

Perfect NIZK with Adaptive Soundness

Masayuki Abe¹ and Serge Fehr²

¹ Information Sharing Platform Laboratories
NTT Corporation, Japan
`abe.masayuki@lab.ntt.co.jp`
² CWI Amsterdam, The Netherlands
`fehr@cwi.nl`

Abstract. The notion of non-interactive zero-knowledge (NIZK) is of fundamental importance in cryptography. Despite the vast attention the concept of NIZK has attracted since its introduction, one question has remained very resistant: Is it possible to construct NIZK schemes for any NP-language with *statistical* or even *perfect* ZK? Groth, Ostrovsky and Sahai recently answered this question in the affirmative. However, in order to achieve *adaptive* soundness, i.e., soundness against dishonest provers who may choose the target statement depending on the common reference string (CRS), their schemes require some restriction to be put upon the statements to be proven, e.g. an a-priori bound on its size.

In this work, we first present a very simple and efficient adaptively-sound perfect NIZK argument system for any NP-language. Besides being the first adaptively-sound statistical NIZK argument for all NP that does not pose any restriction on the statements to be proven, it enjoys a number of additional desirable properties: it allows to re-use the CRS, it can handle *arithmetic* circuits, and the CRS can be set-up very efficiently without the need for an honest party. We then show an application of our techniques in constructing *efficient* NIZK schemes for proving arithmetic relations among committed secrets, whereas previous methods required expensive generic NP-reductions.

The security of the proposed schemes is based on a strong non-standard assumption, an extended version of the so-called Knowledge-of-Exponent Assumption (KEA) over bilinear groups. We give some justification for using such an assumption by showing that the commonly-used approach for proving NIZK *arguments* sound does not allow for adaptively-sound statistical NIZK arguments (unless $\text{NP} \subset \text{P/poly}$). Furthermore, we show that the assumption used in our construction holds with respect to generic adversaries that do not exploit the specific representation of the group elements. We also discuss how to avoid the non-standard assumption in a pre-processing model.

1 Introduction

1.1 Background

NON-INTERACTIVE ZERO-KNOWLEDGE. The notion of non-interactive zero-knowledge (NIZK) captures the problem of proving a statement by just sending one message and without revealing any additional information besides the validity of the statement, provided that a common reference string (CRS) has been properly set up. Since its introduction by Blum, Feldman and Micali in 1988 [6], NIZK has been a fundamental cryptographic primitive used throughout modern cryptography in essential ways.

There is a considerable amount of literature dedicated to NIZK, in particular to the study of which languages allow for what flavor of NIZK proof. As in case of interactive ZK it is known that there cannot be statistical NIZK proofs (i.e., both ZK and soundness are unconditional) for NP-complete languages unless the polynomial hierarchy collapses [23, 2, 33]. Hence, when considering general NP-languages, this only leaves room for a NIZK proof with *computational* ZK or *computational* soundness (where the proof is also called an *argument*), or both. However, in contrast to interactive ZK where it has long been known that both flavors can exist [9, 8, 25], all proposed NIZK proofs or arguments for general NP-languages have computational ZK (see e.g. [6, 21, 7, 30, 34]). Hence the construction of a statistically NIZK (NISZK) argument

has remained an open problem (until very recently, see below). The question of the existence of NISZK arguments is in particular interesting in combination with a result by De Santis *et al.* [34], where they observe that for a strong notion of NIZK, called *same-string* NIZK, soundness can only be computational when considering NP-complete languages (assuming that one-way functions exist).

STATISTICAL NIZK ARGUMENTS. Recently, Groth, Ostrovsky and Sahai proposed an elegant construction for a perfect NIZK (NIPZK) argument for circuit-SAT [27] by using bilinear groups. This shows NIZK *can* come with perfect ZK for any NP-language. However, the scheme only provides security against a non-adaptive dishonest prover who chooses the target instance $x^* \notin L$ (for which it wants to fake a proof) *independent* of the CRS. In an application though, it is likely that the adversary first sees the CRS and then chooses the false statement on which he wants to cheat. Using a counting argument, they argue that under a specific number-theoretic subexponential hardness assumption their scheme is sound against an adaptive dishonest prover if the size of the circuit to be proven is a-priori limited. However, the bound on the size of the circuit is so restrictive that the circuit must be smaller than sublinear in the bit size of the CRS (as discussed in Section 1.3). Though such a restriction is commonly observed in NIZK proof systems and indeed the lower bound in [39] shows that it is inherent, no negative result has been proven for argument systems. Groth *et al.* also proposed a perfect NIZK argument for SAT which is provably secure in Canetti’s Universal Composability (UC) framework [10]. However, besides being much less efficient than their first construction, the scheme still does not guarantee unrestricted security against an adaptive dishonest prover who chooses the target instance $x^* \notin L$ depending on the CRS. For instance, the UC security does not exclude the possibility that a dishonest prover comes up with an accepting proof for the statement “the CRS is invalid or S is true” for an arbitrary false statement S . Since in a real-life execution the CRS is assumed to be valid, this is a convincing argument of the false statement S . Accordingly, the existence of an unrestricted statistical or perfect NIZK argument, which does not pose any restriction on the instances to be proven, is still an open problem.

THE KNOWLEDGE-OF-EXPONENT ASSUMPTION. Informally, the Knowledge-of-Exponent Assumption (KEA) says that for certain groups, given a pair g and $\hat{g} = g^x$ of group elements with unknown discrete-log x , the only way to efficiently come up with another pair A and \hat{A} such that $\hat{A} = A^x$ (for the same x) is by raising g and \hat{g} to some power a : $A = g^a$ and $\hat{A} = \hat{g}^a$. KEA was first introduced and used by Damgård in 1991 [15], and later, together with an extended version (KEA2), by Hada and Tanaka [28]. Recently, Bellare and Palacio [5] showed that KEA2 does not hold, and proposed a new extended version called KEA3 in order to save Hada and Tanaka’s results. KEA3, which we call XKEA for eXtended KEA, says that given *two* pairs (g, \hat{g}) and (h, \hat{h}) with the same unknown discrete-log x , the only way to efficiently come up with another pair A and \hat{A} such that $\hat{A} = A^x$ is by computing $A = g^a h^\alpha$ and $\hat{A} = \hat{g}^a \hat{h}^\alpha$ for some a and α . Assumptions like KEA and XKEA are widely criticized in particular because they do not appear to be “efficiently falsifiable”, as Naor put it [31], though Bellare and Palacio showed that this is not necessarily the case.

1.2 Our Result

Based on XKEA over bilinear groups, we construct an adaptively-sound NIPZK argument for circuit-SAT without any restrictions on the instances to be proven. Besides being the first unrestricted adaptively-sound NISZK argument for any NP-language, the proposed scheme enjoys a number of additional desirable properties: It is *same-string* NIZK, which allows to re-use the

CRS. It is very efficient: the CRS essentially consists of a few group elements, and a proof consists of a few group elements per multiplication gate; this is comparable (if not better) to the first scheme by Groth *et al.*, which is the most efficient general-purpose NIZK scheme known up to date (see the comparison in [27]). Furthermore, our scheme can also be applied to arithmetic circuits over \mathbb{Z}_q for a large prime q whereas known schemes are tailored to binary circuits; this often allows a more compact representation of the statement to be proven. Finally, the CRS does not need to be set-up by a trusted party. It can efficiently be set-up jointly by the prover and the verifier. Finally, our scheme still achieves witness indistinguishability even if the CRS is provided solely by a (possibly dishonest) verifier without any correctness proof. Thus, it can be viewed as a non-interactive witness indistinguishable argument system in the bare public-key model [12], or equivalently a (computationally sound) *zap* [20]. We are not aware of any other NIZK argument or proof that enjoys all these properties.

Based on the techniques developed for the perfect NIZK argument for SAT, we also construct an *efficient* NIPZK argument for arithmetic relations among committed secrets over \mathbb{Z}_q with large prime q . To the best of our knowledge, all known schemes only work for secrets from restricted domains such as \mathbb{Z}_2 and have to rely on generic inefficient reductions to NP-complete problems to handle larger secrets. [17] constructs efficient NIZK argument which allows to prove any arithmetic relations over large field but the CRS can be used only for logarithmic number of times for the sake of soundness. [26] presents elegant NIZK proofs for specific arithmetic relations used for group signatures. Our approach in particular allows for additive and multiplicative relations among secrets committed to by standard Pedersen commitments.

We give two justifications for using such a strong non-standard assumption like XKEA. First, we give some indication that a non-standard assumption is unavoidable for adaptively-sound NISZK arguments. We prove that the common approach for proving computational soundness based on standard assumptions, which has been used for all NIZK arguments (we are aware of), does not allow for statistical ZK unless $\text{NP} \subset \text{P/poly}$ (i.e. unless any NP-problem can be solved by an efficient non-uniform algorithm). And, second, we prove that KEA and XKEA hold in the generic bilinear-group model. This suggests that if there exists an algorithm that breaks, say, KEA in a certain group, then this algorithm must use the specific representation of the elements of that group, and it is likely to fail when some other group (representation) is used. A similar result was independently developed by Dent [19] for non-bilinear groups.

Finally, we discuss how to avoid XKEA in our NIZK arguments by allowing a pre-processing phase. Our scheme allows very efficient pre-processing where the prover only needs to make commitments for random values and prove their knowledge using efficient off-the-shelf zero-knowledge schemes.

1.3 Related Work

In order to make it easier for the reader to position our results, we would like to give a brief discussion about recently proposed NIPZK arguments. In [27] Groth *et al.* presented two schemes for proving circuit satisfiability, where the first one comes in two flavors. Let us name the resulting three schemes by the *non-adaptive*, the *adaptive* and the *UC GOS* scheme. These are the first (and so far only) NISZK arguments proposed in the literature. The non-adaptive GOS scheme is admitted by the authors to be *not* adaptively sound. The adaptive GOS scheme *is* adaptively sound, but it only allows for circuits that are limited in size, and the underlying computational assumption is somewhat non-standard in that it requires that some problem can only be solved with “sub-negligible” probability, like $2^{-\epsilon\kappa^\epsilon \log \kappa} \text{negl}(\kappa)$ where κ is the bit size of the problem instance. The more one relaxes the bound on the size of the circuits, the stronger the underlying

assumption gets in terms of the assumed bound on the success probability of solving the problem; but in any case the size of the circuits are doomed to be sub-linear in the size of the CRS.

Concerning the UC GOS scheme, we first would like to point out that it is of theoretical interest, but it is very inefficient (though poly-time). Furthermore, it has some tricky weak soundness property in that if a dishonest prover should succeed in proving a false statement, then the statement cannot be distinguished from a true one. It is therefore claimed in [27] that the scheme “achieves a weaker, but sufficient, form of adaptive security.” This is true but only if some care is taken with the kind of statements that the (dishonest) prover is allowed to prove; in particular, soundness is only guaranteed if the statement to be proven does not incorporate the CRS. Indeed, the same example that the authors use to reason that their first scheme is not adaptively sound can also be applied to the UC secure scheme: Consider a dishonest prover that comes up with an accepting proof for the statement “the CRS is invalid”, or for a statement like “the CRS is invalid or S is true” where S is an arbitrary false statement. In real-life, where the CRS is guaranteed to be correct, this convinces the verifier of the truth of the false statement S . However, such a prover is *not* ruled out by the UC security: the simulator given in [27] does generate an *invalid* CRS so that the statement in fact becomes true; and thus the proof can obviously be simulated in the ideal-world (when given a corresponding witness, which the simulator has in case of the UC GOS scheme). We stress that this is not a flaw in the UC GOS scheme but it is the UC security definition that does not provide any security guarantees for statements that incorporate the CRS, essentially because in the ideal-life model there is *no* (guaranteed-to-be-correct) CRS.¹ Subtle issues like this one are taken care of in a recent extension of the UC framework [11], which though requires stronger set-up assumptions than a standard CRS for non-trivial two-party protocols.

In conclusion, UC NIZK security provides good enough security under the condition that the statements to be proven do not incorporate the CRS. This is automatically guaranteed in a UC setting, where the statements to be proven must make sense in the ideal-world model, but not necessarily in other settings.

2 Preliminaries

2.1 Notation

We consider uniform probabilistic algorithms (i.e. Turing machines) which take as input (the unary encoding of) a security parameter $\kappa \in \mathbb{N}$ and possibly other inputs and run in deterministic poly-time in κ . We thus always implicitly require the size of the input to be bounded by some polynomial in κ . Adversarial behavior is modeled by *non-uniform* poly-time probabilistic algorithms, i.e., by algorithms which together with the security parameter κ also get some poly-size auxiliary input aux_κ . In order to simplify notation, we usually leave the dependency on κ (and on aux_κ) implicit. By $y \leftarrow \mathcal{A}(x)$, we mean that algorithm \mathcal{A} is executed on input x (and the security parameter κ and, in the non-uniform case, aux_κ) and the output is assigned to y . Similarly, for any finite set S , we use the notation $y \leftarrow S$ to denote that y is sampled uniformly from S , and $y \leftarrow x$ means that the value x is assigned to y .

¹ A minor flaw regarding the UC GOS scheme though is that Groth *et al.* claim the scheme to be *non-malleable*, and their UC NIZK functionality indeed does guarantee non-malleability in that a proof cannot be transformed into a different proof for the same instance without knowing a witness. But it is easy to see that the UC GOS scheme is *not* non-malleable, because the NIZK proof π generated at step 6. in Figure 4 (by using the non-adaptive GOS scheme) is malleable: it uses the NIZK proof from Figure 1 (with h of order n though) which is malleable by raising π_1 and π_3 to some power $s \in \mathbb{Z}_n^*$ and π_2 to power $s^{-1} \pmod{n}$.

For two algorithms \mathcal{A} and \mathcal{B} , we write $\mathcal{B} \circ \mathcal{A}$ for the composed execution of \mathcal{A} and \mathcal{B} , where \mathcal{A} 's output is given to \mathcal{B} as input. Similarly, $\mathcal{A} \parallel \mathcal{B}$ denotes the joint execution of \mathcal{A} and \mathcal{B} on the same input and the same random tape, and we write $(x; y) \leftarrow (\mathcal{A} \parallel \mathcal{B})(w)$ to express that in the joint execution on input w (and the same random tape) \mathcal{A} 's output is assigned to x and \mathcal{B} 's to y . Furthermore, $P[y = \mathcal{A}(x)]$ denotes the probability (taken over the uniformly distributed random tape) that \mathcal{A} outputs y on input x , and we write $P[x \leftarrow \mathcal{B} : \mathcal{A}(x) = y]$ for the (average) probability that \mathcal{A} outputs y on input x when x is output by \mathcal{B} : $P[x \leftarrow \mathcal{B} : \mathcal{A}(x) = y] = \sum_x P[y = \mathcal{A}(x)]P[x = \mathcal{B}]$. We also use natural self-explanatory extensions of this notation.

An *oracle* algorithm \mathcal{A} is an algorithm in the above sense connected to an oracle in that it can write on its own tape an input for the oracle and tell the oracle to execute, and then, in a single step, the oracle processes its input in a prescribed way, and writes its output to the tape. We write $\mathcal{A}^{\mathcal{O}}$ when we consider \mathcal{A} to be connected to the particular oracle \mathcal{O} .

As is common practice, a value $\nu(\kappa) \in \mathbb{R}$, which depends on the security parameter κ , is called *negligible*, denoted by $\nu(\kappa) \leq \text{negl}(\kappa)$ or $\nu \leq \text{negl}$, if $\forall c > 0 \exists \kappa_0 \in \mathbb{N} \forall \kappa \geq \kappa_0 : \nu(\kappa) < 1/\kappa^c$. Furthermore, $\nu(\kappa) \in \mathbb{R}$ is called *noticeable* if $\exists c > 0, \kappa_0 \in \mathbb{N} \forall \kappa \geq \kappa_0 : \nu(\kappa) \geq 1/\kappa^c$.

2.2 Definition

Let $L \subseteq \{0, 1\}^*$ be an NP-language.

Definition 2.1. Consider poly-time algorithms \mathcal{G} , \mathcal{P} and \mathcal{V} of the following form: \mathcal{G} takes the security parameter κ (implicitly treated hereafter) and outputs a common reference string (CRS) Σ together with a trapdoor τ . \mathcal{P} takes as input a CRS Σ and an instance $x \in L$ together with an NP-witness w and outputs a proof π . \mathcal{V} takes as input a CRS Σ , an instance x and a proof π and outputs 1 or 0. The triple $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ is a statistical/perfect NIZK argument for L if the following properties hold.

Completeness: For any $x \in L$ with corresponding NP-witness w

$$P[(\Sigma, \tau) \leftarrow \mathcal{G}, \pi \leftarrow \mathcal{P}(\Sigma, x, w) : \mathcal{V}(\Sigma, x, \pi) = 0] \leq \text{negl}.$$

Soundness: For any non-uniform poly-time adversary \mathcal{P}^*

$$P[(\Sigma, \tau) \leftarrow \mathcal{G}, (x^*, \pi^*) \leftarrow \mathcal{P}^*(\Sigma) : x^* \notin L \wedge \mathcal{V}(\Sigma, x^*, \pi^*) = 1] \leq \text{negl}.$$

Statistical/Perfect Zero-Knowledge (ZK): There exists a poly-time simulator \mathcal{S} such that for any instance $x \in L$ with NP-witness w , and for $(\Sigma, \tau) \leftarrow \mathcal{G}$, $\pi \leftarrow \mathcal{P}(\Sigma, x, w)$ and $\pi_{sim} \leftarrow \mathcal{S}(\Sigma, \tau, x)$, the joint distributions of (Σ, π) and (Σ, π_{sim}) are statistically/perfectly close.

Remark 2.2. The notion of soundness we use here guarantees security against an *adaptive* attacker, which may choose the instance x^* depending on the CRS. We sometimes emphasize this issue by using the term *adaptively-sound*. Note that this is a strictly stronger notion than when the adversary must choose x^* independent of the CRS.

Remark 2.3. In the notion of ZK we use here, \mathcal{P} and \mathcal{S} use the *same* CRS string. In [34], this is called *same-string* ZK. In the context of *statistical* ZK, this notion is *equivalent* (and not only sufficient) to *unbounded* ZK, which captures that the same CRS can be used an unbounded number of times. This is obviously much more desirable compared to the original notion of NIZK, where every proof requires a fresh CRS. In [34], it is shown that there cannot be a *same-string* NIZK *proof* with statistical soundness for a NP-complete language unless there exist no one-way functions. This makes it even more interesting to find out whether there exists a same-string NIZK argument with statistical security on at least one side, namely the ZK side.

2.3 Bilinear Groups and the Hardness Assumptions

We use the standard setting of bilinear groups. Let \mathcal{BGG} be a *bilinear-group generator* that (takes as input the security parameter κ and) outputs $(\mathbb{G}, \mathbb{H}, q, g, e)$ where \mathbb{G} and \mathbb{H} are both groups of prime order q , g is a generator of \mathbb{G} , and e is a non-degenerate bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{H}$, also known as a *pairing*, meaning that $e(g^a, g^b) = e(g, g)^{ab}$ for any $a, b \in \mathbb{Z}_q$ and $e(g, g) \neq 1_{\mathbb{H}}$. Although we use a *symmetric* pairing for simplicity, all constructions in this paper also accommodate to an asymmetric pairing $\mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{H}$, as long as an efficiently computable homomorphism from \mathbb{G}_2 to \mathbb{G}_1 exists, simply by using \mathbb{G}_2 and applying the homomorphism when the pairing needs to be computed. We refer to [24] for a discussion on the existence and efficiency of pairings.

We assume the Discrete-Log Assumption, DLA, that for a random $h \in \mathbb{G}$ it is hard to compute $w \in \mathbb{Z}_q$ with $h = g^w$. In some cases, we also assume the Diffie-Hellman Inversion Assumption, DHIA, which states that, for a random $h = g^w \in \mathbb{G}$, it is hard to compute $g^{1/w}$. Formally, these assumptions for a bilinear-group generator \mathcal{BGG} are stated as follows. In order to simplify notation, we abbreviate the output $(\mathbb{G}, \mathbb{H}, q, g, e)$ of \mathcal{BGG} by pub (for “public parameters”).

Assumption 2.4 (DLA). For every non-uniform poly-time algorithm \mathcal{A}

$$P[pub \leftarrow \mathcal{BGG}, h \leftarrow \mathbb{G}, w \leftarrow \mathcal{A}(pub, h) : g^w = h] \leq \text{negl}.$$

Assumption 2.5 (DHIA). For every non-uniform poly-time algorithm \mathcal{A}

$$P[pub \leftarrow \mathcal{BGG}, h \leftarrow \mathbb{G}, g^{1/w} \leftarrow \mathcal{A}(pub, h) : g^w = h] \leq \text{negl}.$$

Furthermore, we assume XKEA, a variant of the Knowledge-of-Exponent Assumption KEA, (referred to as KEA3 respectively KEA1 in [5]). KEA informally states that given $\hat{g} = g^x \in \mathbb{G}$ with unknown discrete-log x , the only way to efficiently produce a pair $A, \hat{A} \in \mathbb{G}$ such that $\hat{A} = A^x$ for the same x is by choosing some $a \in \mathbb{Z}_q$ and computing $A = g^a$ and $\hat{A} = \hat{g}^a$. XKEA states that given $\hat{g} = g^x \in \mathbb{G}$ as well as another pair h and $\hat{h} = h^x$ with the same unknown discrete-log x , the only way to efficiently produce a pair A, \hat{A} such that $\hat{A} = A^x$ is by choosing $a, \alpha \in \mathbb{Z}_q$ and computing $A = g^a h^\alpha$ and $\hat{A} = \hat{g}^a \hat{h}^\alpha$. Formally, KEA and XKEA are phrased by assuming that for every algorithm which outputs A and \hat{A} as required, there exists an *extractor* which outputs a (and α in case of XKEA) when given the same input and randomness.

Assumption 2.6 (KEA). For every non-uniform poly-time algorithm \mathcal{A} there exists a non-uniform poly-time algorithm $\mathcal{X}_{\mathcal{A}}$, the *extractor*, such that

$$P[pub \leftarrow \mathcal{BGG}, x \leftarrow \mathbb{Z}_q, (A, \hat{A}; a) \leftarrow (\mathcal{A} \parallel \mathcal{X}_{\mathcal{A}})(pub, g^x) : \hat{A} = A^x \wedge A \neq g^a] \leq \text{negl}.$$

Recall that $(A, \hat{A}; a) \leftarrow (\mathcal{A} \parallel \mathcal{X}_{\mathcal{A}})(pub, g^x)$ means that \mathcal{A} and $\mathcal{X}_{\mathcal{A}}$ are executed on the same input (pub, g^x) and the *same random tape*, and \mathcal{A} outputs (A, \hat{A}) whereas $\mathcal{X}_{\mathcal{A}}$ outputs a .

Assumption 2.7 (XKEA). For every non-uniform poly-time algorithm \mathcal{A} there exists a non-uniform poly-time algorithm $\mathcal{X}_{\mathcal{A}}$, the *extractor*, such that

$$P \left[\begin{array}{l} pub \leftarrow \mathcal{BGG}, x \leftarrow \mathbb{Z}_q, h \leftarrow \mathbb{G}, \\ (A, \hat{A}; a, \alpha) \leftarrow (\mathcal{A} \parallel \mathcal{X}_{\mathcal{A}})(pub, g^x, h, h^x) \end{array} : \hat{A} = A^x \wedge A \neq g^a h^\alpha \right] \leq \text{negl}.$$

It is well known that DLA holds for generic groups (see e.g. [38]), i.e., for groups which allow to operate on the group elements only by applying the group operations (multiplication and inversion), but not by making use of the specific representation of the group elements. It is not so hard to see that this result extends to generic bilinear groups \mathbb{G} that come with a pairing $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{H}$, and thus additionally allow to apply the pairing to elements in \mathbb{G} and the group operations to elements in \mathbb{H} . We prove in Section 6 that also KEA and XKEA hold for generic bilinear groups.

We would also like to point out that we only depend on XKEA for “proof-technical” reasons: our perfect NIZK argument still *appears* to be secure even if XKEA should turn out to be false (for the particular generator \mathcal{BGG} used), but we cannot prove it anymore formally. This is in contrast to how KEA and XKEA are used in [28] respectively [5] for 3-round ZK, where there seems to be no simulator anymore as soon as KEA is false.

3 A Perfect NIZK Argument for SAT

3.1 Handling Multiplication Gates

Let $(\mathbb{G}, \mathbb{H}, q, g, e)$ be generated by \mathcal{BGG} , as described in Section 2.3 above. Furthermore, let $h = g^w$ for a random $w \in \mathbb{Z}_q$ which is unknown to anybody. Consider a prover who announces an arithmetic circuit over \mathbb{Z}_q and who wants to prove in NIZK that there is a satisfying input for it. Following a standard design principle, where the prover commits to every input value using Pedersen’s commitment scheme with “basis” g and h as well as to every intermediate value of the circuit when evaluating it on the considered input, the problem reduces to proving the consistency of the multiplication gates in NIZK (whereas the addition gates come for free due to the homomorphic property of Pedersen’s commitment scheme).

Concretely, though slightly informally, given commitments $A = g^a h^\alpha$, $B = g^b h^\beta$ and $C = g^c h^\gamma$ for values a, b and $c \in \mathbb{Z}_q$, respectively, the prover needs to prove in NIZK that $c = a \cdot b$. Note that

$$\begin{aligned} e(A, B) &= e(g^a h^\alpha, g^b h^\beta) = e(g, g)^{ab} e(g, h)^{a\beta + \alpha b} e(h, h)^{\alpha\beta} \quad \text{and} \\ e(C, g) &= e(g^c h^\gamma, g) = e(g, g)^c e(g, h)^\gamma \end{aligned}$$

and hence, if indeed $c = a \cdot b$, then

$$e(A, B)/e(C, g) = e(g, h)^{a\beta + \alpha b - \gamma} e(h, h)^{\alpha\beta} = e(g^{a\beta + \alpha b - \gamma} h^{\alpha\beta}, h). \quad (1)$$

Say that, in order to prove that $c = a \cdot b$, the prover announces $P = g^{a\beta + \alpha b - \gamma} h^{\alpha\beta}$ and the verifier accepts if and only if P is *satisfying* in that

$$e(A, B)/e(C, g) = e(P, h).$$

Then, by the above observations it is immediate that an honest verifier accepts the correct proof of an honest prover. Also, it is quite obvious that a simulator which knows w can “enforce” $c = a \cdot b$ by “cheating” with the commitments, and thus perfectly simulate a satisfying P for the multiplication gate. Note that the simulator needs to know *some* opening of the commitments in order to simulate P ; this though is good enough for our purpose. For completeness, we address this issue again in Section 4 and show a version which allows a full-fledged simulation. Finally, it appears to be hard to come up with a satisfying P unless one can indeed open A, B and C to a, b and c such that $c = a \cdot b$. Concretely, the following holds.

Lemma 3.1. *Given openings of A , B and C to a , b and c , respectively, with $c \neq a \cdot b$, and given an opening of a satisfying P , one can efficiently compute w .*

Proof. Let $P = g^p h^\varpi$ be the given opening of P . Then, inheriting the notation from above,

$$e(A, B)/e(C, g) = e(g^a h^a, g^b h^b)/e(g^c h^c, g) = e(g, g)^{ab-c} e(g, h)^{a\beta+\alpha b-\gamma} e(h, h)^{\alpha\beta}.$$

and

$$e(A, B)/e(C, g) = e(P, h) = e(g^p h^\varpi, h) = e(g, h)^p e(h, h)^\varpi$$

are two different representations of the same element in \mathbb{H} with respect to the “basis” $e(g, g)$, $e(g, h) = e(g, g)^w$, $e(h, h) = e(g, g)^{w^2}$. This allows to compute w by solving a quadratic equation in \mathbb{Z}_q . \square

The need for an opening of P can be circumvented by basing security on DHIA rather than DLA as stated in the following lemma.

Lemma 3.2. *Given openings of A , B and C to a , b and c , respectively, with $c \neq a \cdot b$, and given a satisfying P , one can efficiently compute $g^{1/w}$.*

Proof. For a satisfying P it holds that

$$e(P, h) = e(A, B)/e(C, g) = e(g, g)^{ab-c} e(g, h)^{a\beta+b\alpha-\gamma} e(h, h)^{\alpha\beta}$$

and thus, when $c \neq a \cdot b$ as assumed, the following equalities follow one after the other.

$$\begin{aligned} e(g, g) &= e((P g^{-a\beta-b\alpha+\gamma} h^{-\alpha\beta})^{1/(ab-c)}, h) \\ e(g^{1/w}, g) &= e((P g^{-a\beta-b\alpha+\gamma} h^{-\alpha\beta})^{1/(ab-c)}, g) \\ g^{1/w} &= (P g^{-a\beta-b\alpha+\gamma} h^{-\alpha\beta})^{1/(ab-c)} \end{aligned}$$

Thus, $g^{1/w}$ can be computed from the available information. \square

It remains to argue that a (successful) prover can indeed open all the necessary commitments. This can be enforced as follows. Instead of committing to every value s by $S = g^s h^\sigma$, the prover has to commit to s by $S = g^s h^\sigma$ and $\hat{S} = \hat{g}^s \hat{h}^\sigma$, where $\hat{g} = g^x$ for a random $x \in \mathbb{Z}_q$ and $\hat{h} = h^x$ (with the same x). Note that the same randomness σ is used for computing S and \hat{S} , such that $\hat{S} = S^x$; this can be verified using the bilinear map: $e(\hat{S}, g) = e(S, \hat{g})$. XKEA now guarantees that for every correct double commitment (S, \hat{S}) produced by the prover, he knows (respectively there exists an algorithm that extracts) s and σ such that $S = g^s h^\sigma$.

Based on the above observations, we construct and prove secure an adaptively-sound perfect NIZK argument for circuit-SAT in the next section.

3.2 The Perfect NIZK Scheme

The NIZK scheme for circuit-SAT is given in Figure 1. Note that we assume an arithmetic circuit C over \mathbb{Z}_q (rather than a binary circuit), but of course it is standard to “emulate” a binary circuit by an arithmetic one.

Theorem 3.3. *$(\mathcal{G}, \mathcal{P}, \mathcal{V})$ from Fig. 1 is an adaptively-sound perfect NIZK argument for circuit-SAT, assuming XKEA and DLA.*

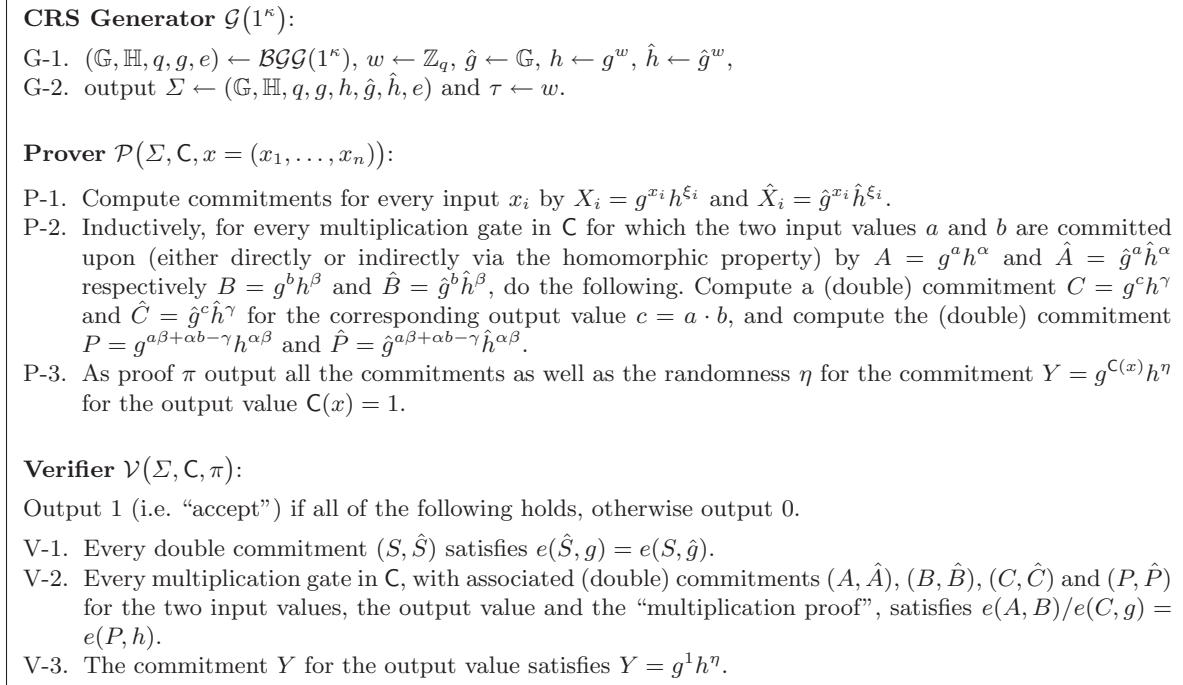


Fig. 1. Perfect NIZK argument for circuit-SAT

Proof. Completeness is straightforward using observation (1). Also, perfect ZK is easy to see. Indeed, the simulator \mathcal{S} can run \mathcal{P} with a default input for x , say $o = (0, \dots, 0)$, and then simply open the commitment Y for the output value $y = C(o)$ (which is likely to be different from 1) to 1 using the trapdoor w . Since Pedersen’s commitment scheme is perfectly hiding, and since P and \hat{P} computed in step P-2. for every multiplication gate are uniquely determined by A , B , and C , it is clear that this simulation is perfectly indistinguishable from a real execution of \mathcal{P} .

It remains to argue soundness. Assume there exists a dishonest poly-time prover \mathcal{P}^* , which on input the CRS Σ outputs a circuit C^* together with a proof π^* such that with non-negligible probability, C^* is not satisfiable but $\mathcal{V}(\Sigma, C^*, \pi^*)$ outputs 1. By XKEA, there exists a poly-time extractor $\mathcal{X}_{\mathcal{P}^*}$ such that when run on the same CRS and the same random tape as \mathcal{P}^* , the extractor $\mathcal{X}_{\mathcal{P}^*}$ outputs the opening information for all (double) commitments in the proof with non-negligible probability.² Concretely, for every multiplication gate and the corresponding commitments A , B , C and P , the extractor $\mathcal{X}_{\mathcal{P}^*}$ outputs $a, \alpha, b, \beta, c, \gamma, \rho, \varpi$ such that $A = g^a h^\alpha$, $B = g^b h^\beta$, $C = g^c h^\gamma$ and $P = g^\rho h^\varpi$. If \mathcal{P}^* succeeds in forging a proof for an unsatisfiable circuit, then there obviously must be an inconsistent multiplication gate with inputs a and b and output $c \neq a \cdot b$. (Note that since addition gates are processed using the homomorphic property, there cannot be an inconsistency in an addition gate.) But this contradicts DLA by Lemma 3.1. \square

Remark 3.4. The NIZK argument from Fig. 1 actually provides ZK even for the case of adaptive corruption of prover and verifier, which is a stronger flavor of ZK than guaranteed by Definition 2.1. It guarantees that \mathcal{S} cannot only perfectly simulate a proof π for any circuit C , but

² As a matter of fact, XKEA guarantees that for every double commitment there exists an efficient extractor that outputs the opening for that particular commitment with non-negligible probability; but running all these extractors together allows to efficiently extract for all commitments simultaneously with non-negligible probability (since the size of C must be polynomially bounded). See Appendix A for a formal proof.

when later given a satisfying input x for \mathbb{C} , it can also provide the randomness that explains how π could have been generated by running \mathcal{P} on witness x .

Remark 3.5. It is reasonable to assume that one can efficiently verify that, for given $(\mathbb{G}, \mathbb{H}, q, g, e)$, \mathbb{G} and \mathbb{H} are groups of order q , g generates \mathbb{G} , and e is a non-degenerate bilinear map. Then, the CRS $\Sigma = (\mathbb{G}, \mathbb{H}, q, g, h, \hat{g}, \hat{h}, e)$ may be generated by the (possibly dishonest) verifier, together with an (interactive) ZK proof of the knowledge of w with $g^w = h$, which can be done very efficiently by using the 4-round ZK proof from [14] for instance. The prover additionally needs to check if $e(\hat{g}, h) = e(g, \hat{h})$. Hence, the set-up of the CRS requires no honest party nor any expensive 2-party (or multi-party) computation. This sort of set-up is captured by the so-called registered public-key model [3]. If the proof of knowledge of w is omitted, so that the verifier only publishes the CRS Σ , then the argument is still witness indistinguishable due to the perfect hiding property of the commitments. Our scheme can thus also be appreciated as a non-interactive witness indistinguishable argument system in the bare public-key model [12], or equivalently as a (computationally sound) *zap* [20].

Remark 3.6. By omitting \hat{P} (and the corresponding verifications), one can obtain a slightly more efficient protocol based on the possibly stronger assumption DHIA instead of DLA. The security can be proven in exactly the same way based on Lemma 3.2 instead of Lemma 3.1. Furthermore, if one is willing to trade XKEA by a new assumption (though of similar flavor, but which can also be proven in the generic model) that the only way to come up with A, B, C and P such that $e(A, B)/e(C, g) = e(P, h)$ is by choosing A, B and C as commitments of a, b and $c = a \cdot b$, respectively, then one can get a NIZK scheme where (not counting the unavoidable commitments A, B and C) the proof for each multiplication gate consists of only 1 group element, P . Note that this requires less communication than using standard interactive ZK techniques in combination with the Fiat-Shamir heuristic [22].

4 Efficient Proof for Relations Among Commitments

We again consider the problem of proving that a Pedersen commitment $C = g^c h^\gamma$ “contains” the product $c = a \cdot b$ of a and b committed to by $A = g^a h^\alpha$ respectively $B = g^b h^\beta$. Recall that the multiplication proof discussed in Section 3.1, consisting of P such that $e(A, B)/e(C, g) = e(P, h)$ (and maybe the corresponding \hat{P}), can only be simulated if the simulator knows *some* openings of A, B and C . This was good enough for the application to NIZK for SAT, as in this case all commitments may be prepared by the simulator. However, for other applications, it might be desirable to have a similarly efficient non-interactive multiplication proof which allows a fully-fledged simulation, i.e., which can be simulated for any *given* A, B and C . We show in this section a modification of the multiplication proof of Section 3.1 which has this property.

The setting is as in the previous section: we assume that a CRS $\Sigma = (\mathbb{G}, \mathbb{H}, q, g, h, \hat{g}, \hat{h}, e)$ has been properly set up and is publicly available. Per default, the prover is required to provide \hat{A}, \hat{B} and \hat{C} for the commitments A, B and C in question, and the verifier should check if $e(\hat{g}, A) = e(g, \hat{A})$ etc., so that the opening of A, B and C can be extracted via XKEA. Note that such \hat{A}, \hat{B} and \hat{C} can be computed from A, B and C and $x = \log_g \hat{g} \in \mathbb{Z}_q$. Thus, the ZK simulator who knows x can simulate \hat{A}, \hat{B} and \hat{C} without knowing the openings of the original commitments. For ease of description, these “hatted” commitments and corresponding verifications are treated implicitly hereafter. We begin with a simple relation for proving that a commitment A can be opened to zero.

Open to Zero ($a = 0$): For $A = g^0 h^\alpha$, the prover publishes $P = g^\alpha$. The verifier accepts if $e(g, A) = e(h, P)$.

It is obvious that an honest verifier accepts the correct proof of an honest prover. ZK is straightforward: the simulator who knows w can compute $P = A^{1/w}$ (without knowing an opening of A). Finally, given an opening (a, α) of A and a satisfying P , i.e., such that $e(g, A) = e(h, P)$, if $a \neq 0$ then one can efficiently compute $g^{1/w}$. This follows from the following equalities:

$$\begin{aligned} e(g, g^a h^\alpha) &= e(h, P) \\ e(g, g^a) &= e(g, P g^{-\alpha})^w \\ e(g, g^{1/w}) &= e(g, (P g^{-\alpha})^{1/a}) \end{aligned}$$

and thus $g^{1/w} = (P g^{-\alpha})^{1/a}$.

The above protocol for opening to zero can be easily applied to show equality ($a = b$) and addition ($a + b = c$) by replacing A in the above protocol with A/B and AB/C , respectively.

We next show a proof system for multiplicative relation $a \cdot b = c$. Recall that the goal is to have a multiplication proof which allows a simulation for any A , B and C given to the simulator.

Multiplication ($ab = c$): The prover publishes $P = (R, S, T)$ such that $R = h^r$, $S = g^r$ for random r and $T = g^{a\beta + \alpha b - \gamma - ar} h^{\alpha\beta - \alpha r}$. The verifier accepts if $e(g, R) = e(h, S)$ and $e(A, B)/e(g, C) = e(A, R)e(h, T)$.

Completeness is verified by seeing that the first verification equation follows from $e(g, R) = e(g, h^r) = e(g^r, h) = e(h, S)$, and the second from

$$e(A, B)/e(g, C) = e(g, g)^{ab-c} e(g, h)^{a\beta + \alpha b - \gamma} e(h, h)^{\alpha\beta}$$

in combination with

$$\begin{aligned} e(A, R) e(h, T) &= e(g^a h^\alpha, h^r) e(h, g^{a\beta + \alpha b - \gamma - ar} h^{\alpha\beta - \alpha r}) \\ &= e(g^a, h^r) e(h^\alpha, h^r) e(h, g^{a\beta + \alpha b - \gamma - ar}) e(h, h^{\alpha\beta - \alpha r}) \\ &= e(g, h)^{a\beta + \alpha b - \gamma} e(h, h)^{\alpha\beta}, \end{aligned}$$

which gives the desired equality if indeed $ab - c = 0$.

Fully-fledged ZK and soundness are captured by following Lemma 4.1 and 4.2, respectively.

Lemma 4.1. *Given A , B and C , one can efficiently simulate random R , S and T such that $e(g, R) = e(h, S)$ and $e(A, B)/e(g, C) = e(A, R)e(h, T)$.*

Proof. Given the trapdoor w , the simulator picks random u and sets $R = B h^u$, $S = R^{1/w}$, and $T = A^{-u} C^{-1/w}$. As in the real proof, (S, R, T) is uniformly distributed subject to the verification equations:

$$e(h, S) = e(h, R^{1/w}) = e(g, R)$$

and

$$\begin{aligned} e(A, R) e(h, T) &= e(A, B h^u) e(h, A^{-u} C^{-1/w}) \\ &= e(A, B) e(A^u, h) e(h, A^{-u}) e(g, C^{-1}) \\ &= e(A, B)/e(g, C). \end{aligned}$$

Thus, the simulation is perfect. □

Lemma 4.2. *Given openings of A , B and C to a , b and c , respectively, with $c \neq a \cdot b$, and given a satisfying $P = (R, S, T)$, one can efficiently compute $g^{1/w}$.*

Proof. We first observe that R and S constitute a proof of zero-opening. Hence we can say that $R = h^r$ for some r . Furthermore, we can extract such r by applying XKEA to R and S since they are in correct relation verified by $e(g, R) = e(h, S)$. Now, for a , α , b , β , c , γ and r , we see the following holds from the second verification equation:

$$\begin{aligned} e(A, B)/e(g, C) &= e(A, R) e(h, T) \\ e(g^a h^\alpha, g^b h^\beta) e(g, g^{-c} h^{-\gamma}) &= e(g^a h^\alpha, R) e(h, T) \\ e(g, g^{ab} h^{a\beta}) e(h, g^{\alpha b} h^{\alpha\beta}) e(g, g^{-c}) e(g^{-\gamma}, h) &= e(g^{ra}, h) e(h, R^\alpha) e(h, T) \\ e(g, g^{ab-c}) &= e(h, g^{-a\beta-\alpha b+ra+\gamma} h^{-\alpha\beta} R^\alpha T) \\ e(g, g^{1/w}) &= e(g, g^{-a\beta-\alpha b+ra+\gamma} h^{-\alpha\beta} R^\alpha T)^{1/(ab-c)} \end{aligned}$$

Therefore, if $ab \neq c$, one can compute $g^{1/w} = (g^{-a\beta-\alpha b+ra+\gamma} h^{-\alpha\beta} R^\alpha T)^{1/(ab-c)}$, which contradicts to DHIA. \square

Now, we need to discuss what kind of NIZK arguments these protocols formally are. The crucial issue stems from the fact that Pedersen’s commitment scheme is unconditionally hiding and thus the language of all triples (A, B, C) which allow an opening with $a \cdot b = c$ is trivial. Therefore, proving a statement among these commitments only makes sense in terms of *proof of knowledge*. By Lemma 4.2, the “knowledge soundness” can be proven by using the extractor of XKEA as knowledge extractor. Accordingly, the quality of the extractor of XKEA immediately affects to the quality of the knowledge extractor. Since XKEA only provides a non-black-box extractor, the best the protocol can achieve is a proof of knowledge characterized by a non-black-box knowledge extractor.

Let \mathcal{R} be a binary poly-time relation, which we allow to depend on $(\kappa$ and) Σ in order to capture schemes that prove something about commitments with “basis” g and h , which are part of the CRS. Let $L_{\mathcal{R}} = \{x \mid \exists w : (x, w) \in \mathcal{R}\}$ be the language characterized by \mathcal{R} .

Definition 4.3. *A NIZK proof of knowledge for \mathcal{R} is a NIZK proof (or argument) for $L_{\mathcal{R}}$ such that additionally for every non-uniform poly-time prover \mathcal{P}^* there exists a non-uniform poly-time extractor $\mathcal{E}_{\mathcal{P}^*}$ such that*

$$P\left[(\Sigma, \tau) \leftarrow \mathcal{G}, (x^*, \pi^*; w^*) \leftarrow (\mathcal{P}^* \parallel \mathcal{E}_{\mathcal{P}^*})(\Sigma) : (x^*, w^*) \notin \mathcal{R} \wedge \mathcal{V}(\Sigma, x^*, \pi^*) = 1\right] \leq \text{negl}.$$

Such NIZK proof of knowledge with non-black-box extractor might be weaker than the one with universal black-box extractor originally defined in [36]. This issue is analogue to black-box vs non-black-box ZK where both definitions are widely accepted. Although a stronger definition is in general favorable, a weaker definition has potential to capture nicer schemes with weaker assumptions or even schemes that are impossible otherwise, but still guarantees sufficient security.

The following now follows immediately from the above lemmas.

Theorem 4.4. *The above scheme gives a perfect NIZK proof of knowledge for the relation $\mathcal{R}_{\text{mult}} = \{((A, B, C), (a, \alpha, b, \beta, \gamma)) \in \mathbb{G}^3 \times \mathbb{Z}_q^5 \mid A = g^a h^\alpha, B = g^b h^\beta, C = g^{ab} h^\gamma\}$ under XKEA and DHIA.*

Finally, we note that all the statements in this section can be strengthened to be based on DLA rather than DHIA by additionally providing $\hat{P} = (\hat{R}, \hat{S}, \hat{T})$, similarly as for the proof of SAT in Section 3.

5 The Difficulty of Achieving Statistical NIZK

DIRECT BLACK-BOX REDUCTION PROOF. Since statistical NIZK can only be achieved with computational soundness, the best we can hope for in proving soundness is to be able to show that a dishonest prover that successfully breaks the soundness property can be used in order to solve some hard problem. Usually, such a soundness proof is done by a black-box reduction, i.e., by specifying an oracle algorithm \mathcal{A} which solves some hard problem when allowed to make black-box calls to a successful dishonest prover \mathcal{P}^* . In this section, we show that a restricted sort of black-box reduction proof, which we call *direct* black-box reduction proof, is not possible for proving the soundness.

For simplicity and better readability, we focus on the case where the underlying hard problem is a decisional problem, where an instance generated by a generator \mathcal{G} is assumed to be hard to distinguish from an instance generated by another generator \mathcal{G}^* . A NIZK argument system $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ has a black-box reduction proof if there exists a polynomial-time oracle algorithm \mathcal{A} such that for any dishonest prover \mathcal{P}^*

$$P[X \leftarrow \mathcal{G} : \mathcal{A}^{\mathcal{P}^*}(X) = 1] \geq q, \text{ and} \quad (2)$$

$$P[X^* \leftarrow \mathcal{G}^* : \mathcal{A}^{\mathcal{P}^*}(X^*) = 1] \leq p \quad (3)$$

where p and q only depend on \mathcal{P}^* 's success probability

$$\text{succ}_{\mathcal{P}^*} = P[\Sigma \leftarrow \mathcal{G}, (x^*, \pi^*) \leftarrow \mathcal{P}^*(\Sigma) : x^* \notin L \wedge \mathcal{V}(\Sigma, x^*, \pi^*) = 1]. \quad (4)$$

To explicitly denote the communication between \mathcal{A} and \mathcal{P}^* , we view \mathcal{A} as a sequence of oracle-free algorithms $\mathcal{A}_1, \dots, \mathcal{A}_n$ where \mathcal{A}_i outputs the i -th query to \mathcal{P}^* and \mathcal{A}_{i+1} takes \mathcal{P}^* 's response as an input. \mathcal{A}_i also passes some state information to \mathcal{A}_{i+1} . The last algorithm \mathcal{A}_n outputs the answer. Again, for simplicity and readability, we restrict ourselves to the case where \mathcal{A} queries \mathcal{P}^* only once. The general case can be treated in a similar though more tedious way. We thus view \mathcal{A} as a pair $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ of oracle-free algorithms and rewrite (2) and (3) as

$$P[(\Sigma, a) \leftarrow \mathcal{A}_1 \circ \mathcal{G}, (x^*, \pi^*) \leftarrow \mathcal{P}^*(\Sigma) : \mathcal{A}_2(x^*, \pi^*, a) = 1] \geq q, \text{ and} \quad (5)$$

$$P[(\Sigma, a) \leftarrow \mathcal{A}_1 \circ \mathcal{G}^*, (x^*, \pi^*) \leftarrow \mathcal{P}^*(\Sigma) : \mathcal{A}_2(x^*, \pi^*, a) = 1] \leq p \quad (6)$$

where a is the state information. We pose the following restrictions on \mathcal{A} and the underlying hard problem.

- I. The underlying hard problem is a standard decisional problem. Precisely, \mathcal{G} and \mathcal{G}^* are poly-time algorithms, and if $\text{succ}_{\mathcal{P}^*}$ is noticeable so is \mathcal{A} 's advantage $\text{adv} = q - p$.
- II. Algorithm \mathcal{A} embeds the instance of the underlying hard problem into the CRS in a “structure-preserving” way. This is to be understood as follows.
 - (a) The composed algorithm $\mathcal{A}_1 \circ \mathcal{G}$ outputs a *valid* Σ whose distribution is statistically close to the one generated by \mathcal{G} , whereas $\mathcal{A}_1 \circ \mathcal{G}^*$ outputs an *invalid* Σ^* such that $P[\mathcal{G} = \Sigma^*] = \text{negl}$.
 - (b) Given some hints from \mathcal{G} and \mathcal{A}_1 , it is possible to efficiently compute a legitimate trapdoor for Σ . Precisely, there exists a poly-time algorithm $\mathcal{T}_{\mathcal{G}, \mathcal{A}_1}$ such that $(\Sigma, a, \tau) \leftarrow \mathcal{T}_{\mathcal{G}, \mathcal{A}_1}$ is statistically close to (Σ, a, τ) generated as $(\Sigma, a) \leftarrow \mathcal{A}_1 \circ \mathcal{G}$ and $(\Sigma', \tau) \leftarrow \mathcal{G}$ conditioned on $\Sigma = \Sigma'$.

Definition 5.1. A NIZK argument system $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ has a direct black-box reduction proof for soundness if its reduction algorithm \mathcal{A} and the underlying hard problem satisfy conditions I. and II.

Obviously, this notion does not capture the most general way of proving a statistical NIZK argument sound, even when restricting to black-box reduction proofs; in particular condition II. might appear quite restrictive at first glance. However, all NIZK arguments (with computational or statistical ZK, adaptive or non-adaptive soundness) we are aware of satisfy II.³ As an example, in the adaptive or non-adaptive GOS scheme (using the naming introduced in Section 1.3), a (valid) CRS Σ is a pair g, h of group elements where g and h both have order $n = pq$ and the trapdoor τ is the discrete-log $\log_g h$, and the underlying computational problem is to distinguish such pairs from pairs where h has order q . The reduction algorithm \mathcal{A} simply forwards the instance g, h untouched to \mathcal{P}^* . Thus, II. is clearly satisfied, where $\mathcal{T}_{\mathcal{G}_{gen}, \mathcal{A}_1}$ simply samples g and h such that it knows $\tau = \log_g h$. For the skeptical reader, we discuss two more examples; the satisfied reader may well jump ahead.

The UC GOS scheme is more complicated, in particular soundness is reduced to several hard problems. However, if one focuses on the reduction to the hardness of distinguishing hk generated by $\mathcal{G}_{gen} = Kstat$ from hk generated by $\mathcal{G}_{gen}^* = Kextract$ (borrowing the notation from [27]), then \mathcal{A} embeds hk (untouched) into the CRS $\Sigma = (hk, ck, pk, \sigma)$ by generating the other parts of Σ as prescribed: $(ck, tk) \leftarrow Kcom$, $(pk, dk) \leftarrow Kpseudo$ and $(\sigma, \tau) \leftarrow S_1$, which in particular allows \mathcal{A} to learn τ , and which is such that Σ is valid if and only if hk is generated by $\mathcal{G}_{gen} = Kstat$. Thus, II. is trivially satisfied.

Similarly, the soundness proof in [34] for the (computational unbounded simulation-sound same-string) NIZK argument in their Theorem 2 is captured: focusing on the reduction to the hardness of distinguishing a random Σ_2 from a pseudorandom $\Sigma_2 = G(d)$ (borrowing the notation from [34]), \mathcal{A} embeds Σ_2 (untouched) into the CRS $\Sigma = (\Sigma_1, \Sigma_2, \Sigma_3)$ by generating Σ_1 (together with τ) and Σ_3 as prescribed. Thus, II. is again trivially satisfied.

In general, it appears to be hard to escape property II: if the reduction algorithm “destroys” the structure of the instance X when embedding it into the CRS, then it appears hard to use \mathcal{P}^* ’s response in order to solve the instance X . But of course, this would be a different approach to try to circumvent the following impossibility result, rather than using a non-standard assumption as we did in Section 3.

IMPOSSIBILITY RESULT. The following theorem states that it is not possible to construct a statistical NIZK argument for an NP-complete language with adaptive soundness (unless $\text{NP} \subset \text{P/poly}$) by means of commonly used reduction techniques, based on standard assumptions.

Theorem 5.2. *There exists no adaptively-sound statistical NIZK argument for an NP-complete language which has a direct black-box reduction proof, unless $\text{NP} \subset \text{P/poly}$.*

Recall, P/poly denotes the family of languages which can be recognized in poly-time by a non-uniform deterministic algorithm. Hence, $\text{NP} \subset \text{P/poly}$ would imply that essentially *any* computational cryptographic scheme is insecure against non-uniform attacks. Also, it is not hard to see that under the reasonable assumption that the correctness of a trapdoor for a CRS can be verified in poly-time by a uniform *deterministic* algorithm, the impossibility result holds even unconditionally.

Proof. Let L be an NP-complete language. Assume a statistical NIZK argument $(\mathcal{G}, \mathcal{P}, \mathcal{V})$ for L which has a direct black-box reduction proof. Hence, there exists a poly-time algorithm \mathcal{A} as described in Definition 5.1. We show how to efficiently recognize L by a non-uniform algorithm.

³ Interestingly, the only scheme we know of which escapes our notion, the adaptive GOS scheme from [27], does not violate II. but I. (in that the underlying problem is assumed to be solvable only with “sub-negligible” probability), and it additionally poses a bound on the size of the instance to be proven.

Fix an arbitrary instance $x \in L$ with witness w and an arbitrary instance x^* not in L . We stress that the following argument does not assume any distribution over the instances and holds for any choice of the instances. The goal will be to construct an efficient non-uniform distinguisher that distinguishes x from x^* for any choice, such that $L \in \text{P/poly}$.

Let \mathcal{S} be a ZK simulator for $(\mathcal{G}, \mathcal{P}, \mathcal{V})$. By definition,

$$P[(\Sigma, \tau) \leftarrow \mathcal{G}, \pi \leftarrow \mathcal{S}(\Sigma, \tau, x) : \mathcal{V}(\Sigma, x, \pi) = 1] \geq 1 - \text{negl} \quad (7)$$

holds. We claim that simulator \mathcal{S} produces accepting proofs even for ill-instance x^* : say

$$P[(\Sigma, \tau) \leftarrow \mathcal{G}, \pi^* \leftarrow \mathcal{S}(\Sigma, \tau, x^*) : \mathcal{V}(\Sigma, x^*, \pi^*) = 1] \geq \frac{1}{2}. \quad (8)$$

If this is not the case then x^* can trivially be recognized as not being in L . Thus, $x^* \notin L$ for which (8) does not hold can be “filtered out” by testing (8) in advance.

Consider a dishonest prover \mathcal{P}^* such that

$$\mathcal{P}^*(\Sigma) = \begin{cases} \tau \leftarrow \mathcal{Q}(\Sigma), \pi^* \leftarrow \mathcal{S}(\Sigma, \tau, x^*), \text{ output } (x^*, \pi^*) & (\text{if } \Sigma \text{ is valid}) \\ \pi \leftarrow \mathcal{P}(\Sigma, x, w), \text{ output } (x, \pi) & (\text{if } \Sigma \text{ is not valid}) \end{cases} \quad (9)$$

where \mathcal{Q} is a (possibly inefficient) algorithm that computes a trapdoor with the correct distribution for a valid CRS Σ .⁴ Since \mathcal{P}^* is a dishonest prover with $\text{succ}_{\mathcal{P}^*} \geq \frac{1}{2}$, and making use of property I, \mathcal{A} satisfies (5) and (6) with a noticeable advantage $\text{adv} = q - p$. By expanding \mathcal{P}^* along with (9) and by using condition IIa, (5) and (6) can be rewritten as

$$P[(\Sigma, a) \leftarrow \mathcal{A}_1 \circ \mathcal{G}_{\text{en}}, \tau \leftarrow \mathcal{Q}(\Sigma), \pi^* \leftarrow \mathcal{S}(\Sigma, \tau, x^*) : \mathcal{A}_2(x^*, \pi^*, a) = 1] \geq p + \text{adv}, \text{ and} \quad (10)$$

$$P[(\Sigma^*, a) \leftarrow \mathcal{A}_1 \circ \mathcal{G}_{\text{en}}^*, \pi \leftarrow \mathcal{P}(\Sigma^*, x, w) : \mathcal{A}_2(x, \pi, a) = 1] \leq p. \quad (11)$$

The next step is to remove the inefficient algorithm \mathcal{Q} . For this, we replace the composed algorithm $\mathcal{A}_1 \circ \mathcal{G}_{\text{en}}$ by $\mathcal{T}_{\mathcal{G}_{\text{en}}, \mathcal{A}_1}$, using property IIb. Equation (10) can now be written as

$$P[(\Sigma, a, \tau) \leftarrow \mathcal{T}_{\mathcal{G}_{\text{en}}, \mathcal{A}_1}, \pi^* \leftarrow \mathcal{S}(\Sigma, \tau, x^*) : \mathcal{A}_2(x^*, \pi^*, a) = 1] \geq p + \text{adv}. \quad (12)$$

Next, in (11), replace $\mathcal{G}_{\text{en}}^*$ with \mathcal{G}_{en} . Due to the hardness of the underlying problem, it follows that

$$P[(\Sigma, a) \leftarrow \mathcal{A}_1 \circ \mathcal{G}_{\text{en}}, \pi \leftarrow \mathcal{P}(\Sigma, x, w) : \mathcal{A}_2(x, \pi, a) = 1] \leq p + \text{negl}. \quad (13)$$

Finally, due to the statistical ZK property, we can replace \mathcal{P} in (13) with simulator \mathcal{S} . In order to create the trapdoor given to \mathcal{S} , we also replace $\mathcal{A}_1 \circ \mathcal{G}_{\text{en}}$ by $\mathcal{T}_{\mathcal{G}_{\text{en}}, \mathcal{A}_1}$, as we did for (12). We now have

$$P[(\Sigma, a, \tau) \leftarrow \mathcal{T}_{\mathcal{G}_{\text{en}}, \mathcal{A}_1}, \pi \leftarrow \mathcal{S}(\Sigma, \tau, x) : \mathcal{A}_2(x, \pi, a) = 1] \leq p + \text{negl}. \quad (14)$$

The above inequalities hold for any $x \in L$ and for any $x^* \notin L$ (for which (8) holds). Comparing (12) and (14), one can see that the composition of the algorithms $\mathcal{T}_{\mathcal{G}_{\text{en}}, \mathcal{A}_1}$, \mathcal{S} and \mathcal{A}_2 constitutes a non-uniform⁵ poly-time probabilistic algorithm for recognizing L . That is, given a challenge instance $x^?$, compute $(\Sigma, a, \tau) \leftarrow \mathcal{T}_{\mathcal{G}_{\text{en}}, \mathcal{A}_1}$ and $\pi \leftarrow \mathcal{S}(\Sigma, \tau, x^?)$, and output $\mathcal{A}_2(x^?, \pi, a)$. Independent repetitions give a large enough success probability. By the NP-completeness, this implies that $\text{NP} \subset \text{BPP/poly} = \text{P/poly}$, where the latter equality was originally shown in [1]. \square

⁴ Such \mathcal{P}^* might be inefficient but it is not a problem since as a black-box reduction, \mathcal{A} should also work for an inefficient \mathcal{P}^* .

⁵ Non-uniform because \mathcal{A} may be non-uniform.

We would like to stress that the above proof does *not* apply to *non*-adaptively-sound NIZK arguments. The reason is that the proof relies on \mathcal{P}^* 's distinguished behavior for choosing x or x^* depending on whether an invalid or valid CRS is given, as in (9).

Also, Theorem 5.2 does not apply to *computational* NIZK arguments. The reason is that in this case a simulated proof need not look like a real one when some side information on CRS Σ is available. Notice that, in (13) and (14), \mathcal{A}_2 is given information a generated together with Σ . This may help \mathcal{A}_2 to distinguish the simulated proof π from the real one if the ZK property is only computational. Hence we cannot argue inequality (14) in such a case.

Furthermore, we would like to emphasize once more that Theorem 5.2 does not apply to the adaptive GOS scheme, nor to the UC GOS scheme. It does not apply to the adaptive GOS scheme because that one uses a non-standard assumption, namely that some decision problem has “sub-negligible” advantage. Tailoring the proof of Theorem 5.2 to such an assumption (and using the *perfect* ZK property) shows that the adaptive GOS scheme implies the existence of a non-uniform algorithm that distinguishes a satisfiable from an unsatisfiable circuit with an advantage larger than exponentially small. But this is no contradiction, since the circuit is required to be sub-linear in size and thus trivially allows such a distinguisher. And Theorem 5.2 does not apply to the UC GOS scheme because in that scheme the dishonest prover is allowed to produce correct proofs for incorrect instances, as long as it is hard to recognize the instance to be incorrect. Tailoring the proof of Theorem 5.2 to such a setting only guarantees a distinguisher for instances that are easy per-se, and thus there is no contradiction.

6 The Security of (X)KEA against Generic Attacks

The notion of a generic algorithm is due to Nechaev [32] and Shoup [38], where it was shown that the discrete-log problem is hard for generic algorithms. Informally, a *generic algorithm* for trying to solve some DL-related problem in a group \mathbb{G} is one that does not exploit and thus does not depend on the actual encoding of the group elements, but only relies on the group structure (and that the encoding is injective). In our context, where \mathbb{G} allows a bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{H}$, we also allow the algorithm to make use of the bilinear map and the group structure in \mathbb{H} .

Formally, a generic algorithm for a bilinear group is an oracle algorithm \mathcal{A} which takes as input a prime q , encodings of elements of \mathbb{Z}_q with respect to a *random* injective encoding function $\sigma : \mathbb{Z}_q \rightarrow S$, and possibly encodings of elements of \mathbb{Z}_q with respect to another random encoding function $\tau : \mathbb{Z}_q \rightarrow T$ (with finite $S, T \subset \{0, 1\}^*$). Furthermore, \mathcal{A} is allowed to make oracle queries in order to compute on encoded group elements: upon a query (`add_in_G`, *sign*, $\sigma(y), \sigma(z)$) the oracle \mathcal{O} replies by $\sigma(y + (-1)^{\text{sign}}z)$ and upon (`add_in_H`, *sign*, $\tau(y), \tau(z)$) by $\tau(y + (-1)^{\text{sign}}z)$, and upon (`pair`, $\sigma(y), \sigma(z)$) the oracle replies by $\tau(y \cdot z)$. Note that the `add`-queries model the group operations in \mathbb{G} and \mathbb{H} , and the `pair`-query models the pairing $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{H}$.

Interestingly, in the literature a generic algorithm \mathcal{A} is typically only allowed to query the oracle on encodings that it has received either as input or as a reply to one of the previous queries, but it is not allowed to take such an encoding and, say, flip the last bit and query the oracle on that encoding. Sometimes (but not always), this is argued by letting the set of encodings (here S and T) be so large that essentially any such query would be invalid anyway. But this also implies that \mathcal{A} cannot sample random group elements without “knowing” their discrete-log. We do not want to make such a restriction, in particular in the context of KEA; even though such a step does not appear to be beneficial, we still feel it should be taken care off in a rigorous analysis. In order to avoid having to deal with invalid encodings, we assume that $S = T = \mathbb{Z}_q$ (actually, the natural representation of \mathbb{Z}_q as strings) and that \mathcal{A} queries \mathcal{O} only on

valid encodings, meaning strings in \mathbb{Z}_q . In some sense this models groups whose elements can be efficiently recognized.

Theorem 6.1. *The assumptions KEA and XKEA over bilinear groups hold with respect to generic algorithms.*

Note that the generic security of KEA in the standard (rather than the bilinear) group setting was concurrently and independently shown by Dent [19].

Proof. Let us first consider KEA. A generic algorithm \mathcal{A} takes as input $\sigma(1)$ and $\sigma(x)$ for a random $x \in \mathbb{Z}_q$, and it should output $\sigma(a)$ and $\sigma(ax)$ for some $a \in \mathbb{Z}_q$. Let m be the (polynomial) number of oracle queries \mathcal{A} makes. It is easy to see that any encoding that \mathcal{A} might use (or receive) in an oracle query or that \mathcal{A} might output is of the form $\sigma(P_k(x, u_1, \dots, u_{2m}))$ respectively $\tau(Q_k(x, u_1, \dots, u_{2m}))$ for multi-variate polynomials $P_k \in \mathbb{Z}_q[X, U_1, \dots, U_{2m}]$ of total degree at most 1 respectively $Q_k \in \mathbb{Z}_q[X, U_1, \dots, U_{2m}]$ of total degree at most 2, and for random (but fixed once and for all P_k and Q_k) pairwise-different $u_1, \dots, u_{2m} \in \mathbb{Z}_q \setminus \{x\}$. Indeed, $\sigma(1)$ and $\sigma(x)$ correspond to the polynomials 1 and X , every encoding that \mathcal{A} uses in a query which is fresh in that it has not been given to \mathcal{A} in a reply (or as input) corresponds to a new variable U_i , and any reply given by the oracle corresponds to a P_k respectively Q_k that is inductively computed as $P_k = P_i \pm P_j$ respectively as $Q_k = Q_i \pm Q_j$ or $Q_k = P_i \cdot P_j$. In particular, it is easy to see that by observing \mathcal{A} 's oracle queries, one can keep track of these polynomials.

We now define the extractor $\mathcal{X}_{\mathcal{A}}$ as follows. $\mathcal{X}_{\mathcal{A}}$ runs \mathcal{A} but keeps track of these polynomials P_k and Q_k ; and if the two polynomials $P_{\text{out}_0}, P_{\text{out}_1} \in \{P_k\}_{k=1\dots m}$ that correspond to the two encodings that \mathcal{A} outputs are of the form $P_{\text{out}_0} = a$ and $P_{\text{out}_1} = aX$, then it outputs a and otherwise 0.

Let us analyze $\mathcal{X}_{\mathcal{A}}$. Obviously, if $\mathcal{X}_{\mathcal{A}}$ fails (in that \mathcal{A} outputs $\sigma(a)$ and $\sigma(ax)$ but $\mathcal{X}_{\mathcal{A}}$ does not output a) then, by the restriction on the degree, $P_{\text{out}_1} \neq X \cdot P_{\text{out}_0}$, whereas P_{out_1} and $X \cdot P_{\text{out}_0}$ coincide when evaluated at (x, u_1, \dots, u_{2m}) . The event that $\mathcal{X}_{\mathcal{A}}$ fails is thus contained in the event \mathcal{E} that at least two distinct polynomials in $\{P_{\text{out}_1}, X \cdot P_{\text{out}_0}\}$, in $\{P_k\}_{k=1\dots m}$ or in $\{Q_k\}_{k=1\dots m}$ evaluate to the same value when applied to (x, u_1, \dots, u_{2m}) . The standard argument for analyzing generic algorithms, using Schwartz' Lemma below, guarantees that the probability of \mathcal{E} is upper bounded by $O(m^2/q)$; since m is polynomial in κ , this proves the claim (for KEA).⁶

The proof for XKEA uses exactly the same reasoning, the only difference is that \mathcal{A} gets four inputs, encodings of $1, x, w, xw \in \mathbb{Z}_q$, which are associated with the polynomials $1, X, W, XW \in \mathbb{Z}_q[X, W, U_1, \dots, U_{2m}]$, and $\mathcal{X}_{\mathcal{A}}$ outputs a, α if P_{out_0} is of the form $P_{\text{out}_0} = a + \alpha W$ (which is the only P_{out_0} which allows $P_{\text{out}_1} = X \cdot P_{\text{out}_0}$). As above we can argue that if $\mathcal{X}_{\mathcal{A}}$ fails then $P_{\text{out}_1}(x, w, u_1, \dots, u_{2m}) = x \cdot P_{\text{out}_0}(x, w, u_1, \dots, u_{2m})$ but $P_{\text{out}_1} \neq X \cdot P_{\text{out}_0}$. Reasoning as above, the probability of this to happen can again be upper bounded by $O(m^2/q)$. \square

Lemma 6.2 (Schwartz [37]). *Let $q \in \mathbb{Z}$ be a prime. For any polynomial P in $\mathbb{Z}_q[X_1, \dots, X_n]$ of total degree at most d , the probability that P vanishes on a uniformly distributed tuple $(x_1, \dots, x_n) \in \mathbb{Z}_q^n$ is at most d/q .*

⁶ To make the argument rigorous, one has to consider a modified ‘‘game’’ where \mathcal{A} is provided with random encodings as long as the corresponding polynomial (rather than its evaluation) is new, and then observe that one can define a joint probability distribution for the original and the modified game which leaves the individual (marginal) distributions intact, and such that \mathcal{E} occurs in the original game if and only if it occurs in the modified one (and thus has the same probability in both cases). In the modified game, however, the polynomials are chosen completely independent of (x, u_1, \dots, u_{2m}) and thus we can apply Schwartz' Lemma.

7 Eliminating XKEA by Pre-Processing

In this section, we briefly discuss the possibility of circumventing XKEA, and to solely rely on standard assumptions, by allowing *pre-processing*. Note that in all of the above results, XKEA is only needed in order to *extract* openings of commitments that are prepared by the possibly dishonest player. Therefore, a possible way to circumvent XKEA is to have all players prepare in a pre-processing phase *random* commitments $U = g^u h^\nu$ in such a way that one can extract the openings (u, ν) of these commitments in the security proof: for instance in the 2-player setting by a standard interactive ZK proof of knowledge (e.g. the 4-round ZK scheme from [14]), or in the multi-player setting with dishonest minority by a simple Pedersen VSS sharing. Then, when the actual NIZK argument needs to be executed, instead of providing for every commitment $S = g^s h^\sigma$ its hatted version $\hat{S} = \hat{g}^s \hat{h}^\sigma$, for every commitment $S = g^s h^\sigma$ the opening $(s + u, \sigma + \nu)$ of SU is provided, where U is a fresh unused random commitment from the pre-processing phase. This then obviously also allows to extract s in the security proof as required. There are some feasibility results about statistical NIZK arguments in the pre-processing model, cf. [35, 29, 16], which rely only on general assumptions but require a complicated pre-processing stage.

Beaver’s pre-processing techniques [4] can be applied in a straightforward way to yield similarly efficient schemes as we do. However, this approach requires the generation of random commitments *with multiplicative relations* in the pre-processing phase, whereas with our techniques purely random commitments, which are potentially easier to prepare, suffice. For instance in the multi-player setting, this is known as the *linear* pre-processing model [18], and when the number of players is small, using the techniques of [13], one can have a once-and-for-all pre-processing stage that allows to produce an *unbounded* number of pseudo-random commitments on the fly.

Acknowledgments

We would like to thank Alexander Dent for useful and interesting remarks regarding our proof of the generic security of KEA and XKEA in the bilinear group setting (Section 6).

References

1. L. M. Adleman. Two theorems on random polynomial time. In *Proceedings of the 19th IEEE Annual Symposium on Foundations of Computer Science*, pages 75–83. IEEE, 1978.
2. W. Aiello and J. Hastad. Statistical zero-knowledge languages can be recognized in two rounds. *Journal of Computer and System Sciences*, 42(3):327–345, 1991.
3. B. Barak, R. Canetti, J. Nielsen, and R. Pass. Universally composable protocols with relaxed set-up assumptions. In *Proceedings of the 45th IEEE Annual Symposium on Foundations of Computer Science*, pages 186–195. IEEE Computer Society, 2004.
4. D. Beaver. Secure multiparty protocols and zero-knowledge proof systems tolerating a faulty minority. *Journal of Cryptology*, 4(2):75–122, 1991.
5. M. Bellare and A. Palacio. The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In M. Franklin, editor, *Advances in Cryptology — CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 273–289. Springer-Verlag, 2004.
6. M. Blum, P. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications. In *Proceedings of the 20th annual ACM Symposium on the Theory of Computing*, pages 103–112. ACM, 1988.
7. M. Blum, A. D. Santis, S. Micali, and G. Persiano. Non-interactive zero-knowledge proof systems. *SIAM Journal on Computing*, 20(6):1084–1118, 1991.
8. G. Brassard, D. Chaum, and C. Crépeau. Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences*, 37(2), 1988.
9. G. Brassard and C. Crépeau. Non-transitive transfer of confidence: A perfect zero-knowledge interactive protocol for SAT and beyond. In *Proceedings of the 27th IEEE Annual Symposium on Foundations of Computer Science*, pages 188–195. IEEE, 1986.

10. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings of the 42nd IEEE Annual Symposium on Foundations of Computer Science*, pages 136–145, 2001.
11. R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally composable security with global setup. In S. Vadhan, editor, *Theory of Cryptography Conference, TCC 2007*, volume 4392 of *Lecture Notes in Computer Science*, pages 61–85. Springer-Verlag, 2007.
12. R. Canetti, O. Goldreich, S. Goldwasser, and S. Micali. Resettable zero-knowledge. In *STOC*, pages 235–244, 2000.
13. R. Cramer, I. Damgård, and Y. Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In J. Kilian, editor, *Theory of Cryptography Conference– TCC'2005*, volume 3378 of *Lecture Notes in Computer Science*, pages 342–362. Springer-Verlag, 2005.
14. R. Cramer, I. Damgård, and P. MacKenzie. Efficient zero-knowledge proofs of knowledge without intractability assumptions. In *Public Key Cryptography, PKC 2000*, volume 1751 of *Lecture Notes in Computer Science*, pages 354–372. Springer-Verlag, 2000.
15. I. Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In J. Feigenbaum, editor, *Advances in Cryptology - CRYPTO '91*, volume 576 of *Lecture Notes in Computer Science*, pages 445–456. Springer-Verlag, 1991.
16. I. Damgård. Non-interactive circuit based proofs and non-interactive perfect zero-knowledge with preprocessing. In *Advances in Cryptology – Eurocrypt '92*, volume 658 of *Lecture Notes in Computer Science*, pages 341–355. Springer-Verlag, 1992.
17. I. Damgård, N. Fazio, and A. Nicolosi. Non-interactive zero-knowledge from homomorphic encryption. In S. Halevi and T. Rabin, editors, *Theory of Cryptography Conference, TCC 2006*, volume 3876 of *Lecture Notes in Computer Science*, pages 41–59. Springer-Verlag, 2006.
18. I. Damgård and Y. Ishai. Constant-round multiparty computation using a black-box pseudorandom generator. In *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 378–394. Springer-Verlag, 2005.
19. A. W. Dent. The hardness of the DHK problem in the generic group model. IACR e-print Archive 2006/156, 2006.
20. C. Dwork and M. Naor. Zaps and their applications. In *Proceedings of the 41st IEEE Annual Symposium on Foundations of Computer Science*, pages 283–293. IEEE, 2000.
21. U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero knowledge proofs based on a single random string. In *Proceedings of the 31st IEEE Annual Symposium on Foundations of Computer Science*, pages 308–317. IEEE, 1990.
22. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *Advances in Cryptology — CRYPTO '86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–199. Springer-Verlag, 1987.
23. L. Fortnow. The complexity of perfect zero-knowledge (extended abstract). In *Proceedings of the 19th annual ACM Symposium on the Theory of Computing*, pages 204–209. ACM, 1987.
24. S. Galbraith, K. Paterson, and N. Smart. Pairings for cryptographers. Technical Report 2006/165, IACR ePrint archive, 2006.
25. O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM*, 1(38):691–729, 1991.
26. J. Groth. Simulation-sound nizk proofs for a practical language and constant size group signatures. In X. Lai and K. Chen, editors, *Advances in Cryptology - ASIACRYPT 2006*, volume 4284 of *Lecture Notes in Computer Science*, pages 444–459. Springer-Verlag, 2006.
27. J. Groth, R. Ostrovsky, and A. Sahai. Perfect non-interactive zero knowledge for NP. In S. Vaudenay, editor, *Advances in Cryptology - EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 339–358. Springer-Verlag, 2006.
28. S. Hada and T. Tanaka. On the existence of 3-round zero-knowledge protocols. In H. Krawczyk, editor, *Advances in Cryptology — CRYPTO '98*, volume 1462 of *Lecture Notes in Computer Science*, pages 354–369. Springer-Verlag, 1998. Full version available from IACR e-print archive 1999/009.
29. J. Kilian, S. Micali, and R. Ostrovsky. Minimum resource zero-knowledge proofs (extended abstract). In G. Brassard, editor, *Advances in Cryptology - CRYPTO '89*, volume 435 of *Lecture Notes in Computer Science*, pages 545–546. Springer-Verlag, 1990.
30. J. Kilian and E. Petrank. An efficient noninteractive zero-knowledge proof system for NP with general assumptions. *Journal of Cryptology*, 11(1):1–27, 1998.
31. M. Naor. On cryptographic assumptions and challenges. In D. Boneh, editor, *Advances in Cryptology - CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 96–109. Springer-Verlag, 2003.
32. V. I. Nechaev. Complexity of a determiniate algorithm for the discrete logarithm. *Mathematical Notes*, 55(2), 1994.

33. R. Pass and A. Shelat. Unconditional characterizations of non-interactive zero-knowledge. In V. Shoup, editor, *Advances in Cryptology - CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 118–134. Springer-Verlag, 2005.
34. A. D. Santis, G. D. Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In J. Kilian, editor, *Advances in Cryptology - CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 566–598. Springer-Verlag, 2001.
35. A. D. Santis, S. Micali, and G. Persiano. Non-interactive zero-knowledge with preprocessing. In S. Goldwasser, editor, *Advances in Cryptology - CRYPTO '88*, volume 403 of *Lecture Notes in Computer Science*, pages 269–282. Springer-Verlag, 1988.
36. A. D. Santis and G. Persiano. Zero-knowledge proofs of knowledge without interaction (extended abstract). In *Proceedings of the 33rd IEEE Annual Symposium on Foundations of Computer Science*, pages 427–436. IEEE, 1992.
37. J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4), 1980.
38. V. Shoup. Lower bounds for discrete logarithms and related problems. In W. Fumy, editor, *Advances in Cryptology — EUROCRYPT '97*, volume 1233 of *Lecture Notes in Computer Science*, pages 256–266. Springer-Verlag, 1997.
39. H. Wee. Lower bounds for non-interactive zero-knowledge. In S. Vadhan, editor, *Theory of Cryptography Conference—TCC'2007*, volume 4392 of *Lecture Notes in Computer Science*, pages 103–117. Springer-Verlag, 2007.

A The Composability of (X)KEA

Lemma A.1. *For every non-uniform poly-time algorithm \mathcal{A} which outputs two poly-size lists $\mathbf{A} = (A_1, \dots, A_r)$ and $\hat{\mathbf{A}} = (\hat{A}_1, \dots, \hat{A}_r)$, KEA implies that there exists a non-uniform poly-time algorithm $\mathcal{X}_{\mathcal{A}}$ which outputs a list $\mathbf{a} = (a_1, \dots, a_r)$ such that*

$$P[\text{pub} \leftarrow \mathcal{BGG}, x \leftarrow \mathbb{Z}_q, (\mathbf{A}, \hat{\mathbf{A}}; \mathbf{a}) \leftarrow (\mathcal{A} \parallel \mathcal{X}_{\mathcal{A}})(\text{pub}, g^x) : \exists i (\hat{A}_i = A_i^x \wedge A_i \neq g^{a_i})] \leq \text{negl},$$

and XKEA implies that there exists a non-uniform poly-time algorithm $\mathcal{X}_{\mathcal{A}}$ which outputs two lists $\mathbf{a} = (a_1, \dots, a_r)$ and $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_r)$ such that

$$P \left[\begin{array}{l} \text{pub} \leftarrow \mathcal{BGG}, x \leftarrow \mathbb{Z}_q, h \leftarrow \mathbb{G}, \\ (\mathbf{A}, \hat{\mathbf{A}}; \mathbf{a}, \boldsymbol{\alpha}) \leftarrow (\mathcal{A} \parallel \mathcal{X}_{\mathcal{A}})(\text{pub}, g^x, h, h^x) \end{array} : \exists i (\hat{A}_i = A_i^x \wedge A_i \neq g^{a_i} h^{\alpha_i}) \right] \leq \text{negl}.$$

Proof. We prove the first part of the claim, the second part can be proven along the same lines. Recall that $r = r(\kappa)$ is polynomial in the security parameter $\kappa \in \mathbb{N}$. Consider the following algorithm \mathcal{B} . \mathcal{B} runs \mathcal{A} (with the same security parameter κ and the same auxiliary input) to produce \mathbf{A} and $\hat{\mathbf{A}}$, it chooses a random $i \in \{1, \dots, r\}$, and it outputs the pair (A_i, \hat{A}_i) . Also, for any fixed integer j , let \mathcal{B}_j be the algorithm that runs \mathcal{A} and outputs the pair (A_j, \hat{A}_j) if $j \leq r(\kappa)$, and otherwise some default output \perp . By KEA, there exists an extractor $\mathcal{X}_{\mathcal{B}}$ for \mathcal{B} which runs in polynomial time t and outputs a_i so that $\hat{A}_i = A_i^x \wedge A_i \neq g^{a_i}$ holds only with negligible probability p . Conditioned on a specific value for i , the probability of $\hat{A}_i = A_i^x \wedge A_i \neq g^{a_i}$ is upper bounded by pr , which is still negligible. Thus, $\mathcal{X}_{\mathcal{B}_j}$, obtained from $\mathcal{X}_{\mathcal{B}}$ by fixing the part of the random tape to j that determines i (if $j \leq r(\kappa)$, otherwise $\mathcal{X}_{\mathcal{B}_j}$ may output \perp), is an extractor for \mathcal{B}_j . Furthermore, all the extractors $\mathcal{X}_{\mathcal{B}_j}$ run in the same polynomial time t .⁷ It follows that running $\mathcal{X}_{\mathcal{B}_1}, \dots, \mathcal{X}_{\mathcal{B}_r}$ requires polynomial time rt , and running them on the same random tape as \mathcal{A} gives a_1, \dots, a_r such that the probability that $\hat{A}_i = A_i^x \wedge A_i \neq g^{a_i}$ for at least one i is upper bounded by pr^2 , which is negligible. \square

⁷ This is the crucial point. If every $\mathcal{X}_{\mathcal{B}_j}$ is allowed to have a different, though still polynomial, running time, like $t_j(\kappa) = \kappa^j$, then the total running time of $\mathcal{X}_{\mathcal{B}_1}, \dots, \mathcal{X}_{\mathcal{B}_r}$ may not be polynomial, e.g. $\kappa^{r(\kappa)}$ in the example.