extractable. Here, the sampler is simply an algorithm taking as input the security parameter and auxiliary input  $z_S \in \{0, 1\}^*$ , and outputting an instance  $x$  together with auxiliary information  $aux \in \{0, 1\}^*$ .

**Definition 1 (Predictable Arguments of Knowledge).** *Let  $\Pi = (\text{Chall}, \text{Resp})$  be a  $\rho$ -round predictable argument for an NP relation  $R$ , with  $\ell$ -bit prover's answer. Consider the properties below.*

**Completeness:** *There exists a negligible function  $\nu : \mathbb{N} \rightarrow [0, 1]$  such that for all sequences  $\{(x_\kappa, w_\kappa)\}_{\kappa \geq 0}$  where  $(x_\kappa, w_\kappa) \in R$ , we have that:*

$$\Pr_{\mathcal{P}, \mathcal{V}} [\langle \mathcal{P}(1^\kappa, x_\kappa, w_\kappa), \mathcal{V}(1^\kappa, x_\kappa) \rangle = \text{rej}] \leq \nu(\kappa).$$

**( $\mathcal{S}, f, \epsilon$ )-Knowledge soundness:** *For all PPT provers  $\mathcal{P}^*$  there exists a PPT extractor  $\mathcal{K}$  such that for all auxiliary inputs  $z_P, z_S \in \{0, 1\}^*$  the following holds. Whenever*

$$p(\kappa) := \Pr_{\mathcal{P}^*, \mathcal{V}, r_S} [\langle \mathcal{P}^*(1^\kappa, aux, x, z_P), \mathcal{V}(x) \rangle = \text{acc} : (x, aux) := \mathcal{S}(1^\kappa, z_S; r_S)] > \epsilon(\kappa)$$

then

$$\Pr_{\mathcal{K}, r_S} \left[ \begin{array}{l} \exists w \text{ s.t. } f(w) = y \\ (x, w) \in R \end{array} : \begin{array}{l} (x, aux) := \mathcal{S}(1^\kappa, z_S; r_S), \\ y \leftarrow_{\$} \mathcal{K}(1^\kappa, x, z_P, z_S, aux) \end{array} \right] \geq p(\kappa) - \epsilon(\kappa).$$

<sup>11</sup> In the description above we let  $\text{Resp}$  output also all previous answers  $a_1, \dots, a_{i-1}$ ; while this is not necessary it can be assumed w.l.o.g. and will simplify the proof of Theorem 1.

We call  $\Pi$  a  $\rho$ -round  $\mathcal{S}$ -PAoK for  $R$ , if  $\Pi$  satisfies completeness and  $(\mathcal{S}, f, \epsilon)$ -knowledge soundness for any efficient computable function  $f$ , and moreover  $\epsilon - 2^{-\rho\ell}$  is negligible. We call  $\Pi$  an  $\mathcal{S}$ -PAoK for  $R$ , if  $\Pi$  is a 1-round  $\mathcal{S}$ -PAoK and we call it a laconic  $\mathcal{S}$ -PAoK if  $\Pi$  is an  $\mathcal{S}$ -PAoK and  $\ell = 1$ . Sometimes we also say that  $\Pi$  is a  $\rho$ -round  $(f, \mathcal{S})$ -PAoK if knowledge soundness holds for a specific function  $f$ .

Consider the dummy sampler  $\mathcal{S}_{\text{dummy}}$  that parses its input  $z_{\mathcal{S}}$  as  $(x, aux)$  and then outputs the pair  $(x, aux)$ . We call  $\Pi$  a  $\rho$ -round  $(f, \epsilon)$ -PAoK for  $R$ , if  $\Pi$  satisfies completeness and  $(\mathcal{S}_{\text{dummy}}, f, \epsilon)$ -knowledge soundness. We say that  $\Pi$  is a  $\rho$ -round PAoK for  $R$ , if  $\Pi$  is a  $\rho$ -round  $\mathcal{S}_{\text{dummy}}$ -PAoK for  $R$ .

The reason why the above definition is parametrized by the function  $f$  instead of considering the relation  $R' = \{(x, y) : \exists w \text{ s.t. } (x, w) \in R \wedge y = f(w)\}$  is that such a relation might not be an NP-relation (as it might be hard to check whether  $\exists w \text{ s.t. } (x, w) \in R \wedge y = f(w)$ ). Our definition, instead, ensures that the honest prover knows  $w$  but we can only extract  $f(w)$ . Also note that, in the above definition, the prover  $\mathcal{P}^*$  takes as input the auxiliary information returned by the sampler.

### 3.1 On Multi-Round PAoK

In this section we show that multi-round PAoK can be squeezed into a one-round PAoK (maintaining knowledge soundness).

Let  $\Pi = (\text{Chall}, \text{Resp})$  be a  $\rho$ -round PAoK. Consider the following protocol between prover  $\tilde{\mathcal{P}}_n$  and verifier  $\tilde{\mathcal{V}}_n$ —let us call it the *collapsed protocol* for future reference—for a parameter  $n \in \mathbb{N}$  to be determined later:

- Repeat the following sub-protocol  $\tilde{\Pi} = (\tilde{\mathcal{P}}, \tilde{\mathcal{V}})$  in parallel for all  $j \in [n]$ :
  - $\tilde{\mathcal{V}}$  runs  $(\mathbf{c}^j, \mathbf{b}^j) \leftarrow \text{Chall}(1^\kappa, x)$ ; let  $\mathbf{c}^j = (c_1^j, \dots, c_\rho^j)$  and similarly  $\mathbf{b}^j = (b_1^j, \dots, b_\rho^j)$ . Then,  $\tilde{\mathcal{V}}$  samples a random index  $i_j^* \leftarrow_{\$} [\rho]$ , and forwards  $(c_1^j, \dots, c_{i_j^*}^j)$  to  $\tilde{\mathcal{P}}$ .
  - $\tilde{\mathcal{P}}$ , given a pair  $(x, w)$  and challenges  $(c_1^j, \dots, c_{i_j^*}^j)$ , computes  $(a_1^j, \dots, a_{i_j^*}^j) \leftarrow_{\$} \text{Resp}(1^\kappa, x, w, c_1^j, \dots, c_{i_j^*}^j)$  and forwards  $(a_1^j, \dots, a_{i_j^*}^j)$  to  $\tilde{\mathcal{V}}$ .
  - $\tilde{\mathcal{V}}$  is said to accept the  $j$ -th parallel execution if and only if  $a_i^j = b_i^j$  for all  $i \in [i_j^*]$
- Return **acc** if and only if all parallel executions are accepting.

We write  $\tilde{\Pi}_n := (\tilde{\mathcal{P}}_n, \tilde{\mathcal{V}}_n)$  for the  $n$ -fold repetition of the sub-protocol  $\tilde{\Pi} = (\tilde{\mathcal{P}}, \tilde{\mathcal{V}})$ . Note that the sub-protocol  $\tilde{\Pi}$  is the one-round protocol (described above) that simply cuts the multi-round protocol  $\Pi$  to a random round. We show the following theorem :

**Theorem 1.** *For any polynomial  $\rho(\cdot)$  and any function  $f$  if  $\Pi$  is a  $\rho(\kappa)$ -round  $f$ -PAoK, then the above defined collapsed protocol  $\tilde{\Pi}_n = (\tilde{\mathcal{P}}_n, \tilde{\mathcal{V}}_n)$  with parameter  $n = \omega(\rho \log \kappa)$  is an  $f$ -PAoK.*

We give an intuition for the proof. For simplicity, assume that  $\Pi$  is a  $\frac{1}{3}$ -PAoK for the relation  $R$ . We claim that the knowledge error of the collapsed protocol is not bigger than  $1 - \frac{2}{3\rho}$ . To see this, consider a prover  $\mathcal{P}^*$  for the original protocol  $\Pi$  which at the  $i$ -th iteration (where  $i \in [\rho]$ ) forwards the challenge  $c_1, \dots, c_i$  to a malicious prover  $\tilde{\mathcal{P}}^*$  for the collapsed protocol. Notice that conditioned on  $i^* = i$  the challenge has exactly the same distribution as a challenge for the collapsed protocol. The prover  $\mathcal{P}^*$  fails if the malicious prover  $\tilde{\mathcal{P}}^*$  of the collapsed protocol answered wrongly at least one of the queries that he received. So if we suppose that  $\tilde{\mathcal{P}}^*$  succeeds with probability strictly bigger than  $1 - \frac{2}{3\rho}$ , then, by the union bound, the failing probability of  $\mathcal{P}^*$  is strictly bounded by  $\frac{2}{3\rho} \cdot \rho$ , therefore  $\mathcal{P}^*$  succeeds with probability strictly bigger than  $\frac{1}{3}$ .

Finally, we can make the knowledge soundness error of the collapsed protocol negligible via parallel repetition. It is important to notice that parallel repetition, in general, does not amplify soundness for argument systems [5,39]. Luckily, it does so (at an exponential rate) in the special case of secret-coin one-round arguments (such as PAoK) [5]. The proof of the above theorem relies on the well-known fact that parallel repetition decreases the (knowledge) soundness error of one-round arguments at an exponential rate.

**Lemma 1 (Theorem 4.1 of [5], adapted to one-round protocols).** *Let  $\Pi = (\mathcal{P}, \mathcal{V})$  be a one-round argument of knowledge and denote by  $\Pi_n = (\mathcal{P}_n, \mathcal{V}_n)$  the one-round protocol that consists of the  $n$ -fold repetition of the initial protocol  $\Pi$ . Suppose  $0 < \alpha, \beta < 1$  and  $n \geq 2$  is an integer. Suppose  $\alpha > (16/\beta) \cdot e^{-\beta \cdot n/128}$ . Then there is an oracle algorithm  $\mathcal{R}$  such that for any prover  $\mathcal{P}^*$ , verifier  $\mathcal{V}$  and input string  $x$ , the following is true: If  $\Pr[\langle \mathcal{P}^*(1^\kappa, x, aux), \mathcal{V}_n(x) \rangle = \text{acc}] \geq 2\alpha$  then  $\Pr[\langle \mathcal{R}^{\mathcal{P}^*}(1^\kappa, x, aux), \mathcal{V}(x) \rangle = \text{acc}] \geq 1 - \beta$ . Furthermore,  $\mathcal{R}^{\mathcal{P}^*}$  runs in time  $\text{poly}(n, |x|, \alpha^{-1})$ .*

*Proof (of Theorem 1).* Let  $\tilde{\mathcal{P}}_n^*$  be a prover for the collapsed protocol such that for some  $x$  and  $z$  succeeds with probability at least  $\kappa^{-c}$  for some constant  $c$ . Let  $\alpha = \frac{1}{2}\kappa^{-c}$  and  $\beta = \frac{1}{2\rho}$ , notice that setting  $n = \omega(\rho \log \kappa)$  the following equation holds for  $\kappa$  big enough:

$$\frac{1}{2}\kappa^{-c} = \alpha > (16/\beta) \cdot e^{-\beta \cdot n/128} = 32\rho \cdot e^{-\omega(\log \kappa)/256}.$$

We can apply Lemma 1 with the parameters  $\alpha$  and  $\beta$  set as above. Therefore, consider a single instance of the sub-protocol  $\tilde{\Pi}$ , the prover  $\mathcal{R}^{\tilde{\mathcal{P}}_n^*}$  succeeds with probability  $1 - \beta = 1 - \frac{1}{2\rho}$ .

We build a prover  $\mathcal{P}^*$  for  $\Pi$  that succeeds with probability  $\frac{1}{2}$ . Specifically, Let  $\tilde{\mathcal{P}}^* := \mathcal{R}^{\tilde{\mathcal{P}}_n^*}$  and let  $\mathcal{P}^*$  interact with the verifier  $\mathcal{V}$  of the multi-round protocol as follow:

1.  $\mathcal{V}$  samples  $(\mathbf{c}, \mathbf{b}) \leftarrow \text{Chall}(1^\kappa, x)$ , where  $\mathbf{c} := (c_1, \dots, c_\rho)$  and  $\mathbf{b} := (b_1, \dots, b_\rho)$ .
2. For all  $i \in [\rho]$  in increasing sequence:
  - Upon input challenge  $c_i$  from the verifier  $\mathcal{V}$ , prover  $\mathcal{P}^*$  runs internally  $\tilde{\mathcal{P}}^*$  on input  $(1^\kappa, x)$  and challenge  $(c_1, \dots, c_i)$ . If  $\tilde{\mathcal{P}}^*$  outputs  $(a_1, \dots, a_i)$ , then  $\mathcal{P}^*$  forwards  $a_i$  to  $\mathcal{V}$ ; otherwise it aborts.

Rewriting explicitly the acceptance probability of  $\tilde{\mathcal{P}}^*$  in the collapsed protocol on  $(x, z)$ :

$$\Pr \left[ \tilde{\mathcal{P}}^*(1^\kappa, x, z, c_1, \dots, c_i) = (b_1, \dots, b_i) : (\mathbf{c}, \mathbf{b}) \leftarrow_{\$} \text{Chall}(1^\kappa, x), i \leftarrow_{\$} [\rho] \right] \geq 1 - \frac{1}{2\rho}.$$

Let  $W_i$  be the event that  $a_i = b_i$  in the interaction between  $\mathcal{P}^*$  and  $\mathcal{V}$  described above. We can write:

$$\begin{aligned} & \Pr[\langle \mathcal{P}^*(1^\kappa, x, z), \mathcal{V}(1^\kappa, x) \rangle = \text{acc}] & (1) \\ &= \Pr[\forall i \in [\rho] : W_i] = 1 - \Pr[\exists i \in [\rho] : \neg W_i] \geq 1 - \sum_{i \in [\rho]} \Pr[\neg W_i] \\ &= 1 - \rho \cdot \mathbb{E}_{i \leftarrow_{\$} [\rho]} \left[ \Pr[\mathcal{P}^*(1^\kappa, x, c_1, \dots, c_i) \neq a_i : (\mathbf{c}, \mathbf{b}) \leftarrow_{\$} \text{Chall}(1^\kappa, x)] \right] \\ &\geq 1 - \left(\frac{1}{2\rho}\right) \cdot \rho = \frac{1}{2}. \end{aligned}$$

where the equations above follow by the definition of average and by our assumption on the success probability of  $\mathcal{P}^*$  on  $(x, z)$ . Notice that for any successful  $\tilde{\mathcal{P}}_n^*$  we can define an extractor that is the same extractor for the machine  $\mathcal{P}^*$  executing  $\tilde{\mathcal{P}}^* = \mathcal{R}^{\tilde{\mathcal{P}}_n^*}$  as a subroutine. Moreover, since  $\tilde{\mathcal{P}}_n^*$  succeeds with probability  $\kappa^{-c}$  then  $\mathcal{P}^*$  runs in polynomial time.

### 3.2 Laconic PAoK

We show that laconic PAoK (where the size of the prover's answer is  $\ell = 1$  bit) are in fact equivalent to PAoK.

**Theorem 2.** *Let  $R$  be an NP relation. If there exists a PAoK for  $R$  then there exists a laconic PAoK for  $R$ .*

The proof of the theorem, which appears in the full version of the paper [21], relies on the Goldreich-Levin Theorem [26, Theorem 2.5.2]. Here is the intuition. Let  $(\mathcal{P}, \mathcal{V})$  be a PAoK with  $\ell$ -bit answers. We define a new PAoK  $(\mathcal{P}', \mathcal{V}')$  where the verifier  $\mathcal{V}'$  runs  $\mathcal{V}$  in order to obtain a pair  $(c, b)$ , samples randomness  $r$ , and defines the new predicted answer to be the inner product between  $b$  and  $r$ . Given challenge  $(c, r)$  the prover  $\mathcal{P}'$  simply runs  $\mathcal{P}$  in order to obtain  $a$  and defines the answer to be the inner product between  $a$  and  $r$ . Knowledge soundness follows by the Goldreich-Levin theorem.

## 4 Constructing PAoK

We explore constructions of PAoK. For space reasons we defer to the full version of the paper [21] the constructions based on Extractable Witness Encryption [24,31] and on Extractable Hash-Proof Systems [42].

In Section 4.1, we focus on constructing PAoK for so-called random self-reducible relations. In particular, we show that, for such relations, a fully-extractable PAoK can be obtained by generically leveraging a PAoK for a (much weaker) specific sampler (which depends on the random self-reducible relation).

In Section 4.2, we show that a PAoK for a specific sampler can be obtained generically by using a differing-input obfuscator [6] for a related (specific) circuit sampler.

#### 4.1 PAoK for Random Self-Reducible Languages

We construct a PAoK for languages that are random self-reducible. Roughly speaking, a random self-reducible language is a language for which average-case hardness implies worst-case hardness. Random self-reducibility is a very natural property, with many applications in cryptography (see, e.g., [3,40,37]). Informally a function is random self-reducible if, given an algorithm that computes the function on random inputs, one can compute the function on any input. When considering  $NP$  relations, one has to take a little more care while defining random self-reducibility. We say that  $\mathcal{O}_R(\cdot)$  is an *oracle* for the relation  $R$ , if on any input  $x \in L_R$  we have that  $(x, \mathcal{O}_R(x)) \in R$ .

**Definition 2 (Self-reducible relation).** *An NP-relation  $R$  for a language  $L$  is random self-reducible if there exists a tuple of PPT algorithms  $\mathcal{W} := (\mathcal{W}_{\text{smp}}, \mathcal{W}_{\text{cmp}}, \mathcal{W}_{\text{inv}})$  such that, for any oracle  $\mathcal{O}_R$  for the relation  $R \subseteq X \times W$ , the following holds.*

**Correctness.** *For any  $x \in L_R$  and for any  $r \in \{0, 1\}^{p(|x|)}$ , let  $x' := \mathcal{W}_{\text{smp}}(x; r)$  and  $w := \mathcal{W}_{\text{cmp}}(x, w'; r)$  where  $w' \leftarrow \mathcal{O}_R(x')$ . Then  $(x, w) \in R$ .*

**Witness re-constructability.** *For any  $x \in L_R$  and for any  $r \in \{0, 1\}^{p(|x|)}$ , let  $x' := \mathcal{W}_{\text{smp}}(x; r)$  and  $w := \mathcal{W}_{\text{cmp}}(x, w'; r)$  where  $w' \leftarrow \mathcal{O}_R(x')$ , and define  $w'' := \mathcal{W}_{\text{inv}}(x, w; r)$ . Then  $(x', w'') \in R$ .*

**Uniformity.** *For any  $x$  the output of  $\mathcal{W}_{\text{smp}}(x)$  is uniformly distributed over  $X$ .*

We call the tuple of algorithms  $\mathcal{W}$  an *average-to-worst-case (AWC) reduction with witness re-constructibility*.

Notice that the reduction  $\mathcal{W}$  has access to a “powerful” oracle that produces a witness for a randomized instance, and uses such witness to compute a witness for the original instance. Moreover, for any fixed instance the function can be easily inverted. The witness re-constructibility property is not standard in the context of random self reducibility, however we note that it holds for many interesting random self-reducible relationships (e.g., for the case of discrete logarithm).

**Theorem 3.** *Let  $R$  be an NP-relation which has AWC reduction  $\mathcal{W} = (\mathcal{W}_{\text{smp}}, \mathcal{W}_{\text{cmp}}, \mathcal{W}_{\text{inv}})$  with witness re-constructability. If there exists a  $(\mathcal{W}_{\text{smp}}, \epsilon)$ -PAoK for the relation  $R$ , then there exists an  $\epsilon$ -PAoK for  $R$ .*

The proof of the above theorem is deferred to the full version of the paper [21]. Here we discuss some intuition. Let  $\Pi' := (\text{Chall}', \text{Resp}')$  be a PAoK (w.r.t. the sampler  $\mathcal{W}_{\text{smp}}$ ) for  $R$ , the idea of the construction is to map the input instance  $x$  into a random instance  $x'$  using  $\mathcal{W}_{\text{smp}}(x; r)$  for a random  $r$ , then sample a challenge using the algorithm  $\text{Chall}'$  on input instance  $x'$  and additionally send the prover the auxiliary information  $r$  needed to compute a valid witness  $w'$  for

$x'$ . The response algorithm first computes the valid witness  $w'$  for the instance  $x'$  using  $\mathcal{W}_{\text{inv}}$  and then answers the challenge. Let  $\Pi$  be the PAoK described above, given a prover  $\mathcal{P}^*$  for  $\Pi$  we need to define a knowledge extractor  $\mathcal{K}$ . The point is that  $\mathcal{P}^*$  can equivalently be seen as a prover for  $\Pi'$  where instances are sampled using  $\mathcal{W}_{\text{smp}}(x; \cdot)$ . For this scenario the knowledge soundness of  $\Pi$  provides a knowledge extractor  $\mathcal{K}'$ , and such an extractor can output a valid witness for a uniformly sampled instance. This is where we use the random self-reducibility property. The extractor  $\mathcal{K}'$ , in fact, can be seen as an oracle for the relation  $R$  that with noticeable probability produces a valid witness for a uniformly chosen instance. Therefore, using the AWC reduction  $\mathcal{W}$  with oracle access to  $\mathcal{K}'$  we can reconstruct a valid witness for the instance  $x$ .

## 4.2 PAoK for a Specific Sampler

We use the framework for obfuscation proposed by Bellare *et al.* in [6]. A circuit sampling algorithm is a PPT algorithm  $\mathcal{S} = \{\mathcal{S}_\kappa\}_{\kappa \in \mathbb{N}}$  whose output is distributed over  $\mathcal{C}_\kappa \times \mathcal{C}_\kappa \times \{0, 1\}^{p(\kappa)}$ , for a class of circuit  $\mathcal{C} = \{\mathcal{C}_\kappa\}_{\kappa \in \mathbb{N}}$  and a polynomial  $p$ . We assume that for every  $C_0, C_1 \in \mathcal{C}_\kappa$  it holds that  $|C_0| = |C_1|$ . Given any class of samplers  $\mathbf{S}$  for a class of circuits  $\mathcal{C}$  consider the following definition:

**Definition 3 (S-Obfuscator).** *A PPT algorithm Obf is an S-obfuscator for the parametrized collection of circuits  $\mathcal{C} = \{\mathcal{C}_\kappa\}_{\kappa \in \mathbb{N}}$  if the following requirements are met.*

- **Correctness:**  $\forall \kappa, \forall C \in \mathcal{C}_\kappa, \forall x : \Pr[C'(x) = C(x) : C' \leftarrow_{\mathcal{S}} \text{Obf}(1^\kappa, C)] = 1$ .
- **Security:** *For every sampler  $\mathcal{S} \in \mathbf{S}$ , for every PPT (distinguishing) algorithm  $\mathcal{D}$ , and every auxiliary inputs  $z_D, z_S \in \{0, 1\}^*$ , there exists a negligible function  $\nu : \mathbb{N} \rightarrow [0, 1]$  such that for all  $\kappa \in \mathbb{N}$ :*

$$\left| \Pr \left[ \mathcal{D}(C', aux, z_D, z_S) = 1 : \begin{array}{c} (C_0, C_1, aux) \leftarrow_{\mathcal{S}} \mathcal{S}(1^\kappa, z_S), \\ C' \leftarrow_{\mathcal{S}} \text{Obf}(1^\kappa, C_0) \end{array} \right] - \Pr \left[ \mathcal{D}(C', aux, z_D, z_S) = 1 : \begin{array}{c} (C_0, C_1, aux) \leftarrow_{\mathcal{S}} \mathcal{S}(1^\kappa, z_S), \\ C' \leftarrow_{\mathcal{S}} \text{Obf}(1^\kappa, C_1) \end{array} \right] \right| \leq \nu(\kappa),$$

where the probability is over the coins of  $\mathcal{S}$  and Obf.

Abusing the notation, given a circuit sampler  $\mathcal{S}$ , we say that Obf is an  $\mathcal{S}$ -obfuscator if it is an  $\{\mathcal{S}\}$ -obfuscator. It is easy to see that the above definition allows to consider various flavours of obfuscation as a special case (including indistinguishability and differing-input obfuscation [4]). In particular, we say that a circuit sampler is differing-input if for any PPT adversary  $\mathcal{A}$  and any auxiliary input  $z_S \in \{0, 1\}^*$  there exists a negligible function  $\nu : \mathbb{N} \rightarrow [0, 1]$  such that the following holds:

$$\Pr \left[ C_0(x) \neq C_1(x) : \begin{array}{c} (C_0, C_1, aux) \leftarrow_{\mathcal{S}} \mathcal{S}(1^\kappa, z_S) \\ x \leftarrow \mathcal{A}(C_0, C_1, aux, z_S) \end{array} \right] \leq \nu(\kappa).$$

Let  $\mathbf{S}^{\text{diff}}$  be the class of all differing-input samplers; it is clear that an  $\mathbf{S}^{\text{diff}}$ -obfuscator is equivalent to a differing-input obfuscator.

Consider the following construction of a PAoK  $\Pi = (\text{Chall}, \text{Resp})$  for a relation  $R$ .

- Upon input  $(1^\kappa, x)$  algorithm  $\text{Chall}(1^\kappa, x)$  outputs  $c := \text{Obf}(C_{x,b})$  where  $b \leftarrow_{\$} \{0, 1\}^\kappa$  and  $C_{x,b}$  is the circuit that hard-wires  $x$  and  $b$  and, upon input a value  $w$ , it returns  $b$  if and only if  $(x, w) \in R$  (and  $\perp$  otherwise).
- Upon input  $(1^\kappa, x, w, c)$ , algorithm  $\text{Resp}(1^\kappa, x, w, c)$  executes  $a := c(w)$  and outputs  $a$ .

Given an arbitrary instance sampler  $\mathcal{S}$ , let  $\text{CS}[\mathcal{S}]$  be the circuit samplers that sample randomness  $r' := r\|b$ , execute  $(x, aux) := \mathcal{S}(1^\kappa, z_{\mathcal{S}}; r)$ , and output the tuple  $(C_{x,b}, C_{x,\perp}, aux\|b)$ . We prove the following result, whose proof appears in the full version of the paper [21].

**Theorem 4.** *Let  $\mathcal{S}$  be an arbitrary instance sampler and  $\mathbf{S}^{\text{diff}}$  and  $\text{CS}[\mathcal{S}]$  be as above. If  $\text{CS}[\mathcal{S}] \in \mathbf{S}^{\text{diff}}$  and  $\text{Obf}$  is a  $\text{CS}[\mathcal{S}]$ -obfuscator, then the protocol  $\Pi$  described above is an  $\mathcal{S}$ -PAoK for the relation  $R$ .*

By combining Theorem 4 together with Theorem 3 we get the following corollary.

**Corollary 1.** *Let  $R$  be a random self-reducible NP-relation which is witness reconstructible and has AWC reduction  $\mathcal{W} = (\mathcal{W}_{\text{smp}}, \mathcal{W}_{\text{cmp}}, \mathcal{W}_{\text{inv}})$ . If there exists a  $\text{CS}[\mathcal{W}_{\text{smp}}]$ -obfuscator and  $\text{CS}[\mathcal{W}_{\text{smp}}] \in \mathbf{S}^{\text{diff}}$  then there exists a PAoK for  $R$ .*

## 5 On Zero Knowledge

One can easily verify that PAoK are always honest-verifier zero-knowledge, since the answer to a (honest) challenge from the verifier can be predicted without knowing a valid witness.

It is also not too hard to see that in general PAoK may not be witness indistinguishable (more details in the full version of the paper [21]).

Furthermore, we note that PAoK in the plain model can be zero-knowledge only for trivial languages. The reason is that predictable arguments have inherently deterministic provers and, as shown by Goldreich and Oren [29, Theorem 4.5], the zero-knowledge property for such protocols is achievable only for languages in  $BPP$ .

In this section we show how to circumvent this impossibility using setup assumptions. In particular, we show how to transform any PAoK into another PAoK additionally satisfying the zero-knowledge property (without giving up on predictability). We provide two solutions. The first one is in the common random string (CRS) model,<sup>12</sup> while the second one is in the non-programmable random oracle (NPRO) model.

<sup>12</sup> This model is sometimes also known as the Uniform Random String (URS) model.

## 5.1 Compiler in the CRS Model

We start by recalling the standard notion of zero-knowledge interactive protocols in the CRS model. Interactive protocols in the CRS model are defined analogously to interactive protocols in the plain model (cf. Section 2), with the only difference that at setup a uniformly random string  $\omega \leftarrow_{\$} \{0, 1\}^\ell$  is sampled and both the prover and the verifier additionally take  $\omega$  as input. For space reasons, the definition of zero-knowledge protocols in the CRS model is given in the full version of the paper [21].

*The compiler.* Our first compiler is based on a NIZK-PoK system (see full version for the formal definition). Let  $\Pi = (\text{Chall}, \text{Resp})$  be a PAoK for a relation  $R$ , and assume that  $\text{Chall}$  uses at most  $\rho(|x|, \kappa)$  random bits for a polynomial  $\rho$ . Let  $\mathcal{NIZK} = (\ell, \text{Prove}, \text{Ver})$  be a NIZK for the relation

$$R_{\text{chal}} = \{((c, x), r) : \exists b \text{ s.t. } (c, b) := \text{Chall}(1^\kappa, x; r)\}.$$

Consider the following one-round PAoK  $\Pi' = (\text{Chall}', \text{Resp}')$  in the CRS model.

- At setup a uniform CRS  $\omega \leftarrow_{\$} \{0, 1\}^\ell$  is sampled.
- Algorithm  $\text{Chall}'$  takes as input  $(1^\kappa, \omega, x)$  and proceeds as follows:
  1. Sample random tape  $r \leftarrow_{\$} \{0, 1\}^\rho$ .
  2. Generate a proof  $\pi \leftarrow_{\$} \text{Prove}(\omega, (c, x), r)$  for  $((c, x), r) \in R_{\text{chal}}$ .
  3. Output  $c' := (c, \pi)$ .
- Algorithm  $\text{Resp}'$  takes as input  $(1^\kappa, \omega, x, w, c')$  and proceeds as follows:
  1. Parse  $c' := (c, \pi)$ ; in case  $\text{Ver}(\omega, (c, x), \pi) = 0$  return  $\perp$ .
  2. Output  $b' := \text{Resp}(1^\kappa, x, w, c)$ .

Roughly speaking, in the above construction the verifier sends the challenge  $c$  together with a NIZK-PoK  $\pi$  that  $c$  is “well formed” (i.e., there exist random coins  $r$  such that the verifier of the underlying PAoK with coins  $r$  returns a pair  $(c, b)$ ); the prover answers only in case the proof  $\pi$  is correct. We show the following result, whose proof appears in the full version of the paper [21].

**Theorem 5.** *Let  $\Pi$  be a PAoK for the relation  $R \in NP$  and let  $\mathcal{NIZK}$  be a NIZK-PoK for the relation  $R_{\text{chal}}$ . Then the protocol  $\Pi'$  is a ZK-PAoK in the CRS model.*

The knowledge soundness of  $\Pi'$  follows almost directly from the zero-knowledge property of  $\mathcal{NIZK}$  and from the knowledge soundness of  $\Pi$ . In fact, one can consider a mental experiment where the verifier generates a simulated proof  $\pi$  instead of a real one. This proof does not carry any information about the randomness but it is indistinguishable from a real one. A successful prover in the real world is still successful in this mental experiment and, therefore, we reduced to the knowledge soundness of  $\Pi$ . The zero-knowledge of  $\Pi'$  follows from the fact that PAoK are honest-verifier zero-knowledge, and from the knowledge soundness of  $\mathcal{NIZK}$ . In particular, given a maliciously generated challenge  $(c^*, \pi^*)$ , the simulator can use the knowledge extractor of  $\mathcal{NIZK}$  on  $\pi^*$ , extract a valid witness  $r^*$ , and then produce a valid answer.

## 5.2 Compiler in the NPRO Model

We start by recalling the definition of zero-knowledge in the NPRO model, for interactive protocols. Recall that a NPRO is weaker than a programmable random oracle. Intuitively, in the NPRO model the simulator can observe the verifier's queries to the hash function, but is not allowed to program the behaviour of the hash function. The definition below is adapted from Wee [41].

**Definition 4 (Zero-knowledge protocol in the NPRO model).** *Let  $(\mathcal{P}, \mathcal{V})$  be an interactive protocol for an NP relation  $R$ . We say that  $(\mathcal{P}, \mathcal{V})$  satisfies the zero-knowledge property in the NPRO model if for every PPT malicious verifier  $\mathcal{V}^*$  there exists a PPT simulator  $\mathcal{Z}$  and a negligible function  $\nu : \mathbb{N} \rightarrow [0, 1]$  such that for all PPT distinguishers  $\mathcal{D}$ , all  $(x, w) \in R$ , and all auxiliary inputs  $z \in \{0, 1\}^*$ , the following holds:*

$$\Delta(\Pi, \mathcal{Z}, \mathcal{V}^*) := \max_{\mathcal{D}, z} \left| \Pr [\mathcal{D}^H(x, \tau, z) = 1 : \tau \leftarrow (\mathcal{P}^H(x, w) \leftrightarrow \mathcal{V}^{*H}(x, z))] \right. \\ \left. - \Pr [\mathcal{D}^H(x, \tau, z) = 1 : \tau \leftarrow \mathcal{Z}^H(x, z)] \right| \leq \nu(|x|).$$

*The compiler.* Let  $\Pi = (\text{Chall}, \text{Resp})$  be a PAoK for a relation  $R$  with  $\ell$ -bit prover's answer, and assume that  $\text{Chall}$  uses at most  $\rho(|x|, \kappa)$  random bits for a polynomial  $\rho$ . Let  $H$  be a random oracle with output length  $\rho(\kappa)$ . Consider the following derived one-round PAoK  $\Pi' = (\text{Chall}', \text{Resp}')$ .

- Algorithm  $\text{Chall}'$  takes as input  $(1^\kappa, x)$  and proceeds as follows:
  1. Sample a random tag  $t_1 \leftarrow_{\$} \{0, 1\}^\rho$  and compute  $r := H(t_1)$ .
  2. Run  $(c, b) := \text{Chall}(1^\kappa, x; r)$ .
  3. Define  $t_2 := H(b)$ , and set the challenge to  $c' := (c, t)$  where  $t := t_1 \oplus t_2$ .
- Algorithm  $\text{Resp}'$  takes as input  $(x, w, c)$  and proceeds as follows:
  1. Parse  $c' := (c, t)$  and run  $a \leftarrow_{\$} \text{Resp}(1^\kappa, x, w, c)$ .
  2. Define  $t_1 := t \oplus H(a)$ , and check whether  $(c, a) = \text{Chall}(1^\kappa, x; H(t_1))$ . If this is the case, output  $a$  and otherwise output  $\perp$ .

The main idea behind the above construction is to force the malicious verifier to follow the underlying protocol  $\Pi$ ; in order to do so we generate the challenge feeding the algorithm  $\text{Chall}$  with the uniformly random string  $H(t_1)$ . What we need now is to both make able the prover to check that the verifier followed the algorithm  $\text{Chall}$  and to maintain soundness. Unfortunately, since PAoK are private-coin protocols, we can't simply make the verifier output  $t_1$ ; what we do instead is to one-time pad the value with the value  $t_2$  which is computable only knowing the answer. We show the following result:

**Theorem 6.** *If  $\Pi$  is a PAoK with  $\ell$ -bit prover's answer for the relation  $R$ , and  $\ell = \omega(\log \kappa)$ , then the protocol  $\Pi'$  is a ZK-PAoK in the NPRO model.*

To prove soundness we show that  $t = t_1 \oplus t_2$  is essentially uniformly random if the prover does not know  $b$ : this explains why we need  $\ell = \omega(\log \kappa)$ , otherwise a malicious prover could just brute force the right value of  $b$  and check for

consistency. Note that here we are leveraging the power of the random oracle model, that allows us to produce polynomially-long pseudorandomness from unpredictability. To prove zero-knowledge we note that a simulator can look into the random-oracle calls made by the malicious verifier while running it. Given the output  $(c^*, t^*)$  produced by the malicious verifier two cases can happen:

- The simulator finds an oracle call  $t'$  that “explains” the challenge  $c^*$ , namely  $(c^*, b) = \text{Chall}(1^\kappa, x; H(t'))$ ; in this case the simulator just outputs  $b$ . We argue that the simulator produces an indistinguishable view because the protocol  $\Pi$  has overwhelming completeness.
- The simulator does not find any  $t'$  that explains the challenge. Then it outputs  $\perp$ . Let  $b'$  be the answer that the real prover would compute using the algorithm  $\text{Resp}$ . We argue that the malicious verifier can find a challenge  $(c^*, t^*)$  that passes the check, namely  $(c^*, b') = \text{Chall}(1^\kappa, x; H(H(b') \oplus t^*))$  only with negligible probability. Therefore the real prover would output  $\perp$  as well, and so the views are indistinguishable.

*Proof (of Theorem 6).* Completeness follows readily from the completeness of the underlying PAoK.

We proceed to prove knowledge soundness of  $\Pi'$ . Given a prover  $\mathcal{P}'^*$  for  $\Pi'$  that makes the verifier accept with probability  $p(\kappa)$ , we define a prover  $\mathcal{P}^*$  for  $\Pi$  that is successful with probability  $p(\kappa)/Q(\kappa)$  where  $Q$  is a polynomial that upper bounds the number of oracle calls made by  $\mathcal{P}'^*$  to the NPRO  $H$ . Prover  $\mathcal{P}^*$  proceeds as follow:

1. Upon input  $(1^\kappa, c, z)$ , set  $c' := (c, t)$  for uniformly random  $t \leftarrow_{\$} \{0, 1\}^\rho$  and run  $\mathcal{P}'^*(1^\kappa, c', z)$ . Initialize counter  $j$  to  $j := 1$ ,  $\mathcal{Q} := \emptyset$ , and pick a uniformly random index  $i^* \leftarrow_{\$} [Q(\kappa)]$ .
2. Upon input a random oracle query  $x$  from  $\mathcal{P}'^*$ , pick  $y \leftarrow_{\$} \{0, 1\}^\rho$  and add the tuple  $(x, y, j)$  to  $H$ . If  $j = i^*$ , then output  $x$  and stop. Otherwise set  $j \leftarrow j + 1$  and forward  $y$  to  $\mathcal{P}'^*$ .
3. In case  $\mathcal{P}'^*$  aborts or terminates, output  $\perp$  and stop.

Without loss of generality we can assume that the prover  $\mathcal{P}'^*$  does not repeat random oracle queries, and that before outputting an answer  $a^*$ , it checks that  $(c, a^*) := \text{Chall}(1^\kappa, x; H(t \oplus H(a^*)))$ . We now analyse the winning probability of  $\mathcal{P}^*$ . Let  $a$  be the correct answer corresponding to the challenge  $c$ . Observe that the view produced by  $\mathcal{P}^*$  is exactly the same as the real view (i.e., the view that  $\mathcal{P}'^*$ , with access to the random oracle, expects from an execution with the verifier  $\mathcal{V}'$  from  $\Pi'$ ), until  $\mathcal{P}'^*$  queries  $H$  with the value  $a$ . In this case, in fact,  $\mathcal{P}'^*$  expects to receive a tag  $t_2$  such that  $(c, a) := \text{Chall}(1^\kappa, x; H(t \oplus t_2))$ . We can

write,

$$\begin{aligned}
& \Pr[\mathcal{P}^*(1^\kappa, c, z) \text{ returns } a] \\
&= \Pr[(a, *, i^*) \in \mathcal{Q}] \\
&= \Pr[a \text{ is the } i^*\text{-th query to } H \wedge a = \mathcal{P}'^*(1^\kappa, c', z)] \\
&= \Pr[a \text{ is the } i^*\text{-th query to } H \mid a = \mathcal{P}^*(1^\kappa, c', z)] \Pr[a = \mathcal{P}^*(1^\kappa, c', z)] \\
&\geq 1/Q(\kappa) \cdot p(\kappa).
\end{aligned} \tag{2}$$

Notice that in Eq.(2) the two probabilities are taken over two different probability spaces, namely the view provided by  $\mathcal{P}'^*$  to the prover  $\mathcal{P}^*$  together with  $i^*$  on the left hand side and the view that  $\mathcal{P}'^*$  would expect in an execution with a honest prover together with the index  $i^*$  in the right hand side. Knowledge soundness of  $\Pi'$  follows.

We now prove the zero-knowledge property. Upon input  $(1^\kappa, x, z)$  the simulator  $\mathcal{Z}$  proceeds as follows:

1. Execute algorithm  $\mathcal{V}^*(1^\kappa, x, z)$  and forward all queries to  $H$ ; let  $\mathcal{Q}$  be the set of queries made by  $\mathcal{V}^*$ .
2. Eventually  $\mathcal{V}^*$  outputs a challenge  $c^* = (c'^*, t^*)$ . Check if there exist  $(a^*, t_1^*) \in \mathcal{Q}$  such that  $(c'^*, a^*) = \text{Chall}(1^\kappa, x; H(t_1^*))$  and  $t^* = t_1^* \oplus H(a^*)$ . Output the transcript  $\tau := (c^*, a^*)$ . If no such pair is found, output  $(c^*, \perp)$ .

Let  $r'$  be the randomness used by the prover. For any challenge  $c$ , instance  $x$  and witness  $w$ , we say that  $r$  is *good* for  $c$  w.r.t.  $x, w, r'$  if  $(c, a) = \text{Chall}(1^\kappa, x; r) \wedge a = \text{Resp}(1^\kappa, x, w, c; r')$ . By completeness, the probability that  $r$  is not good, for  $r \leftarrow_s \{0, 1\}^\rho$ , is negligible. Therefore by letting *Good* be the event that  $\mathcal{V}^*$  queries  $H$  only on inputs that output good randomness for some  $c$ , by taking a union bound over all queries we obtain

$$\Pr[\text{Good}] \geq 1 - Q(\kappa) \cdot \nu'(\kappa) \geq 1 - \nu(\kappa), \tag{3}$$

for negligible functions  $\nu, \nu' : \mathbb{N} \rightarrow [0, 1]$ .

From now on we assume that the event *Good* holds; notice that this only modifies by a negligible factor the distinguishing probability of the distinguisher  $\mathcal{D}$ .

We proceed with a case analysis on the possible outputs of the simulator and the prover:

- The second output of  $\mathcal{Z}$  is  $a \neq \perp$ , whereas the second output of  $\mathcal{P}$  is  $\perp$ . Conditioning on  $\mathcal{Z}$ 's second output being  $a \neq \perp$ , we get that the challenge  $c$  is *well formed*, namely,  $c$  is in the set of all possible challenges for the instance  $x$  and security parameter  $1^\kappa$ . On the other hand, the fact that  $\mathcal{P}$  outputs  $\perp$  means that either algorithm *Resp* aborted or the check in step 2 of the description of  $\Pi'$  failed. However, neither of the two cases can happen unless event *Good* does not happen. Namely, if *Resp* outputs  $\perp$  the randomness  $H(t^* \oplus H(a))$  is not good for  $c$  (w.r.t.  $x, w, r'$ ), and therefore *Resp* must have output  $a$  which, together with  $t^*$ , would pass the test in step 2 by definition of  $\mathcal{Z}$ . It follows that this case happens only with negligible probability.

- The second output returned by  $\mathcal{Z}$  is  $\perp$ , whereas  $\mathcal{P}$ 's second output is  $a \neq \perp$ . Conditioning on  $\mathcal{Z}$ 's second output being  $\perp$ , we get that  $\mathcal{V}^*$  made no queries  $(a^*, t_1^*)$  such that  $(c, a^*) = \text{Chall}(1^\kappa, x; t_1^*)$  and  $t_1^* = H(t^* \oplus H(a^*))$ . In such a case, there exists a negligible function  $\nu : \mathbb{N} \rightarrow [0, 1]$  such that:

$$\begin{aligned} & \Pr[(c, a) = \text{Chall}(1^\kappa, x; H(t^* \oplus H(a^*)))] \\ & \leq \Pr[t_1^* := (H(a) \oplus t) \in \mathcal{Q} \vee \text{Chall}(1^\kappa, x; H(t_1^*)) = (c, a)] \\ & \leq Q \cdot 2^{-\rho} + 2^{-\gamma} + \epsilon \leq \nu(\kappa), \end{aligned} \tag{4}$$

where  $2^\gamma$  is the size of the challenge space. Notice that by overwhelming completeness and  $\ell = \omega(\log \kappa)$ , it follows that  $\gamma = \omega(\log \kappa)$ .

- Both  $\mathcal{Z}$ 's and  $\mathcal{P}$ 's second output are not  $\perp$ , but they are different. This event cannot happen, since we are conditioning on *Good*.

Combining Eq.(3) and Eq.(4) we obtain that  $\Delta(\Pi, \mathcal{Z}, \mathcal{V}^*)$  is negligible, as desired.

*On RO-dependent auxiliary input.* Notice that Definition 4 does not allow the auxiliary input to depend on the random oracle. Wee [41] showed that this is necessary for one-round protocols, namely zero-knowledge w.r.t. RO-dependent auxiliary input is possible only for trivial languages. This is because the result of [29] relativizes.

In a similar fashion, for the case of multi-round protocols, one can show that also the proof of [29, Theorem 4.5] relativizes. It follows that the assumption of disallowing RO-dependent auxiliary input is necessary also in our case.

## 6 Predictable ZAPs

We recall the concept of ZAP introduced by Dwork and Naor [17]. ZAPs are two-message (i.e., one-round) protocols in which:

- (i) The first message, going from the verifier to the prover, can be fixed “once and for all,” and is independent of the instance being proven;
- (ii) The verifier’s message consists of public coins.

Typically a ZAP satisfies two properties. First, it is witness indistinguishable meaning that it is computationally hard to tell apart transcripts of the protocols generated using different witnesses (for a given statement). Second, the protocol remains sound even if the statement to be proven is chosen after the first message is fixed.

In this section we consider the notion of Predictable ZAP (PZAP). With the terminology “ZAP” we want to stress the particular structure of the argument system we are interested in, namely a one-round protocol in which the first message can be fixed “once and for all.” However, there are a few important differences between the notion of ZAPs and PZAPs. First off, PZAPs cannot be public coin, because the predictability requirement requires that the verifier uses private coins. Second, we relax the privacy requirement and allow PZAPs

not to be witness indistinguishable; notice that, in contrast to PZAPs, ZAPs become uninteresting in this case as the prover could simply forward the witness to the verifier. Third, ZAPs are typically only computationally sound, whereas we insist on knowledge soundness.

More formally, a PZAP is fully specified by a tuple of PPT algorithms  $\Pi = (\text{Chall}, \text{Resp}, \text{Predict})$  as described below:

1.  $\mathcal{V}$  samples  $(c, \vartheta) \leftarrow_s \text{Chall}(1^\kappa)$  and sends  $c$  to  $\mathcal{P}$ .
2.  $\mathcal{P}$  samples  $a \leftarrow_s \text{Resp}(1^\kappa, x, w, c)$  and sends  $a$  to  $\mathcal{V}$ .
3.  $\mathcal{V}$  computes  $b := \text{Predict}(1^\kappa, \vartheta, x)$  and outputs  $\text{acc}$  iff  $a = b$ .

Notice that, in contrast to the syntax of PAoK, now the verifier runs two algorithms  $\text{Chall}, \text{Predict}$ , where  $\text{Chall}$  is independent of the instance  $x$  being proven, and  $\text{Predict}$  uses the trapdoor  $\vartheta$  and the instance  $x$  in order to predict the prover's answer.

Care needs to be taken while defining (knowledge) soundness for PZAPs. In fact, observe that while the verification algorithm needs private coins, in many practical circumstances the adversary might be able to infer the outcome of the verifier, and thus learn one bit of information about the verifier's private coins. For this reason, as we aim to constructing argument systems where the first message can be re-used, we enhance the adversary with oracle access to the verifier in the definition of soundness.

**Definition 5 (Predictable ZAP).** *Let  $\Pi = (\text{Chall}, \text{Resp}, \text{Predict})$  be as specified above, and let  $R$  be an NP relation. Consider the properties below.*

**Completeness:** *There exists a negligible function  $\nu : \mathbb{N} \rightarrow [0, 1]$  such that for all  $(x, w) \in R$ :*

$$\Pr_{c, \vartheta} [\text{Predict}(1^\kappa, \vartheta, x) \neq \text{Resp}(1^\kappa, x, w, c) : (c, \vartheta) \leftarrow \text{Chall}(1^\kappa)] \leq \nu(\kappa).$$

**(Adaptive) Knowledge soundness with error  $\epsilon$ :** *For all PPT provers  $\mathcal{P}^*$  making polynomially many queries to its oracle, there exists a PPT extractor  $\mathcal{K}$  such that for any auxiliary input  $z \in \{0, 1\}^*$  the following holds. Whenever*

$$p_z(\kappa) := \Pr \left[ \begin{array}{l} (c, \vartheta) \leftarrow_s \text{Chall}(1^\kappa), \\ (x, a) \leftarrow_s \mathcal{P}^{*\mathcal{V}(1^\kappa, \vartheta, \cdot, \cdot)}(c, z) \text{ where } |x| = \kappa, \\ b := \text{Predict}(1^\kappa, \vartheta, x). \end{array} \right] > \epsilon(\kappa),$$

*we have*

$$\Pr \left[ \begin{array}{l} (c, \vartheta) \leftarrow_s \text{Chall}(1^\kappa), \\ (x, w) \in R : (x, a) \leftarrow_s \mathcal{P}^{*\mathcal{V}(1^\kappa, \vartheta, \cdot, \cdot)}(c, z) \text{ where } |x| = \kappa, \\ w \leftarrow_s \mathcal{K}(1^\kappa, x, z, \mathcal{Q}). \end{array} \right] \geq p_z(\kappa) - \epsilon(\kappa).$$

*In the above equations, we denote by  $\mathcal{V}(1^\kappa, \vartheta, \cdot, \cdot)$  the oracle machine that upon input a query  $(x, a)$  computes  $b := \text{Predict}(1^\kappa, \vartheta, x)$  and outputs 1 iff  $a = b$ ; we also write  $\mathcal{Q}$  for the list  $\{(x_i, a_i), d_i\}$  of oracle queries (and answers to these queries) made by  $\mathcal{P}^*$ .*

Let  $\ell$  be the size of the prover's answer, we call  $\Pi$  a predictable ZAP (PZAP) for  $R$  if  $\Pi$  satisfies completeness and adaptive knowledge soundness with error  $\epsilon$ , and moreover  $\epsilon - 2^{-\ell}$  is negligible. In case knowledge soundness holds provided that no verification queries are allowed, we call  $\Pi$  a weak PZAP.

The definition of *laconic* PZAPs is obtained as a special case of the above definition by setting  $\ell = 1$ . Note, however, that in this case we additionally need to require that the value  $x$  returned by  $\mathcal{P}^*$  is not contained in  $\mathcal{Q}$ .<sup>13</sup>

In the full version of the paper [21] we show a construction of PZAP based on any extractable witness pseudo-random function (Ext-WPRF), a primitive recently introduced in [43].

## 6.1 On Weak PZAP versus PZAP

We investigate the relation between the notions of weak PZAP and PZAP. On the positive side, we show that weak PZAP for  $NP$  can be generically leveraged to PZAP for  $NP$  in a generic (non-black-box) manner. On the negative side, we show an impossibility result ruling out a broad class of black-box reductions from weak PZAP to PZAP. Both results assume the existence of one-way functions.

**From Weak PZAP to PZAP.** We show the following result:

**Theorem 7.** *Under the assumption that non-interactive zero-knowledge proof of knowledge systems for  $NP$  and non-interactive computationally-hiding commitment schemes exist, weak PZAP for  $NP$  imply PZAP for  $NP$ .*

Before coming to the proof, let us introduce some useful notation. Given a set  $I \subseteq \{0, 1\}^\kappa$ , we will say that  $I$  is *bit-fixing* if there exists a string in  $x \in \{0, 1, \star\}^\kappa$  such that  $I_x = I$  where  $I_x := \{y \in \{0, 1\}^\kappa : \forall i \in [\kappa], (x_i = y_i \vee x_i = \star)\}$  is the set of all  $\kappa$ -bit strings matching  $x$  in the positions where  $x$  is equal to 0/1. The symbol  $\star$  takes the role of a special “don't care” symbol. Notice that there is a bijection between the set  $\{0, 1, \star\}^\kappa$  and the family of all bit-fixing sets contained in  $\{0, 1\}^\kappa$ ; in particular, for any  $I \subseteq \{0, 1\}^\kappa$  there exists a unique  $x \in \{0, 1, \star\}^\kappa$  such that  $I = I_x$  (and viceversa). Therefore, in what follows, we use  $x$  and  $I_x$  interchangeably. We also enforce the empty set to be part of the family of all bit-fixing sets, by letting  $I_\perp = \emptyset$  (corresponding to  $x = \perp$ ).

We now give some intuition for the proof of Theorem 7. The proof is divided in two main steps. In the first step, we define three algorithms (Gen, Sign, Verify). Roughly speaking, such a tuple constitutes a special type of signature scheme where the key generation algorithm Gen additionally takes as input a bit-fixing set  $I$  and returns a secret key that allows to sign messages  $m \notin I$ . There are two main properties we need from such a signature scheme: (i) The verification key and any set of polynomially many (adaptively chosen) signature queries do not

<sup>13</sup> This is necessary, as otherwise a malicious prover could query both  $(x, 0)$  and  $(x, 1)$ , for  $x \notin L$ , and succeed with probability 1.

reveal any information on the set  $I$ ; (ii) It should be hard to forge signatures on messages  $m \in I$ , even when given the set  $I$  and the secret key corresponding to  $I$ . A variation of such a primitive, with a few crucial differences, already appeared in the literature under the name of functional signatures [11].<sup>14</sup> Fix now some  $NP$ -relation  $R$ . In the second step of the proof, we consider an augmented  $NP$ -relation where the witness of an instance  $(x, VK)$  is either a witness  $w$  for  $(x, w) \in R$ , or a valid signature of  $x$  under  $VK$ . We then construct a PZAP based on a weak PZAP and on a NIZK-PoK for the above augmented  $NP$ -relation.

The reduction from weak PZAP to PZAP uses a partitioning technique, similar to the one used to prove unforgeability of several signature schemes (see, e.g., [10,34,35,12,20]). Intuitively, we can set the reduction in such a way that by sampling a random bit-fixing set  $I$  all the verification queries made by a succeeding prover for the PZAP will not be in  $I$  with good probability (and therefore such queries can be dealt with using knowledge of the signature key corresponding to  $I$ ); this holds because the prover has no information on the set  $I$ , as ensured by property (i) defined above. On the other hand, the challenge  $x^*$  output by the prover will be contained in the set  $I$ , which will allow the reduction to break the weak PZAP. Here, is where we rely on property (ii) described above, so that the reduction is not able to forge a signature for  $x^*$ , and thus the extracted witness  $w^*$  must be a valid witness for  $(x^*, w^*) \in R$ .

*Proof (of Theorem 7).* Let  $\text{Com}$  be a computationally hiding commitment scheme with message space  $\{0, 1, \star\}^\kappa \cup \{\perp\}$ . Consider the following relation:

$$R_{\text{com}} := \{(m, \text{com}), (x, r) : \text{com} = \text{Com}(x; r) \wedge m \notin I_x\}.$$

Let  $\mathcal{NIZK} = (\ell, \text{Prove}, \text{Ver})$  be a NIZK-PoK for the relation  $R_{\text{com}}$ . We define the following tuple of algorithms ( $\text{Gen}, \text{Sign}, \text{Verify}$ ).

- Algorithm  $\text{Gen}$  takes as input the security parameter and a string  $x \in \{0, 1, \star\}^\kappa \cup \{\perp\}$ , samples  $\omega \leftarrow_{\$} \{0, 1\}^{\ell(\kappa)}$ , and defines  $\text{com} := \text{Com}(x; r)$  for some random tape  $r$ . It then outputs  $VK := (\omega, \text{com})$  and  $SK := (\omega, x, r)$ .
- Algorithm  $\text{Sign}$  takes as input a secret key  $SK$  and a message  $m$ , and outputs  $\sigma := \pi \leftarrow_{\$} \text{Prove}(\omega, (m, \text{com}), (x, r))$ .
- Algorithm  $\text{Verify}$  takes as input a verification key  $VK$  and a pair  $(m, \sigma)$ , parses  $VK := (\omega, \text{com})$ , and outputs the same as  $\text{Ver}(\omega, (m, \text{com}), \sigma)$ .

The lemmas below show two main properties of the above signature scheme.

**Lemma 2.** *For any PPT distinguisher  $\mathcal{D}$ , and any bit-fixing set  $I \subseteq \{0, 1\}^\kappa$ , there exists a negligible function  $\nu : \mathbb{N} \rightarrow [0, 1]$  such that:*

$$\begin{aligned} & \left| \Pr[\mathcal{D}^{\text{Sign}(SK_I, \cdot)}(VK, I) : (VK, SK_I) \leftarrow_{\$} \text{Gen}(1^\kappa, I)] \right. \\ & \quad \left. - \Pr[\mathcal{D}^{\text{Sign}(SK, \cdot)}(VK, I) : (VK, SK) \leftarrow_{\$} \text{Gen}(1^\kappa, \perp)] \right| \leq \nu(\kappa), \end{aligned}$$

<sup>14</sup> On a high level, the difference is that functional signatures allow to generate punctured signature keys, whereas our signature scheme allows to puncture the message space.

where  $\mathcal{D}$  is not allowed to query its oracle on messages  $m \in I$ .

*Proof.* We consider a series of hybrid experiments, where each hybrid is indexed by a bit-fixing set  $I$  and outputs the view of a distinguisher  $\mathcal{D}$  taking as input a verification key and the set  $I$ , while given oracle access to a signing oracle.

**Hybrid  $\mathcal{H}_1^I$ :** The first hybrid samples  $(VK, SK) \leftarrow_{\$} \text{Gen}(1^\kappa, I)$  and runs the distinguisher  $\mathcal{D}$  upon input  $(VK, I)$  and with oracle access to  $\text{Sign}(SK, \cdot)$ .

**Hybrid  $\mathcal{H}_2^I$ :** Let  $\mathcal{Z}$  be the simulator of the underlying NIZK-PoK. The second hybrid samples  $(\tilde{\omega}, \vartheta) \leftarrow_{\$} \mathcal{Z}_0(1^\kappa)$  and defines  $com := \text{Com}(x; r)$  (for random tape  $r$ ) and  $\tilde{VK} = (\tilde{\omega}, com)$ . It then runs the distinguisher  $\mathcal{D}$  upon input  $(\tilde{VK}, I)$ , and answers its oracle queries  $m$  by returning  $\tilde{\sigma} \leftarrow_{\$} \mathcal{Z}_1(\vartheta, (m, com))$ .

The two claims below imply the statement of Lemma 2.

*Claim.* For all bit-fixing sets  $I$ , we have  $\{\mathcal{H}_1^I\}_{\kappa \in \mathbb{N}} \stackrel{c}{\approx} \{\mathcal{H}_2^I\}_{\kappa \in \mathbb{N}}$ .

*Proof (of claim).* The only difference between the two experiments is in the way the verification key is computed and in how the signature queries are answered. In particular, the second experiment replaces the CRS with a simulated CRS and answers signature queries by running the ZK simulator of the NIZK. Note that the commitment  $com$  has the same distribution in both experiments.

Clearly, given any distinguisher that tells apart the two hybrids for some set  $I$  we can derive a distinguisher contradicting the unbounded zero-knowledge property of the NIZK. This concludes the proof.

*Claim.* Let  $I_\perp := \emptyset$ . For all bit-fixing sets  $I$ , we have  $\{\mathcal{H}_2^I\}_{\kappa \in \mathbb{N}} \stackrel{c}{\approx} \{\mathcal{H}_2^{I_\perp}\}_{\kappa \in \mathbb{N}}$ .

*Proof (of claim).* Given a PPT distinguisher  $\mathcal{D}$  telling apart  $\mathcal{H}_2^I$  and  $\mathcal{H}_2^{I_\perp}$ , we construct a PPT distinguisher  $\mathcal{D}'$  that breaks computational hiding of the commitment scheme. Distinguisher  $\mathcal{D}'$  is given as input a value  $com'$  which is either a commitment to  $I$  or a commitment to  $I_\perp$ . Thus,  $\mathcal{D}'$  simply emulates the view for  $\mathcal{D}$  but uses  $com'$  instead of  $com$ .

The claim follows by observing that in case  $com'$  is a commitment to  $I$  the view generated by  $\mathcal{D}'$  is identical to that in hybrid  $\mathcal{H}_2^I$ , whereas in case  $com'$  is a commitment to  $I_\perp$  the view generated by  $\mathcal{D}'$  is identical to that in hybrid  $\mathcal{H}_2^{I_\perp}$ . Hence,  $\mathcal{D}'$  retains the same advantage as  $\mathcal{D}$ , a contradiction.

**Lemma 3.** For any PPT forger  $\mathcal{F}$ , and for any bit-fixing set  $I$ , there exists a negligible function  $\nu : \mathbb{N} \rightarrow [0, 1]$  such that the following holds:

$$\Pr \left[ m^* \in I \wedge \text{Verify}(VK, m^*, \sigma^*) = 1 : \begin{array}{l} (m^*, \sigma^*) \leftarrow_{\$} \mathcal{F}(I, r), \\ (VK, SK_I) := \text{Gen}(1^\kappa, I; r) \end{array} \right] \leq \nu(\kappa).$$

*Proof.* We rely on the knowledge soundness property of the NIZK-PoK and on the binding property of the commitment scheme. By contradiction, assume that

there exists a PPT forger  $\mathcal{F}$ , a bit-fixing set  $I_x$ , and some polynomial  $p(\cdot)$ , such that for infinitely many values of  $\kappa \in \mathbb{N}$

$$\Pr \left[ m^* \in I_x \wedge \text{Ver}(\omega, (m^*, \text{com}), \sigma^*) = 1 : \begin{array}{l} r \leftarrow_{\$} \{0, 1\}^*, \omega \leftarrow_{\$} \{0, 1\}^\ell \\ \text{com} \leftarrow_{\$} \text{Com}(x; r) \\ (m^*, \sigma^*) \leftarrow_{\$} \mathcal{F}(\omega, r, I_x) \end{array} \right] \geq 1/p(\kappa).$$

Consider the following adversary  $\mathcal{B}$  attacking the binding property of the commitment scheme:

- i) Upon input  $1^\kappa$ , run  $(\tilde{\omega}, \vartheta) \leftarrow_{\$} \mathcal{K}_0(1^\kappa)$ ;
- ii) Obtain  $(m^*, \sigma^*) \leftarrow_{\$} \mathcal{F}(\tilde{\omega}, r, I_x)$  for some  $x \in \{0, 1, \star\}^\kappa$  and  $r \leftarrow_{\$} \{0, 1\}^*$ ;
- iii) Extract  $(x', r') \leftarrow_{\$} \mathcal{K}_1(\tilde{\omega}, \vartheta, (m^*, \text{com}), \sigma^*)$ , where  $\text{com} = \text{Com}(x; r)$ ;
- iv) Output  $(x, r), (x', r')$  and  $m$  (as an auxiliary output).

By relying on the knowledge soundness property of the NIZK-PoK, and using the fact that the forger outputs an accepting proof with non-negligible probability, we obtain:

$$\begin{aligned} & \Pr[\mathcal{B} \text{ wins}] \\ &= \Pr[\text{com} = \text{Com}(x'; r') \wedge (x, r) \neq (x', r') : ((x, r), (x', r')), m) \leftarrow_{\$} \mathcal{B}(1^\kappa)] \\ &\geq \Pr \left[ \begin{array}{l} \text{com} = \text{Com}(x'; r'), \\ m \notin I_{x'}, \\ m \in I_x \end{array} : ((x, r), (x', r')), m) \leftarrow_{\$} \mathcal{B}(1^\kappa) \right] - \nu(\kappa) \\ &\geq \Pr \left[ \begin{array}{l} m^* \in I_x, \\ \text{Ver}(\omega, (m^*, \text{com}), \sigma^*) = 1 : \begin{array}{l} r \leftarrow_{\$} \{0, 1\}^*, \omega \leftarrow_{\$} \{0, 1\}^\ell \\ \text{com} \leftarrow_{\$} \text{Com}(x; r) \\ (m^*, \sigma^*) \leftarrow_{\$} \mathcal{F}(\omega, r, I_x) \end{array} \end{array} \right] \\ &\geq 1/p(\kappa) - \nu(\kappa), \end{aligned}$$

for some negligible function  $\nu(\cdot)$ . The first inequality uses the fact that the condition  $(m \notin I_{x'}) \wedge m \in I_x$  implies  $I_x \neq I_{x'}$  (and thus  $x \neq x'$ ), and thus is sufficient for violating the binding property. This concludes the proof.

We can now explain how to transform a weak PZAP for  $NP$  into a PZAP for  $NP$ . Let  $R$  be an  $NP$ -relation. Consider the following derived relation:

$$R' = \{((x, VK), w) : (x, w) \in R \vee \text{Verify}(VK, x, w) = 1\}.$$

Clearly,  $R'$  is in  $NP$ , so let  $\Pi = (\text{Chall}, \text{Resp}, \text{Predict})$  be a weak PZAP for  $R'$ . Define the following PZAP  $\Pi' = (\text{Chall}', \text{Resp}', \text{Predict}')$  for the relation  $R$ .

- Algorithm  $\text{Chall}'$  takes as input  $(1^\kappa, x)$  and proceeds as follows:
  - Run  $(c, \vartheta) \leftarrow_{\$} \text{Chall}(1^\kappa)$ .
  - Sample  $(VK, SK) \leftarrow_{\$} \text{Gen}(1^\kappa, \perp)$ , and let the challenge be  $c' := (c, VK)$  and the trapdoor be  $\vartheta' = (\vartheta, VK)$ .
- Algorithm  $\text{Resp}'$  takes as input  $(1^\kappa, x, w, c')$ , parses  $c' := (c, VK)$ , and outputs  $a := \text{Resp}(1^\kappa, (x, VK), w, c)$ .

- Algorithm  $\text{Predict}'$  takes as input  $1^\kappa, \vartheta', x$ , parses  $\vartheta' := (\vartheta, VK)$ , and outputs  $b := \text{Predict}(\vartheta, (x, VK))$ .

The lemma below concludes the proof of Theorem 7.

**Lemma 4.** *Let  $\Pi$  and  $\Pi'$  be as above. If  $\Pi$  is a weak PZAP for  $R'$ , then  $\Pi'$  is a PZAP for  $R$ .*

*Proof.* Given a prover  $\mathcal{P}^*$  for  $\Pi'$ , we construct a prover  $\mathcal{P}_\alpha$  for  $\Pi'$  for a parameter  $\alpha \in [\kappa]$  to be determined later. The description of  $\mathcal{P}_\alpha$  follows.

- Upon input challenge  $c$ , choose  $s \in \{0, 1, \star\}^\kappa$  in such a way that  $\alpha := |\{i \in [\kappa] : s_i = \star\}|$ . Sample  $(VK, SK_I) \leftarrow_s \text{Gen}(1^\kappa, I)$  for  $I := I_s$ , and forward the challenge  $c' := (c, VK)$  to  $\mathcal{P}^*$ .
- Upon input a verification query  $(x_i, a_i)$  from  $\mathcal{P}^*$  behave as follows:
  - In case  $x_i \in I$ , stop simulating  $\mathcal{P}^*$ , pick a random  $x^* \leftarrow_s \{0, 1\}^\kappa \setminus I$ , and return the instance  $(x^*, VK)$  and answer  $a^* := \text{Resp}(1^\kappa, c, (x^*, VK), \text{Sign}(SK_I, x^*))$ .
  - In case  $x_i \notin I$ , compute  $\sigma \leftarrow_s \text{Sign}(SK_I, x_i)$  and answer the verification query with 1 iff  $a = \text{Resp}(1^\kappa, c, (x, VK), \sigma)$ .
- Whenever  $\mathcal{P}^*$  outputs  $(x^*, a^*)$ , if  $x^* \in I$  output  $((x^*, VK), a^*)$ . Else pick a random  $x^* \leftarrow_s \{0, 1\}^\kappa \setminus I$  and return the instance  $(x^*, VK)$  and answer  $a^* := \text{Resp}(1^\kappa, c, (x^*, VK), \text{Sign}(SK_I, x^*))$ .

We define the extractor for  $\Pi'$  (w.r.t. the relation  $R$ ) to be the same as the extractor  $\mathcal{K}$  for  $\Pi$  (w.r.t. the relation  $R'$ ). It remains to bound the probability that  $\mathcal{K}$  output a valid witness for the relation  $R$ .

Let  $\text{Good}$  be the event that  $x^* \in I$  and all the  $x_i$ 's corresponding to  $\mathcal{P}^*$ 's verification queries are such that  $x_i \notin I$ . Moreover, let  $\text{Ext}_R$  (resp.  $\text{Ext}_{R'}$ ) be the event that  $(x, w) \in R$  (resp.  $((x, VK), w) \in R'$ ) where  $w$  comes from running the extractor  $\mathcal{K}$  in the definition of PZAP. We can write:

$$\begin{aligned}
\Pr[\text{Ext}_R] &\geq \Pr[\text{Ext}_R \wedge \text{Good}] & (5) \\
&\geq \Pr[\text{Ext}_{R'} \wedge \text{Good}] - \nu(\kappa) \\
&\geq \Pr[\text{Ext}_{R'}] - \Pr[\neg \text{Good}] - \nu(\kappa) \\
&\geq (\Pr[\mathcal{P}' \text{ succeeds}] - \nu'(\kappa)) - \Pr[\neg \text{Good}] - \nu(\kappa), & (6)
\end{aligned}$$

for negligible functions  $\nu(\cdot), \nu'(\cdot)$ . Here, Eq. (5) holds because of Lemma 3, whereas Eq. (6) follows by knowledge soundness of  $\Pi$ .

Observe that, by definition of  $\mathcal{P}_\alpha$ , the success probability when we condition on the event  $\text{Good}$  not happening is overwhelming (this is because in that case  $\mathcal{P}_\alpha$  just computes a valid signature, and thus it succeeds with overwhelming probability by completeness of  $\Pi$ ), therefore:

$$\Pr[\mathcal{P}_\alpha \text{ succeeds}] \geq \Pr[\mathcal{P}_\alpha \text{ succeeds} | \text{Good}] \cdot \Pr[\text{Good}] + (1 - \nu''(\kappa)) \Pr[\neg \text{Good}],$$

for some negligible function  $\nu''(\cdot)$ . Combining the last two equations, we obtain that there exists a negligible function  $\nu'''(\cdot)$  such that:

$$\Pr[\text{Ext}_R] \geq \Pr[\mathcal{P}_\alpha \text{ succeeds} | \text{Good}] \cdot \Pr[\text{Good}] - \nu'''(\kappa).$$

We analyse the probability that  $\mathcal{P}_\alpha$  succeeds conditioning on *Good* and the probability of event *Good* separately. We claim that the first term is negligibly close to the success probability of  $\mathcal{P}^*$ . In fact, when the event *Good* happens, by Lemma 2, the view generated by  $\mathcal{P}_\alpha$  is indistinguishable from the view in the knowledge soundness definition of PZAP.

As for the second term, again by Lemma 2, it is not hard to see that it is negligibly close to  $(1 - 2^{-\kappa+\alpha})^Q \cdot 2^{-\kappa+\alpha}$ , where  $Q$  is an upper bound for the number of verification queries made by the prover. Since when  $2^{-\kappa+\alpha} := 1 - Q/(Q + 1)$ , then  $(1 - 2^{-\kappa+\alpha})^Q \cdot 2^{-\kappa+\alpha} \geq 1/e$ , it suffices to set  $\alpha := \kappa + \log(1 - Q/(Q + 1))$  to enforce that the probability of *Good* is noticeable. This concludes the proof.

**Ruling-Out Challenge-Passing Reductions.** We show an impossibility result ruling out a broad class of black-box reductions from weak laconic PZAP to laconic PZAP. This negative result holds for so-called “challenge-passing” black-box reductions, which simply forward their input to the inner prover of the PZAP protocol.

**Theorem 8.** *Assume that pseudo-random generators exist, and let  $\Pi$  be laconic weak PZAP for NP. There is no challenge-passing black-box reduction from weak knowledge soundness to knowledge soundness of  $\Pi$ .*

The impossibility exploits the fact that oracle access to the verifier  $\mathcal{V}^*$  in the adaptive-knowledge soundness of laconic PZAP is equivalent to oracle access to a succeeding prover for the same relation. Consider the relation of pseudo-random string produced by a PRG  $G$ . The adversary can query the reduction with either a valid instance, namely a value  $x$  such that  $x = G(s)$  for  $s \leftarrow_s \{0, 1\}^\kappa$ , or an invalid instance  $x \leftarrow_s \{0, 1\}^{\kappa+1}$ . Notice that, since the reduction is black-box the two instances are indistinguishable, therefore a good reduction must be able to answer correctly both kind of instances. This allows us to use the reduction itself as a succeeding prover. We refer the reader to the full version of the paper [21] for the formal proof.

## 7 Conclusion and Open Problems

We initiated the study of Predictable Arguments of Knowledge (PAoK) systems for NP. Our work encompasses a full characterization of PAoK (showing in particular that they can without loss of generality assumed to be extremely laconic), provides several constructions of PAoK (highlighting that PAoK are intimately connected to witness encryption and program obfuscation), and studies PAoK with additional properties (such as zero-knowledge and Predictable ZAP).

Although, the notions of PAoK and Ext-WE are equivalent, we think that they give two different points of view on the same object. Ultimately, this can only give more insights.

There are several interesting questions left open by our work. First, one could try to see whether there are other ways (beyond the ones we explored in

the paper) how to circumvent the implausibility result of [23]. For instance it remains open if full-fledged PAoK for  $NP$  exist in the random oracle model.

Second, while it is impossible to have PAoK that additionally satisfy the zero-knowledge property in the plain model—in fact, we were able to achieve zero-knowledge in the CRS model and in the non-programmable random oracle model—such a negative result does not apply to witness indistinguishability. Hence, it would be interesting to construct PAoK that are additionally witness indistinguishable in the plain model. An analogous question holds for PZAP.

Third, we believe the relationship between the notions of weak PZAP (where the prover is not allowed any verification query) and PZAP deserves further study. Our impossibility result for basing PZAP on weak PZAP in a black-box way, in fact, only rules out very basic types of reductions (black-box, and challenge-passing), and additionally only works for laconic PZAP. It remains open whether the impossibility proof can be extended to rule-out larger classes of reductions for non-laconic PZAP, or if the impossibility can somehow be circumvented using non-black-box techniques.

## References

1. Hamza Abusalah, Georg Fuchsbauer, and Krzysztof Pietrzak. Offline witness encryption. *IACR Cryptology ePrint Archive*, 2015:838, 2015.
2. Prabhanjan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. Differing-inputs obfuscation and applications. *IACR Cryptology ePrint Archive*, 2013:689, 2013.
3. Dana Angluin and David Lichtenstein. Provable security of cryptosystems: A survey. Technical Report TR-288, Yale University, October 1983.
4. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6, 2012.
5. Mihir Bellare, Russell Impagliazzo, and Moni Naor. Does parallel repetition lower the error in computationally sound protocols? In *FOCS*, pages 374–383, 1997.
6. Mihir Bellare, Igors Stepanovs, and Stefano Tessaro. Poly-many hardcore bits for any one-way function and a framework for differing-inputs obfuscation. In *ASIACRYPT*, pages 102–121, 2014.
7. Mihir Bellare, Igors Stepanovs, and Brent Waters. New negative results on differing-inputs obfuscation. In *EUROCRYPT, Part II*, pages 792–821, 2016.
8. Fabrice Benhamouda, Olivier Blazy, Céline Chevalier, David Pointcheval, and Damien Vergnaud. New techniques for sphfs and efficient one-round PAKE protocols. In *CRYPTO*, pages 449–475, 2013.
9. Nir Bitansky, Ran Canetti, and Shai Halevi. Leakage-tolerant interactive protocols. In *TCC*, pages 266–284, 2012.
10. Dan Boneh and Xavier Boyen. Secure identity based encryption without random oracles. In *CRYPTO*, pages 443–459, 2004.
11. Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In *PKC*, pages 501–519, 2014.
12. Elette Boyle, Gil Segev, and Daniel Wichs. Fully leakage-resilient signatures. *J. Cryptology*, 26(3):513–558, 2013.

13. Ran Canetti, Shai Halevi, and Michael Steiner. Hardness amplification of weakly verifiable puzzles. In *TCC*, pages 17–33, 2005.
14. Kai-Min Chung and Feng-Hao Liu. Parallel repetition theorems for interactive arguments. In *TCC*, pages 19–36, 2010.
15. Kai-Min Chung and Rafael Pass. Tight parallel repetition theorems for public-coin arguments using KL-divergence. In *TCC*, pages 229–246, 2015.
16. David Derler and Daniel Slamanig. Practical witness encryption for algebraic languages and how to reply an unknown whistleblower. *IACR Cryptology ePrint Archive*, 2015:1073, 2015.
17. Cynthia Dwork and Moni Naor. ZAPs and their applications. *SIAM J. Comput.*, 36(6):1513–1543, 2007.
18. Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. In *Innovations in Computer Science*, pages 434–452, 2010.
19. Antonio Faonio and Jesper Buus Nielsen. Fully leakage-resilient codes. *IACR Cryptology ePrint Archive*, 2015:1151, 2015.
20. Antonio Faonio, Jesper Buus Nielsen, and Daniele Venturi. Mind your coins: Fully leakage-resilient signatures with graceful degradation. In *ICALP*, pages 456–468, 2015.
21. Antonio Faonio, Jesper Buus Nielsen, and Daniele Venturi. Predictable arguments of knowledge. Cryptology ePrint Archive, Report 2015/740, 2015. <http://eprint.iacr.org/2015/740>.
22. Sebastian Faust, Pratyay Mukherjee, Jesper Buus Nielsen, and Daniele Venturi. A tamper and leakage resilient von Neumann architecture. In *PKC*, pages 579–603, 2015.
23. Sanjam Garg, Craig Gentry, Shai Halevi, and Daniel Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In *CRYPTO*, pages 518–535, 2014.
24. Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In *STOC*, pages 467–476, 2013.
25. Sanjam Garg, Rafail Ostrovsky, Ivan Visconti, and Akshay Wadia. Resettable statistical zero knowledge. In *TCC*, pages 494–511, 2012.
26. Oded Goldreich. *The Foundations of Cryptography - Volume 1, Basic Techniques*. Cambridge University Press, 2001.
27. Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Inf. Process. Lett.*, 67(4):205–214, 1998.
28. Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *STOC*, pages 25–32, 1989.
29. Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *J. Cryptology*, 7(1):1–32, 1994.
30. Oded Goldreich, Salil P. Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1-2):1–53, 2002.
31. Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. How to run Turing machines on encrypted data. In *CRYPTO*, pages 536–553, 2013.
32. Iftach Haitner. A parallel repetition theorem for any interactive argument. *SIAM J. Comput.*, 42(6):2487–2501, 2013.
33. Johan Håstad, Rafael Pass, Douglas Wikström, and Krzysztof Pietrzak. An efficient parallel repetition theorem. In *TCC*, pages 1–18, 2010.
34. Dennis Hofheinz and Eike Kiltz. Programmable hash functions and their applications. In *CRYPTO*, pages 21–38, 2008.

35. Tal Malkin, Isamu Teranishi, Yevgeniy Vahlis, and Moti Yung. Signatures resilient to continual leakage on memory and computation. In *TCC*, pages 89–106, 2011.
36. Jesper Buus Nielsen, Daniele Venturi, and Angela Zottarel. On the connection between leakage tolerance and adaptive security. In *PKC*, pages 497–515, 2013.
37. Tatsuaki Okamoto and Kazuo Ohta. Divertible zero knowledge interactive proofs and commutative random self-reducibility. In *EUROCRYPT*, pages 134–148, 1989.
38. Rafael Pass and Muthuramakrishnan Venkitasubramaniam. An efficient parallel repetition theorem for Arthur-Merlin games. In *STOC*, pages 420–429, 2007.
39. Krzysztof Pietrzak and Douglas Wikström. Parallel repetition of computationally sound protocols revisited. *J. Cryptology*, 25(1):116–135, 2012.
40. Martin Tompa and Heather Woll. Random self-reducibility and zero knowledge interactive proofs of possession of information. In *FOCS*, pages 472–482, 1987.
41. Hoeteck Wee. Zero knowledge in the random oracle model, revisited. In *ASIACRYPT*, pages 417–434, 2009.
42. Hoeteck Wee. Efficient chosen-ciphertext security via extractable hash proofs. In *CRYPTO*, pages 314–332, 2010.
43. Mark Zhandry. How to avoid obfuscation using witness PRFs. In *TCC*, pages 421–448, 2016.