Efficiency Limitations for Σ -Protocols for Group Homomorphisms*

Endre Bangerter¹, Jan Camenisch², and Stephan Krenn³

¹ Bern University of Applied Sciences, Biel-Bienne, Switzerland
`endre.bangerter@bfh.ch`

² IBM Research — Zurich, Rüschlikon, Switzerland
`jca@zurich.ibm.com`

³ Bern University of Applied Sciences, Biel-Bienne, Switzerland, and
University of Fribourg, Fribourg, Switzerland
`stephan.krenn@bfh.ch`

Abstract. Efficient zero-knowledge proofs of knowledge for group homomorphisms are essential for numerous systems in applied cryptography. Especially, Σ -protocols for proving knowledge of discrete logarithms in known and hidden order groups are of prime importance. Yet, while these proofs can be performed very efficiently within groups of known order, for hidden order groups the respective proofs are far less efficient.

This paper shows strong evidence that this efficiency gap cannot be bridged. Namely, while there are efficient protocols allowing a prover to cheat only with negligibly small probability in the case of known order groups, we provide strong evidence that for hidden order groups this probability is bounded below by $1/2$ for all efficient Σ -protocols not using common reference strings or the like.

We prove our results for a comprehensive class of Σ -protocols in the generic group model, and further strengthen them by investigating certain instantiations in the plain model.

Keywords: Generic Group Model, Σ -Protocols, Proofs of Knowledge, Error Bounds.

1 Introduction

A *Zero-Knowledge Proof of Knowledge (ZK-PoK)* is a two party protocol between a prover and a verifier enabling the prover to convince the verifier that he knows some secret value, without the verifier being able to learn anything about it. More precisely, in a ZK-PoK an honest prover can always convince the verifier, while no malicious prover (not knowing the secret) can do so with a probability larger than some threshold value (the *knowledge error*).

Fundamental results show that there are ZK-PoK for all languages in $\mathcal{NP}$ [2]. Yet, the respective protocols are of theoretical interest only, because executing

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-3-642-11799-2_36](https://doi.org/10.1007/978-3-642-11799-2_36)

* This work was in part funded by the European Community's Seventh Framework Programme (FP7) under grant agreements no. 216499 and 216483. We also refer to the full version of this paper [3].

them once is either computationally and communicationally too expensive for real world use, or enables the prover to cheat with a high probability. In the latter case, the protocols have to be repeated numerous times to reduce the knowledge error (remember that r repetitions of a ZK-PoK with knowledge error κ result in a protocol with knowledge error κ^r), and thus they become inefficient again.

A (group) homomorphism is a mapping between two groups $\mathcal{G}$ and $\mathcal{H}$ satisfying $\phi(a+b) = \phi(a) \cdot \phi(b)$ for all $a, b \in \mathcal{G}$. Proving knowledge of a preimage under a homomorphism (i.e., of w satisfying $x = \phi(w)$) can often be done very efficiently by using the so-called Σ^ϕ -protocol (i.e., the Schnorr [3] or Guillou/Quisquater [4] protocol generalized to arbitrary homomorphisms [5,6,7]). This protocol consists of three messages being exchanged: the prover chooses r at random from the domain of the homomorphism, and sends the *commitment* $t := \phi(r)$ to the verifier. The verifier then chooses a random *challenge* c from a predefined challenge set $\mathcal{C}$, and sends it to the prover, who computes its response $s := r + c \cdot w$. The verifier now accepts the proof, if and only if $\phi(s) = x^c \cdot t$. Standard techniques [8] allow one to transform this protocol into non-interactive versions or so called signatures of knowledge.

The Σ^ϕ -protocol is a very efficient proof of knowledge for many proof goals existing in cryptography (e.g., knowledge of a discrete logarithm in a known order group, or of the plaintext encrypted in a Paillier ciphertext). The reason is that for the respective homomorphisms, a negligibly small knowledge error can be obtained in a *single run* of the Σ^ϕ -protocol. Yet, the situation is different for the important class of exponentiation homomorphisms with hidden order co-domain (e.g., $\phi(\cdot) : \mathbb{Z} \rightarrow \mathbb{Z}_n^* : a \mapsto g^a$, where g is a generator of the quadratic residues modulo n). Such homomorphisms play an important role for many cryptographic applications, e.g., [9,10,11,12,13,14,15,16], including *Direct Anonymous Attestation (DAA)* [17], and the *identity mixer (idemix)* anonymous credential system [18]. In this case, the Σ^ϕ -protocol is only known to be a PoK with knowledge error $1/2$, and hence must be repeated sequentially to get a sufficiently small knowledge error (e.g., 80 sequential repetitions are required to obtain a knowledge error of $1/2^{80}$). The resulting computational and communication costs are much too high for many practical applications.

A number of authors have tried to overcome the above problem by proposing alternative protocols for exponentiation homomorphisms with hidden order co-domain [5,19,20,21,22,23]. All these protocols build on a basic idea put forth by Fujisaki and Okamoto [22], and we thus call them *FO-based* henceforth. Unfortunately, none of these FO-based protocols is fully satisfactory, neither from a practical nor from a theoretical point of view:

- One run of any FO-based protocol is much more expensive than running the Σ^ϕ -protocol once. Moreover, if only standard complexity assumptions (i.e. the Strong RSA Assumption [22]) are made, a recent analysis has revealed that in many cases FO-based protocols are even more expensive than the sequential repetition of the Σ^ϕ -protocol with knowledge error $1/2$ [20].
- The FO-based protocols in [5,19,20,21,22] make use of a common reference string, which is either issued by a trusted third party or generated in an

expensive interactive setup phase. Yet, the presence of common reference strings reduces the modularity, and thus increases the complexity of the security analysis of larger applications (as discussed, e.g., in [23,24,25]). The security proofs for the protocols in [5,19] additionally assume the existence of ideal hash functions, and thus only hold true in the random oracle model¹.

Because of these disadvantages, the natural question arises *whether it is necessary to use FO-based protocols at all?* After all, the possibilities of Σ -protocols have not yet been explored thoroughly, and it could be possible that a novel, cleverly designed Σ -protocol or even the existing Σ^ϕ -protocol could be used to overcome the current efficiency limitations. (We note that the latter could be quite possible, if one could find a new knowledge extractor working for the Σ^ϕ -protocol with a suitably chosen challenge set that allows one to obtain a small knowledge error in a single execution of the protocol.).

Contribution and Results. In this paper we are aiming at answering this question. We provide ample evidence suggesting that the known minimal knowledge error of the Σ^ϕ -protocol cannot be underrun, neither by a better knowledge extractor for the Σ^ϕ -protocol nor by any other Σ -protocol. In particular, our results indicate that using Σ -protocols the knowledge error of $1/2$ cannot be decreased for exponentiation homomorphisms with hidden order co-domain.

More precisely, we first consider PoK based on Σ -protocols in the generic group model. That is, Σ -protocols where prover, verifier, and knowledge extractor are generic algorithms that can only access the homomorphism and its domain and co-domain through an oracle. We then show that there are lower bounds on the knowledge error for (almost) arbitrary Σ -protocols. These lower bounds on the knowledge error in turn imply efficiency limitations for most possible protocol instances. Roughly, these follow by the fact that a PoK with a large knowledge error needs to be repeated sequentially to reduce the knowledge error, which results in a high computational and communicational overhead. Within the generic group model our efficiency analysis shows that the existing Σ^ϕ -protocol is *optimal* and there cannot be another, more efficient Σ -protocol.

We further complement our results by proving lower bounds on the knowledge error of the Σ^ϕ -protocol in the plain model. First, for homomorphisms of the form $w \mapsto w^e$ in RSA groups we show that $1/d$ is a lower bound on the knowledge error, where d is the smallest divisor of e . Then, we show that for exponentiation homomorphisms with hidden order co-domain, $1/2$ is a lower bound on the knowledge error for all knowledge extractors structurally related to the only one currently known. These results are in accord with those in the generic model and again suggest that the knowledge error that is currently known to be achievable and the associated efficiency limitations cannot be underrun.

Finally, we note that our results do not rule out entirely the possibility to obtain efficient PoK using Σ -protocols. On the one hand, we describe a large number of cases (i.e., instances of Σ -protocols) where this is indeed impossible,

¹ For completeness, we note that while the protocol in [23] yields ZK-PoK in the plain model, it is by far too inefficient for practical usage.

indicating that there are inherent efficiency limitations for Σ -protocols. On the other hand, the cases that are not covered by our results also seem to be valuable, since they provide cues for protocol designers on how it could be possible to conceive novel Σ -protocols that overcome current efficiency limitations.

Related Work. Given the abundant usage of Σ -protocols, very little work on their theoretical foundations has been done. Shoup [26] shows that the knowledge error of $1/2$ for homomorphisms of the form $\phi(w) = w^{2^t}$ in RSA groups cannot be improved. One of our results in the plain model extends this to arbitrary exponents. Further, parts of our results are based on unpublished results of one of the authors [5]. Apart from this we are not aware of any other work on efficiency limitations of Σ -protocols. Yet, technically we make use of generic group proof techniques devised by Shoup [26] as well as the extension of these techniques to groups of hidden order by Damgård/Koprowski [27].

The generic group model goes back to Nechaev and Shoup [28,29]. It has been extensively used since then to provide evidence for the security of various cryptographic systems, e.g., [27,28,29,30,31,32,33,34,35,36,37,38]. The model is often criticized, because of the risk of lulling a user in a false sense of security. Indeed, there are cases where information only available in the plain model (i.e., obtained from encoding specific properties of the group) can be used to break a system which was proved secure in the generic model [39,40]. Yet, the implications of these observations are different for all the systems cited above than for our results. All the proofs in the former case are used to give evidence for the security of a cryptographic system. Thus, if any of them does not hold true in the plain model, the security of the according system can be flawed, resulting in dire consequences for all applications using the respective scheme. In contrast to this, we use the generic group model in a more conservative way. Namely, we show efficiency limitations on the efficiency of a cryptographic primitive. Thus, if our results do not hold true in the plain model this means that the efficiency of the scheme can be increased, but the security of the scheme is not affected by any means.

We finally remark that our results do not conflict with those in [41]. The authors there show how to build efficient Σ -protocols for certain exponentiation homomorphisms with hidden order co-domain. Yet, their approach is not generic, but rather uses certain properties of the homomorphism at hand. Further, only very few proofs of practical interest can be performed with their technique.

Structure of this Document. In §2 we recap the basic definitions, and introduce the notion of lower bounds and the class of Σ -protocols for which our results hold true. In §3 we then formulate our main result in the generic group model. This result is strengthened in §4, where we give results in the plain model. We finally conclude and point out some open problems in §5.

2 Preliminaries

In §2.1 we give a short introduction to ZK-PoK and briefly discuss the Σ^ϕ -protocol in §2.2. Then, in §2.3 we introduce the notion of lower bounds on the

knowledge error of a protocol. In §2.4 we recap the generic group model we are working in, and finally describe the class of protocols for which our results in the generic group model hold true in §2.5.

2.1 Zero-Knowledge Proofs of Knowledge

After having defined ZK-PoK We recall the widely accepted definition of zero-knowledge proofs of knowledge (ZK-PoK) [42,43]. We by $(P(w), V)(x)$ denote a two party protocol between a prover P and a verifier V with common input x and private input w to P .

Definition 1 (Computational Proof of Knowledge [42,43]). A computational proof of knowledge for a binary relation $\mathcal{R}$ with knowledge error $\kappa(\cdot) : \mathbb{N} \rightarrow [0, 1]$ is a two party protocol $(P(w), V)(x)$, satisfying the following two conditions:

Completeness: The verifier always accepts the proof, if $(x, w) \in \mathcal{R}$.

Soundness: There exists a polynomial $\text{poly}(\cdot)$, and a probabilistic algorithm M (the knowledge extractor) with input x and rewindable black-box access to the prover, such that the following holds true. For every probabilistic polynomial-time (PPT) prover P^* that can make V accept the proof with probability $\varepsilon(x) > \kappa(x)$, M outputs w' satisfying $(x, w') \in \mathcal{R}$ in expected time at most

$$t^+(\varepsilon, \kappa, x) := \frac{\text{poly}(\|x\|)}{\varepsilon(x) - \kappa(\|x\|)},$$

where access to P^* counts as one step only.

The computational aspect of this definition, i.e., the restriction of P^* to be a PPT algorithm, is of importance for our results, as it (almost) allows us to stay in the standard complexity class of PPT algorithms. This issue will also be discussed in §2.3.

A proof of knowledge (PoK) is called *honest verifier zero knowledge (HVZK)*, if no verifier following the protocol is able to gain any information about the secret value w except that it satisfies the stated relation. For a formal description we refer to [43]. There are well known techniques to transform HVZK protocols into protocols which are zero-knowledge also against maliciously behaving verifiers [8].

2.2 The Σ^ϕ -Protocol in Hidden-Order Groups

Most practical applications using ZK-PoK make use of the Σ^ϕ -protocol explained in §1. This allows one to prove knowledge of a preimage w of a public value x under some group homomorphism $\phi(\cdot) : \mathcal{G} \rightarrow \mathcal{H}$. If $\phi(\cdot)$ is an exponentiation homomorphism with hidden order co-domain, e.g., $\phi(\cdot) : \mathbb{Z} \rightarrow \mathbb{Z}_n^* : a \mapsto g^a$ for some RSA modulus n , the domain of the homomorphism is infinite. To circumvent the problem of drawing random values from an infinite set in P 's first step, the random choice $r \in_R \mathcal{G} = \mathbb{Z}$ is substituted by $r \in_R \mathcal{G}' = \{-\Delta w, \dots, \Delta w\}$

with $(\Delta w - \text{ord } \mathcal{G}) / \text{ord } \mathcal{G}$ being negligibly small. The rest of the protocol remains unchanged. This approach can be generalized also to the case $\mathcal{G} = \mathbb{Z}^u$ for some integer u . For more details see, e.g., [5, 23].

It is well known that the Σ^ϕ -protocol is a PoK with knowledge error $1/2$ for exponentiation homomorphisms with hidden order co-domain. For homomorphisms with a co-domain of known order v , and power homomorphisms $(w_1, w_2) \mapsto \psi(w_1) \cdot w_2^e$, the protocol is known to have a knowledge error of $1/d$, where d is the smallest prime dividing v in the former, respectively e in the latter case [6].

2.3 Lower Bounds of the Knowledge Error

Let us now introduce the notion of *lower bounds*, which is a key to our results stated in the following. Intuitively, β is a lower bound of the knowledge error of a protocol, if for this protocol it is not possible to achieve any knowledge error *smaller than or equal to* β :

Definition 2 (Lower Bound). *A function $\beta(\cdot) : \mathbb{N} \rightarrow [0, 1]$ is called a lower bound on the knowledge error of the protocol (P, V) for a binary relation $\mathcal{R}$, if (P, V) is not a computational proof of knowledge for $\mathcal{R}$ for any $\kappa'(\cdot) : \mathbb{N} \rightarrow [0, 1]$ with $\kappa'(\cdot) \leq \beta(\cdot)$.*

An alternative but equivalent characterization is that of $\beta(\cdot)$ being a lower bound if and only if (P, V) is not a computational PoK with knowledge error $\beta(\cdot)$ for the given relation.

All our results on lower bounds are proven by showing that the conditions of the following theorem are satisfied.

Theorem 3 (Sufficient Conditions for Lower Bounds). *Let (P, V) be a two-party protocol, let $\mathcal{R}$ be a binary relation, and let $\beta(\cdot) : \mathbb{N} \rightarrow [0, 1]$ be a function. Then $\beta(\cdot)$ is a lower bound on the knowledge error of (P, V) for $\mathcal{R}$, if the following two conditions are satisfied:*

Uniformity: *There are a polynomial $\text{poly}(\cdot)$ and PPT algorithms P^* and D such that $\varepsilon(x) - \beta(\|x\|) \geq 1/\text{poly}(\|x\|)$ holds for all sufficiently long x generated by D , where $\varepsilon(x)$ is the probability that P^* makes V accept on common input x .*

Hardness: *For all expected PPT algorithms M having rewindable black-box access to P^* , the probability that M outputs a w' with $(x, w') \in \mathcal{R}$ is negligible.*

From the uniformity condition and Definition 1 it follows that any hypothetical knowledge extractor must be an expected PPT algorithm. This is important, as in our results we show that the hardness condition has to be satisfied by showing that otherwise the respective knowledge extractor could be used to break a cryptographic standard assumption, which is typically defined against PPT attackers. Still, we will have to adopt these assumptions in a natural way. As the standard definition of PoK allows the knowledge extractor to be an *expected* time

algorithm [42,43], we have to generalize the class of attackers the cryptographic assumption holds against to *expected* PPT algorithms as well. Yet, we believe that this generalization is reasonable as by Markov's inequality we see that an expected PPT algorithm may only run super-polynomially long for a small fraction of its executions.

2.4 The Generic Group Model and Groups of Hidden Order

Our main result holds in the generic group model, which we briefly recap next.

The *generic group model* is used to analyze the complexity of problems by considering algorithms in groups whose representation does not reveal any information to the algorithm. That is, such an algorithm must not exploit encoding dependent properties of the group, but is restricted to only use group operations. The hardness of a problem in the generic model is a necessary but not sufficient condition for a problem to be hard in the plain model [39,40].

Various formalizations of this model have been proposed [28,29,35,44]. They all have in common that an algorithm does not get the concrete group description, but only handles to group elements (e.g., via random encodings [29] or indices to elements [35]). Further, the algorithm gets access to an oracle. To evaluate a group operation, the algorithm inputs the handles of elements and the operation to perform to the oracle, which then returns the handle of the result. Similarly, a homomorphism $\phi(\cdot) : \mathcal{G} \rightarrow \mathcal{H}$ has to be evaluated through an oracle.

We call an algorithm a *generic homomorphism algorithm* for $\phi(\cdot) : \mathcal{G} \rightarrow \mathcal{H}$, if, through an oracle $\mathcal{O}^{\phi(\cdot)}$, it might perform the following operations.

- + : Evaluation of the group operation within $\mathcal{G}$ or $\mathcal{H}$,
- − : inverting an element within $\mathcal{G}$ or $\mathcal{H}$,
- $\stackrel{?}{=}$: testing the equality of two elements from the same group,
- $\in_{\mathbf{R}}$: choosing a group element uniformly at random within $\mathcal{G}$ and $\mathcal{H}$, and
- $\phi(\cdot)$: evaluating the homomorphism on arbitrary elements $a \in \mathcal{G}$.

When proving our results, we show that any generic algorithm, acting as hypothetical knowledge extractor for a knowledge error smaller than the stated lower bounds, must fail with overwhelming probability. We therefore describe next which operations such an algorithm may perform.

Definition 4 (Generic Black-Box Algorithm). *A generic black-box algorithm is a generic homomorphism algorithm for $\phi(\cdot)$ with oracle $\mathcal{O}^{\phi(\cdot)}$, which additionally has rewindable black-box access to $\mathbf{P}^*$. That is, it can (i) execute $\mathbf{P}^*$, (ii) choose the random inputs of $\mathbf{P}^*$, and (iii) repeatedly reset $\mathbf{P}^*$. Resetting $\mathbf{P}^*$ does not reset $\mathcal{O}^{\phi(\cdot)}$.*

We remark that the black-box property of such an algorithm is exactly the same as for a knowledge extractor according to Definition 1.

Groups of Hidden Order. In the following we will be interested in group homomorphisms with hidden order co-domain (resp., image). Intuitively this means

that the order of the co-domain (image of $\phi(\cdot)$, denoted by $\text{Im } \phi(\cdot)$) cannot be computed with non-negligible probability. More precisely, using the formalization of Damgård/Koprowski [27], we let π be the largest prime dividing the order of the co-domain ($\text{Im } \phi(\cdot)$), and let $\alpha(\pi)$ denote the maximal probability that π occurs when $\phi(\cdot)$ is chosen randomly from a predefined finite set of homomorphisms. Then $\phi(\cdot)$ is said to have a *hidden order co-domain (image)*, if $\alpha(\pi)$ is negligibly small.

2.5 Generic Σ -Protocols

We call the class of protocols for which our results hold true generic Σ -protocols. Informally, this class consists of almost all HVZK Σ -protocols of the following form. The prover is allowed to compute and send arbitrary elements obtained from generic homomorphism algorithms in both moves. The verifier may send multiple randomly chosen challenges in its first move, and use an arbitrary generic algorithm to decide whether to accept or to reject the proof.

Definition 5 (Generic (Group) Σ -Protocols). *Let $a_{ij}, b_{ij}, d_i, e_i, f_i, g_i$ be integer coefficients, let $\{(b_{11}, \dots, b_{1l}), \dots, (b_{n1}, \dots, b_{nl})\}$ be linearly independent over the integers, and let $C_1, \dots, C_p \subseteq \mathbb{Z}$ be arbitrary finite sets. Let further $\text{Verify}(\cdot, \dots, \cdot)$ be a generic homomorphism algorithm, and let the verifier always accept for an honest prover. We then call an HVZK two party protocol a generic (group) Σ -protocol for a homomorphism $\phi(\cdot) : \mathcal{G} \rightarrow \mathcal{H}$, if it has the form depicted in Fig. 1.*

It can easily be seen that this class covers the existing Σ^ϕ -protocol as well as the parallel execution of multiple instantiations thereof. Yet, a much broader set of protocols is covered by the class of generic Σ -protocols.

We make two minor remarks on this definition. First, the required linear independence can often be inferred from the HVZK property. Namely, if the

Fig. 1. Structure of a generic Σ -protocol for a homomorphism $\phi : \mathcal{G} \rightarrow \mathcal{H}$

vectors were not linearly independent the verifier could compute a multiple of w , and using Shamir's trick [5] could thus often compute the secret. Second, the definition of generic homomorphism algorithms also allows to draw random choices in the co-domain of the homomorphism. The above definition allows to draw random choices in the image by drawing $r \in_R \mathcal{G}$ and computing $\phi(r)$.

3 Efficiency Limitations in the Generic Group Model

In this section we describe lower bounds on the knowledge error for generic Σ -protocols with generic black-box algorithms as knowledge extractors. From these lower bounds we infer efficiency limitations for ZK-PoK using Σ -protocols.

In the statement of our results we refer to the notion of *expected PPT pseudo random functions*. Such functions are defined just as pseudo random functions (cf., e.g., [43]), except for one minor modification. Namely, we require that no *expected* PPT algorithm can distinguish such a function from a truly random one (usually one considers only *strict* PPT distinguishers). See §2.3 for a brief discussion why we resort to expected PPT time assumptions.

We are now ready to formulate our main result in the generic group model.

Theorem 6 (Lower Bounds in the Generic Group Model). *Let be given an arbitrary but fixed polynomial $\text{poly}(\cdot)$, a homomorphism $\phi(\cdot) : \mathcal{G} \rightarrow \mathcal{H}$ with hidden order image [2], and $x \in \mathcal{H}$, for which knowledge of a preimage under $\phi(\cdot)$ shall be proven. Consider a generic Σ -protocol as in Definition 5, and let q be the number of responses sent by the prover in its second step, i.e., $q := n - n^- + m - m^-$. Assuming that expected PPT pseudo random functions exist, the knowledge error of this protocol in the generic group model is lower bounded by*

$$\frac{1}{2^{\min(p,q)}} - \frac{1}{\text{poly}(\|x\|)}.$$

Let us briefly discuss the relevance and implications of this result.

- Our results indicates that a knowledge error of $1/2$ is an inherent limitation of the Σ^ϕ -protocol for homomorphisms with hidden order co-domain, which especially cover exponentiation homomorphisms in RSA groups.
- The best known technique to decrease the knowledge error is to repeat the Σ^ϕ -protocol, sequentially or in parallel. In either case, the number of elements sent by the prover and the verifier increases by the number of repetitions. Our results show that at least for the second and third move, i.e., the challenges sent by V and the responses sent by P , this growth cannot be avoided.

Put differently, Theorem 6 shows that the number p of challenges, and the number q of responses are the key parameters determining the size of the knowledge error. This implies that the strategy of repeating the Σ^ϕ -protocol parallelly is optimal concerning the second and third move of the protocol.

² Note that this is a stronger requirement than the requirement that the co-domain has hidden order. Yet, typically these two properties accompany each other.

- Finally, a protocol designer can deduce from Theorem 6 how an alternative for the Σ^ϕ -protocol must not look like. Namely, it must either not be a generic Σ -protocol, or the protocol must have a non-generic knowledge extractor, which uses particulars of the homomorphism.

3.1 Generalization to Other Classes of Homomorphisms

So far we have considered homomorphisms with hidden order co-domain. Yet, in practice this information is sometimes available and could potentially be used to decrease the lower bounds on the knowledge error.

More generally, we thus consider the class of *special* homomorphisms next. A homomorphism $\phi(\cdot) : \mathcal{G} \rightarrow \mathcal{H}$ is called special, if for every $x \in \mathcal{H}$ a pair $(u, v) \in \mathcal{G} \times \mathbb{Z} \setminus \{0\}$ satisfying $\phi(u) = x^v$ can be computed efficiently. The pair (u, v) is called *pseudo-preimage* of x under $\phi(\cdot)$. Besides homomorphisms with known order co-domain, also power homomorphisms are known to be special.

We model this property by adding one more query to the oracle $\mathcal{O}^{\phi(\cdot)}$, i.e., we allow a generic homomorphism algorithm to request a pseudo-preimage under $\phi(\cdot)$ for arbitrary elements from the co-domain of the homomorphism. We then obtain the following lemma:

Lemma 7 (Lower Bounds for Special Homomorphisms)

- (i) For power homomorphisms $(w_1, w_2) \mapsto \psi(w_1) \cdot w_2^e$ with hidden order co-domain, Theorem 6 can be generalized to a lower bound of $\frac{1}{d^{\min(p,q)}} - \frac{1}{\text{poly}(\|x\|)}$, where d is the smallest prime dividing e .
- (ii) For arbitrary homomorphisms with a co-domain of known order v , Theorem 6 generalizes literally with a lower bound of $\frac{1}{d^{\min(p,q)}} - \frac{1}{\text{poly}(\|x\|)}$, where d is the smallest prime dividing v , if v has a super-polynomially large prime factor.

Note that no such generalization is suitable for exponentiation homomorphisms with hidden order co-domain, as they are not known to be special. Analogue observations as for Theorem 6 on the implications of this lemma hold. Especially, in the generic group model the known knowledge error of the Σ^ϕ -protocol cannot be underrun for special homomorphisms.

Examples for homomorphisms of Case (i) are those used in the RSA, Paillier, and Damgård/Jurik encryption schemes [45, 46, 47, 48]. The ZK-PoK for these homomorphisms was introduced by Guillou/Quisquater [4]. Case (ii) covers the homomorphisms underlying the ElGamal encryption scheme [49] and the ZK-PoK for it was proposed by Schnorr [3].

3.2 Proof of Theorem 6

The remainder of this section is now dedicated to proving the theorem. We therefore recap the following lemma introduced by Damgård/Koprowski.

Lemma 8 (Lemma 3 of [27]). *Let $E := a_1X_1 + \dots + a_uX_u \in \mathbb{Z}[X_1, \dots, X_u]$ be a non-zero polynomial, and let $z \geq |a_i|$ for all i . Let further $\mathcal{G}$ be a group of hidden order, and $x_1, \dots, x_u \in_R \mathcal{G}$. For any positive A , we then have*

$$\Pr[a_1x_1 + \dots + a_ux_u = 0] \leq \frac{1}{A} + (\log_2 z + A)\alpha(\pi).$$

Proof (of Theorem 6 – Sketch). The proof is structured as follows. We describe a prover P^* for which we show that it satisfies the conditions of Theorem 3. We will see that the uniformity condition holds true by definition. For the hardness condition we simulate the behavior of P^* in the additive subgroup of a suitable polynomial ring. We then estimate the success probability of this simulated game and the error made when making this simulation.

We start with describing a malicious prover P^* . This cheating prover essentially behaves like the honest prover, except that it does not answer all challenges but only certain ones. Depending on whether $p \leq q$ or not, the set $\mathcal{C}'$ of answered challenges is defined as follows:

- $p \leq q$: For $i = 1, \dots, p$, let $\bar{c}_i \in \{0, 1\}$ such that at least half of the elements of $\mathcal{C}_i$ have the same parity as $\bar{c}_i$. Then $\mathcal{C}' := \{(c_1, \dots, c_p) \in \mathcal{C} \mid c_i \equiv \bar{c}_i \pmod{2}\}$.
 $q < p$: We define $\mathcal{C}'$ as a subset of $\mathcal{C}$, which has a cardinality of at least $\#\mathcal{C}/2^q$, and all $(c_1, \dots, c_p), (c'_1, \dots, c'_p) \in \mathcal{C}'$ satisfy the following q equations for all $j = m^- + 1, \dots, m$ and all $k = n^- + 1, \dots, n$:

$$\sum g_{ji}c_i \equiv \sum g_{ji}c'_i \pmod{2} \quad \text{and} \quad \sum e_{ki}c_i \equiv \sum e_{ki}c'_i \pmod{2}$$

We next describe P^* . We therefore make the random input $\zeta = (\zeta_1, \dots, \zeta_l)$ to the prover explicit, and let $\rho(\cdot)$ be a pseudo random function.

- (i) It sets $r'_i := \rho(\zeta_i)$ for $i = 1, \dots, l$, and using these random elements, it behaves just as an honest prover.
- (ii) If $c_i \in \mathcal{C}'$, P^* behaves like an honest prover, using $(r'_{l^-+1}, \dots, r'_l)$ as random elements. Otherwise it halts.

The **uniformity property** of Theorem 3 is obviously satisfied, as the prover answers a fraction of at least $1/2^{\min(p,q)}$ of all challenges, and makes the verifier accept (because the verifier would accept for an honest prover).

Let us now turn towards the **hardness property**. We say that a generic black-box algorithm succeeds, if after v steps it outputs the handle corresponding to a preimage of x under $\phi(\cdot)$. Now, instead of letting the knowledge extractor interact with P^* and the oracle $\mathcal{O}^{\phi(\cdot)}$, we play the following game. We substitute $\mathcal{G}$ and $\mathcal{H}$ by the following subgroups of the polynomial rings over the indeterminates $W, O_{ij}, R_{ij}, T_{ij}$:

$$\begin{aligned} \mathcal{G}' &:= \langle W, O_{11}, \dots, O_{1l}, \dots, O_{v1}, \dots, O_{vl}, R_{11}, \dots, R_{1m}, \dots, R_{v1}, \dots, R_{vm} \rangle \\ \mathcal{H}' &:= \langle \mathcal{G}', T_{11}, \dots, T_{1n}, \dots, T_{v1}, \dots, T_{vn} \rangle. \end{aligned}$$

Accordingly, the oracle $\mathcal{O}'^{\phi(\cdot)}$ now performs its computations within $\mathcal{G}'$ and $\mathcal{H}'$.

The prover P^* is adopted as described next. It maintains a list L , which is initially empty, and sets $u := 0$. On random input ζ , it performs the following steps:

- (i) For each ζ_i , it checks whether there is a pair $(\zeta_{ji}, \bar{R}_{ji})$ with $\zeta_i = \zeta_{ji}$ in L . If so, it sets $\hat{R}_i := \bar{R}_{ji}$. Otherwise, it increases u by 1 (but at most once in each run), sets $\hat{R}_i := R_{ui}$, and adds $(\zeta_i, \hat{R}_i)$ to L . Then it sends

$$\left(\left(\sum a_{ji} \cdot \hat{R}_i + f_j \cdot W \right)_{j=1}^{m^-}, \left(\sum b_{ki} \cdot \hat{R}_i + d_k \cdot W \right)_{k=1}^{n^-} \right)$$

to V . Former are marked as elements of $\mathcal{G}'$, latter as elements of $\mathcal{H}'$.

- (ii) If $c_i \in \mathcal{C}'$, P^* analogously computes its response according to the protocol. Otherwise, if $c \notin \mathcal{C}'$, P^* halts.

By $\mathbf{r}$ we denote an element from the set of from which the oracle and by the generator of the input to the protocol draw their random choices, i.e.,

$$\mathbf{r} \in \left\{ (\phi(\cdot), x, w, \rho, o, t) \mid \phi(\cdot) : \mathcal{G} \rightarrow \mathcal{H} \text{ has hidden order co-domain,} \right. \\ \left. x = \phi(w), \rho(\cdot) \text{ pseudo random, } o \in \mathcal{G}^{v \times l}, t \in \mathcal{H}^{v \times m} \right\}$$

We then define the following two mappings. By $\iota_{\mathcal{G}'}^{\mathbf{r}}(\cdot)$ we denote the evaluation homomorphism from $\mathcal{G}'$ into $\mathcal{G}$. That is, by $\iota_{\mathcal{G}'}^{\mathbf{r}}(E)$ we denote the element in $\mathcal{G}$ which results when all indeterminates in E are substituted in the following way:

$$W \mapsto w \quad O_{ij} \mapsto o_{ij} \quad R_{ij} \mapsto r'_{ij}.$$

In absolute analogy we let $\iota_{\mathcal{H}'}^{\mathbf{r}}(\cdot)$ be the evaluation homomorphism from $\mathcal{H}'$ into $\mathcal{H}$. That is, the substitution is given by:

$$W \mapsto \phi(w) \quad O_{ij} \mapsto \phi(o_{ij}) \quad R_{ij} \mapsto \phi(r'_{ij}) \quad T_{ij} \mapsto t_{ij}.$$

We observe that for all $E \in \mathcal{G}'$ we have $\phi(\iota_{\mathcal{G}'}^{\mathbf{r}}(E)) = \iota_{\mathcal{H}'}^{\mathbf{r}}(E)$.

During its computation the generic black-box algorithm maintains a list of elements $E_i \in \mathcal{G}'$ respectively $F_i \in \mathcal{H}'$. We say that the algorithm wins this modified game, if one of the following two cases occurs. In case (a), the algorithm finds a preimage of x under $\phi(\cdot)$, while in case (b) there is a pair $i \neq j$ satisfying the following. For a randomly chosen $\mathbf{r}$, we either have $E_i \neq E_j$ and $\iota_{\mathcal{G}'}^{\mathbf{r}}(E_i - E_j) = 0$, or $F_i \neq F_j$ and $\iota_{\mathcal{H}'}^{\mathbf{r}}(F_i - F_j) = 0$.

Observing that the behavior of this game and the actual interaction between the algorithm and the real oracle are indistinguishable as long as the above game is not won, we get that the success probability of the generic black-box algorithm is upper bounded by the probability that the algorithm wins the game [50].

Case (a). Finding a preimage means to compute E_i such that $\phi(\iota_{\mathcal{G}'}^{\mathbf{r}}(E_i)) = x$. Using the observation that we always have $\iota_{\mathcal{H}'}^{\mathbf{r}}(W) = x$ this means to find an E_i such that $\iota_{\mathcal{H}'}^{\mathbf{r}}(E_i - X) = 0$. By introspection of how the E_i 's are computed, and by

using the linear independency of the vectors $\{(b_{11}, \dots, b_{1l}), \dots, (b_{n1}, \dots, b_{nl})\}$, one can show that $W \neq E_i$ for all i .

Let $K := K(\mathcal{C}, a_{ji}, b_{ki}, g_{ji}, e_{ji}, f_j, d_k)$ be an integer such that K is larger than the absolute values of all coefficients occurring in the definition of the examined generic Σ -protocol. Using that $E_i \neq W$ and Lemma 8, and noting that after v oracle queries for E_i 's and F_j 's each, all coefficients are smaller than $2^v \cdot K$, we get

$$Pr[(a)] \leq \frac{1}{A} + (v + \log_2 K + A)\alpha(\pi) \quad \text{for all } A \in \mathbb{Z}.$$

Case (b). Using K as before, and observing that there are at most v different E_i 's and F_j 's each, we get by a similar argument that the probability for (b) is bounded by

$$Pr[(b)] \leq v^2 \left(\frac{1}{A} + (v + \log_2 K + A)\alpha(\pi) \right) \quad \text{for all } A \in \mathbb{Z}.$$

We here assumed that $\phi(\cdot)$ is surjective, and that $\rho(\cdot)$ is a truly random function. The former can easily be seen to be just a technical issue to ease presentation, and the latter yields only a negligible error as $\rho(\cdot)$ is pseudo random by definition.

Demonstration of Hardness Condition. The overall probability that the algorithm wins the game described above is hence limited by

$$Pr[(a)] + Pr[(b)] \leq (v^2 + 1) \left(\frac{1}{A} + (v + \log_2 K + A)\alpha(\pi) \right) \quad \text{for all } A \in \mathbb{Z}.$$

for a fixed choice of $\mathbf{r}$. We now set the so far arbitrary value of A to $A := \sqrt{1/\alpha(\pi)}$ such that both $1/A$ and $A \cdot \alpha(\pi)$ are negligible, and observe that K and $\alpha(\pi)$ are independent from $\mathbf{r}$. Using now that for the hardness condition to be satisfied we only need to consider generic black-box algorithms the expected number v of steps of which is polynomially bounded, and computing the expectation value over all choices of $\mathbf{r}$, we get that the success probability of the generic black-box algorithm is negligible. $\square$

4 Lower Bounds for the Σ^ϕ -Protocol in the Plain Model

As pointed out by Dent and Fischlin [39,40], restrictions proven in the generic model do not necessarily hold true in the plain model as well. In this section we thus confirm our results obtained in the generic model by showing the existence of lower bounds in the plain model. That is, we provide evidence that for exponentiation homomorphisms with hidden order co-domain, and for power homomorphisms of the form $\phi(\cdot) : \mathcal{H} \rightarrow \mathcal{H} : w \mapsto w^e$, no smaller knowledge error than in the generic model can be reached in the plain model. The results only hold for the Σ^ϕ -protocol, and not for the entire class of generic Σ -protocols.

The following results are based on a generalization of the Root Assumption [47], which we call the *Expected Root Assumption*. We say that the Expected

Root Assumption holds for a group $\mathcal{H}$ if there exists no expected PPT algorithm that on input a random element $h \in_R \mathcal{H}$ and $e \geq 2$ outputs an e^{th} root of h with non-negligible probability. In contrast to the standard Root Assumption, we here also require that no *expected* PPT algorithm has a noticeable success probability. This requirement naturally arises from the fact that the definition of PoK only restricts the *expected* running time of the knowledge extractor, cf. §2.3.

4.1 Lower Bounds for Power Homomorphisms

We first consider the Σ^ϕ -protocol for power homomorphisms of the form $\phi_P(\cdot) : \mathcal{H} \rightarrow \mathcal{H} : w \mapsto w^e$. This is a generalization of the protocol proposed Guillou/Quisquater [4]. We generalize the result from Shoup [26] from exponents of the form $e = 2^t$ to arbitrary values of e .

In the following we use the following notation. For a set $\mathcal{S}$ and $r \in \mathbb{Z}$, we define $\text{Div}(\mathcal{S}, r)$ to be all multiples of r within $\mathcal{S}$, i.e., $\text{Div}(\mathcal{S}, r) := \{s : s \in \mathcal{S}, r|s\}$.

Theorem 9 (Bounds for Power Homomorphisms). *Let $\text{poly}(\cdot)$ be an arbitrary but fixed polynomial. Then for every power homomorphism $\phi_P(\cdot) : \mathcal{H} \rightarrow \mathcal{H} : w \mapsto w^e$ with $e \geq 2$, the knowledge error of the Σ^ϕ -protocol for $\phi_P(\cdot)$ is lower bounded by*

$$\max_{2 \leq r \leq e, r|e} \frac{\#\text{Div}(\mathcal{C}, r)}{\#\mathcal{C}} - \frac{1}{\text{poly}(\|x\|)},$$

if the Expected Root Assumption is satisfied for $\mathcal{H}$ and $\gcd(e, \text{ord } \mathcal{H}) = 1$.

Note here that, if $\mathcal{H}$ is an RSA group, i.e., $\mathcal{H} = \mathbb{Z}_n^*$ for a composite modulus n of unknown factorization, the condition $\gcd(e, \text{ord } \mathcal{H}) = 1$ is always satisfied.

We stress that, if the challenge set $\mathcal{C}$ is an integer interval, the theorem implies a lower bound which is equal to the smallest knowledge error that is currently known to be achievable:

Corollary 10. *Let the conditions of Theorem 9 be satisfied, and let the challenge set be an integer interval (i.e., $\mathcal{C} = \{a, \dots, b\}$ for some $a, b \in \mathbb{Z}$). Let d denote the smallest divisor of e . Then knowledge error of the Σ^ϕ -protocol is bounded from below by*

$$\frac{1}{d} - \frac{1}{\text{poly}(\|x\|)}.$$

Theorem 9 becomes meaningless if all elements of $\mathcal{C}$ are co-prime (e.g., if all elements of $\mathcal{C}$ are primes), as it then implies a lower bound of 0. Though, the result is still relevant when seen in connection with Theorem 6. Namely, while the latter states that any hypothetical knowledge extractor has to use encoding specific properties of the homomorphism $\phi_P(\cdot)$, the former further restricts the situations where the generic result could potentially be violated in the plain model. In summary, the existence of an extractor underrunning the limitation of $1/d$ seems unlikely.

4.2 Lower Bounds for Exponentiation Homomorphisms

For exponentiation homomorphisms $\phi_E(\cdot) : \mathbb{Z} \rightarrow \mathcal{H} : w \mapsto h^w$ with hidden order co-domain $\mathcal{H}$, the Σ^ϕ -protocol is only known to be a PoK with knowledge error $1/2$. In this section we show that (if existing at all) any knowledge error achieving a smaller knowledge error in this case would require fundamentally new insights to the Σ^ϕ -protocol.

Although being used for numerous different homomorphisms, essentially only one knowledge extractor is known for the Σ^ϕ -protocol. This standard knowledge extractor works as described next. In a first phase, it is given rewindable black-box access to the prover, and extracts a *pseudo preimage* (u, v) , i.e., a pair satisfying $v \neq 0$ and $x^v = \phi_E(u)$, cf. §3.1. Then, in a second phase in which the extractor does not have access to the prover any more, it computes a preimage of x from this pseudo preimage. We call knowledge extractors working this way *pseudo preimage based*. We show that no such knowledge extractor can underrun a knowledge error of $1/2$ for the Σ^ϕ -protocol and exponentiation homomorphisms with hidden order co-domain.

Let us introduce some notation: for a set $\mathcal{S}$ of integers, we write $\text{Diff}(\mathcal{S})$ for the set of all possible absolute values of differences between different elements of $\mathcal{S}$, i.e., $\text{Diff}(\mathcal{S}) := \{|s_1 - s_2| : s_1 \neq s_2 \in \mathcal{S}\}$. We further say that an integer d and a set $\mathcal{S}$ are co-prime, if $\gcd(d, s) = 1$ for all $s \in \mathcal{S}$.

Theorem 11 (Bounds for Exponentiation Homomorphisms). *Let $\text{poly}(\cdot)$ be an arbitrary but fixed polynomial. Then for every exponentiation homomorphisms $\phi_E(\cdot) : \mathbb{Z} \rightarrow \mathcal{H}' : w \mapsto h^w$, with $h \in \mathcal{H}'$, the knowledge error of the Σ^ϕ -protocol for $\phi_E(\cdot)$ is lower bounded by*

$$\frac{1}{2} - \frac{1}{\text{poly}(\|x\|)},$$

against pseudo preimage based knowledge extractors, if the following conditions are satisfied. The co-domain $\mathcal{H}'$ is a large subgroup of $\mathcal{H}$ (i.e., $\#\mathcal{H}'/\#\mathcal{H}$ is not negligible), the Expected Root Assumption is satisfied for $\mathcal{H}$, and $\text{ord } \mathcal{H}'$ and $\text{Diff}(\mathcal{C})$ are co-prime.

We remark that this result can straightforwardly be generalized to homomorphisms of the form $\phi_M(\cdot) : \mathcal{G}^r \rightarrow \mathcal{H} : (w_1, \dots, w_r) \mapsto h_1^{w_1} \dots h_r^{w_r}$.

In practice the conditions of this theorem are most often satisfied. For instance consider the case where $\mathcal{H} = \mathbb{Z}_n^*$ for a safe RSA modulus n , i.e., $n = (2p + 1) \cdot (2q + 1)$, where $p, q, (2p + 1)$, and $(2q + 1)$ are primes. Then $\mathcal{H}'$ is usually given by the set of quadratic residues modulo n , and we have $\#\mathcal{H}'/\#\mathcal{H} = 1/4$. Further, $\text{ord } \mathcal{H}' = p \cdot q$, and hence any challenge set $\mathcal{C}$ only containing elements smaller than p, q will satisfy the condition of $\text{Diff}(\mathcal{C})$ and $\text{ord } \mathcal{H}'$ being co-prime.

Although this result only considers pseudo preimage based knowledge extractors, it is still relevant for the following reason. Together with the results in the generic group model in §3, Theorem 11 implies that a knowledge extractor for exponentiation homomorphisms with hidden order co-domain must neither be generic nor pseudo preimage based. Thus, if possible at all, substantially new

insights were required to underrun the restriction of $1/2$ in this case. According to current knowledge, we doubt the existence of such an extractor. We thus believe that for reaching a small knowledge error in the case of exponentiation homomorphisms with hidden order co-domain, either running the Σ^ϕ -protocol repeatedly or employing an FO-based protocol cannot be avoided.

5 Conclusion

We have introduced the class of generic Σ -protocols, and have shown that in the generic group model a knowledge error of $1/2^n$ (where n is the minimum of the number of challenges and responses sent in the protocol) is inherent to any of these protocols for homomorphisms with hidden order co-domain. We further generalized this result to special homomorphisms as well, covering essentially all homomorphisms being used in cryptography. Especially, those underlying various crypto systems fall into this class [45,46,47,48,49]. We then confirmed our results for the Σ^ϕ -protocol and certain homomorphisms in the plain model as well.

Besides pointing out these limitations, our results also give insights in how these restrictions could be overcome. Namely, any Σ -protocol overcoming these bounds must either be substantially different from the Σ^ϕ -protocol (i.e., it must not be a generic Σ -protocol), or it must have a non-generic knowledge extractor.

The former seems to be hard to achieve without using auxiliary constructions resulting from a common reference string as done in [5,20,21], because the class of generic Σ -protocols does not leave much design options for other Σ -protocols to look like. Yet, the latter also is unlikely, because of our results in the plain model. Thus, although being riddled with various limitations from a theoretical point of view, FO-based protocols [5,19,20,21,22,23] using common reference strings seem to be inevitable for many real systems.

References

1. Bangerter, E., Camenisch, J., Krenn, S.: Efficiency limitations for Σ -protocols for group homomorphisms. Cryptology ePrint Archive, Report 2009/595 (2009)
2. Goldreich, O., Micali, S., Wigderson, A.: Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. Journal of the ACM 38, 691–729 (1991); Preliminary version in 27th FOCS (1986)
3. Schnorr, C.: Efficient signature generation by smart cards. Journal of Cryptology 4, 161–174 (1991)
4. Guillou, L., Quisquater, J.J.: A “paradoxical” identity-based signature scheme resulting from zero-knowledge. In: Goldwasser, S. (ed.) CRYPTO 1988. LNCS, vol. 403, pp. 216–231. Springer, Heidelberg (1990)
5. Bangerter, E.: Efficient Zero-Knowledge Proofs of Knowledge for Homomorphisms. PhD thesis, Ruhr-University Bochum (2005)
6. Cramer, R.: Modular Design of Secure yet Practical Cryptographic Protocols. PhD thesis, CWI and University of Amsterdam (1997)

7. Maurer, U.: Unifying zero-knowledge proofs of knowledge. In: Preneel, B. (ed.) AFRICACRYPT 2009. LNCS, vol. 5580, pp. 272–286. Springer, Heidelberg (2009)
8. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987)
9. Boudot, F.: Efficient proofs that a committed number lies in an interval. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 431–444. Springer, Heidelberg (2000)
10. Camenisch, J.: Better privacy for trusted computing platforms (extended abstract). In: Samarati, P., Ryan, P.Y.A., Gollmann, D., Molva, R. (eds.) ESORICS 2004. LNCS, vol. 3193, pp. 73–88. Springer, Heidelberg (2004)
11. Camenisch, J., Michels, M.: Proving in zero-knowledge that a number is the product of two safe primes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 107–122. Springer, Heidelberg (1999)
12. Camenisch, J., Shoup, V.: Practical verifiable encryption and decryption of discrete logarithms. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 126–144. Springer, Heidelberg (2003)
13. Lipmaa, H.: On diophantine complexity and statistical zero-knowledge arguments. In: Lai, C.-S. (ed.) ASIACRYPT 2003. LNCS, vol. 2894, pp. 398–415. Springer, Heidelberg (2003)
14. Song, D.X.: Practical forward secure group signature schemes. In: ACM CCS 2001, pp. 225–234. ACM, New York (2001)
15. Tang, C., Liu, Z., Wang, M.: A verifiable secret sharing scheme with statistical zero-knowledge. Cryptology ePrint Archive, Report 2003/222 (2003)
16. Tsang, P.P., Wei, V.K., Chan, T.K., Au, M.H., Liu, J.K., Wong, D.S.: Separable linkable threshold ring signatures. In: Canteaut, A., Viswanathan, K. (eds.) INDOCRYPT 2004. LNCS, vol. 3348, pp. 384–398. Springer, Heidelberg (2004)
17. Brickell, E., Camenisch, J., Chen, L.: Direct anonymous attestation. In: ACM CCS 2004, pp. 132–145. ACM, New York (2004)
18. Camenisch, J., Herreweghen, E.V.: Design and implementation of the idemix anonymous credential system. In: ACM CCS 2002, pp. 21–30. ACM, New York (2002)
19. Bangerter, E., Camenisch, J., Maurer, U.: Efficient proofs of knowledge of discrete logarithms and representations in groups with hidden order. In: Vaudenay, S. (ed.) PKC 2005. LNCS, vol. 3386, pp. 154–171. Springer, Heidelberg (2005)
20. Bangerter, E., Krenn, S., Sadeghi, A.R., Schneider, T., Tsay, J.K.: Design and implementation of efficient zero-knowledge proofs of knowledge. In: SPEED-CC 2009 (2009)
21. Damgård, I., Fujisaki, E.: A statistically-hiding integer commitment scheme based on groups with hidden order. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 125–142. Springer, Heidelberg (2002)
22. Fujisaki, E., Okamoto, T.: Statistical zero knowledge protocols to prove modular polynomial relations. In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 16–30. Springer, Heidelberg (1997)
23. Camenisch, J., Kiayias, A., Yung, M.: On the portability of Generalized Schnorr Proofs. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 425–442. Springer, Heidelberg (2009)
24. Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited. *Journal of the ACM* 51, 557–594 (2004); Preliminary version in STOC (1998)

25. Pass, R.: On deniability in the common reference string and random oracle model. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 316–337. Springer, Heidelberg (2003)
26. Shoup, V.: On the security of a practical identification scheme. In: Maurer, U.M. (ed.) EUROCRYPT 1996. LNCS, vol. 1070, pp. 344–353. Springer, Heidelberg (1996)
27. Damgård, I., Koprowski, M.: Generic lower bounds for root extraction and signature schemes in general groups. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 256–271. Springer, Heidelberg (2002)
28. Nechaev, V.I.: Complexity of a determinate algorithm for the discrete logarithm. *Mathematical Notes* 55, 165–172 (1994); Translated from *Matematicheskie Zametki* 55(2), 91–101 (1994)
29. Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 256–266. Springer, Heidelberg (1997)
30. 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)
31. Aggarwal, D., Maurer, U.: Breaking RSA generically is equivalent to factoring. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 36–53. Springer, Heidelberg (2009)
32. Boneh, D., Boyen, X., Goh, E.J.: Hierarchical identity based encryption with constant size ciphertext. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg (2005)
33. Brown, D.: Generic groups, collision resistance, and ECDSA. *Cryptology ePrint Archive*, Report 2002/026 (2002)
34. Dent, A.W.: The hardness of the DHK problem in the generic group model. *Cryptology ePrint Archive*, Report 2006/156 (2006)
35. Maurer, U.: Index search, discrete logarithms, and Diffie-Hellman. In: *Number-theoretic cryptography workshop*. Mathematical Sciences Research Institute, Berkeley (2000)
36. Maurer, U., Wolf, S.: Lower bounds on generic algorithms in groups. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 72–84. Springer, Heidelberg (1998)
37. Schnorr, C., Jakobsson, M.: Security of signed elgamal encryption. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 73–89. Springer, Heidelberg (2000)
38. Smart, N.P.: The exact security of ECIES in the generic group model. In: Honary, B. (ed.) *Cryptography and Coding* 2001. LNCS, vol. 2260, pp. 73–84. Springer, Heidelberg (2001)
39. Dent, A.W.: Adapting the weaknesses of the random oracle model to the generic group model. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 100–109. Springer, Heidelberg (2002)
40. Fischlin, M.: A note on security proofs in the generic model. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 458–469. Springer, Heidelberg (2000)
41. Cramer, R., Damgård, I.: On the amortized complexity of zero-knowledge protocols. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 177–191. Springer, Heidelberg (2009)
42. Bellare, M., Goldreich, O.: On defining proofs of knowledge. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 390–420. Springer, Heidelberg (1993)
43. Goldreich, O.: *Foundations of Cryptography – Basic Tools*. Cambridge University Press, Cambridge (2001)

44. Babai, L., Szemerédi, E.: On the complexity of matrix group problems I. In: IEEE FOCS 1984, pp. 229–240 (1984)
45. Damgård, I., Jurik, M.: A generalisation, a simplification and some applications of Paillier's probabilistic public-key system. In: Kim, K.-c. (ed.) PKC 2001. LNCS, vol. 1992, pp. 119–136. Springer, Heidelberg (2001)
46. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg (1999)
47. Rivest, R., Shamir, A., Adleman, L.: A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM* 21, 120–126 (1978)
48. Takagi, T.: Fast RSA-type cryptosystem modulo $p^k q$. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 318–326. Springer, Heidelberg (1998)
49. Gamal, T.E.: A public key cryptosystem and a signature scheme based on discrete logarithms. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 10–18. Springer, Heidelberg (1985)
50. Shoup, V.: OAEP reconsidered. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 239–259. Springer, Heidelberg (2001)