



Single-to-Multi-theorem Transformations for Non-interactive Statistical Zero-Knowledge

Marc Fischlin^(✉) and Felix Rohrbach

Cryptoplexity, Technische Universität Darmstadt, Darmstadt, Germany
{marc.fischlin,felix.rohrbach}@cryptoplexity.de
<http://www.cryptoplexity.de>

Abstract. Non-interactive zero-knowledge proofs or arguments allow a prover to show validity of a statement without further interaction. For non-trivial statements such protocols require a setup assumption in form of a common random or reference string (CRS). Generally, the CRS can only be used for one statement (single-theorem zero-knowledge) such that a fresh CRS would need to be generated for each proof. Fortunately, Feige, Lapidot and Shamir (FOCS 1990) presented a transformation for any non-interactive zero-knowledge proof system that allows the CRS to be reused any polynomial number of times (multi-theorem zero-knowledge). This FLS transformation, however, is only known to work for either computational zero-knowledge or requires a structured, non-uniform common reference string.

In this paper we present FLS-like transformations that work for non-interactive statistical zero-knowledge arguments in the common *random* string model. They allow to go from single-theorem to multi-theorem zero-knowledge and also preserve soundness, for both properties in the adaptive and non-adaptive case. Our first transformation is based on the general assumption that one-way permutations exist, while our second transformation uses lattice-based assumptions. Additionally, we define different possible soundness notions for non-interactive arguments and discuss their relationships.

Keywords: Non-interactive arguments · Statistical zero-knowledge · Soundness · Transformation · One-way permutation · Lattices · Dual-mode commitments

1 Introduction

In a non-interactive proof for a language \mathcal{L} the prover P shows validity of some theorem $x \in \mathcal{L}$ via a proof π based on a common string crs chosen by some external setup procedure. The common requirements are completeness—that the honest prover is able to convince the verifier V for true statements x —and soundness—that the verifier will not accept false statements $x \notin \mathcal{L}$ from

malicious provers. Blum et al. [5] showed that such non-interactive proofs can also be zero-knowledge [22], saying that a simulator can create a proof π on behalf of P if it has the ability to place some trapdoor information in crs.

1.1 Flavors of Non-interactive Zero-Knowledge

Non-interactive zero-knowledge protocols come in many variations:

- If the prover is computationally unbounded then one speaks of a NIZK *proof system* whereas in *arguments* or *argument systems* the prover runs in polynomial time [8].
- Zero-knowledge may be *computational* (NICZK) or *statistical* (NISZK) or even *perfect* (NIPZK). Note that non-interactive statistical (or perfect) zero-knowledge for NP requires that the prover is computationally bounded, unless the polynomial hierarchy collapses [31].
- The common string crs may be uniformly distributed over all bit strings of a certain length, in which case one speaks of the *common random string* or, less frequently, of the *uniform reference string* model. In any other case the string may have more structure and one calls it a *common reference string* or, sometimes, also *public parameter* model. In this work, we will focus on the case where the crs is uniformly distributed.

Another important aspect is the question of when malicious parties choose their challenge statement x . Both zero-knowledge and soundness come in an adaptive and in a non-adaptive version. The adaptive versions say that the adversary may choose the statement x after having seen the common reference string. For zero-knowledge this means that the simulator must prepare crs independently of x and then find a valid proof π after learning a maliciously chosen $x \in \mathcal{L}$. Adaptive soundness says that the malicious prover P^* first receives crs and then tries to find a false statement $x \notin \mathcal{L}$ with a convincing proof π .

Remarkably, for soundness one usually merely distinguishes between non-adaptive and adaptive notions. But there are also different ways how to capture the fact that a malicious prover P^* needs to succeed for an invalid statement $x \notin \mathcal{L}$. Either one assumes that the prover only outputs invalid statements, thus excluding some adversaries, or one penalizes the prover and declares it to lose if it chooses some $x \in \mathcal{L}$.¹ The penalizing definition implies the exclusive one. We note that Arte and Bellare [3], in a concurrent work, have proposed a similar distinction between exclusive and penalizing soundness.

Both notions, exclusive and penalizing soundness, already appeared implicitly in the literature, e.g., the work by Blum et al. [7] gives both an adaptive and a non-adaptive soundness definition in the exclusive setting. Indeed, non-adaptive soundness in the literature is often cast in this style. In contrast, for adaptive soundness nowadays one often encounters the penalizing variant. It seems,

¹ We use here the terminology from [4] for the comparable scenario of admissible decryption queries in chosen-ciphertext security.

however, that the adaptive/exclusive version is already sufficient for many applications, e.g., to build universally composable NIZK protocols [26]. We discuss this in more detail in Sect. 3 when defining the different versions.

1.2 From Single-theorem to Multi-theorem Proofs

In this work we focus on another important property of NIZK, namely, if the crs can be used only once (*bounded* or *single-theorem*) or is applicable for many proofs (*unbounded* or *multi-theorem*). The latter is of course preferable, and indeed Feige et al. [17, 18] show how to generally turn single-theorem NIZK proofs and arguments into multi-theorem zero-knowledge protocols. We call this the FLS-transformation.

The idea of the FLS-transformation is to augment the common random string by an extra uniformly distributed portion crs^{aux} and let the prover for this NP-language show that “ $x \in \mathcal{L}$ or crs^{aux} is the output of a pseudorandom generator”. This allows the simulator to create this part crs^{aux} pseudorandomly and use the generator’s seed as a witness for simulating the or-proof. If the original proof is zero-knowledge, then it is also witness indistinguishable [19], and then one cannot distinguish or-proofs generated by the genuine prover with the witness for x from proofs created by the simulator with the witness for crs^{aux} .

Soundness, on the other hand, is not affected because a random string crs^{aux} is not pseudorandom, except with exponentially small probability. Hence, for invalid x the “or” of the statements $x \notin \mathcal{L}$ or “ crs^{aux} is pseudorandom” would not be satisfied either with overwhelming probability. This implies that a prover would still need to break soundness of the or-protocol.

The FLS-transformation, per se, is only known to work for non-interactive *computational* zero-knowledge. The reason is that the pseudorandom string crs^{aux} of the zero-knowledge simulator is only computationally indistinguishable from a truly random string. There exists a folklore “dual version” of the FLS-transformation for non-interactive perfect (and therefore also statistical) zero-knowledge, where the crs contains a pseudorandom value by construction. But this transformation requires a structured, non-uniformly chosen crs, whereas we are interested in the setting of common *random* strings. For completeness, we provide a formal description of that folklore result along our terminology in the eprint version [20].

It is thus unclear if the FLS-transformation can be used equally smoothly for statistical zero-knowledge in the common random string model. For example, Peikert and Shiehian [32] recently presented a statistical zero-knowledge argument for NP based on LWE in the common random string model, which is only zero-knowledge for a single theorem. They therefore asked whether there is an FLS-like transformation to achieve multi-theorem zero-knowledge in the statistical case.

1.3 Known NISZK Constructions

There are only a few known constructions of NISZK and NIPZK protocols for the general class NP. Groth et al. [25, 26] were the first to give a NIPZK argument for NP based on specific number-theoretic constructions over bilinear groups. Their protocol achieves multi-theorem adaptive zero-knowledge, but only non-adaptive/exclusive soundness (although this can be extended to some limited form of adaptive soundness, called adaptive culpable soundness). It is cast in the common reference string model.

Abe and Fehr [1] later showed how to achieve NIPZK arguments for NP under some form of the knowledge-of-exponent assumption. Their protocol achieves adaptive multi-theorem zero-knowledge and is adaptively sound (in the penalizing setting). This protocol is again in the common reference string model.

Sahai and Waters [34] show how to build NIPZK arguments for NP based on indistinguishability obfuscation and one-way functions. Their solution is adaptive multi-theorem zero-knowledge and non-adaptively/exclusively sound. It is designed in the common reference string model.

Peikert and Shiehian [32] constructed NISZK arguments for NP based on the LWE assumption. Their construction is based on the NIZK framework of Canetti et al. [9, 10] as well as Holmgren and Lombardi [27] which, among others, constructs a non-adaptively/exclusively sound NISZK argument for NP in the common random string model. Their protocol is adaptively zero-knowledge for single theorems. The instantiation of Peikert and Shiehian [32] uses the LWE assumption to implement the primitives and inherits the characteristics of the solutions in [9, 10, 27].

An interesting observation, based on [11, Footnote 13], is that one should be able to show adaptive soundness for the constructions in [9] when using the exclusive notion. Noteworthy, Canetti et al. [11] merely claim non-adaptive soundness, because for the adaptive version they switch to the penalizing variant. They detail why this notion cannot be achieved with the current construction, and the point touches precisely the difference between penalizing and exclusive soundness. Reverting to adaptive/exclusive soundness, the construction may satisfy this weaker level. This gives the interesting twist that the solution by Peikert and Shiehian [32] may already be adaptively/exclusively sound, such that our transformation lifts it from single-theorem to multi-theorem (adaptive) zero-knowledge.

Libert et al. [29] recently showed how to build *designated-verifier* statistical zero-knowledge arguments based on the (kernel) k -linear assumption, and how this construction can also be turned into a public verifiable NISZK argument. Their public verifiable construction achieves multi-theorem zero-knowledge and non-adaptive/exclusive soundness in the common reference string model.

In another work, Libert et al. [28] achieve multi-theorem zero-knowledge in the common *random* string model. Their protocol provides non-adaptive/non-uniform soundness, i.e., where one quantifies over all inputs $x \notin \mathcal{L}$ and the crs is chosen as part of the experiment. We will later argue that in the non-

adaptive case this notion is equivalent to non-adaptive/exclusive and to non-adaptive/penalizing soundness for non-uniform provers (Fig. 1).

1.4 Our Results

In this work we show multiple FLS-SZK-transformations which preserve statistical zero-knowledge. Moreover, they allow to preserve non-adaptive or adaptive zero-knowledge and also inherit the adaptive security of soundness (in the exclusive variant). In detail, we show:

- For statistical zero-knowledge we show how to transform any single-theorem zero-knowledge NISZK argument for NP-languages into one which is a multi-theorem zero-knowledge NISZK argument in the common random string model. This requires only the existence of one-way permutations².
- For perfect zero-knowledge we show that our transformation can be augmented to preserve perfect zero-knowledge. This, however, comes at the cost of having a zero-knowledge simulator which runs in expected polynomial-time.
- Finally, we show that we can build a transformation for statistical zero-knowledge from the Learning with Errors (LWE) assumption in the common

Work	Soundness	CRS uniform?	ZK	Required
FLS* [17,18]	adaptive/ penalizing	✓	computational	PRGs
folklore*	adaptive/ exclusive	✗	perfect	PRGs
Groth et al. [25,26]	non-adaptive/ exclusive	✗	perfect	bilinear groups
Abe and Fehr [1]	adaptive/ penalizing	✗	perfect	knowledge-of-exponent
Sahai and Waters [34]	non-adaptive/ exclusive	✗	perfect	iO
Libert et al. [29]	non-adaptive/ exclusive	✗	statistical	k-linear
Libert et al. [28]	non-adaptive/ non-uniform	✓	statistical	LWE
this work*	adaptive/ exclusive	✓	statistical	OWP or LWE
	adaptive/ exclusive	✓	perfect	OWP+ expected simulation

Fig. 1. Comparison of different multi-theorem NIZK schemes. The entries marked with * are actually transformations for the single-to-multi-theorem cases.

² Note that we define one-way permutations as one-way functions that are 1-1 and length-preserving, not as a family of such functions.

random string model. This transformation, in contrast to the construction by Libert et al. [28], even works for *adaptively* sound NISZK arguments. This fits in nicely with the recent construction of statistical zero-knowledge arguments based on LWE [32].

- Additionally, we define and discuss the different soundness properties for non-interactive arguments and analyze their relationship. In particular, we show that in the non-adaptive case, the notions of exclusive, penalizing, and non-uniform soundness are all equivalent when considering non-uniform provers.

Our techniques for the constructions based on general assumptions uses a “dual” version of the original FLS-transformation. That is, instead of building the or-language for crs^{aux} being pseudorandom, we use that crs^{aux} is *not* pseudorandom. Since this is in general a coNP -language we need to make sure that it is also in NP . We achieve this by using the Blum-Micali-Yao pseudorandom generator [6,35] based on one-way permutations and hardcore bits, which lies in $\text{NP} \cap \text{coNP}$. Soundness for our dual FLS-transformation then follows since we can let the malicious prover run on a pseudorandom string crs^{aux} instead, since this is indistinguishable for the efficient prover in an argument. Then the or of the two statements, $x \in \mathcal{L}$ or crs^{aux} is not pseudorandom, is again not satisfied.

The construction based on LWE is inspired by a primitive called dual-mode commitment scheme, i.e., a commitment which can be either perfectly-binding or statistically-hiding, based on the choice of how to generate the public key. The public keys for both modes are computationally indistinguishable. We note that the usefulness of such dual-mode commitments for non-interactive zero-knowledge is well known, starting with the work by Groth et al. [25] where this technique was called parameter switching, to recent efforts like the construction of Libert et al. [29]. Most times, however, the solutions work over certain structures and yield arguments in the common reference string model.

Here, we use a construction of Gorbunov et al. [23] to build these dual-mode commitments where the (statistically-hiding) public key and a commitment can be chosen as uniform bit strings. As in the FLS-transformation we extend the CRS by a public key string pk and a random commitment string c and extend the language to “ $x \in \mathcal{L}$ or c is a commitment to 1”. For the simulator, we choose our public key to be statistically-hiding. In our construction, a statistically-hiding public key will be statistically close to a uniformly random string and indeed generate a commitment to the value 1.

However, for the soundness game we exchange the public key pk for a perfectly-binding one and change the commitment to 0, thereby forcing the malicious prover to prove x to be in \mathcal{L} . We emphasize that we only switch between these modes and merely require computational indistinguishability of the different types of public keys. In particular, we do not need to rely on the SIS assumption as considered in [23] but, as pointed out in [13], the LWE assumption suffices. Indeed, one could directly use Regev’s LWE encryption scheme [33] which also supports a statistically-hiding, lossy mode.

1.5 Squeezing in into Possibility and Impossibility Results

There are some known impossibility results for statistical and perfect zero-knowledge arguments. Strictly speaking, these results do not infringe with our results here, since we show how to *transform* statistical zero-knowledge arguments (from single to multiple theorems) but do not give constructions. Still, one may wonder if the combination of our transformations with the impossibility results have any implications on potential constructions.

Abe and Fehr [1] were the first to show that NISZK arguments cannot be proven to be adaptively sound via so-called direct black-box reductions, unless the language is in P/poly . One property which such direct reductions has is that one can use an efficient alternative to the crs generator which in addition outputs the simulator's trapdoor information (property II.(b) in [1]). Our construction, however, bypasses this property because for the soundness proof it generates a bad crs which does not have a trapdoor. In this sense, our technique indicates that the notion of direct black-box reductions may be too restrictive.

Pass [30], using similar ideas and techniques as [1], shows that *adaptive* statistical and perfect zero-knowledge arguments with *adaptive* soundness cannot be based on hard primitives via black-box reductions. How does the result of Pass [30] match our results? First we remark that our NIPZK is indeed *adaptively* sound and *adaptively* zero-knowledge. But the simulator only runs in polynomial time averaged over its internal randomness. Such simulators escape the results in [30].

Yet, the most striking difference between the results in [1,30] and our transformations lies in the distinct notions of adaptive soundness. We show that our transformations preserve adaptive/exclusive soundness. Opposite to that, the impossibility results of [1,30] rely on the ability of the malicious prover to occasionally output theorems $x \in \mathcal{L}$. Put differently, they rule out the stronger form of adaptive/penalizing sound arguments, whereas we argue that adaptive/exclusive soundness is preserved. As remarked above, however, adaptive/exclusive sound arguments may still be sufficient for applications.

1.6 Concurrent Work

As mentioned earlier, Arte and Bellare [3] have touched upon the issue of different soundness notions in non-interactive proofs as well. Their starting point are dual-mode systems in which the common reference string can be generated in two modes, and in how far such systems allow for transference of security properties in the different modes. Our work instead focuses on the transformations for multi-theorem statistical zero-knowledge arguments.

Arte and Bellare define notions of penalizing and exclusive soundness, called SND-P and SND-E, with which our adaptive notions for soundness coincide (for efficient provers).³ Remarkably, they show a separating example of their exclusive and penalizing soundness notion in the adaptive case, under the decisional

³ Strictly speaking, their notion of exclusiveness allows for a negligible error which could be integrated in our notion as well.

Diffie-Hellman assumption. This example applies to our notions in the adaptive setting as well. We complement this result by showing that the notions are equivalent in the non-adaptive case, assuming non-uniform provers.

Another notable difference between the two works lies in the applications of the different soundness notions. Arte and Bellare discuss the example of the Bellare-Goldwasser signature scheme where penalizing soundness is required and exclusive soundness is insufficient. We argue along the implication of culpability that exclusive soundness may suffice in many settings.

2 Preliminaries

An NP-relation \mathcal{R} consists of pairs (x, ω) of theorems and witnesses where the length of witness is polynomially bounded in the length of the theorem, and where one can efficiently decide membership. More formally, there exists a polynomial-time Turing machine $M_{\mathcal{R}}$ and a polynomial $p_{\mathcal{R}}$ such that

$$\mathcal{R} = \{(x, \omega) \mid |\omega| \leq p_{\mathcal{R}}(|x|) \wedge M_{\mathcal{R}}(x, \omega) = 1\}.$$

The induced language $\mathcal{L}_{\mathcal{R}}$ is given by

$$\mathcal{L}_{\mathcal{R}} = \{x \in \{0, 1\}^* \mid \exists \omega : (x, \omega) \in \mathcal{R}\}.$$

2.1 Non-interactive Arguments

A non-interactive argument or proof system for an NP-relation is now a protocol in which the setup algorithm **Setup** generates a common string crs which the prover **P** then uses to generate a proof π for the input (x, ω) . The verifier **V** then checks this proof against crs and x only. There are some length restrictions, of course, namely that the length of the theorem x determines the length of the common string. In particular, we assume that there is a polynomial p_{Setup} such that $\text{crs} \in \{0, 1\}^{p_{\text{Setup}}(n)}$ for any $\text{crs} \xleftarrow{\$} \text{Setup}(1^n)$. Let $\mathcal{R}(1^n) = \{(x, \omega) \in \mathcal{R} \mid |x| = n\}$ and $\mathcal{L}_{\mathcal{R}}(1^n) = \{x \in \mathcal{L}_{\mathcal{R}} \mid |x| = n\}$ denote the restriction of inputs of the relation and language with length $|x| = n$ such that the length of the common string for such inputs is given by $p_{\text{Setup}}(n)$. Note that the verifier can easily check that $|x|$ matches the security parameter n such that we can assume that this is always the case.

We note that the string crs generated by **Setup** may be uniformly distributed, in which case we speak of a common random string. It may have a different distribution, in which case we call it a common reference string. In particular, we see a common random string as a special case of a common reference string.

The usual completeness notion of non-interactive arguments and proofs asks that the verifier **V** accepts genuine proofs π generated by the prover **P** for input $x \in \mathcal{L}_{\mathcal{R}}$. Soundness, on the hand, demands that the verifier does not accept false proofs generated by a malicious prover \mathbf{P}^* for inputs $x \notin \mathcal{L}_{\mathcal{R}}$. As explained in the introduction there are various possibilities to define soundness, which we will discuss in Sect. 3, and just use one example of the possible definitions here.

Definition 1 (Non-interactive Argument). A non-interactive argument for an NP-relation \mathcal{R} (in the common reference string model) is a triple of probabilistic polynomial-time algorithms $\Pi = (\text{Setup}, P, V)$ satisfying the completeness and soundness condition:

(Perfect) Completeness: For every $n \in \mathbb{N}$, every $(x, \omega) \in \mathcal{R}(1^n)$, every $\text{crs} \xleftarrow{\$} \text{Setup}(1^n)$, every $\pi \xleftarrow{\$} P(1^n, x, \omega, \text{crs})$ we have that $V(1^n, x, \pi, \text{crs}) = 1$ with probability 1.

(Non-adaptive/Exclusive) Soundness: For every (possibly malicious) probabilistic polynomial-time prover P^* outputting only $x \notin \mathcal{L}_{\mathcal{R}}$ there exists a negligible function $\epsilon(n)$ such that for every $n \in \mathbb{N}$ we have

$$\Pr[V(1^n, x, \pi, \text{crs}) = 1] \leq \epsilon(|x|),$$

where the probability is over $(x, st) \xleftarrow{\$} P^*(1^n)$, $\text{crs} \xleftarrow{\$} \text{Setup}(1^n)$, as well as $\pi \xleftarrow{\$} P^*(1^n, st, \text{crs})$, and V 's randomness.

We say that the argument is in the common random string model if $\text{Setup}(n)$ outputs uniformly distributed strings over $\{0, 1\}^{p_{\text{Setup}}(n)}$ for every $n \in \mathbb{N}$.

2.2 Zero-knowledge

We next define zero-knowledge with the usual notion of a simulator ZKSim . In the non-interactive setting this algorithm has the advantage to choose the common string crs to simulate proofs. In the bounded case the distinguisher only gets to see a single proof for a chosen theorem, where the proof is either genuine or fabricated by the simulator. We simultaneously define the single-theorem and multi-theorem case where the distinguisher learns one or many (genuine or simulated) proofs. We first define both cases in the adaptive setting where the distinguisher selects the theorems in dependence of the common string and of previous proofs and in the non-adaptive case where the distinguisher chooses the statement(s) in advance. We stress that we are interested in statistical zero-knowledge here such that the distinguisher is unbounded, except that it can only ask for polynomially many proofs. We also allow the simulator to run in *expected* polynomial time in specially marked cases.

Definition 2 (Statistical and Perfect Zero Knowledge). Let \mathcal{R} be an NP-relation and let $\Pi = (\text{Setup}, P, V)$ be a non-interactive argument for \mathcal{R} . The argument is zero-knowledge if it satisfies one of the following properties:

Non-adaptive Multi-theorem Zero-knowledge: For any unbounded algorithm D there exists a probabilistic algorithm ZKSim , the simulator, running in (expected) polynomial time, such that the advantage

$$\text{Adv}_{\Pi, \text{ZKSim}, D}^{\text{naSZK}}(1^n) := \Pr \left[\text{Expt}_{\Pi, \text{ZKSim}, D}^{\text{naSZK}}(1^n) = 1 \right] - \frac{1}{2}$$

is negligible for polynomially bounded q , where experiment $\text{Expt}_{\Pi, \text{ZKSim}, D}^{\text{naSZK}}(1^n)$ is defined in Fig. 2. If the advantage of any such D is always 0 then the argument is called perfect zero-knowledge.

<u>Expt_{II,ZKSim,D}^{naSZK}(1ⁿ):</u>	<u>Expt_{(Setup,P,V),ZKSim,D}^{aSZK}(1ⁿ):</u>
1 $b \xleftarrow{\$} \{0, 1\}$	1 $b \xleftarrow{\$} \{0, 1\}, q \leftarrow 0, \text{st}_D \leftarrow \perp$
2 $(\text{st}_D, x_1, \omega_1, \dots, x_q, \omega_q) \xleftarrow{\$} D(1^n)$	2 $\text{crs}_0 \xleftarrow{\$} \text{Setup}(1^n)$
3 $\text{crs}_0 \xleftarrow{\$} \text{Setup}(1^n)$	3 $(\text{crs}_1, \text{st}_{\text{ZKSim}}) \xleftarrow{\$} \text{ZKSim}(1^n)$
4 $(\text{crs}_1, \text{st}_{\text{ZKSim}}) \xleftarrow{\$} \text{ZKSim}(1^n)$	4 repeat
5 for $i = 1..q$ do	5 $q \leftarrow q + 1$
6 if $(x_i, \omega_i) \in \mathcal{R}$ then	6 $(\text{st}_D, x, \omega) \xleftarrow{\$} D(1^n, \text{st}_D, \text{crs}_b)$
7 $\pi_{i,0} \xleftarrow{\$} P(1^n, x_i, \omega_i, \text{crs}_0)$	7 if $(x, \omega) \in \mathcal{R}$ then
8 $\pi_{i,1} \xleftarrow{\$} \text{ZKSim}(1^n, \text{st}_{\text{ZKSim}}, x_i)$	8 $\pi_0 \xleftarrow{\$} P(1^n, x, \omega, \text{crs}_0)$
9 else $\pi_{i,0} \leftarrow \pi_{i,1} \leftarrow \perp$	9 $\pi_1 \xleftarrow{\$} \text{ZKSim}(1^n, \text{st}_{\text{ZKSim}}, x)$
10 $d \xleftarrow{\$} D(1^n, \text{st}_D, \pi_{1,b}, \dots, \pi_{q,b}, \text{crs}_b)$	10 else $\pi_0 \leftarrow \pi_1 \leftarrow \perp$
11 return $b = d$	11 $(\text{st}_D, \text{cont}, d) \xleftarrow{\$} D(1^n, \text{st}_D, \pi_b)$
	12 until $\text{cont} = \text{false}$
	13 return $b = d$

Fig. 2. Non-adaptive and adaptive statistical zero-knowledge experiments.

Adaptive Multi-theorem Zero Knowledge: For any unbounded algorithm D there exists a probabilistic algorithm ZKSim , the simulator, running in (expected) polynomial time, such that the advantage

$$\text{Adv}_{II, \text{ZKSim}, D}^{\text{aSZK}}(1^n) := \Pr \left[\text{Expt}_{II, \text{ZKSim}, D}^{\text{aSZK}}(1^n) = 1 \right] - \frac{1}{2}$$

is negligible for polynomially bounded q , where experiment $\text{Expt}_{II, \text{ZKSim}, D}^{\text{aSZK}}(1^n)$ is defined in Fig. 2. If the advantage of any such D is always 0 then the argument is called perfect zero-knowledge.

The argument is single-theorem zero-knowledge of the corresponding type if the property holds for $q = 1$.

Definition 3 (Statistical Witness Indistinguishability). Let \mathcal{R} be an NP-relation. A non-interactive argument $\Pi = (\text{Setup}, P, V)$ for \mathcal{R} is called statistical witness indistinguishable (NISWI) if it satisfies one of the following properties:

Non-adaptive Multi-theorem Witness Indistinguishability: For any unbounded algorithm D the advantage

$$\text{Adv}_{II, D}^{\text{naSWI}}(1^n) := \Pr \left[\text{Expt}_{II, D}^{\text{naSWI}}(1^n) = 1 \right] - \frac{1}{2}$$

is negligible for polynomially bounded q , where the experiment $\text{Expt}_{II, D}^{\text{naSWI}}(1^n)$ is defined in Fig. 3. If the advantage of any such D is always 0 then the argument is called perfect witness indistinguishable.

$\text{Expt}_{\Pi, D}^{\text{naSWI}}(1^n)$:	$\text{Expt}_{\Pi, D}^{\text{aSWI}}(1^n)$:
1 $b \stackrel{\$}{\leftarrow} \{0, 1\}$	1 $b \stackrel{\$}{\leftarrow} \{0, 1\}, \mathbf{q} \leftarrow 0, \text{st}_D \leftarrow \perp$
2 $(\text{st}_D, (x_i, \omega_{i,0}, \omega_{i,1})_{i=1..q}) \stackrel{\$}{\leftarrow} D(1^n)$	2 $\text{crs} \stackrel{\$}{\leftarrow} \text{Setup}(1^n)$
3 $\text{crs} \stackrel{\$}{\leftarrow} \text{Setup}(1^n)$	3 repeat
4 for $i = 1..q$ do	4 $\mathbf{q} \leftarrow \mathbf{q} + 1$
5 if $(x_i, \omega_{i,0}) \in \mathcal{R} \wedge (x_i, \omega_{i,1}) \in \mathcal{R}$	5 $(\text{st}_D, x, \omega_0, \omega_1) \stackrel{\$}{\leftarrow} D(1^n, \text{st}_D, \text{crs}_b)$
6 $\pi_{i,0} \stackrel{\$}{\leftarrow} P(1^n, x_i, \omega_{i,0}, \text{crs})$	6 if $(x, \omega_0) \in \mathcal{R} \wedge (x, \omega_1) \in \mathcal{R}$
7 $\pi_{i,1} \stackrel{\$}{\leftarrow} P(1^n, x_i, \omega_{i,1}, \text{crs})$	7 $\pi_0 \stackrel{\$}{\leftarrow} P(1^n, x, \omega_0, \text{crs})$
8 else $\pi_{i,0} \leftarrow \pi_{i,1} \leftarrow \perp$	8 $\pi_1 \stackrel{\$}{\leftarrow} P(1^n, x, \omega_1, \text{crs})$
9 $d \stackrel{\$}{\leftarrow} D(1^n, \text{st}_D, \pi_{1,b}, \dots, \pi_{q,b}, \text{crs})$	9 else $\pi_0 \leftarrow \pi_1 \leftarrow \perp$
10 return $b = d$	10 $(\text{st}_D, \text{cont}, d) \stackrel{\$}{\leftarrow} D(1^n, \text{st}_D, \pi_b)$
	11 until $\text{cont} = \text{false}$
	12 return $b = d$

Fig. 3. Non-adaptive and adaptive statistical witness indistinguishability experiments.

Adaptive Multi-theorem Witness Indistinguishability: *For any unbounded algorithm D the advantage*

$$\text{Adv}_{\Pi, D}^{\text{aSWI}}(1^n) := \Pr \left[\text{Expt}_{\Pi, D}^{\text{aSWI}}(1^n) = 1 \right] - \frac{1}{2}$$

is negligible for polynomially bounded q , where the experiment $\text{Expt}_{\Pi, D}^{\text{aSWI}}(1^n)$ is defined in Fig. 3. If the advantage of any such D is always 0 then the argument is called perfect witness indistinguishable.

The argument is single-theorem witness indistinguishable of the corresponding type if the property holds for $q = 1$.

2.3 From Single-Theorem Zero-Knowledge to Multi-Theorem Witness Indistinguishability

We repeat here the well known fact that zero-knowledge implies witness indistinguishability, and that witness indistinguishability is closed under repetitions [19]. We state the results here for sake of completeness and according to our terminology in the statistical setting.

Lemma 1. *Any adaptive resp. non-adaptive single-theorem NISZK argument is also an adaptive resp. non-adaptive single-theorem NISWI argument.*

Proof (Sketch). We only argue the adaptive case; the non-adaptive case follows analogously. We can perform a game hop starting with the witness-indistinguishability experiment $\text{Expt}_{\Pi, D}^{\text{aSWI}}(1^n)$. In this hop we replace the CRS and both proofs π_0 and π_1 in each iteration by simulated ones, all created by the simulator ZKSim

without knowledge of the witnesses ω_0 and ω_1 but using the same trapdoor. Note that we can view the proofs in the WI experiment as two sequentially requested proofs in the ZK experiment, such that the SZK property ensures that this hop is statistically indistinguishable. (In the non-adaptive case we would split each entry $(x_i, \omega_{i,0}, \omega_{i,1})$ in D 's initial choice into two entries $(x_i, \omega_{i,0})$ and $(x, \omega_{i,1})$.)

But now both proofs π_0 and π_1 are created without the specific witness, and since the simulator does not update its state for giving proofs, the order in which the proofs are computed is irrelevant. In this case the bit b is perfectly hidden from the distinguisher such that the advantage in predicting b is 0. \square

Lemma 2. *Any adaptive resp. non-adaptive single-theorem NISWI argument is also an adaptive resp. non-adaptive multi-theorem NISWI argument.*

Proof (Sketch). We again only discuss the adaptive case since the non-adaptive case follows analogously. The proof follows by a hybrid argument. For this we reduce the multi-theorem distinguisher D to a bounded one D_1 which only makes one query. Let $Q(n)$ be a polynomial upper bound on the number of queries q which D makes. The bounded distinguisher D_1 initially picks an index $i \xleftarrow{\$} \{1, 2, \dots, Q(n)\}$ and then internally runs in the first stage (Line 5) the distinguisher D up to the i -th query $(st_D, x, \omega_0, \omega_1)$. All requested proofs up to this step are computed internally by D_1 via P and the left witness, and returned to D . The i -th query is then computed externally, and D_1 then hands the proof back to D . In the final steps till halting, D_1 computes the remaining proofs for ω_1 , and eventually returns D 's decision bit d unchanged.

It can be shown that the advantage of the bounded distinguisher D_1 is at most a factor $Q(n)$ larger than the one of D . Since $Q(n)$ is polynomial, the difference is negligible. \square

3 Soundness of Non-interactive Arguments

Soundness of a non-interactive argument assures that a (computationally-bound) malicious prover is unable to convince the verifier of a false statement. Commonly, soundness is defined in two variants: Adaptive soundness, which allows the (possibly malicious) prover P^* to choose the statement to prove x before seeing the common random string crs , and non-adaptive soundness, in which the prover P^* has to decide on the statement x before the common random string crs is generated.

Remarkably, there is another dimension of definitional choice for soundness which often goes unnoticed in the literature. This dimension refers to the question how we measure success of the malicious prover. Clearly, the malicious prover should not make the verifier accept for a statement x not in the language. But there are two possibilities to capture the non-membership requirement. One is to disallow P^* to output $x \in \mathcal{L}$ at all. The other one is to declare P^* to lose if it picks $x \in \mathcal{L}$. Following the work of Bellare et al. [4] about the question how to deal with inadmissible decryption queries in CCA-secure encryption schemes, we call the former stipulation of P^* outputting only $x \notin \mathcal{L}$ *exclusive*, because it

excludes certain adversaries. The latter is called *penalizing* as it punishes P^* if it chooses $x \in \mathcal{L}$.

3.1 Soundness Definitions

In total, we define five soundness notions: adaptive vs. non-adaptive, and exclusive vs. penalizing, as well as a non-uniform variant that only exists for the non-adaptive case. We typically speak of non-adaptive/exclusive and adaptive/penalizing soundness etc. to distinguish the different types. Figure 4 provides an overview. It is also easy to see that adaptive soundness implies non-adaptive soundness in both settings, and penalizing soundness implies exclusive soundness in any of the other dimensions. The latter is easy to see because any malicious prover P^* breaking exclusive soundness must output $x \notin \mathcal{L}$ such that this prover also satisfies the winning condition in the penalizing setting. In this chapter, we highlight the further connections between these definitions and their implications.

The difference between exclusive and penalizing soundness may appear to be insignificant. Indeed, for non-interactive *proofs* it is folklore to show that the weakest one of the five notions, non-adaptive/exclusive soundness, implies the strongest one, adaptive/penalizing soundness. See for instance [21]. This may explain why today’s literature mostly distinguishes between the (exclusive) non-adaptive notion and the (penalizing) adaptive notion. An exception is the seminal paper by Blum et al. [7] which defines the adaptive version according to the exclusive dimension (without using our terminology here, of course). We emphasize, however, that the equivalence of all notions is not known to hold for non-interactive *arguments*.

Is a more fine-grained distinction between exclusive and penalizing soundness in arguments necessary? We argue that it is. Roughly, the difference is that in the exclusive case the malicious prover (and any other party) knows that its output is not in the language, in the penalizing case even the prover may itself be oblivious about this. This is an important ingredient in Pass’ impossibility result to build adaptive sound and adaptive statistical zero-knowledge arguments based on black-box reductions [30]. The result crucially relies on the malicious prover choosing a (random or pseudorandom) statement for which it does not know the status. In other words, this impossibility result rules out the strongest form of adaptive/penalizing soundness.

We next argue that the weaker form of adaptive/exclusive soundness is very relevant. It is easy to see that this notion implies a slightly weaker notion of adaptive/*culpable* soundness [26]. This notion is similar to our definition of adaptive/exclusive soundness, but also requires the malicious prover to output an efficiently verifiable witness (denoted ω_{guilt} in [26]) that the statement x is *not* in the language \mathcal{L} . Our exclusive notion asks P^* to output $x \notin \mathcal{L}$. We prove the implication that adaptive/exclusive yields adaptive/culpable soundness formally in Sect. 3.3.

The noteworthy fact is that adaptive/culpable soundness suffices for many applications. One of the most important ones is the possibility to derive

<i>non-adaptive</i>	<i>adaptive</i>
<p><u>exclusive soundness:</u></p> <ol style="list-style-type: none"> 1 // only P^* outputting $x \notin \mathcal{L}$ 2 $(x, st_{P^*}) \stackrel{\\$}{\leftarrow} P^*(1^n)$ 3 $crs \stackrel{\\$}{\leftarrow} \text{Setup}(1^n)$ 4 $\pi \stackrel{\\$}{\leftarrow} P^*(1^n, crs, st_{P^*})$ 5 return $V(1^n, x, \pi, crs)$ <p><u>penalizing soundness:</u></p> <ol style="list-style-type: none"> 1 $(x, st_{P^*}) \stackrel{\\$}{\leftarrow} P^*(1^n)$ 2 $crs \stackrel{\\$}{\leftarrow} \text{Setup}(1^n)$ 3 $\pi \stackrel{\\$}{\leftarrow} P^*(1^n, crs, st_{P^*})$ 4 return $V(1^n, x, \pi, crs) \wedge x \notin \mathcal{L}$ <p><u>non-uniform soundness:</u></p> <ol style="list-style-type: none"> 1 // for all $x \notin \mathcal{L}$ 2 $crs \stackrel{\\$}{\leftarrow} \text{Setup}(1^n)$ 3 $\pi \stackrel{\\$}{\leftarrow} P^*(1^n, crs, x)$ 4 return $V(1^n, x, \pi, crs)$ 	<p><u>exclusive soundness:</u></p> <ol style="list-style-type: none"> 1 // only P^* outputting $x \notin \mathcal{L}$ 2 3 $crs \stackrel{\\$}{\leftarrow} \text{Setup}(1^n)$ 4 $(x, \pi) \stackrel{\\$}{\leftarrow} P^*(1^n, crs)$ 5 return $V(1^n, x, \pi, crs)$ <p><u>penalizing soundness:</u></p> <ol style="list-style-type: none"> 1 2 $crs \stackrel{\\$}{\leftarrow} \text{Setup}(1^n)$ 3 $(x, \pi) \stackrel{\\$}{\leftarrow} P^*(1^n, crs)$ 4 return $V(1^n, x, \pi, crs) \wedge x \notin \mathcal{L}$

Fig. 4. Different notions of soundness.

universally composable NIZK argument [26]. Other applications include correctness proofs for shuffles [14, 15, 24] or for e-voting [12]. Since adaptive/exclusive soundness implies adaptive/culpable soundness, any protocol satisfying the exclusive notion is also applicable in such settings.

We can now define arguments with the different soundness properties:

Definition 4 (Soundness of non-interactive Arguments). *A non-interactive argument for an NP-relation \mathcal{R} (in the common reference string model) is a triple of probabilistic polynomial-time algorithms $\Pi = (\text{Setup}, P, V)$ satisfying the completeness as well as at least one of the soundness conditions:*

Non-adaptive/Exclusive Soundness: *For every (possibly malicious) probabilistic polynomial-time prover P^* outputting only $x \notin \mathcal{L}_{\mathcal{R}}$ there exists a negligible function $\epsilon(n)$ such that for every $n \in \mathbb{N}$ we have*

$$\Pr[V(1^n, x, \pi, crs) = 1] \leq \epsilon(|x|),$$

where the probability is over $(x, st) \xleftarrow{\$} P^*(1^n)$, $crs \xleftarrow{\$} \text{Setup}(1^n)$, as well as $\pi \xleftarrow{\$} P^*(1^n, st, crs)$, and V 's randomness.

Non-adaptive/Penalizing Soundness: For every (possibly malicious) probabilistic polynomial-time prover P^* there exists a negligible function $\epsilon(n)$ such that for every $n \in \mathbb{N}$ we have

$$\Pr[V(1^n, x, \pi, crs) = 1 \wedge x \notin \mathcal{L}_{\mathcal{R}}] \leq \epsilon(|x|),$$

where the probability is over $(x, st) \xleftarrow{\$} P^*(1^n)$, $crs \xleftarrow{\$} \text{Setup}(1^n)$, as well as $\pi \xleftarrow{\$} P^*(1^n, st, crs)$, and V 's randomness.

Adaptive/Exclusive Soundness: For every (possibly malicious) probabilistic polynomial-time prover P^* outputting only $x \notin \mathcal{L}_{\mathcal{R}}$ there exists a negligible function $\epsilon(n)$ such that for every $n \in \mathbb{N}$ we have

$$\Pr[V(1^n, x, \pi, crs) = 1] \leq \epsilon(|x|),$$

where the probability is over $crs \xleftarrow{\$} \text{Setup}(1^n)$, $(x, \pi) \xleftarrow{\$} P^*(1^n, crs)$, and V 's randomness.

Adaptive/Penalizing Soundness: For every (possibly malicious) probabilistic polynomial-time prover P^* there exists a negligible function $\epsilon(n)$ such that for every $n \in \mathbb{N}$ we have

$$\Pr[V(1^n, x, \pi, crs) = 1 \wedge x \notin \mathcal{L}_{\mathcal{R}}] \leq \epsilon(|x|),$$

where the probability is over $crs \xleftarrow{\$} \text{Setup}(1^n)$, $(x, \pi) \xleftarrow{\$} P^*(1^n, crs)$, and V 's randomness.

Non-adaptive/Non-uniform Soundness: For every (possibly malicious) probabilistic polynomial-time prover P^* there exists a negligible function $\epsilon(n)$ such that for every $n \in \mathbb{N}$ and every $x \notin \mathcal{L}_{\mathcal{R}}$ with $|x| = n$, we have

$$\Pr[V(1^n, x, \pi, crs) = 1 \wedge x \notin \mathcal{L}_{\mathcal{R}}] \leq \epsilon(|x|),$$

where the probability is over $crs \xleftarrow{\$} \text{Setup}(1^n)$, and $\pi \xleftarrow{\$} P^*(1^n, x, crs)$, and V 's randomness.

3.2 Equivalence of the Non-adaptive Soundness Notions

We now show that the non-adaptive soundness definitions are all equivalent if we allow the malicious provers to be non-uniform:

Theorem 1. For non-uniform (malicious) provers, a non-interactive argument $\Pi = (\text{Setup}, P, V)$ has non-adaptive/exclusive soundness iff it has non-adaptive/non-uniform soundness, and has non-adaptive/non-uniform soundness iff it has non-adaptive/penalizing soundness.

Proof. Non-adaptive/exclusive soundness follows directly from non-adaptive/penalizing soundness, therefore we only show that non-adaptive/non-uniform soundness follows from non-adaptive/exclusive soundness and that non-adaptive/penalizing soundness follows from non-adaptive/non-uniform soundness.

We start by showing non-adaptive/non-uniform soundness follows from non-adaptive/exclusive soundness. Let $\Pi = (\text{Setup}, P, V)$ be the non-interactive argument in question. Assume that there exists a successful malicious prover $P_{na/nu}^*$ against the non-adaptive/non-uniform soundness, i.e., for any negligible function $\epsilon(n)$ there exists an $x \notin \mathcal{L}$ such that

$$\Pr \left[V(\text{crs}, x, P_{na/nu}^*(\text{crs}, x)) \right] > \epsilon(|x|),$$

where the probability is over $\text{crs} \xleftarrow{\$} \text{Setup}(1^n)$, as well as $P_{na/nu}^*$'s and V 's randomness. We can now construct a malicious prover $P_{na/ex}^*$ against non-adaptive/exclusive soundness as follows: We define the first-stage algorithm $P_{na/ex,1}^*(1^n)$ to choose $x \notin \mathcal{L}$ of length n non-uniformly, such that $P_{na/nu}^*$'s success probability is maximized. The state st is left empty. Further, the second-stage algorithm $P_{na/ex,2}^*$ merely calls $P_{na/nu}^*$ internally, ignoring the state st . Then, the success probability of $P_{na/ex}^*$ is at least as large as the one of $P_{na/nu}^*$ and thus non-negligible.

Next, we show that non-adaptive/penalizing soundness follows from non-adaptive/non-uniform soundness. Assume that there exists a successful malicious prover $P_{na/pn}^*$ against the non-adaptive/penalizing soundness, i.e., for any negligible function ϵ there exists an $n \in \mathbb{N}$ such that

$$\Pr [V(\text{crs}, x, \pi) = 1 \wedge x \notin \mathcal{L}] > \epsilon(n),$$

where the probability is over $(x, st) \xleftarrow{\$} P_{na/pn,1}^*(1^n)$, $\text{crs} \xleftarrow{\$} \text{Setup}(1^n)$, $\pi \xleftarrow{\$} P_{na/pn,2}^*$ as well as V 's internal randomness.

We can now construct a malicious prover $P_{na/nu}^*$ against non-adaptive/non-uniform soundness as follows: For each input length n , we fix the pair (\bar{x}, \bar{st}) , $\bar{x} \in \{0, 1\}^n$, $\bar{x} \notin \mathcal{L}$, on which $P_{na/pn,2}^*$'s success probability is maximized (we bound the length of \bar{st} by $P_{na/pn,1}^*$'s running time). Next we define $P_{na/nu}^*$ as follows: On input x , $P_{na/nu}^*$ checks whether x equals \bar{x} , and if that is the case, it internally calls $P_{na/pn,2}^*(\text{crs}, \bar{x}, \bar{st})$ to generate a proof. Otherwise, $P_{na/nu}^*$ returns an empty proof. Note that we use the non-uniformity to save the sequence of (\bar{x}, \bar{st}) for each input length. It is again easy to see that this prover is indeed a successful malicious prover against non-adaptive/non-uniform soundness. \square

For adaptive soundness, Arte and Bellare [3] showed that there exists a protocol that provides *adaptive/exclusive* soundness but not *adaptive/penalizing* soundness. This indicates that a NISZK protocol with *adaptive/exclusive* soundness might indeed be achievable, compared to one with *adaptive/penalizing* soundness, for which Pass [30] showed a black-box impossibility result.

3.3 Exclusive Soundness Implies Culpable Soundness

In this section we show that adaptive/exclusive soundness implies the notion of adaptive/culpable soundness of [26]. We first recall the definition of culpable soundness (according to our terminology). For an NP-relation \mathcal{R} let $\mathcal{R}_{\text{guilt}}$ be an NP-relation for the complement of $\mathcal{L}_{\mathcal{R}}$, i.e., $x \notin \mathcal{L}_{\mathcal{R}}$ means that there is a polynomial size ω_{guilt} such that $(x, \omega_{\text{guilt}}) \in \mathcal{R}_{\text{guilt}}$. Note that the relation $\mathcal{R}_{\text{guilt}}$ is efficiently verifiable as an NP-relation (and $\mathcal{L}_{\mathcal{R}}$ is therefore in co-NP).

Definition 5 (Adaptive/Culpable Soundness). *A non-interactive argument (Setup, P, V) for an NP-relation \mathcal{R} (in the common reference string model) has adaptive culpable soundness if for any PPT algorithm P_{culp}^* there exists a negligible function ϵ such that*

$$\Pr[V(1^n, x, \pi, \text{crs}) = 1 \wedge (x, \omega_{\text{guilt}}) \in \mathcal{R}_{\text{guilt}}] \leq \epsilon(n),$$

where the probability is over $\text{crs} \xleftarrow{\$} \text{Setup}(1^n)$, $(x, \pi, \omega_{\text{guilt}}) \xleftarrow{\$} P_{\text{culp}}^*(1^n, \text{crs})$, and V 's internal randomness.

Proposition 1. *A non-interactive argument (Setup, P, V) for an NP-relation \mathcal{R} (in the common reference string model) which has a corresponding relation $\mathcal{R}_{\text{guilt}}$ and is adaptive/exclusive sound is also adaptive/culpable sound.*

Proof. Assume that we have a successful prover P_{culp}^* against culpable soundness. We construct a malicious prover P_{ex}^* against exclusive soundness as follows. P_{ex}^* receives as input crs and forwards this to P_{culp}^* which, then, outputs $(x, \pi, \omega_{\text{guilt}})$. Our prover P_{ex}^* checks in polynomial time if $(x, \omega_{\text{guilt}}) \in \mathcal{R}_{\text{guilt}}$. If not it immediately outputs \perp , else it returns (x, π) .

Note that since we interpret outputs \perp as $\perp \notin \mathcal{L}_{\mathcal{R}}$ our prover P_{ex}^* only outputs values not in the language. It is thus an admissible attacker against exclusive soundness. Furthermore, P_{culp}^* can only win for $x \notin \mathcal{L}_{\mathcal{R}}$ such that only outputting (x, π) for those x cannot decrease the success probability. This yields that P_{ex}^* has the same success probability as P_{culp}^* . \square

4 Constructions Based on General Assumptions

4.1 Multi-theorem NISZK Based on One-Way Permutations

Our approach uses the same idea as in [17] of having crs^{aux} , but we apply it in a dual way. That is, we use a language saying that crs^{aux} is *not* pseudorandom. Since this is in general a coNP-relation we use the Blum-Micali-Yao [6, 35] generator for one-way permutations,

$$G(s) = f^{|s|}(s) \| \text{hb}(s) \| \text{hb}(f(s)) \| \dots \| \text{hb}(f^{|s|-1}(x)),$$

where s is the seed of length $|s| = n$, f is a one-way permutation, $f^i(s)$ the i -fold iteration of f for input s , and hb is a hardcore bit for f . Proving that a string

crs^{aux} is *not* in the range of G is easy if one presents the unique seed s such that the first bits are equal to $f^{|s|}(s)$ and that the remaining bits are not the hardcore bits.

For our simulator we can thus generate a perfectly distributed common random string by picking s randomly, computing $G(s)$, and randomly flipping the hardcore bits:

$$\text{crs}^{\text{aux}} \leftarrow G(s) \oplus 0^{|s|} \| t$$

where each bit $t_i \stackrel{\$}{\leftarrow} \{0, 1\}$ in $t = t_1 \| \dots \| t_{|s|}$ is chosen uniformly and independently. Unless all t_i 's are 0 —which happens with probability $2^{-|s|}$ — this gives the simulator a witness for crs^{aux} not being pseudorandom in form of s, t . If $t = 0^{|s|}$ then we let the simulator abort. This unlikely event of all t_i 's being 0 causes our simulator to be statistical zero-knowledge instead of being perfect zero-knowledge.

For the malicious prover in the soundness game we will hand over a pseudorandom string $G(s)$ instead of a truly random one. For the bounded prover this is computationally indistinguishable. But then the prover does not have a witness for the or-part and would thus need to break soundness of the other protocol part for $x \notin \mathcal{L}_{\mathcal{R}}$. This step preserves any exclusive soundness notion but not penalizing soundness, because we need to be able to detect diverging success behavior of the prover in the two cases (which we may not necessarily be able to in the penalizing setting since we cannot check if x is in the language or not).

Below we formally define the augmented language $\mathcal{L}_{\mathcal{R}}^{\text{or}}$ as

$$\mathcal{L}_{\mathcal{R}}^{\text{or}} = \left\{ (x, y) \mid \exists \omega : (x, \omega) \in \mathcal{R} \vee \exists s, t \in \{0, 1\}^{\lfloor |y|/2 \rfloor} : y = G(s) \oplus 0^{|s|} \| t, t \neq 0^{|s|} \right\}$$

and the corresponding relation \mathcal{R}^{or} accordingly. Note that this is an NP-relation such that, if we have any single-theorem statistical NIZK for general NP-relations, then we also have an multi-theorem statistical witness-indistinguishable argument for this relation \mathcal{R}^{or} .

For pseudorandomness of G we consider for any probabilistic polynomial-time algorithm \mathcal{D} the probability that $\mathcal{D}(1^n, y_{b'}) = b'$ where the probability is taken over $b' \stackrel{\$}{\leftarrow} \{0, 1\}$, $y_0 \leftarrow G(s)$ for $s \stackrel{\$}{\leftarrow} \{0, 1\}^n$, $y_1 \stackrel{\$}{\leftarrow} \{0, 1\}^{2n}$. Let $\text{Adv}_{G, \mathcal{D}}^{\text{PRG}}(1^n) := \Pr[\mathcal{D}(1^n, y_{b'}) = b'] - \frac{1}{2}$ be \mathcal{D} 's advantage. We say that G is a pseudorandom generator if for any probabilistic polynomial-time algorithm \mathcal{D} this advantage is negligible. Note that the Blum-Micali-Yao generator based on a one-way permutation f achieves this property.

Construction 2 (SZK-FLS-Transformation). *Let \mathcal{R} be an NP-relation. Let f be a one-way permutation and $\Pi^{\text{or}} = (\text{Setup}^{\text{or}}, P^{\text{or}}, V^{\text{or}})$ be a multi-theorem non-interactive statistical witness-indistinguishable argument for the NP-relation \mathcal{R}^{or} . We construct a multi-theorem non-interactive statistical zero knowledge argument $\Pi = (\text{Setup}, P, V)$ for \mathcal{R} as follows (see also Fig. 5):*

$\text{Setup}(1^n)$	$P(1^n, x, \omega, \text{crs})$	$V(1^n, x, \pi, \text{crs})$
$\text{crs}^{\text{or}} \stackrel{\$}{\leftarrow} \text{Setup}^{\text{or}}(1^n)$	$\parallel \text{crs} = \text{crs}^{\text{or}} \parallel \text{crs}^{\text{aux}}$	$\parallel \text{crs} = \text{crs}^{\text{or}} \parallel \text{crs}^{\text{aux}}$
$\text{crs}^{\text{aux}} \stackrel{\$}{\leftarrow} \{0, 1\}^{2n}$	$\pi^{\text{or}} \stackrel{\$}{\leftarrow} P^{\text{or}}((x, \text{crs}^{\text{aux}}), \omega, \text{crs}^{\text{or}})$	$d \stackrel{\$}{\leftarrow} V^{\text{or}}((x, \text{crs}^{\text{aux}}), \pi, \text{crs}^{\text{or}})$
$\text{crs} \leftarrow \text{crs}^{\text{or}} \parallel \text{crs}^{\text{aux}}$	$\pi \leftarrow \pi^{\text{or}}$	return d
return crs	return π	

Fig. 5. SZK-FLS-Transformation for multi-theorem NISZK argument (additional input 1^n omitted for P^{or} and V^{or} for space reasons).

CRS: We define the sampling algorithm $\text{Setup}(1^n)$ for the common random string crs for our construction as

$$\text{Setup}(1^n) = \text{Setup}^{\text{or}}(1^n) \parallel U_{2n},$$

where U_{2n} is the uniform distribution on all $2n$ -bit strings.

Prover: The prover P , receiving 1^n , $\text{crs} = \text{crs}^{\text{or}} \parallel \text{crs}^{\text{aux}}$, x and ω (for \mathcal{R}) as input, uses $(x, \text{crs}^{\text{aux}})$ and ω for the augmented relation \mathcal{R}^{or} and computes a witness-indistinguishable proof π^{or} for this NP-relation using the string crs^{or} .

Verifier: The verifier V receives 1^n , $\text{crs} = \text{crs}^{\text{or}} \parallel \text{crs}^{\text{aux}}$, x , and a proof π^{or} for \mathcal{R}^{or} . The verifier accepts iff $V^{\text{or}}(1^n, (x, \text{crs}^{\text{aux}}), \pi^{\text{or}}, \text{crs}^{\text{or}})$ accepts.

Theorem 3. Let \mathcal{R} be an NP-relation. Assuming that $\Pi^{\text{or}} = (\text{Setup}^{\text{or}}, P^{\text{or}}, V^{\text{or}})$ is a non-interactive statistical single-theorem zero-knowledge argument for \mathcal{R}^{or} and that f is a one-way permutation, the non-interactive argument system $\Pi = (\text{Setup}, P, V)$ in Construction 2 is a multi-theorem statistical zero-knowledge argument. Furthermore, if the underlying protocol Π^{or} is (non-adaptively resp. adaptively) exclusively sound, then so is the derived protocol Π ; if Π^{or} is adaptive resp. non-adaptive zero-knowledge, then so is Π .

Proof. (Perfect) Completeness: Note that the verifier V accepts a genuine proof $\pi^{\text{or}} \stackrel{\$}{\leftarrow} P(1^n, x, \omega, \text{crs})$ for original data $\text{crs} = \text{crs}^{\text{or}} \parallel \text{crs}^{\text{aux}} \stackrel{\$}{\leftarrow} \text{Setup}(1^n)$ and $x \in \mathcal{L}_{\mathcal{R}}$ if and only if V^{or} accepts π^{or} for $(x, \text{crs}^{\text{aux}})$ under crs^{or} . The latter is always true since $x \in \mathcal{L}_{\mathcal{R}}$ such that the pair $(x, \text{crs}^{\text{aux}})$ of the or-relation is also in $\mathcal{L}_{\mathcal{R}^{\text{or}}}$, the output of P is given by the output of P^{or} for valid input, and the verifier V^{or} accepts genuine proofs of P^{or} .

Non-adaptive/Exclusive Soundness: Assume that Π^{or} is non-adaptively/exclusively sound. Our argument to show that Π , too, has this property is as follows. We will first substitute the “real” common random string by one in which the augmented component crs^{aux} is always in the range of the pseudorandom generator G . This will be indistinguishable for the bounded prover P^* such that P^* outputs a valid proof with roughly equal probability for pseudorandom G . In this step we exploit the property of non-adaptive/exclusive soundness that

$x \notin \mathcal{L}_{\mathcal{R}}$ is chosen before crs . But then the or-language does not have a witness for either part, such that the malicious prover would have to break (non-adaptive) exclusive soundness of the protocol for \mathcal{R}^{or} .

More formally, let crs be a CRS generated as described above and crs_G an artificial CRS generated as

$$\text{crs}_G \leftarrow \text{Setup}^{\text{or}}(1^n) || G(s),$$

where s is chosen uniformly from $\{0, 1\}^n$. In a first game hop we argue that a successful malicious prover P^* for such a CRS is almost as successful as for a genuine one, that is,

$$\Pr[V(1^n, x, \pi, \text{crs}) = 1] \approx \Pr[V(1^n, x, \pi, \text{crs}_G) = 1]$$

are negligibly close, where the probability is over $(x, \text{st}) \xleftarrow{\$} P^*(1^n)$, $\text{crs} \xleftarrow{\$} \text{Setup}(1^n)$ and $\pi \xleftarrow{\$} P^*(1^n, \text{st}, \text{crs})$ and V 's randomness in the first case, and accordingly over $(x, \text{st}) \xleftarrow{\$} P^*(1^n)$, $\text{crs}_G \xleftarrow{\$} \text{Setup}^{\text{or}}(1^n) || G(s)$, $\pi \xleftarrow{\$} P^*(1^n, \text{st}, \text{crs}_G)$ and V 's randomness in the second case.

We show the indistinguishability by defining a distinguisher \mathcal{D} against the pseudorandom generator G . For security parameter n the distinguisher receives a string $y \in \{0, 1\}^{2n}$ as input, either picked uniformly at random, or being the output of the pseudorandom generator. The distinguisher then invokes the prover and verifier to decide:

$$\begin{array}{l} \mathcal{D}(1^n, y) \\ \hline (x, \text{st}) \xleftarrow{\$} P^*(1^n) \\ \text{crs}^{\text{or}} \xleftarrow{\$} \text{Setup}^{\text{or}}(1^n) \\ \text{crs} \leftarrow \text{crs}^{\text{or}} || y \\ \pi \xleftarrow{\$} P^*(1^n, \text{st}, \text{crs}) \\ \mathbf{return} V(1^n, x, \pi, \text{crs}) \end{array}$$

We claim that the distinguishing advantage bounds the difference between the two games, where G_0 is the original soundness game (with output 1 indicating that P^* has won) and G_1 describes the game where we use the artificial string crs_G instead. Since the two games correspond syntactically to the cases that the distinguisher receives a random y resp. a pseudorandom y we get:

$$\Pr[G_0(1^n) = 1] - \Pr[G_1(1^n)] \leq 2 \cdot \text{Adv}_{G, \mathcal{D}}^{\text{PRG}}(1^n).$$

Next we turn the malicious prover P^* in G_1 against non-adaptive/exclusive soundness against the unbounded scheme Π into one of the same type for the augmented scheme Π^{or} . Note that we are guaranteed that P^* always outputs $x \notin \mathcal{L}_{\mathcal{R}}$ by assumption. Our prover P_{or}^* against Π^{or} works as follows:

$P_{\text{or}}^*(1^n)$	$P_{\text{or}}^*(1^n, \text{st}_{\text{or}}, \text{crs}^{\text{or}})$
$(x, \text{st}) \xleftarrow{\$} P^*(1^n)$	$\parallel \text{st}_{\text{or}} = (\text{st}, \text{crs}^{\text{aux}})$
$s \xleftarrow{\$} \{0, 1\}^n$	$\text{crs} \leftarrow \text{crs}^{\text{or}} \parallel \text{crs}^{\text{aux}}$
$\text{crs}^{\text{aux}} \leftarrow G(s)$	$\pi \xleftarrow{\$} P^*(1^n, \text{st}, \text{crs})$
$\text{st}_{\text{or}} \leftarrow (\text{st}, \text{crs}^{\text{aux}})$	return π
return $((x, \text{crs}^{\text{aux}}), \text{st}_{\text{or}})$	

We first observe that, if P^* always outputs $x \notin \mathcal{L}_{\mathcal{R}}$, then our prover P_{or}^* always outputs $(x, \text{crs}^{\text{aux}}) \notin \mathcal{L}_{\mathcal{R}}^{\text{or}}$. This holds as the string crs^{aux} is pseudorandom such that neither condition of the or-language is satisfied. In addition, P_{or}^* is efficient. Hence, P_{or}^* is also an admissible attacker against non-adaptive/exclusive soundness, this time against $\mathcal{L}_{\mathcal{R}}^{\text{or}}$.

We conclude that, by the soundness of Π^{or} , the success probability of prover P_{or}^* must be negligible. But because P_{or}^* has the same success probability as P^* in G_1 it follows that the winning probability of P^* in G_1 must also be negligible. Since this success probability is negligibly close to the one of P^* in G_0 by the pseudorandomness of G , we derive that P^* success probability against our derived protocol Π must be negligible.

Adaptive/Exclusive Soundness: The proof in the adaptive case follows exactly as in the non-adaptive case. Only this time P^* chooses $x \notin \mathcal{L}_{\mathcal{R}}$ after seeing crs . But both the distinguisher \mathcal{D} against the pseudorandomness \mathcal{D} , as well as the prover P_{or}^* against soundness, can assemble the common random string before P^* selects x . It follows as before that the probability of P_{or}^* against adaptive/exclusive soundness of Π^{or} and thus the one of P^* against Π must be negligible.

Zero Knowledge: The simulator ZKSim works as follows: On input 1^n it first generates $\text{crs} = \text{crs}^{\text{or}} \parallel \text{crs}^{\text{aux}}$, where $\text{crs}^{\text{or}} \xleftarrow{\$} \text{Setup}^{\text{or}}(1^n)$ and crs^{aux} is sampled as

$$\text{crs}^{\text{aux}} \leftarrow G(s) \oplus 0^{|s|} \parallel t$$

for s, t chosen uniformly from $\{0, 1\}^n$. Note that since f is a permutation this CRS has the same distribution as a truly random string. If $t = 0^{|s|}$ then the simulator immediately aborts. Else it outputs crs as the common random string and (s, t) as state st_{ZKSim} . When receiving a (valid) theorem $x \in \mathcal{L}_{\mathcal{R}}$ the simulator runs the prover P^{or} for \mathcal{R}^{or} on input $1^n, (x, \text{crs}^{\text{aux}}), \text{crs}^{\text{or}}$ and witness (s, t) to generate a proof π^{or} . The state remains unchanged.

By assumption, Π^{or} is single-theorem statistical zero knowledge (either adaptively or non-adaptively secure). Further, by Lemma 1 it is single-theorem statistical witness indistinguishable, and by Lemma 2 also multi-theorem statistical witness indistinguishable for the same level of adaptiveness. Therefore, whenever ZKSim is able to find a valid $t \neq 0^{|s|}$, the statistical distance between genuine proofs by P^{or} (for witness ω) and proofs by ZKSim resp. P^{or} (with witness (s, t)) is

given by a negligible term $\epsilon(n)$ for any distinguisher requesting at most q proofs. As ZKSim fails to derive $t \neq 0^{|s|}$ with probability 2^{-n} , the overall statistical distance is therefore at most $\epsilon(n) + 2^{-n}$ and thus negligible. Thus, $\Pi = (\text{Setup}, \text{P}, \text{V})$ is multi-theorem statistical zero knowledge. We note that the protocol inherits the notion of zero-knowledge adaptiveness from Π^{or} . \square

We remark that the transformation also preserves adaptive/culpable soundness. For this notion the distinguisher against the pseudorandom generator in the soundness part can check efficiently if the prover's choice x is in the language or not with the help of the witness ω_{guilt} which the prover needs to output, too.

4.2 Adaptive Perfect Zero-Knowledge Under Expected Poly-Time

The construction in the previous section displays a small error in the simulation, even if we would start with a perfect zero-knowledge or witness-indistinguishable argument. The reason is that our simulator may not generate a valid pair (s, t) with $t \neq 0^{|s|}$. However, to preserve perfect zero-knowledge the simulator cannot simply discard such bad pairs, else outputs of the form $G(s)$ would not be hit (while a uniformly chosen string may actually be in the range of G).

The solution in the single-theorem case is to use the fact that the event of picking bad t 's is very unlikely, namely, 2^{-n} . We will now decrease the probability further such that we can safely search for the actual witness ω for the x part in this rare case, without violating polynomial run time on the average. For this let $p_{\mathcal{R}}$ denote the polynomial which bounds the witness length of relation \mathcal{R} . Then we use a pseudorandom generator $G(s)$ as before, but we iterate the one-way permutation f for $p_{\mathcal{R}}(n)$ steps. Now the probability of picking some input $(s, t) \in \{0, 1\}^n \times \{0, 1\}^{p_{\mathcal{R}}(n)}$ with $t = 0^{p_{\mathcal{R}}(n)}$ is $2^{-p_{\mathcal{R}}(n)}$. Given that this happens we let the simulator (later, after having obtained the input x) search through all potential witnesses $w \in \{0, 1\}^{\leq p_{\mathcal{R}}(n)}$ and each time check in polynomial time $q_{\mathcal{R}}(n)$ if $(x, w) \in \mathcal{R}$. The run time of the simulator for the exhaustive search is then bounded from above by $2 \cdot 2^{p_{\mathcal{R}}(n)} \cdot q_{\mathcal{R}}(n)$. But since this step is only executed with probability at most $2^{-p_{\mathcal{R}}(n)}$ the overall run time of the simulator remains polynomial in expectation.

If we assume that the original argument system Π^{or} is perfectly witness indistinguishable for non-adaptively chosen statements, then the derived protocol is perfectly zero-knowledge, with as simulator running in expected polynomial time and holding either a witness s, t for the auxiliary part or a witness for x to compute the proof. As in the statistical case, the protocol still preserves non-adaptive/exclusive or adaptive/exclusive soundness.

The next step is to extend the above idea to multiple theorems. If we have polynomial many statements x_1, \dots, x_q then we would have to search for all witnesses to simulate the proofs if $t = 0 \dots 0$. But the time to search for all these witnesses by brute force is additive and requires at most $2q \cdot q_{\mathcal{R}}(n) \cdot 2^{p_{\mathcal{R}}(n)}$ many steps. Hence, the expected run time is still polynomial.

We finally remark that our simulator only attains the simple notion of expected polynomial where we average the number of steps over the randomness

of the algorithm. It is not known if one can modify the simulator to achieve more robust notions, such as Levin's average-time complexity.

5 A Lattice-based Construction

The main drawbacks of the previous constructions based on general assumptions is that they are not directly applicable to lattice-based problems because they require a one-way permutation. In this section we therefore present a multi-theorem extension in the common random string using dual-mode commitments, based on the Learning-With-Errors (LWE) assumption.

5.1 Dual-mode Commitment Schemes Based on Lattices

A (non-interactive) commitment scheme consists of a probabilistic polynomial-time algorithm to generate a public key and another probabilistic polynomial-time algorithm which allows to commit to a message under a public key. The scheme can be statistically-hiding (and computationally-binding), or it can be perfectly-binding (and computationally-hiding). A dual-mode scheme has now two key generation algorithms, one for the statistically-hiding and one for the perfectly-binding case. Furthermore, the output of the two key generation algorithms is computationally indistinguishable. To preserve statistical zero-knowledge we make the additional assumption that the public key output in hiding mode is close to uniform:

Definition 6 (Dual-mode Commitment Scheme). *A non-interactive commitment scheme $\Gamma = (\text{Gen}_H, \text{Gen}_B, \text{Com})$ is called a dual-mode commitment scheme if,*

Statistically-Hiding Mode: *The scheme $(\text{Gen}_H, \text{Com})$ is a statistically-hiding, computationally-binding commitment scheme. Furthermore, the output of Gen_H is statistically close to the uniform distribution.*

Perfectly-Binding Mode: *The scheme $(\text{Gen}_B, \text{Com})$ is a perfectly-binding, computationally-hiding commitment scheme.*

Indistinguishability of Modes: *The random variables Gen_H and Gen_B are computationally indistinguishable.*

Note that for a dual commitment scheme, it suffices to show that the scheme is statistically-hiding in the hiding mode, perfectly-binding in the binding mode, and that the modes are computationally indistinguishable. The complementary property of the corresponding mode (with computational guarantees) follows immediately.

For the dual-mode commitments, we will use (a stripped-off version of) the two homomorphic trapdoor functions defined by Gorbunov et al. [23]. As pointed out in [13], these two trapdoor functions give rise to a dual-mode commitment scheme. It has been shown in [13] that it can be used together with a non-interactive witness-indistinguishable proof system for bounded distance decoding to build non-interactive *designated-verifier* computational zero-knowledge

arguments. We will describe this dual-mode commitment scheme now in detail and provide proof sketches based on the security proofs in [23].

The construction of the commitment scheme in [23] itself is based on the SIS problem [2], stating that for parameters $n, m = \text{poly}(n), q$ and β_{SIS} it is hard to find a short non-zero integer vector u (of length at most β_{SIS}) to a given random $n \times m$ -matrix A over \mathbb{Z}_q such that $Au = 0$. The noteworthy property is that there is also a method to generate an $n \times m$ matrix A over \mathbb{Z}_q together with a trapdoor in a secure way. This is implemented by an algorithm TrapGen , taking $1^n, 1^m$ and q as input. Furthermore, there exists an algorithm $\text{Sam}(1^m, 1^m, q)$ which outputs a “small” matrix $U \in \mathbb{Z}_q^{m \times m}$. As discussed in [23] it holds that A generated by $\text{TrapGen}(1^n, 1^m, q)$ is statistically close to uniform, and that A and $A \cdot U$ (sampled according to Sam) are statistically close to A and a uniform matrix V' . The final ingredient is a fixed and easy to compute matrix $G \in \mathbb{Z}_q^{n \times m}$ for the given parameter which allows us to build the commitment scheme. We can then commit to a value $x \in \mathbb{Z}_q$ for matrix A by computing $A \cdot U + x \cdot G$. Note that since $A \cdot U$ is statistically close to a uniform matrix V' we obtain that x is statistically hidden.

We note that we do not take advantage of the trapdoor property here in our construction, but instead sample a uniform matrix A (in the hiding mode). Moreover, as pointed out in [13], the SIS assumption is not necessary either. The LWE assumption suffices for our purpose, since we only need that the mode switching is computationally indistinguishable. Indeed, the same could be already accomplished with Regev’s encryption scheme [33] where one can alter to a lossy mode. We describe the dual-commitment scheme more formally in the following constructions:

Construction 4 (Hiding-mode Commitment Scheme).

Key Generation Gen_H : We sample $A \in \mathbb{Z}_q^{n \times m}$ uniformly and set $\text{pk} \leftarrow A$.

Commitment Com : For input pk and $x \in \mathbb{Z}_q$, we sample $U \leftarrow \text{Sam}(1^m, 1^m, q)$ and return $\text{Com}(\text{pk}, x; U) = \text{pk} \cdot U + x \cdot G$. To open the commitment, we reveal x and U (or the randomness used to sample U).

Proposition 2. *Construction 4 is a statistically hiding commitment scheme.*

Proof. As shown in [23], we have that the following two tuples are statistically close:

$$(\text{pk}, x, \text{pk} \cdot U + x \cdot G) \equiv_s (\text{pk}, x, V')$$

where $U \leftarrow \text{Sam}(1^m, 1^m, q)$ and $V' \leftarrow \mathbb{Z}_q^{n \times m}$, i.e., the commitment is statistically indistinguishable from a random matrix. This holds for public keys generated by TrapGen and, since that algorithm’s output is close to uniform, also for the random matrix A . \square

Next we recall from [23] how we can switch to a perfectly-binding mode by assuming the hardness of LWE. This problem states that given a matrix A and $As + e$ for a small error vector e sampled from a distribution χ , recovering s is hard [33].

Construction 5 (Binding-mode Commitment Scheme).

Key Generation Gen_B : We sample $A' \leftarrow \mathbb{Z}_q^{(n-1) \times m}$ uniformly and $s' \xleftarrow{\$} \mathbb{Z}_q^{n-1}$ and set

$$pk \leftarrow \begin{pmatrix} A' \\ s' A' + e \end{pmatrix},$$

where e is a short “noise vector” sampled from χ .

Commitment Com: The commitment is identical to the one in Construction 4.

Proposition 3. Construction 5 is a perfectly binding commitment scheme.

Proof. To show this construction is perfectly binding, it suffices to show that we can uniquely recover x using s . Indeed, if we know s' , we can set $s = (-s', 1)$ and $z = (0, \dots, 0, r)$ and calculate

$$s(pk \cdot U + x \cdot G)G^{-1}(z) = e \cdot U \cdot G^{-1}(z) + x \cdot \langle s, z \rangle = x \cdot r + e'.$$

Note that G^{-1} is a polynomial-time algorithm whose existence is guaranteed by Lemma 2.2 in [23]. For correctly chosen parameters r and e , this lets us recover x uniquely. Now, as s does not depend on x or U , if for two pairs (x, U) and (x', U')

$$pk \cdot U + x \cdot G = pk \cdot U' + x' \cdot G,$$

holds, then we have $x = x'$. □

Proposition 4. Assuming the $\text{LWE}(q, \chi)$ -assumption holds, Constructions 4 and 5 together form a dual-mode commitment scheme.

Proof. We start by showing that the public keys of both schemes are computationally indistinguishable. First, note that all but the last column of matrix A are generated uniformly random (or statistically close to that) for both public keys. Therefore, the problem is equivalent to distinguish between $A's + e$ and v' given A' , where $v' \in \mathbb{Z}_q^n$ is a uniformly random vector and s and e are sampled as described in the scheme. However, this is exactly the decisional LWE problem. By our assumption, the two public keys are therefore indistinguishable.

We have not yet shown that Construction 4 is computationally binding and that Construction 5 is computationally hiding. However, we argue this follows directly from what we have shown already. Assume Construction 5 would not be computationally hiding, i.e., there exists an adversary that, given a public key pk , can distinguish between a commitment to x and x' with notable advantage. However, in this case, we can use this adversary to distinguish the public keys of the schemes, as Construction 4 is statistically hiding and no adversary with notable advantage can exist here.

Similarly, assume that Construction 4 is not computationally binding. Then, there exists an adversary that, given a public key pk , can generate a commitment c that opens to two values x and x' with non-negligible probability. However, as Construction 5 is perfectly binding, we can use such an adversary to distinguish between public keys of the two schemes. □

5.2 SZK-FLS-Transformation Based on Lattices

We will now define our multi-theorem transformation based on the dual-mode commitment scheme in the previous section. As before, we will use the FLS-type transform, therefore we only need to define a sampling algorithm for the auxiliary CRS crs^{aux} and an augmented or-relation \mathcal{R}^{or} for this string.

The sampling algorithm $\text{Setup}^{\text{aux}}$ to generate crs^{aux} will just generate uniformly random values representing a public key pk and a commitment c :

$$\text{crs}^{\text{aux}} = (\text{pk}, c) \leftarrow U_{nmq} \times U_{nmq}.$$

Note that a random public key corresponds to the hiding-mode public key.

Technically the public key and the commitment in crs^{aux} are matrices over \mathbb{Z}_q , and not uniform strings as required by the common random string model. However, we can generate random elements in \mathbb{Z}_q from uniform strings by interpreting a random string of length $|q| + n$ as an integer and mapping it to the residue mod q . The statistical distance to a uniform element from \mathbb{Z}_q is then exponentially small. We stress that we can also go “backwards” with this technique. Given a random value $v \in \mathbb{Z}_q$ we can add a random multiple $i \cdot q$ to v for $i \xleftarrow{\$} \{0, 1, \dots, 2^{n-1}\}$ to get an (almost) uniform $|q| + n$ bit string which would map to v again. Hence, from now on we switch between random matrices from \mathbb{Z}_q and uniformly random string whenever convenient.

Our relation will now ask for a given public key pk of the commitment scheme and commitment c , both found in the common random string, if there is a matrix $U \leftarrow \text{Sam}(1^m, 1^m, q)$ resp. randomness u such that $U = \text{Sam}(1^m, 1^m, q; u)$, such that the commitment opens to 1:

$$((\text{pk}, c), u) \in \mathcal{R}^{\text{or}} \iff U = \text{Sam}(1^m, 1^m, q; u) \wedge c = \text{Com}(\text{pk}, 1; U).$$

Given these two properties we can now use the same construction as for the one-way permutation, only that we use the relation above and the sampler $\text{Setup}^{\text{aux}}$ to generate crs^{aux} . In fact the construction is otherwise identical to the one in Fig. 5:

Construction 6 (SZK-FLS-Dual-Mode-Transformation). *Let \mathcal{R} be an NP-relation. Further, let $\Gamma = (\text{Gen}_H, \text{Gen}_B, \text{Com})$ be a non-interactive dual-mode commitment scheme and suppose that $\Pi^{\text{or}} = (\text{Setup}^{\text{or}}, P^{\text{or}}, V^{\text{or}})$ be a multi-theorem non-interactive statistical witness-indistinguishable argument for the NP-relation \mathcal{R}^{or} . We construct a multi-theorem non-interactive statistical zero knowledge argument $\Pi = (\text{Setup}, P, V)$ for \mathcal{R} as in Fig. 5 with the following exception:*

CRS: *We define the sampling algorithm $\text{Setup}(1^n)$ for the common random string crs for our construction as*

$$\text{Setup}(1^n) = \text{Setup}^{\text{or}}(1^n) \parallel \text{Setup}^{\text{aux}}(1^n).$$

The prover algorithm P and verifier algorithm V are as before.

Theorem 7. *Let \mathcal{R} be an NP-relation. Assuming that $\Pi^{or} = (\text{Setup}^{or}, P^{or}, V^{or})$ is a non-interactive statistical single-theorem zero-knowledge argument for \mathcal{R}^{or} and that $\Gamma = (\text{Gen}_H, \text{Gen}_B, \text{Com})$ is a dual-mode non-interactive commitment scheme, the non-interactive argument $\Pi = (\text{Setup}, P, V)$ in Construction 6 is a multi-theorem statistical zero-knowledge argument. Furthermore, if the underlying protocol Π^{or} is (non-adaptively resp. adaptively) exclusively sound, then so is the derived protocol Π ; if Π^{or} is adaptive resp. non-adaptive zero-knowledge, then so is Π .*

Proof. The proof is very close to the one of Theorem 3 such that we only sketch the main differences here.

(Perfect) Completeness: It follows as in the one-way permutation case that the honest verifier accepts proofs generated by P for $x \in \mathcal{L}_{\mathcal{R}}$.

Exclusive Soundness: To show exclusive soundness (in the non-adaptive or adaptive case) we first switch the auxiliary string to a randomly sampled binding key $\text{pk} \stackrel{\$}{\leftarrow} \text{Gen}_B(1^n)$ and a 0-commitment $\text{Com}(\text{pk}, 0; U)$, instead of using uniformly random values. Note that we can use two game hops to show that this is computationally indistinguishable from genuine common random strings. In the first hop we replace the random key component in crs^{aux} by a key $\text{pk} \stackrel{\$}{\leftarrow} \text{Gen}_H(1^n)$, which is even statistically close. Then we replace the random commitment component in crs^{aux} by a random commitment to 0, $\text{Com}(\text{pk}, 0; U)$. This is again statistically indistinguishable.

And finally we switch to a binding key $\text{pk} \stackrel{\$}{\leftarrow} \text{Gen}_B(1^n)$ and a 0-commitment under this key. This is computationally indistinguishable by the indistinguishability of the dual-mode key generation. (The additional 0-commitment can be computed easily given a hiding or binding key.) This is where we again use exclusive soundness to turn a malicious prover into a distinguisher against the dual-mode scheme, analogously to the distinguisher against the pseudorandomness of the generator in the one-way permutation case.

We now have an auxiliary string which contains a binding key and a 0-commitment, such that the or-part in the \mathcal{R}^{or} cannot be satisfied. It follows now as before that soundness of the constructed protocol follows from the soundness of the original non-interactive argument.

Zero-Knowledge: For adaptive multi-theorem zero-knowledge we remark that the simulator ZKSIm can create the key part in the auxiliary string as a hiding key $\text{pk} \stackrel{\$}{\leftarrow} \text{Gen}_H(1^n)$ and the commitment part as a 1-commitment under pk . Since the key pk and the 1-commitment are statistically close to a uniform strings, the simulator's string crs^{aux} is statistically close to a uniform string. For this string crs^{aux} the simulator can use the randomness of the commitment as a witness. The remaining steps in the proof are identical to the ones in the proof of Theorem 3. \square

6 Conclusion

We have shown how to apply the idea of the FLS transformation also for statistical zero-knowledge arguments. Let us highlight two important aspects of our transformations.

First, our transformations based on one-way permutations and on lattices work in the common *random* string model and does not require any structure of the CRS. Common *reference* strings have the inherent disadvantage that they have some structure and that one needs to trust the party which generates the string. A prominent example is the discussion about the trustworthiness of the Zcash reference string and follow-up suggestions to use common random strings instead, e.g., [16]. Of course, a party generating a common *random* string may also impose some trust assumption, as our lattice-based solution with the implicit trapdoor generation algorithm shows. But several measures to thwart attacks can be implemented much easier than for structured strings. This includes the computation of the string as the output of a hash function, or by xoring common random strings from several sources.

The other aspect we would like to emphasize that our transformations preserve adaptive security for both zero-knowledge and soundness. This does not conflict with black-box impossibility result for such statistical zero-knowledge arguments [1, 30], because in the course of showing adaptive soundness we have, in passing, encountered a possibility to bypass the impossibility results. A key observation is that one may be able to achieve adaptive soundness and zero-knowledge if one switches to the notion of exclusive soundness. This adaptive/exclusive soundness implies adaptive/culpable soundness and thus suffices for many practical applications.

Acknowledgments. We thank the anonymous reviewers for valuable comments. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1119 – 236615297.

References

1. Abe, M., Fehr, S.: Perfect NIZK with adaptive soundness. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 118–136. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-70936-7_7
2. Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: 28th ACM STOC, pp. 99–108. ACM Press, May 1996
3. Arte, V., Bellare, M.: Dual-Mode NIZKs: possibility and impossibility results for property transfer. In: Bhargavan, K., Oswald, E., Prabhakaran, M. (eds.) INDOCRYPT 2020. LNCS, vol. 12578, pp. 859–881. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-65277-7_38
4. Bellare, M., Hofheinz, D., Kiltz, E.: Subtleties in the definition of IND-CCA: when and how should challenge decryption be disallowed? *J. Cryptol.* **28**(1), 29–48 (2015)
5. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications (extended abstract). In: 20th ACM STOC. pp. 103–112. ACM Press, May 1988

6. Blum, M., Micali, S.: How to generate cryptographically strong sequences of pseudo-random bits. *SIAM J. Comput.* **13**(4), 850–864 (1984). <https://doi.org/10.1137/0213053>
7. Blum, M., Santis, A.D., Micali, S., Persiano, G.: Noninteractive zero-knowledge. *SIAM J. Comput.* **20**(6), 1084–1118 (1991). <https://doi.org/10.1137/0220068>
8. Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.* **37**(2), 156–189 (1988). [https://doi.org/10.1016/0022-0000\(88\)90005-0](https://doi.org/10.1016/0022-0000(88)90005-0)
9. Canetti, R., et al.: Fiat-Shamir: from practice to theory. In: Charikar, M., Cohen, E. (eds.) 51st ACM STOC, pp. 1082–1090. ACM Press, June 2019
10. Canetti, R., Chen, Y., Reyzin, L., Rothblum, R.D.: Fiat-Shamir and correlation intractability from strong KDM-secure encryption. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10820, pp. 91–122. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78381-9_4
11. Canetti, R., Lombardi, A., Wichs, D.: Fiat-Shamir: from practice to theory, Part II (NIZK and correlation intractability from circular-secure FHE). *Cryptology ePrint Archive*, Report 2018/1248 (2018). <https://eprint.iacr.org/2018/1248>
12. Chaidos, P., Groth, J.: Making sigma-protocols non-interactive without random oracles. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 650–670. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46447-2_29
13. Couteau, G., Hofheinz, D.: Designated-verifier pseudorandom generators, and their applications. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11477, pp. 562–592. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17656-3_20
14. Fauzi, P., Lipmaa, H.: Efficient culpably sound NIZK shuffle argument without random oracles. In: Sako, K. (ed.) CT-RSA 2016. LNCS, vol. 9610, pp. 200–216. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-29485-8_12
15. Fauzi, P., Lipmaa, H., Siim, J., Zając, M.: An efficient pairing-based shuffle argument. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10625, pp. 97–127. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70697-9_4
16. Fauzi, P., Meiklejohn, S., Mercer, R., Orlandi, C.: Quisquis: a new design for anonymous cryptocurrencies. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11921, pp. 649–678. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34578-5_23
17. Feige, U., Lapidot, D., Shamir, A.: Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In: 31st FOCS, pp. 308–317. IEEE Computer Society Press, October 1990
18. Feige, U., Lapidot, D., Shamir, A.: Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM J. Comput.* **29**(1), 1–28 (1999). <https://doi.org/10.1137/S0097539792230010>
19. Feige, U., Shamir, A.: Witness indistinguishable and witness hiding protocols. In: 22nd ACM STOC, pp. 416–426. ACM Press, May 1990
20. Fischlin, M., Rohrbach, F.: Single-to-multi-theorem transformations for non-interactive statistical zero-knowledge. *Cryptology ePrint Archive*, Report 2020/1204 (2020). <https://eprint.iacr.org/2020/1204>
21. Goldreich, O.: *Foundations of Cryptography*, vol. 1. Cambridge University Press, Cambridge (2006)
22. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. *SIAM J. Comput.* **18**(1), 186–208 (1989). <https://doi.org/10.1137/0218012>

23. Gorbunov, S., Vaikuntanathan, V., Wichs, D.: Leveled fully homomorphic signatures from standard lattices. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC, pp. 469–477. ACM Press, June 2015
24. Groth, J., Lu, S.: A non-interactive shuffle with pairing based verifiability. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 51–67. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-76900-2_4
25. Groth, J., Ostrovsky, R., Sahai, A.: Perfect non-interactive zero knowledge for NP. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 339–358. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_21
26. Groth, J., Ostrovsky, R., Sahai, A.: New techniques for noninteractive zero-knowledge. *J. ACM* **59**(3), 111–1135 (2012). <https://doi.org/10.1145/2220357.2220358>
27. Holmgren, J., Lombardi, A.: Cryptographic hashing from strong one-way functions (or: One-way product functions and their applications). In: Thorup, M. (ed.) 59th FOCS, pp. 850–858. IEEE Computer Society Press, October 2018
28. Libert, B., Nguyen, K., Passelègue, A., Titiu, R.: Simulation-sound arguments for LWE and applications to KDM-CCA2 security. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020. LNCS, vol. 12491, pp. 128–158. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64837-4_5
29. Libert, B., Passelègue, A., Wee, H., Wu, D.J.: New constructions of statistical NIZKs: dual-mode DV-NIZKs and more. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12107, pp. 410–441. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45727-3_14
30. Pass, R.: Unprovable security of perfect NIZK and non-interactive non-malleable commitments. *Comput. Complex.* **25**(3), 607–666 (2016). <https://doi.org/10.1007/s00037-016-0122-2>
31. Pass, R., Shelat, A.: Unconditional characterizations of non-interactive zero-knowledge. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 118–134. Springer, Heidelberg (2005). https://doi.org/10.1007/11535218_8
32. Peikert, C., Shiehian, S.: Noninteractive zero knowledge for NP from (plain) learning with errors. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11692, pp. 89–114. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26948-7_4
33. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC. pp. 84–93. ACM Press, May 2005
34. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (ed.) 46th ACM STOC, pp. 475–484. ACM Press, May/June 2014
35. Yao, A.C.C.: Theory and applications of trapdoor functions (extended abstract). In: 23rd FOCS, pp. 80–91. IEEE Computer Society Press, November 1982