



The Algebraic Group Model and its Applications

Georg Fuchsbauer¹, Eike Kiltz², and Julian Loss²(✉)

¹ Inria, ENS, CNRS, PSL, Paris, France
`georg.fuchsbauer@ens.fr`

² Ruhr University Bochum, Bochum, Germany
`{eike.kiltz,julian.loss}@rub.de`

Abstract. One of the most important and successful tools for assessing hardness assumptions in cryptography is the Generic Group Model (GGM). Over the past two decades, numerous assumptions and protocols have been analyzed within this model. While a proof in the GGM can certainly provide some measure of confidence in an assumption, its scope is rather limited since it does not capture group-specific algorithms that make use of the representation of the group.

To overcome this limitation, we propose the Algebraic Group Model (AGM), a model that lies in between the Standard Model and the GGM. It is the first restricted model of computation covering group-specific algorithms yet allowing to derive simple and meaningful security statements. To prove its usefulness, we show that several important assumptions, among them the Computational Diffie-Hellman, the Strong Diffie-Hellman, and the interactive LRSW assumptions, are equivalent to the Discrete Logarithm (DLog) assumption in the AGM. On the more practical side, we prove tight security reductions for two important schemes in the AGM to DLog or a variant thereof: the BLS signature scheme and Groth's zero-knowledge SNARK (EUROCRYPT 2016), which is the most efficient SNARK for which only a proof in the GGM was known. Our proofs are quite simple and therefore less prone to subtle errors than those in the GGM.

Moreover, in combination with known lower bounds on the Discrete Logarithm assumption in the GGM, our results can be used to derive lower bounds for all the above-mentioned results in the GGM.

Keywords: Algebraic algorithms · Generic group model
Security reductions · Cryptographic assumptions

1 Introduction

Starting with Nechaev [Nec94] and Shoup [Sho97], much work has been devoted to studying the computational complexity of problems with respect to generic group algorithms over cyclic groups [BL96, MW98, Mau05]. At the highest level, generic group algorithms are algorithms that do not exploit any special structure

of the representation of the group elements and can thus be applied in any cyclic group. More concretely, a generic algorithm may use only the abstract group operation and test whether two group elements are equal. This property makes it possible to prove information-theoretic lower bounds on the running time for generic algorithms. Such lower bounds are of great interest since for many important groups, in particular for elliptic curves, no helpful exploitation of the representation is currently known.

The class of generic algorithms encompasses many important algorithms such as the baby-step giant-step algorithm and its generalization for composite-order groups (also known as Pohlig-Hellman algorithm [HP78]) as well as Pollard’s rho algorithm [Pol78]. However, part of the common criticism against the generic group model is that many algorithms of practical interest are in fact not generic. Perhaps most notably, index-calculus and some factoring attacks fall outside the family of generic algorithms, as they are applicable only over groups in which the elements are represented as *integers*.

Another example is the “trivial” discrete logarithm algorithm over the additive group \mathbb{Z}_p , which is the identity function.

With this motivation in mind, a number of previous works considered extensions of the generic group model [Riv04, LR06, AM09, JR10]. Jager and Rupp [JR10] considered assumptions over groups equipped with a bilinear map $e: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_3$, where \mathbb{G}_1 and \mathbb{G}_2 are modeled as generic groups, and \mathbb{G}_3 is modeled in the Standard Model. (This is motivated by the fact that in all practical bilinear groups, \mathbb{G}_1 and \mathbb{G}_2 are elliptic curves whereas \mathbb{G}_3 is a sub-group of a finite field). However, none of these models so far capture algorithms that can freely exploit the representation of the group. In this work, we propose a restricted model of computation which does exactly this.

1.1 Algebraic Algorithms

Let \mathbb{G} be a cyclic group of prime order p . Informally, we call an algorithm \mathbf{A}_{alg} *algebraic* if it fulfills the following requirement: whenever \mathbf{A}_{alg} outputs a group element $\mathbf{Z} \in \mathbb{G}$, it also outputs a “representation” $\vec{z} = (z_1, \dots, z_t) \in \mathbb{Z}_p^t$ such that $\mathbf{Z} = \prod_i \mathbf{L}_i^{z_i}$, where $\vec{\mathbf{L}} = (\mathbf{L}_1, \dots, \mathbf{L}_t)$ is the list of all group elements that were given to \mathbf{A}_{alg} during its execution so far.

Such algebraic algorithms were first considered by Boneh and Venkatesan [BV98] in the context of straight-line programs computing polynomials over the ring of integers \mathbb{Z}_n , where $n = pq$. Later, Paillier and Vergnaud [PV05] gave a more formal and general definition of algebraic algorithms using the notion of an *extractor algorithm* which efficiently computes the representation \vec{z} .

In our formalization of algebraic algorithms, we distinguish group elements from all other parameters at a *syntactical level*, that is, other parameters must not depend on any group elements. This is to rule out pathological exploits of the model, see below.

While this class of algebraic algorithms certainly captures a much broader class of algorithms than the class of generic algorithms (e.g., index-calculus algorithms), it was first noted in [PV05] that the class of algebraic algorithms actually *includes* the class of generic algorithms.

Algebraic algorithms have mostly been studied so far in the context of proving *impossibility results* [BV98, Cor02, PV05, BMV08, GBL08, AGO11, KMP16], i.e., to disprove the existence of an algebraic security reduction between two cryptographic primitives (with certain good parameters). Only quite recently, a small number of works have considered the idea of proving statements with respect to *algebraic adversaries* [ABM15, BFW16].

1.2 Algebraic Group Model

We propose the *algebraic group model* (AGM) — a computational model in which all adversaries are modeled as algebraic. In contrast to the GGM, the AGM does not allow for proving information-theoretic lower bounds on the complexity of an algebraic adversary. Similar to the Standard Model, in the AGM one proves security implications via reductions. Specifically, $H \Rightarrow_{\text{alg}} G$ for two primitives H and G means that every algebraic adversary \mathbf{A}_{alg} against G can be transformed into an algebraic adversary \mathbf{B}_{alg} against H with (polynomially) related running times and success probabilities. It follows that if H is secure against algebraic adversaries, so is G . While algebraic adversaries have been considered before (see above), to the best of our knowledge, our work is the first to provide a clean and formal framework for security proofs with respect to algebraic adversaries. We elaborate further on our model below.

CONCRETE SECURITY IMPLICATIONS IN THE AGM. Indeed, one can exploit the algebraic nature of an adversary in the AGM to obtain stronger security implications than in the Standard Model. The first trivial observation is that the classical *knowledge of exponent assumption*¹ [Dam92] holds by definition in the AGM.

We are able to show that several important computational assumptions are in fact equivalent to the Discrete Logarithm assumption over prime-order groups in the AGM, including the following:

- Diffie-Hellman assumption [DH76]
- (Interactive) strong Diffie-Hellman assumption [ABR01]
- (Interactive) LRSW assumption [LRSW99, CL04].

The significance of the Strong Diffie-Hellman Assumption comes from its equivalence to the IND-CCA security of Hashed ElGamal encryption (also known as Diffie-Hellman Integrated Encryption Standard) in the random oracle model [ABR01]. The LRSW assumption (named its authors [LRSW99]) is of importance since it is equivalent to the (UF-CMA) security of Camenisch-Lysyanskaya

¹ The knowledge of exponent assumption states that for every algorithm \mathbf{A} that, given g and $\mathbf{X} = g^x$, outputs (\mathbf{A}, \mathbf{B}) with $\mathbf{B} = \mathbf{A}^x$, there exists an extractor algorithms that, given the same input, outputs a satisfying $(\mathbf{A}, \mathbf{B}) = (g^a, \mathbf{X}^a)$.

(CL) signatures [CL04]. CL signatures are a central building block for anonymous credentials [CL04, BCL04, BCS05], group signatures [CL04, ACHdM05], e-cash [CHL05], unclonable functions [CHK+06], batch verification [CHP07], and RFID encryption [ACdM05]. Via our results, the security of all these schemes is implied by the discrete logarithm assumption in the AGM.

Our result can be interpreted as follows. Every algorithm attacking one of the above-mentioned problems and schemes must solve the standard discrete logarithm problem directly, unless the algorithm relies on inherently non-algebraic operations. In particular, powerful techniques such as the index-calculus algorithms do not help in solving these problems any better than they do for solving the discrete logarithm problem directly.

Moreover, we show the *tight* equivalence of the security of the following schemes to the underlying hardness assumptions in the AGM:

- IND-CCA1 (aka lunchtime) security of the standard ElGamal Encryption to a parametrized variant of Decisional Diffie-Hellman assumption where in addition to g^x, g^y the adversary receives g^{x^2}, \dots, g^{x^q} , where q is the maximal number of decryption queries.
- The UF-CMA security of the BLS signature scheme [BLS04] to the discrete logarithm problem in the random oracle model. Previous reductions non-tightly reduced from the CDH problem, with a tightness loss linear in the number of signing queries. This loss is known to be inherent [Cor02, KK12], even in the random oracle model.
- The security of the so far most efficient zero-knowledge SNARK scheme by Groth [Gro16] to a parametrized variant of the discrete logarithm problem, where in addition to g^x the adversary receives $g^{x^2}, \dots, g^{x^{2n-1}}$, where n is the degree of the quadratic arithmetic programs. The only previous proof of the security of this scheme is in the generic group model.

RELATION TO THE GENERIC GROUP MODEL. The AGM is stronger (in the sense that it puts more restrictions on the attackers) than the Standard Model, but weaker than the GGM. In spite of this, all of our reductions are purely generic algorithms. As mentioned above, any generic algorithm can be modeled within the AGM. In particular, combining arbitrary generic operations with algebraic ones will yield an algebraic algorithm. This suggests the following idea. Let H and G be two computational problems and let A_{alg} be an algebraic algorithm that solves problem G . If we can convert A_{alg} by means of a generic reduction algorithm R_{gen} into an algorithm B_{alg} for problem H , then clearly, B_{alg} is also an algebraic algorithm. However, we obtain an even stronger statement for free: Namely, if A_{gen} is a generic algorithm solving G , then B_{gen} is a generic algorithm solving H . This means that results in the AGM directly carry over to the GGM.

For this reason, we believe that our model offers an alternative, perhaps simpler method of proving the hardness of computational problems within the GGM. This applies in particular to interactive assumptions, which can be rather difficult to analyze in the GGM. For example, we prove that the discrete logarithm assumption implies the LRSW assumption in the AGM. As the discrete logarithm assumption holds in the GGM, we instantly obtain that the

LRSW assumption holds in the GGM. The first (rigorous) proof of the LRSW assumption within the GGM was presented in the work of [BFF+14] (the original work [LRSW99] provided only a proof sketch), but was derived from a more general theorem and proven using an automated proof verification tool. We hope that our proof can offer some additional insight over the proof of [BFF+14]. Another example is our tight equivalence of the IND-CCA1 security of ElGamal and our parametrized variant of the Decisional Diffie-Hellman (DDH) assumption in the algebraic group model. Together with the known generic $\sqrt{p/q}$ attack on ElGamal [BG04] for certain primes p (see also [Che06]), our result proves the tight generic bound $\tilde{\Theta}(\sqrt{p/q})$ on the complexity of breaking IND-CCA1 security of ElGamal in the GGM.

We also remark that proofs in the AGM have an inherently different interpretation than proofs in the GGM. To analyze the hardness of an assumption in the GGM, one must explicitly augment the model by any functionality that is offered by the structure of the group. As a simple example, let us consider a group \mathbb{G} which is equipped with a symmetric bilinear map $e: \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. The bilinear map can be modeled in the GGM via an oracle. However, it is not clear whether e can be used to gather even further information about the elements of \mathbb{G} . Though it is widely believed that this is not the case, a proof in the GGM provides no answer to this question, because the GGM itself is based on the conjecture that e does not offer any functionality beyond a bilinear map. In contrast, the AGM captures any such exploit without the need of having to model it explicitly and considers the relation between two problems instead of their *individual hardness*. This means that if one can reduce H to G in the AGM and H is conjectured to remain hard with respect to algebraic algorithms, even when given e , then also G remains hard. No similar statement can be inferred in the GGM. Thus, the AGM allows for a more fine grained assessment of the hardness of computational problems than the GGM.

The gap between the two models becomes even more apparent if one considers structural properties of \mathbb{G} which cannot be meaningfully modeled as an oracle in the GGM. As an example, consider the Jacobi symbol, which was shown to be generically hard to compute in [JS09]. Indeed, it was left as an open problem in [AM09] to re-examine the equivalence of factoring and breaking the RSA assumption if an additional oracle for the Jacobi symbol were given. Though their results are stated in the *generic ring model* rather than the GGM, it seems they are similarly confronted with the issue of explicitly modeling such an oracle.

LIMITATIONS OF THE AGM. As already noted, one of the main benefits of our model over the GGM is the ability to reason about algorithms that arbitrarily exploit the structure of the group. So which algorithms are not covered in this manner? Obviously, outputting an obliviously sampled group element (with unknown representation) is forbidden. This coincides with the GGM of Maurer [Mau05] and which also excludes the possibility of obliviously sampling a random group element. For this reason, our model is strictly weaker than the one from [Mau05] in the sense that any security reduction derived in Maurer's GGM also holds in the AGM. In contrast, the GGM defined by Shoup [Sho97]

does allow for such a sampling process. Similar to Maurer’s GGM, we can allow obviously sampling a random group element \mathbf{X} through an additional oracle $\mathcal{O}()$ that can be called during the execution of \mathbf{A}_{alg} . By definition, the outputs of $\mathcal{O}()$ are added to the list $\vec{\mathbf{L}}$. We have thus argued that both versions of the GGM (i.e., the ones by Maurer and Shoup) are strictly stronger than the AGM. Also note that simulating $\mathcal{O}()$ to \mathbf{A}_{alg} as part of a reduction is straight-forward and always possible; the reduction simply samples r and returns g^r to the adversary. As the reduction knows r , adding $\mathcal{O}()$ to an experiment does not change it and is completely without loss of generality. From a practical point of view, it seems that generating and outputting a random group element without knowing a representation is generally not of much help. We therefore believe that the AGM captures most algorithms of practical interest.

1.3 Related Work and Open Questions

We have already mentioned the semi generic group model (SGGM) [JR10] as related work, but we discuss here some key differences of their model to ours in more detail. First, the SGGM is a very restrictive model in the sense that the class of problems it captures is limited. The main theorem of [JR10] (Theorem 3) holds only for pairing-based computational problems in which the output consists of a single element in either one of the base groups. In contrast, the AGM does not require a pairing group setting and thus applies to a much broader class of computational problems. Second, by extending the AGM to pairing groups, we are able to model all three groups as algebraic and reason again about a broader class of problems, in which the output can also consist of elements in the target group. To extend the AGM to the pairing setting, we allow the algebraic adversary to compute any element in the target group by applying the pairing to elements in the respective base groups.

Dent [Den02] shows that the generic group model as proposed Shoup [Sho97] inherits the known weaknesses in the random oracle model [CGH98]. Thus, there exist schemes which can be proven secure in Shoup’s GGM, but are pathologically insecure when viewed in the standard model. An interesting open question is whether the AGM bears similar weaknesses. A promising line of research related to this question has recently been initiated by Bitansky et al. [BCPR16]. Namely, they show that indistinguishability obfuscation (iO) implies the existence of non-extractable one-way functions. If these non-extractable one-way functions were furthermore algebraic (such as the knowledge of exponent assumption [Dam92]), then this would invalidate the AGM (under the assumption that iO exists).

Another promising direction for future research is to prove further reductions between common computational assumptions in the AGM. In particular, it would be interesting to classify different such assumptions within the AGM, for example along the lines of work [SS01, Kil01, Boy08, JR15, CM14, MRV16, GG17].

We leave it as an open problem to come up with a meaningful formalization of the AGM for *decisional assumptions*. At a technical level, the main difficulty in this task arises from the fact that an algorithm, i.e., distinguisher, in a decisional problem is asked to output a *bit* rather than a group element. Therefore, such an

algorithm is trivially considered algebraic in our framework. It would therefore be interesting to develop a model which captures the algebraic properties of such algorithms in more detail.

A further potential for follow-up work would be to investigate whether it is possible to automate proofs in the AGM. Indeed, for the case of the GGM this has been considered in [BFF+14, ABS16] and it would be interesting to see if similar automated tools can be derived for the AGM.

Finally, we remark that all of our results require prime-order groups and do not yet extend to the setting of pairing groups. When generalizing our results to composite-order groups, we expect to encounter the following technical difficulty: Given, e.g., an equation of the form $ax \equiv_n b$, where n is composite, there might be (exponentially) many solutions for the unknown x in case $\gcd(a, n) > 1$. This interferes with the proof strategies presented in this work and requires a more involved analysis. In fact, proving a reduction from the discrete logarithm problem to the CDH problem in the AGM for group orders containing multiple prime factors (e.g., $n = p^2$) is excluded by [MW98]. Hardness bounds in the GGM for composite-order groups have been considered in [Sho97, MW98, Mau05]. Generalizing the GGM to pairing groups has been the subject, e.g., of the works of [Boy08, KSW08, RLB+08]. Extending the AGM to either one of these regimes is an interesting line of research for future work.

2 Algebraic Algorithms

ALGORITHMS. We denote by $s \stackrel{\$}{\leftarrow} S$ the uniform sampling of the variable s from the (finite) set S . All our algorithms are probabilistic (unless stated otherwise) and written in uppercase letters A, B . To indicate that algorithm A runs on some inputs (x_1, \dots, x_n) and returns y , we write $y \stackrel{\$}{\leftarrow} A(x_1, \dots, x_n)$. If A has access to an algorithm B (via oracle access) during its execution, we write $y \stackrel{\$}{\leftarrow} A^B(x_1, \dots, x_n)$.

SECURITY GAMES. We use a variant of (code-based) *security games* [BR04, Sho04]. In game \mathbf{G}_{par} (defined relative to a set of parameters par), an adversary A interacts with a challenger that answers oracle queries issued by A . It has a main procedure and (possibly zero) oracle procedures which describe how oracle queries are answered. We denote the output of a game \mathbf{G}_{par} between a challenger and an adversary A via \mathbf{G}_{par}^A . A is said to *win* if $\mathbf{G}_{par}^A = 1$. We define the *advantage* of A in \mathbf{G}_{par} as $\text{Adv}_{par, A}^{\mathbf{G}} := \Pr [\mathbf{G}_{par}^A = 1]$ and the running time of \mathbf{G}_{par}^A as $\text{Time}_{par, A}^{\mathbf{G}}$.

cdh _{\mathcal{G}} ^A	cdh _{\mathcal{G}} ^{A_{alg}}
00 $x, y \xleftarrow{\$} \mathbb{Z}_p$	00 $x, y \xleftarrow{\$} \mathbb{Z}_p$
01 $(\mathbf{X}, \mathbf{Y}) := (g^x, g^y)$	01 $(\mathbf{X}, \mathbf{Y}) := (g^x, g^y)$
02 $\mathbf{Z} \xleftarrow{\$} \mathbf{A}(\mathbf{X}, \mathbf{Y})$	02 $[\mathbf{Z}]_{\vec{z}} \xleftarrow{\$} \mathbf{A}_{\text{alg}}(\mathbf{X}, \mathbf{Y})$
03 Return $(\mathbf{Z} = g^{xy})$	03 Return $(\mathbf{Z} = g^{xy})$

Fig. 1. Left: Algebraic game **cdh** relative to group description $\mathcal{G} = (\mathbb{G}, g, p)$ and adversary \mathbf{A} . All group elements are written in bold, uppercase letters. **Right:** Algebraic game **cdh** relative to group description $\mathcal{G} = (\mathbb{G}, g, p)$ and algebraic adversary \mathbf{A}_{alg} . The algebraic adversary \mathbf{A}_{alg} additionally returns a representation $\vec{z} = (a, b, c)$ of \mathbf{Z} such that $\mathbf{Z} = g^a \mathbf{X}^b \mathbf{Y}^c$.

SECURITY REDUCTIONS. Let \mathbf{G}, \mathbf{H} be security games. We write $\mathbf{H}_{\text{par}} \xrightarrow{(\Delta_\varepsilon, \Delta_t)} \mathbf{G}_{\text{par}}$ if there exists an algorithm \mathbf{R} (called $(\Delta_\varepsilon, \Delta_t)$ -reduction) such that for all algorithms \mathbf{A} , algorithm \mathbf{B} defined as $\mathbf{B} := \mathbf{R}^{\mathbf{A}}$ satisfies

$$\text{Adv}_{\text{par}, \mathbf{B}}^{\mathbf{H}} \geq \frac{1}{\Delta_\varepsilon} \cdot \text{Adv}_{\text{par}, \mathbf{A}}^{\mathbf{G}}, \quad \text{Time}_{\text{par}, \mathbf{B}}^{\mathbf{H}} \leq \Delta_t \cdot \text{Time}_{\text{par}, \mathbf{A}}^{\mathbf{G}}.$$

2.1 Algebraic Security Games and Algorithms

We consider *algebraic security games* $\mathbf{G}_{\mathcal{G}}$ for which we set *par* to a fixed group description $\mathcal{G} = (\mathbb{G}, g, p)$, where \mathbb{G} is a cyclic group of prime order p generated by g . In algebraic security games, we syntactically distinguish between elements of group \mathbb{G} (written in bold, uppercase letters, e.g., \mathbf{A}) and all other elements, which must not depend on any group elements. As an example of an algebraic security game, consider the Computational Diffie-Hellman game $\text{cdh}_{\mathcal{G}}^{\mathbf{A}}$, depicted in Fig. 1 (left).

We now define algebraic algorithms. Intuitively, the only way for an algebraic algorithm to output a new group element \mathbf{Z} is to derive it via group multiplications from known group elements.

Definition 1. (*Algebraic algorithm*) An algorithm \mathbf{A}_{alg} executed in an algebraic game $\mathbf{G}_{\mathcal{G}}$ is called algebraic if for all group elements \mathbf{Z} that \mathbf{A}_{alg} outputs (i.e., the elements in bold uppercase letters), it additionally provides the representation of \mathbf{Z} relative to all previously received group elements. That is, if $\vec{\mathbf{L}}$ is the list of group elements $\mathbf{L}_0, \dots, \mathbf{L}_m \in \mathbb{G}$ that \mathbf{A}_{alg} has received so far (w.l.o.g. $\mathbf{L}_0 = g$), then \mathbf{A}_{alg} must also provide a vector \vec{z} such that $\mathbf{Z} = \prod_i \mathbf{L}_i^{z_i}$. We denote such an output as $[\mathbf{Z}]_{\vec{z}}$.

REMARKS ON OUR MODEL. Algebraic algorithms were first considered in [BV98, PV05], where they are defined using an additional extractor algorithm which computes for an output group element a representation in basis $\vec{\mathbf{L}}$. We believe that our definition gives a simpler and cleaner definition of algebraic algorithms. If one assumes that the extractor algorithm has constant running time, then our definition is easily seen to be equivalent to theirs. Indeed, this

view makes sense for algorithms in the GGM since the representation \vec{z} trivially follows from the description of the algorithm. However, if running the extractor algorithm imposes some additional cost, then this will clearly affect the running times of our reductions. If the cost of the extractor is similar to that of the solver adversary, then reductions in our model that neither call an algebraic solver multiple times nor receive from it a non-constant amount of group elements (along with their representations) will remain largely the same in both models.

For the inputs to algebraic adversaries we syntactically distinguish group elements from other inputs and require that the latter not depend on any group elements. This is necessary to rule out pathological cases in which an algorithm receives “disguised” group elements and is forced to output an algebraic representation of them (which it might not know). To illustrate the issue, consider an efficient algorithm A , which on input $X' := \mathbf{X} \parallel \perp$ returns \mathbf{X} , where \mathbf{X} is a group element, but X' is not. If A is algebraic then it must return a representation of \mathbf{X} in g (the only group element previously seen), which would be the discrete logarithm of \mathbf{X} .

Allowing inputs of form X' while requiring algorithms to be algebraic leads to contradictions. (E.g., one could use A_{alg} to compute discrete logarithms: given a challenge $\mathbf{X} = g^x$, run $[\mathbf{X}]_x \stackrel{s}{\leftarrow} A_{\text{alg}}(\mathbf{X} \parallel \perp)$ and return x .) We therefore demand that non-group-element inputs must not depend on group elements. (Note that if A_{alg} 's input contains \mathbf{X} explicitly then it can output $[\mathbf{X}]_{(0,1)}$ with a valid representation of \mathbf{X} relative to $\vec{\mathbf{L}} = (g, \mathbf{X})$.)

Finally, we slightly abuse notation and let an algebraic algorithm also represent output group elements as combinations of previous *outputs*. This makes some of our proofs easier and is justified since all previous outputs must themselves have been given along with an according representation. Therefore, one can always recompute a representation that depends only on the initial inputs to the algebraic algorithm.

INTEGRATING WITH RANDOM ORACLES IN THE AGM. As mentioned above, an algorithm A that samples (and outputs) a group element \mathbf{X} obliviously, i.e., without knowing its representation, is not algebraic. This appears to be problematic if one wishes to combine the AGM with the Random Oracle Model [BR93]. However, group elements output by the random oracle are included by definition in the list $\vec{\mathbf{L}}$. This means that for any such element, a representation is trivially available to A_{alg} .

2.2 Generic Security Games and Algorithms

Generic algorithms A_{gen} are only allowed to use generic properties of group \mathcal{G} . Informally, an algorithm is generic if it works regardless of what group it is run in. This is usually modeled by giving an algorithm indirect access to group elements via abstract handles. It is straight-forward to translate all of our algebraic games into games that are syntactically compatible with generic algorithms accessing group elements only via abstract handles.

We say that winning algebraic game \mathbf{G}_G is (ε, t) -hard in the generic group model if for every generic algorithm A_{gen} it holds that

$$\mathbf{Time}_{G, A_{\text{gen}}}^G \leq t \implies \mathbf{Adv}_{G, A_{\text{gen}}}^G \leq \varepsilon.$$

We remark that usually in the generic group model one considers group operations (i.e., oracle calls) instead of the running time. In our context it is more convenient to measure the running time instead, assuming every oracle call takes one unit time.

As an important example, consider the algebraic Discrete Logarithm Game \mathbf{dlog}_G in Fig. 2 which is $(t^2/p, t)$ -hard in the generic group model [Sho97, Mau05].

We assume that a generic algorithm A_{gen} additionally provides the representation of \mathbf{Z} relative to all previously received group elements, for all group elements \mathbf{Z} that it outputs. This assumption is w.l.o.g. since a generic algorithm can only obtain new group elements by multiplying two known group elements; hence it always knows a valid representation. This way, every generic algorithm is also an algebraic algorithm.

Furthermore, if B_{gen} is a generic algorithm and A_{alg} is an algebraic algorithm, then $B_{\text{alg}} := B_{\text{gen}}^{A_{\text{alg}}}$ is also an algebraic algorithm. We refer to [Mau05] for more on generic algorithms.

2.3 Generic Reductions Between Algebraic Security Games

Let \mathbf{G}_G and \mathbf{H}_G be two algebraic security games. We write $\mathbf{H}_G \xrightarrow{(\Delta_\varepsilon, \Delta_t)}_{\text{alg}} \mathbf{G}_G$ if there exists a generic algorithm R_{gen} (called generic $(\Delta_\varepsilon, \Delta_t)$ -reduction) such that for every algebraic algorithm A_{alg} , algorithm B_{alg} defined as $B_{\text{alg}} := R_{\text{gen}}^{A_{\text{alg}}}$ satisfies

$$\mathbf{Adv}_{G, B_{\text{alg}}}^H \geq \frac{1}{\Delta_\varepsilon} \cdot \mathbf{Adv}_{G, A_{\text{alg}}}^G, \quad \mathbf{Time}_{G, B_{\text{alg}}}^H \leq \Delta_t \cdot \mathbf{Time}_{G, A_{\text{alg}}}^G.$$

Note that we deliberately require reduction R_{gen} to be generic. Hence, if A_{alg} is algebraic, then $B_{\text{alg}} := R_{\text{gen}}^{A_{\text{alg}}}$ is algebraic; if A_{alg} is generic, then $B_{\text{alg}} := R_{\text{gen}}^{A_{\text{alg}}}$ is generic. If one is only interested in algebraic adversaries, then it suffices to require reduction R_{gen} to be algebraic. But in that case one can no longer infer that $B_{\text{alg}} := R_{\text{gen}}^{A_{\text{alg}}}$ is generic in case A_{alg} is generic.

COMPOSING INFORMATION-THEORETIC LOWER BOUNDS WITH REDUCTIONS IN THE AGM. The following lemma explains how statements in the AGM carry over to the GGM.

Lemma 1. *Let \mathbf{G}_G and \mathbf{H}_G be algebraic security games such that $\mathbf{H}_G \xrightarrow{(\Delta_\varepsilon, \Delta_t)}_{\text{alg}} \mathbf{G}_G$ and winning \mathbf{H}_G is (ε, t) -hard in the GGM. Then, \mathbf{G}_G is $(\varepsilon \cdot \Delta_\varepsilon, t/\Delta_t)$ -hard in the GGM.*

Proof. Let A_{gen} be a generic algorithm playing in game \mathbf{G}_G . Then by our premise there exists a generic algorithm $B_{\text{alg}} = R_{\text{gen}}^{A_{\text{alg}}}$ such that

$$\mathbf{Adv}_{G, B_{\text{alg}}}^H \geq \frac{1}{\Delta_\varepsilon} \cdot \mathbf{Adv}_{G, A_{\text{alg}}}^G, \quad \mathbf{Time}_{G, B_{\text{alg}}}^H \leq \Delta_t \cdot \mathbf{Time}_{G, A_{\text{alg}}}^G.$$

\mathbf{dlog}_G^A	$\mathbf{lc-dh}_G^A$	$\mathbf{sq-dh}_G^A$
00 $x \xleftarrow{\$} \mathbb{Z}_p$	00 $x, y \xleftarrow{\$} \mathbb{Z}_p$	00 $x \xleftarrow{\$} \mathbb{Z}_p$
01 $\mathbf{X} := g^x$	01 $(\mathbf{X}, \mathbf{Y}) := (g^x, g^y)$	01 $\mathbf{X} := g^x$
02 $z \xleftarrow{\$} A(\mathbf{X})$	02 $(\mathbf{Z}, u, v, w) \xleftarrow{\$} A(\mathbf{X}, \mathbf{Y})$	02 $\mathbf{Z} \xleftarrow{\$} A(\mathbf{X})$
03 Return ($z = x$)	03 Return ($\mathbf{Z} = g^{ux^2+vx y+wy^2}$ $\wedge (u \neq 0 \vee v \neq 0 \vee w \neq 0)$)	03 Return ($\mathbf{Z} = g^{x^2}$)

Fig. 2. Discrete Logarithm Game \mathbf{dlog} , Square Diffie-Hellman Game $\mathbf{sq-dh}$, and Linear Combination Diffie-Hellman Game $\mathbf{lc-dh}$ relative to group \mathcal{G} and adversary A .

Assume $\mathbf{Time}_{\mathcal{G}, A_{\text{alg}}}^G \leq t/\Delta_t$; then $\mathbf{Time}_{\mathcal{G}, B_{\text{alg}}}^H \leq \Delta_t \cdot \mathbf{Time}_{\mathcal{G}, A_{\text{alg}}}^G \leq t$. Since winning \mathbf{H}_G is (ε, t) -hard in the GGM, it follows that

$$\varepsilon \geq \mathbf{Adv}_{\mathcal{G}, B_{\text{alg}}}^H \geq \frac{1}{\Delta_\varepsilon} \cdot \mathbf{Adv}_{\mathcal{G}, A_{\text{alg}}}^G$$

and thus $\varepsilon \cdot \Delta_\varepsilon \geq \mathbf{Adv}_{\mathcal{G}, A_{\text{alg}}}^G$, which proves that \mathbf{G}_G is $(\varepsilon\Delta_\varepsilon, t/\Delta_t)$ -hard in the GGM. \square

3 The Diffie-Hellman Assumption and Variants

In this section we consider some variants of the standard Diffie-Hellman assumption [DH76] and prove them to be equivalent to the discrete logarithm assumption (defined via algebraic game \mathbf{dlog}_G of Fig. 2) in the Algebraic Group Model.

3.1 Computational Diffie-Hellman

Consider the Square Diffie-Hellman Assumption [MW99] described in algebraic game $\mathbf{sq-dh}_G$ and the Linear Combination Diffie-Hellman Assumption described in algebraic game $\mathbf{lc-dh}_G$ (both in Fig. 2), which will be convenient for the proof of Theorem 2.

As a warm-up we now prove that the Discrete Logarithm assumption is tightly equivalent to the Diffie-Hellman, the Square Diffie-Hellman, and the Linear Combination Diffie-Hellman Assumption in the Algebraic Group Model. The equivalence of the Square Diffie-Hellman and Diffie-Hellman problems was previously proven in [MW99, BDZ03].

Theorem 1. $\mathbf{dlog}_G \xrightarrow{(1,1)}_{\text{alg}} \{\mathbf{cdh}_G, \mathbf{sq-dh}_G\}$ and $\mathbf{dlog}_G \xrightarrow{(3,1)}_{\text{alg}} \mathbf{lc-dh}_G$.

Proof. Let A_{alg} be an algebraic adversary executed in game $\mathbf{sq-dh}_G$; cf. Fig. 3.

As A_{alg} is an algebraic adversary, it returns a solution \mathbf{Z} together with a representation $(a, b) \in \mathbb{Z}_p^2$ such that

$$\mathbf{Z} = g^{x^2} = g^a (g^x)^b. \quad (1)$$




Fig. 3. Algebraic adversary \mathcal{A}_{alg} playing in $\mathbf{sq-dh}_G$.

We now show how to construct a generic reduction \mathcal{R}_{gen} that calls \mathcal{A}_{alg} exactly once such that for $\mathcal{B}_{\text{alg}} := \mathcal{R}_{\text{gen}}^{\mathcal{A}_{\text{alg}}}$ we have

$$\mathbf{Adv}_{\mathcal{G}, \mathcal{B}_{\text{alg}}}^{\text{dlog}} = \mathbf{Adv}_{\mathcal{G}, \mathcal{A}_{\text{alg}}}^{\mathbf{sq-dh}}.$$

\mathcal{R}_{gen} works as follows. On input a discrete logarithm instance \mathbf{X} , it runs \mathcal{A}_{alg} on \mathbf{X} . Suppose \mathcal{A}_{alg} is successful. Equation (1) is equivalent to the quadratic equation $x^2 - bx - a \equiv_p 0$ with at most two solutions in x . (In general such equations are not guaranteed to have a solution but since the representation is valid and \mathcal{A}_{alg} is assumed to be correct, there exists at least one solution for x .) \mathcal{R}_{gen} can test which one (out of the two) is the correct solution x by testing against $\mathbf{X} = g^x$. Moreover, it is easy to see that \mathcal{R}_{gen} only performs generic group operations and is therefore generic. Hence, $\mathcal{B}_{\text{alg}} := \mathcal{R}_{\text{gen}}^{\mathcal{A}_{\text{alg}}}$ is algebraic, which proves

$$\text{dlog}_G \xrightarrow{(1,1)}_{\text{alg}} \mathbf{sq-dh}_G.$$

The statement $\text{dlog}_G \xrightarrow{(1,1)}_{\text{alg}} \mathbf{cdh}_G$ follows, since given an adversary against \mathbf{cdh}_G (see Fig. 1), we can easily construct an adversary against $\mathbf{sq-dh}_G$ that runs in the same time and has the same probability of success (given $\mathbf{X} = g^x$, sample $r \stackrel{\$}{\leftarrow} \mathbb{Z}_p$, run the \mathbf{cdh} adversary on $(\mathbf{X}, \mathbf{X}^r)$, obtain \mathbf{Z} and return $\mathbf{Z}^{\frac{1}{r}}$).

It remains to show that $\mathbf{sq-dh}_G \xrightarrow{(3,1)}_{\text{alg}} \mathbf{lc-dh}_G$. Given an algebraic solver \mathcal{C}_{alg} executed in game $\mathbf{lc-dh}_G$, we construct an adversary \mathcal{A}_{alg} against $\mathbf{sq-dh}_G$ as follows: On input $\mathbf{X} = g^x$, \mathcal{A}_{alg} samples $r \stackrel{\$}{\leftarrow} \mathbb{Z}_p$ and computes either (\mathbf{X}, g^r) , (g^r, \mathbf{X}) , or $(\mathbf{X}, \mathbf{X}^r)$ each with probability 1/3. Note that this instance is correctly distributed. It then runs \mathcal{C}_{alg} on the resulting tuple $(\mathbf{X}_1, \mathbf{X}_2)$ and receives (\mathbf{Z}, u, v, w) together with (a, b, c) s.t. $\mathbf{Z} = g^a \mathbf{X}_1^b \mathbf{X}_2^c$. If $u \neq 0$, then the choice $\mathbf{X}_1 = \mathbf{X}$, $\mathbf{X}_2 = g^r$ yields $\mathbf{Z} = g^{u x^2 + v x r + w r^2}$, from which g^{x^2} can be computed as $g^{x^2} = (\mathbf{Z} \mathbf{X}^{-v r} g^{-w r^2})^{\frac{1}{u}}$. Clearly, \mathcal{A}_{alg} is able to compute an algebraic representation of g^{x^2} from the values (a, b, c) and thus is algebraic itself. The cases $v \neq 0, w \neq 0$ follow in a similar fashion. \square

Corollary 1. \mathbf{cdh}_G and $\mathbf{sq-dh}_G$ are $(t^2/p, t)$ -hard in the generic group model and $\mathbf{lc-dh}_G$ is $(3t^2/p, t)$ -hard in the generic group model.

For the subsequent sections and proofs, we will not make explicit the reduction algorithm \mathcal{R}_{gen} every time (as done above).

3.2 Strong Diffie-Hellman

Consider the Strong Diffie-Hellman Assumption [ABR01] described via game $\mathbf{sdh}_{\mathcal{G}}$ in Fig. 4. We now prove that the Discrete Logarithm Assumption (non-tightly) implies the Strong Diffie-Hellman Assumption in the Algebraic Group Model. We briefly present the main ideas of the proof. The full proof of Theorem 2 can be found in the full version of our paper [FKL17]. Let \mathbf{A}_{alg} be an algebraic adversary playing in $\mathbf{sdh}_{\mathcal{G}}$ and let $\mathbf{Z} = g^z$ denote the Discrete Logarithm challenge. We show an adversary \mathbf{B}_{alg} against $\mathbf{dlog}_{\mathcal{G}}$ that simulates $\mathbf{sdh}_{\mathcal{G}}$ to \mathbf{A}_{alg} . \mathbf{B}_{alg} appropriately answers \mathbf{A}_{alg} 's queries to the oracle $\mathbf{O}(\cdot, \cdot)$ by using the algebraic representation of the queried elements provided by \mathbf{A}_{alg} . Namely, when $(\mathbf{Y}', \mathbf{Z}')$ is asked to the oracle, \mathbf{B}_{alg} obtains vectors \vec{b}, \vec{c} such that $\mathbf{Y}' = g^{b_1} \mathbf{X}^{b_2} \mathbf{Y}^{b_3}$ and $\mathbf{Z}' = g^{c_1} \mathbf{X}^{c_2} \mathbf{Y}^{c_3}$. As long as $b_2 = b_3 = 0$, \mathbf{B}_{alg} can answer all of \mathbf{A}_{alg} 's queries by checking whether $\mathbf{X}^{b_1} = \mathbf{Z}'$. On the other hand, if $b_2 \neq 0$ or $b_3 \neq 0$, then \mathbf{B}_{alg} simply returns 0. Informally, the simulation will be perfect unless \mathbf{A}_{alg} manages to compute a valid solution to $\mathbf{lc-dh}_{\mathcal{G}}$. All of these games can be efficiently simulated by \mathbf{B}_{alg} , as we have shown in the previous section.

$\mathbf{sdh}_{\mathcal{G}}^{\mathbf{A}}$ 00 $x, y \xleftarrow{\$} \mathbb{Z}_p$ 01 $(\mathbf{X}, \mathbf{Y}) := (g^x, g^y)$ 02 $\mathbf{Z} \xleftarrow{\$} \mathbf{A}^{\mathbf{O}(\cdot, \cdot)}(\mathbf{X}, \mathbf{Y})$ 03 Return $(\mathbf{Z} = g^{xy})$	$\mathbf{O}(\mathbf{Y}', \mathbf{Z}')$ 04 Return $(\mathbf{Z}' = (\mathbf{Y}')^x)$
--	---

Fig. 4. Strong Diffie-Hellman Game \mathbf{sdh} relative to \mathcal{G} and adversary \mathbf{A} .

Theorem 2. $\mathbf{dlog}_{\mathcal{G}} \xrightarrow[\text{alg}]{(4q,1)} \mathbf{sdh}_{\mathcal{G}}$, where q is the maximum number of queries to oracle $\mathbf{O}(\cdot, \cdot)$ in $\mathbf{sdh}_{\mathcal{G}}$.

Corollary 2. $\mathbf{sdh}_{\mathcal{G}}$ is $(\frac{t^2}{4pq}, t)$ -hard in the generic group model.

4 The LRSW Assumption

The interactive LRSW assumption [LRSW99, CL04] is defined via the algebraic security game \mathbf{lrsw} in Fig. 5.

We prove that the LRSW assumption is (non-tightly) implied by the Discrete Logarithm Assumption in the Algebraic Group Model. We give a high-level sketch of the main ideas here; the full proof of Theorem 3 can be found in the full version of our paper [FKL17]. Let \mathbf{A}_{alg} be an algebraic adversary playing in $\mathbf{lrsw}_{\mathcal{G}}$ and let $\mathbf{Z} = g^z$ denote the Discrete Logarithm challenge. We construct an adversary \mathbf{B}_{alg} against $\mathbf{dlog}_{\mathcal{G}}$, which simulate $\mathbf{lrsw}_{\mathcal{G}}$ to \mathbf{A}_{alg} by embedding the value of z in one of three possible ways. Namely, it either sets $\mathbf{X} := \mathbf{Z}$ or $\mathbf{Y} := \mathbf{Z}$, or it chooses a random the query by \mathbf{A}_{alg} to the oracle $\mathbf{O}(\cdot)$ in $\mathbf{lrsw}_{\mathcal{G}}$ to embed the

$\text{lrsw}_{\mathcal{G}}^{\mathbf{A}}$ 00 $Q := \emptyset$ 01 $x, y \xleftarrow{\$} \mathbb{Z}_p$ 02 $\mathbf{X} := g^x, \mathbf{Y} := g^y$ 03 $(m^*, \mathbf{A}^*, \mathbf{B}^*, \mathbf{C}^*) \xleftarrow{\$} \mathbf{A}^{\mathcal{O}(\cdot)}(\mathbf{X}, \mathbf{Y})$ 04 Return $(m^* \notin Q \wedge m^* \neq 0 \wedge \mathbf{A}^* \neq 1$ $\quad \wedge \mathbf{B}^* = (\mathbf{A}^*)^y \wedge \mathbf{C}^* = (\mathbf{A}^*)^{m^*xy+x})$	$\mathcal{Q}(m_j)$ //For query j 05 $r_j \xleftarrow{\$} \mathbb{Z}_p$; 06 $\mathbf{A}_j := g^{r_j}$ 07 $\mathbf{B}_j := g^{r_j y}$ 08 $\mathbf{C}_j := g^{r_j m_j xy + r_j x}$ 09 $Q := Q \cup \{m_j\}$ 10 Return $(\mathbf{A}_j, \mathbf{B}_j, \mathbf{C}_j)$
--	--

Fig. 5. Game lrsw relative to \mathcal{G} and adversary \mathbf{A} .

value of z . These behaviours correspond in our proof to the adversaries $\mathbf{C}_{\text{alg}}, \mathbf{D}_{\text{alg}}$, and \mathbf{E}_{alg} , respectively. After obtaining a solution $(m^*, [\mathbf{A}^*]_{\vec{a}}, [\mathbf{B}^*]_{\vec{b}}, [\mathbf{C}^*]_{\vec{c}})$ on a fresh value $m^* \neq 0$ from \mathbf{A}_{alg} , the adversaries use the algebraic representations $\vec{a}, \vec{b}, \vec{c}$ obtained from \mathbf{A}_{alg} to suitably rewrite the values of $\mathbf{A}^*, \mathbf{C}^*$. They then make use of the relation $(\mathbf{A}^*)^{(xm^*y+x)} = \mathbf{C}^*$ to obtain an equation mod p , which in turn gives z .

Theorem 3. $\text{dlog}_{\mathcal{G}} \xrightarrow{(6q,1)}_{\text{alg}} \text{lrsw}_{\mathcal{G}}$, where $q \geq 6$ is the maximum number of queries to $\mathcal{O}(\cdot)$ in $\text{lrsw}_{\mathcal{G}}$.

Corollary 3. $\text{lrsw}_{\mathcal{G}}$ is $(t, \frac{t^2}{6pq})$ -hard in the generic group model.

5 ElGamal Encryption

In this section we prove that the IND-CCA1 (aka. lunchtime security) of the ElGamal encryption scheme (in its abstraction as a KEM) is implied by a parametrized (“ q -type”) variant of the Decision Diffie-Hellman Assumption in the Algebraic Group Model.

ADVANTAGE FOR DECISIONAL ALGEBRAIC SECURITY GAMES. We parameterize a *decisional* algebraic game \mathbf{G} (such as the game in Fig. 7) with a parameter bit b . We define the advantage of adversary \mathbf{A} in \mathbf{G} as

$$\text{Adv}_{\text{par}, \mathbf{A}}^{\mathbf{G}} := |\Pr [\mathbf{G}_{\text{par}, 0}^{\mathbf{A}} = 1] - \Pr [\mathbf{G}_{\text{par}, 1}^{\mathbf{A}} = 1]|.$$

We define $\text{Time}_{\text{par}, \mathbf{A}_{\text{alg}}}^{\mathbf{G}}$ independently of the parameter bit b , i.e., we consider only adversaries that have the same running time in both games $\mathbf{G}_{\text{par}, 0}, \mathbf{G}_{\text{par}, 1}$. In order to cover games that define the security of schemes (rather than assumptions), instead of $\text{par} = \mathcal{G}$, we only require that \mathcal{G} be *included* in par . Let $\mathbf{G}_{\text{par}}, \mathbf{H}_{\text{par}}$ be decisional algebraic security games. As before, we write $\mathbf{H}_{\text{par}} \xrightarrow{(\Delta_\varepsilon, \Delta_t)}_{\text{alg}} \mathbf{G}_{\text{par}}$ if there exists a generic algorithm \mathbf{R}_{gen} (called generic $(\Delta_\varepsilon, \Delta_t)$ -reduction) such that for algebraic algorithm \mathbf{B}_{alg} defined as $\mathbf{B}_{\text{alg}} := \mathbf{R}_{\text{gen}}^{\mathbf{A}_{\text{alg}}}$, we have

$$\text{Adv}_{\text{par}, \mathbf{B}_{\text{alg}}}^{\mathbf{H}} \geq \frac{1}{\Delta_\varepsilon} \cdot \text{Adv}_{\text{par}, \mathbf{A}_{\text{alg}}}^{\mathbf{G}}, \quad \text{Time}_{\text{par}, \mathbf{B}_{\text{alg}}}^{\mathbf{H}} \leq \Delta_t \cdot \text{Time}_{\text{par}, \mathbf{A}_{\text{alg}}}^{\mathbf{G}}.$$

$\mathbf{ind}\text{-cca1}_{\text{KEM}, \text{par}, b}^A$	$\text{Dec}(C)$	$\text{Enc}()$ // <i>One time</i>
00 $(pk, sk) \xleftarrow{\$} \text{Gen}(\text{par})$	// <i>Before Enc is called</i>	05 $(K_0^*, C^*) \xleftarrow{\$} \text{Enc}(pk)$
01 $b' \xleftarrow{\$} \text{A}^{\text{Dec}, \text{Enc}}(pk)$	03 $K \xleftarrow{\$} \text{Dec}(C, sk)$	06 $K_1^* \xleftarrow{\$} \mathcal{K}$
02 Return b'	04 Return K	07 Return (K_b^*, C^*)

Fig. 6. IND-CCA1 Game $\mathbf{ind}\text{-cca1}$ relative to KEM $\text{KEM} = (\text{Gen}, \text{Enc}, \text{Dec})$, parameters par , and adversary A .

KEY ENCAPSULATION MECHANISMS. A *key encapsulation mechanism* (KEM for short) $\text{KEM} = (\text{Gen}, \text{Enc}, \text{Dec})$ is a triple of algorithms together with a symmetric-key space \mathcal{K} . The randomized *key generation algorithm* Gen takes as input a set of parameters, par , and outputs a public/secret key pair (pk, sk) . The *encapsulation algorithm* Enc takes as input a public key pk and outputs a key/ciphertext pair (K, C) such that $K \xleftarrow{\$} \mathcal{K}$. The deterministic *decapsulation algorithm* Dec takes as input a secret key sk and a ciphertext C and outputs a key $K \in \mathcal{K}$ or a special symbol \perp if C is invalid. We require that KEM be *correct*: For all possible pairs (K, C) output by $\text{Enc}(pk)$, we have $\text{Dec}(sk, C) = K$. We formalize IND-CCA1 security of a KEM via the games (for $b = 0, 1$) depicted in Fig. 6.

In the following, we consider the ElGamal KEM EG defined in Fig. 8. We also consider a stronger variant of the well-known Decisional Diffie-Hellman (DDH) assumption that is parametrized by an integer q . In the q -DDH game, defined in Fig. 7, the adversary receives, in addition to (g^x, g^y) , the values g^{x^2}, \dots, g^{x^q} .

Lemma 2. [Che06] For $q < p^{1/3}$, $\mathbf{q}\text{-ddh}_{\mathcal{G}}$ is $(\frac{t^2 q}{p \log p}, t)$ -hard in the generic group model.

The proof of the following theorem can be found in the full version of our paper [FKL17]. In the proof, we consider the algebraic games depicted in Fig. 9.

Theorem 4. $\mathbf{ind}\text{-cca1}_{\text{EG}, \mathcal{G}} \xleftrightarrow{\text{alg}}^{(1,1)} \mathbf{q}\text{-ddh}_{\mathcal{G}}$, where $q - 1$ is the maximal number of queries to $\text{Dec}(\cdot)$ in $\mathbf{ind}\text{-cca1}_{\text{EG}, \mathcal{G}}$.

Corollary 4. For $q < p^{1/3}$, $\mathbf{ind}\text{-cca1}_{\text{EG}, \mathcal{G}}$ is $(\frac{t^2 q}{p \log p}, t)$ -hard in the generic group model, where $q - 1$ is the maximal number of queries to $\text{Dec}(\cdot)$ in $\mathbf{ind}\text{-cca1}_{\text{EG}, \mathcal{G}}$.

6 Tight Reduction for the BLS Scheme

For this section, we introduce the notion of groups \mathbb{G} equipped with a symmetric, (non-degenerate) bilinear map $e: \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$, where \mathbb{G}_T denotes the so-called *target group*. We now set $\mathcal{G} = (p, \mathbb{G}, \mathbb{G}_T, g, e)$ (Fig. 9).

SIGNATURE SCHEMES. A *signature scheme* $\text{SIG} = (\text{SIGGen}, \text{SIGSig}, \text{SIGVer})$ is a triple of algorithms. The randomized *key generation algorithm* SIGGen takes as input a set of parameters, par , and outputs a public/secret key pair (pk, sk) . The randomized *signing algorithm* SIGSig takes as input a secret key sk and a message

$\mathbf{q\text{-ddh}}_{\mathcal{G},b}^A$ 00 $x, r, z \xleftarrow{\$} \mathbb{Z}_p$ 01 $b' \xleftarrow{\$} A(g^x, g^{x^2}, \dots, g^{x^q}, g^r, g^{xr+zb})$ 02 Return b'
--

Fig. 7. q -Decisional Diffie-Hellman Game $\mathbf{q\text{-ddh}}$ relative to \mathcal{G} and adversary A .

$\mathbf{Gen}(\mathcal{G})$ 00 $x \xleftarrow{\$} \mathbb{Z}_p$ 01 $\mathbf{X} := g^x$ 02 Return $(pk, sk) := (\mathbf{X}, x)$	$\mathbf{Enc}(pk) :$ 03 $r \xleftarrow{\$} \mathbb{Z}_p$ 04 $\mathbf{C} := g^r$ 05 $\mathbf{K} := \mathbf{X}^r$ 06 Return (\mathbf{K}, \mathbf{C})	$\mathbf{Dec}(\mathbf{C}, sk) :$ 07 If $\mathbf{C} \notin \mathcal{G}$ 08 Return \perp 09 $\tilde{\mathbf{K}} := \mathbf{C}^x$ 10 Return $\tilde{\mathbf{K}}$
--	---	--

Fig. 8. ElGamal KEM $\mathbf{EG} = (\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$

$\mathbf{ind\text{-}cca1}_{\mathbf{EG}, \mathcal{G}, 0}^A$ 00 $x \xleftarrow{\$} \mathbb{Z}_p$ 01 $\mathbf{X} := g^x$ 02 $b' \xleftarrow{\$} A_{\text{alg}}^{\mathbf{Dec}, \mathbf{Enc}}(\mathbf{X})$ 03 Return b'	$\mathbf{ind\text{-}cca1}_{\mathbf{EG}, \mathcal{G}, 1}^A$	$\mathbf{Dec}([\mathbf{C}]_{\bar{x}})$ // Before Enc is called 04 $\mathbf{K} := \mathbf{C}^x$ 05 Return \mathbf{K}	$\mathbf{Enc}() \quad // \text{One time}$ 06 $r \xleftarrow{\$} \mathbb{Z}_p$ 07 $\mathbf{C}^* := g^r$ 08 $\mathbf{K}^* := \mathbf{X}^r$ 09 $\tilde{\mathbf{K}}^* \xleftarrow{\$} \mathcal{K}$ 10 Return $(\tilde{\mathbf{K}}^*, \mathbf{C}^*)$
---	--	---	---

Fig. 9. Games $\mathbf{ind\text{-}cca1}_{\mathbf{EG}, \mathcal{G}, 0}^A$ and $\mathbf{ind\text{-}cca1}_{\mathbf{EG}, \mathcal{G}, 1}^A$ with algebraic adversary A_{alg} . The boxed statement is only executed in $\mathbf{ind\text{-}cca1}_{\mathbf{EG}, \mathcal{G}, 1}^A$.

m in the message space \mathcal{M} and outputs a signature σ in the signature space \mathcal{S} . The deterministic *signature verification algorithm* \mathbf{SIGVer} takes as input a public key pk , a message m , and a signature σ and outputs $b \in \{0, 1\}$. We require that \mathbf{SIG} be *correct*: For all possible pairs (pk, sk) output by \mathbf{SIGGen} , and all messages $m \in \mathcal{M}$, we have $\Pr[\mathbf{SIGVer}(pk, m, \mathbf{SIGSig}(m, sk)) = 1] = 1$. We formalize *unforgeability under chosen message attacks* for \mathbf{SIG} via game $\mathbf{uf\text{-}cma}_{\mathbf{SIG}, \text{par}}$ depicted in Fig. 10.

In the following, we show how in the AGM with a random oracle, the security of the BLS signature scheme [BLS04], depicted in Fig. 11, can be tightly reduced to the discrete logarithm problem. Boneh, Lynn and Shacham [BLS04] only prove a loose reduction to the CDH problem. In the AGM we manage to improve the quality of the reduction, leveraging the fact that a forgery comes with a representation in the basis of all previously answered random oracle and signature queries. We embed a discrete logarithm challenge in either the secret key or inside the random oracle queries—a choice that remains hidden from the adversary. Depending on the adversary’s behavior we always solve the discrete logarithm challenge in one of the cases.

$\mathbf{uf-cma}_{\text{SIG}, \text{par}}^{\text{A}}$ 00 $(pk, sk) \xleftarrow{\$} \text{SIGGen}$ 01 $Q := \emptyset$ 02 $(m^*, \sigma^*) \xleftarrow{\$} \text{A}^{\text{O}(\cdot)}$ 03 Return $(m^* \notin Q \wedge \text{SIGVer}(m^*, \sigma^*))$	$\underline{\text{Q}}(m)$ 04 $Q := Q \cup \{m\}$ 05 $\sigma \xleftarrow{\$} \text{SIGSig}(m, sk)$ 06 Return σ
--	---

Fig. 10. Game $\mathbf{uf-cma}$ defining (existential) unforgeability under chosen-message attacks for signature scheme SIG , parameters par and adversary A .

$\underline{\text{BLSGen}}(\mathcal{G})$ 00 $x \xleftarrow{\$} \mathbb{Z}_p$ 01 $\mathbf{X} := g^x$ 02 $sk := x$ 03 $pk := \mathbf{X}$ 04 Return (pk, sk)	$\underline{\text{BLSSig}}(m)$ 05 $\Sigma := H(m)^x$ 06 Return Σ	$\underline{\text{BLSVer}}(m, \Sigma)$ 07 Return $(e(H(m), \mathbf{X}) = e(\Sigma, g))$
--	---	--

Fig. 11. Boneh, Lynn and Shacham's signature scheme $\text{BLS}_{\mathcal{G}}$. Here, H is a hash function that is modeled as a random oracle.

Theorem 5. $\text{dlog}_{\mathcal{G}} \xrightarrow{(4,1)}_{\text{alg}} \mathbf{uf-cma}_{\text{BLS}, \mathcal{G}}$ in the random oracle model.

Proof. Let A_{alg} be an algebraic adversary playing in $\mathbf{G} := \mathbf{uf-cma}_{\text{BLS}, \mathcal{G}}^{\text{A}_{\text{alg}}}$, depicted in Fig. 12.

As A_{alg} is an algebraic adversary, at the end of \mathbf{G} , it returns a forgery Σ^* on a message $m^* \notin Q$ together with a representation $\vec{a} = (\hat{a}, a', \bar{a}_1, \dots, \bar{a}_q, \tilde{a}_1, \dots, \tilde{a}_q)$ s.t.

$$\Sigma^* = H(m^*)^x = g^{\hat{a}} \mathbf{X}^{a'} \prod_{i=1}^q \mathbf{H}_i^{\bar{a}_i} \prod_{i=1}^q \Sigma_i^{\tilde{a}_i}. \quad (2)$$

Here, the representation is split (from left to right) into powers of the generator g , the public key \mathbf{X} , all of the answers to hash queries $\mathbf{H}_i, i \in [q]$, and the signatures $\Sigma_i, i \in [q]$, returned by the signing oracle. In the following, we denote the discrete logarithm of $H(m^*)$ w.r.t. basis g as r^* and for all $i \in [q]$, we denote

$\underline{\mathbf{G}}^{\text{A}_{\text{alg}}}$ 00 $x \xleftarrow{\$} \mathbb{Z}_p$ 01 $\mathbf{X} := g^x$ 02 $Q := \emptyset$ 03 $(m^*, [\Sigma^*]_{\vec{a}}) \xleftarrow{\$} \text{A}_{\text{alg}}^{\text{O}(\cdot), \text{H}(\cdot)}(\mathbf{X})$ 04 Return $(m^* \notin Q \wedge \Sigma^* = \text{H}(m^*)^x)$	$\underline{\text{Q}}(m_i)$ 05 $Q := Q \cup \{m_i\}$ 06 $\Sigma_i \leftarrow \text{H}(m_i)^x$ 07 Return Σ_i	$\underline{\text{H}}(m_i)$ 08 $\mathbf{H}_i \leftarrow H(m_i)$ 09 Return \mathbf{H}_i
---	---	--

Fig. 12. Game $\mathbf{G} = \mathbf{uf-cma}_{\text{BLS}, \mathcal{G}}^{\text{A}_{\text{alg}}}$ relative to adversary A_{alg} .

the discrete logarithm of $H(m_i)$ as r_i . Equation (2) is thus equivalent to

$$xr^* \equiv_p x(a' + \sum_i r_i \tilde{a}_i) + (\hat{a} + \sum_i r_i \tilde{a}_i). \tag{3}$$

In the full version of our paper we show how to efficiently compute x from Eq. (3). Again, the main idea is to perform a case distinction over the cases where the coefficient of x is zero, or non-zero, respectively. \square

Corollary 5. *uf-cma_{BLS,G} is $(t, \frac{t^2}{4p})$ -hard in the generic group model with a random oracle.*

7 Groth’s Near-Optimal zk-SNARK

In order to cover notions such as knowledge soundness, which are defined via games for two algorithms, we generalize the notion of algebraic games and reductions between them. We write $\mathbf{G}_{par}^{A,X}$ to denote that A and X play in \mathbf{G}_{par} and define the advantage $\mathbf{Adv}_{par,A,X}^{\mathbf{G}} := \Pr[\mathbf{G}_{par}^{A,X} = 1]$ and the running time $\mathbf{Time}_{par,A,X}^{\mathbf{G}}$ as before. To capture definitions that require that for every A there exists some X (which has black-box access to A) such that $\mathbf{Adv}_{par,A,X}^{\mathbf{G}}$ is small, we define algebraic reductions for games \mathbf{G}_{par} of this type as follows.

We write $\mathbf{H}_{par} \xrightarrow{(\Delta_x, \Delta_t)}_{\text{alg}} \mathbf{G}_{par}$ if there exist generic algorithms R_{gen} and S_{gen} such that for all algebraic algorithms A_{alg} we have

$$\mathbf{Adv}_{par,B_{\text{alg}}}^{\mathbf{H}} \geq \frac{1}{\Delta_\varepsilon} \cdot \mathbf{Adv}_{par,A_{\text{alg}},X_{\text{alg}}}^{\mathbf{G}}, \quad \mathbf{Time}_{par,B_{\text{alg}}}^{\mathbf{H}} \leq \Delta_t \cdot \mathbf{Time}_{par,A_{\text{alg}},X_{\text{alg}}}^{\mathbf{G}}.$$

with B_{alg} defined as $B_{\text{alg}} := R_{\text{gen}}^{A_{\text{alg}}}$ and X_{alg} defined as $X_{\text{alg}} := S_{\text{gen}}^{A_{\text{alg}}}$.

THE q -DISCRETE LOGARITHM ASSUMPTION. We define a parametrized (“ q -type”) variant of the DLog assumption via the algebraic security game q -**dlog** in Fig. 13. We will show that Groth’s [Gro16] scheme, which is the most efficient SNARK system to date, is secure under q -DLog in the algebraic group model.




Fig. 13. q -Discrete Logarithm Game q -**dlog** relative to \mathcal{G} and adversary A .

NON-INTERACTIVE ZERO-KNOWLEDGE ARGUMENTS OF KNOWLEDGE. Groth [Gro16] considers proof systems for satisfiability of arithmetic circuits, which consist of addition and multiplication gates over a finite field \mathbb{F} . As a tool, Gennaro et al. [GGPR13] show how to efficiently convert any arithmetic circuit into a quadratic arithmetic program (QAP) R , which is described by \mathbb{F} ,

integers $\ell \leq m$ and polynomials $u_i, v_i, w_i \in \mathbb{F}[X]$, for $0 \leq i \leq m$, and $t \in \mathbb{F}[X]$, where the degrees of u_i, v_i, w_i are less than the degree n of t . (The relation R can also contain additional information *aux*.)

A QAP R defines the following binary relation of statements ϕ and witnesses w , where we set $a_0 := 1$:

$$R = \left\{ (\phi, w) \mid \begin{array}{l} \phi = (a_1, \dots, a_\ell) \in \mathbb{F}^\ell, w = (a_{\ell+1}, \dots, a_m) \in \mathbb{F}^{m-\ell} \\ (\sum_{i=0}^m a_i u_i(X)) \cdot (\sum_{i=0}^m a_i v_i(X)) \equiv \sum_{i=0}^m a_i w_i(X) \pmod{t(X)} \end{array} \right\}$$

A non-interactive argument system for a class of relations \mathcal{R} is a 3-tuple $\text{SNK} = (\text{Setup}, \text{Prv}, \text{Vfy})$ of algorithms. **Setup** on input a relation $R \in \mathcal{R}$ outputs a common reference string *crs*; prover algorithm **Prv** on input *crs* and a statement/witness pair $(\phi, w) \in R$ returns an argument π ; Verification **Vfy** on input *crs*, ϕ and π returns either 0 (reject) or 1 (accept). We require SNK to be *complete*, i.e., for all *crs* output by **Setup**, all arguments for true statements produced by **Prv** are accepted by **Vfy**.

Knowledge soundness requires that for every adversary A there exists an extractor X_A that extracts a witness from any valid argument output by A . We write $(y; z) \stackrel{s}{\leftarrow} (A \parallel X_A)(x)$ when A on input x outputs y and X_A on the same input (including A 's coins) returns z . Knowledge soundness is defined via game $\text{knw-snd}_{\text{SNK}, R}^{A, X_A}$ in Fig. 14.

$\text{knw-snd}_{\text{SNK}, R}^{A, X_A}$ 00 $\text{crs} \stackrel{s}{\leftarrow} \text{Setup}(R)$ 01 $((\phi, \pi); w) \stackrel{s}{\leftarrow} (A \parallel X_A)(R, \text{crs})$ 02 Return $((\phi, w) \notin R$ $\quad \wedge \text{Vfy}(R, \text{crs}, \phi, \pi) = 1)$	$\mathbf{k-snd-aff}_{\text{NILP}, R}^{X, A}$ 03 $\vec{\sigma} \stackrel{s}{\leftarrow} \text{LinSetup}(R)$ 04 $(\phi, P) \stackrel{s}{\leftarrow} A(R)$ 05 $w \stackrel{s}{\leftarrow} X(R, \phi, P)$ 06 Return $(P \in \mathbb{F}^{\mu \times \kappa} \wedge (\phi, w) \notin R$ $\quad \wedge \text{LinVfy}(R, \vec{\sigma}, \phi, P\vec{\sigma}) = 1)$
---	--

Fig. 14. **Left:** Knowledge soundness game knw-snd relative to $\text{SNK} = (\text{Setup}, \text{Prv}, \text{Vfy})$, adversary A and extractor X_A . **Right:** Knowledge soundness game $\mathbf{k-snd-aff}$ relative to $\text{NILP} = (\text{LinSetup}, \text{PrfMtrx}, \text{Test})$, extractor X and *affine* adversary A (right).

Zero-knowledge for SNK requires that arguments do not leak any information besides the truth of the statement. It is formalized by demanding the existence of a simulator which on input a trapdoor (which is an additional output of **Setup**) and a true statement ϕ returns an argument which is indistinguishable from an argument for ϕ output by **Prv** (see [Gro16] for a formal definition).

A (preprocessing) *succinct argument of knowledge (SNARK)* is a knowledge-sound non-interactive argument system whose arguments are of size polynomial in the security parameter and can be verified in polynomial time in the security parameter and the length of the statement.

NON-INTERACTIVE LINEAR PROOFS OF DEGREE 2. NILPs (in Groth’s [Gro16] terminology) are an abstraction of many SNARK constructions introduced by Bitansky et al. [BCI+13]. We only consider NILPs of degree 2 here. Such a system NILP is defined by three algorithms as follows. On input a quadratic arithmetic program R , LinSetup returns $\vec{\sigma} \in \mathbb{F}^\mu$ for some μ . On input R , ϕ and w , algorithm PrfMtrx generates a matrix $P \in \mathbb{F}^{\kappa \times \mu}$ (where κ is the proof length). And on input R and ϕ , Test returns matrices $T_1, \dots, T_\eta \in \mathbb{F}^{\mu + \kappa}$. The last two algorithms implicitly define a prover and a verification algorithm as follows:

- $\vec{\pi} \stackrel{\$}{\leftarrow} \text{LinPrv}(R, \vec{\sigma}, \phi, w)$: run $P \stackrel{\$}{\leftarrow} \text{PrfMtrx}(R, \phi, w)$; return $\vec{\pi} := P\vec{\sigma}$.
- $b \stackrel{\$}{\leftarrow} \text{LinVfy}(R, \vec{\sigma}, \phi, \vec{\pi})$: $(T_1, \dots, T_\eta) \stackrel{\$}{\leftarrow} \text{Test}(R, \phi)$; return 1 iff for all $1 \leq k \leq \eta$:

$$(\vec{\sigma}^\top \mid \vec{\pi}^\top) T_k (\vec{\sigma}^\top \mid \vec{\pi}^\top)^\top = 0. \quad (4)$$

Let $T_k =: (t_{k,i,j})_{1 \leq i,j \leq \mu + \kappa}$; w.l.o.g. we assume that $t_{k,i,j} = 0$ for all k and $i \in \{\mu + 1, \dots, \mu + \kappa\}$ and $j \in \{1, \dots, \mu\}$.

We require a NILP to satisfy *statistical knowledge soundness against affine prover strategies*, which requires the existence of an (efficient) extractor X that works for all (unbounded) adversaries A . Whenever A returns a proof matrix P which leads to a valid proof $P\vec{\sigma}$ for a *freshly sampled* $\vec{\sigma}$, X can extract a witness from P . The notion is defined via game $\mathbf{k}\text{-snd}\text{-aff}_{\text{NILP}, R}^{X,A}$ in Fig. 14.

NON-INTERACTIVE ARGUMENTS FROM NILPs. From a NILP for a quadratic arithmetic program over a finite field $\mathbb{F} = \mathbb{Z}_p$ for some prime p , one can construct an argument system over a bilinear group $\mathcal{G} = (p, \mathbb{G}, g, e)$. We thus consider QAP relations R of the form

$$R = (\mathcal{G}, \mathbb{F} = \mathbb{Z}_p, \ell, \{u_i(X), v_i(X), w_i(X)\}_{i=0}^m, t(X)), \quad (5)$$

and define the *degree* of R as the degree of n of $t(X)$.

$\text{Setup}(R)$ 00 $g \stackrel{\$}{\leftarrow} \mathbb{G}$ 01 $\vec{\sigma} \stackrel{\$}{\leftarrow}$ $\text{LinSetup}(R)$ 02 Return $\vec{\Sigma} := \langle \vec{\sigma} \rangle$	$\text{Prv}(R, \vec{\Sigma}, \phi, w)$ 03 $P \stackrel{\$}{\leftarrow} \text{PrfMtrx}(R, \phi, w)$ 04 Parse $P = (p_{i,j})_{i,j}$ 05 For $i = 1 \dots \kappa$: 06 $\Pi_i := \prod_{j=1}^m \Sigma_j^{p_{i,j}}$ 07 $\pi := (\Pi_1, \dots, \Pi_\kappa)$ 08 Return π // Note that $\pi := \langle P\vec{\sigma} \rangle$	$\text{Vfy}(R, \vec{\Sigma}, \phi, \vec{\Pi})$ 09 $T_1, \dots, T_\eta \stackrel{\$}{\leftarrow} \text{Test}(R, \phi)$ 10 Return 1 iff for all $1 \leq \ell \leq \eta$: $0 = \prod_{i=1}^m \prod_{j=1}^m e(\Sigma_i, \Sigma_j)^{t_{\ell,i,j}}$ $\cdot \prod_{i=1}^m \prod_{j=1}^\kappa e(\Sigma_i, \Pi_j)^{t_{\ell,i,m+j}}$ $\cdot \prod_{i=1}^\kappa \prod_{j=1}^\kappa e(\Pi_i, \Pi_j)^{t_{\ell,m+i,m+j}}$ // This evaluates (6) in the exponent
--	--	--

Fig. 15. Argument system (Setup, Prv, Vfy) from a NILP (LinSetup, PrfMtrx, Test).

The construction of $\text{SNK} = (\text{Setup}, \text{Prv}, \text{Vfy})$ from $\text{NILP} = (\text{LinSetup}, \text{PrfMtrx}, \text{Test})$ is given in Fig. 15, where we write $\langle \vec{x} \rangle$ for $(g^{x_1}, \dots, g^{x_{|\vec{x}|}})$. **Setup** first samples a random group generator g and then embeds the NILP CRS “in the exponent”. Using group operations, **Prv** computes LinPrv in the exponent, and using the pairing, **Vfy** verifies LinVfy in the exponent.

GROTH’S NEAR-OPTIMAL SNARK FOR QAPS. Groth [Gro16] obtains his SNARK system by constructing a NILP for QAPs and then applying the conversion in Fig. 15. Recall that R , as in (5), defines a language of statements $\phi = (a_1, \dots, a_\ell) \in \mathbb{F}^\ell$ with witnesses of the form $w = (a_{\ell+1}, \dots, a_m) \in \mathbb{F}^{m-\ell}$ such that (with $a_0 := 1$):

$$\left(\sum_{i=0}^m a_i u_i(X) \right) \cdot \left(\sum_{i=0}^m a_i v_i(X) \right) = \sum_{i=0}^m a_i w_i(X) + h(X)t(X) \quad (6)$$

for some $h(X) \in \mathbb{F}[X]$ of degree at most $n-2$. Groth’s NILP is given in Fig. 16.

Theorem 6 ([Gro16, Theorem 1]). *The construction in Fig. 16 in a NILP with perfect completeness, perfect zero-knowledge and statistical knowledge soundness against affine prover strategies.*

Groth embeds his NILP in *asymmetric* bilinear groups, which yields a more efficient SNARK. He then shows that the scheme is knowledge-sound in the generic group model for *symmetric* bilinear groups (which yields a stronger result, as the adversary is more powerful than in asymmetric groups). Since we aim at strengthening the security statement, we also consider the symmetric-group variant (which is obtained by applying the transformation in Fig. 15). We now show how from an algebraic adversary breaking knowledge soundness one can construct an adversary against the q -DLog assumption.

<p>LinSetup(R)</p> <p>00 $\alpha, \beta, \gamma, \delta, \tau \xleftarrow{\\$} \mathbb{F}^*$</p> <p>01 $\vec{\sigma} := \left(1, \alpha, \beta, \gamma, \delta, \left\{ \tau^i \right\}_{i=0}^{n-1}, \left\{ \frac{1}{\gamma} (\beta u_i(\tau) + \alpha v_i(\tau) + w_i(\tau)) \right\}_{i=0}^{\ell}, \left\{ \frac{1}{\delta} (\beta u_i(\tau) + \alpha v_i(\tau) + w_i(\tau)) \right\}_{i=\ell+1}^m, \left\{ \frac{1}{\delta} (\tau^i t(\tau)) \right\}_{i=0}^{n-2} \right)$</p> <p>02 Return $\vec{\sigma}$</p>
<p>PrfMtrx(R, ϕ, w)</p> <p>03 Let $h(X)$ be as in (8)</p> <p>04 $r, s \xleftarrow{\\$} \mathbb{F}$</p> <p>05 Return $P \in \mathbb{F}^{3 \times (m+2n+4)}$ s.t. $P\vec{\sigma} = (A, B, C)$ with</p> <p>06 $A := \alpha + \sum_{i=0}^m a_i u_i(\tau) + r\delta$</p> <p>07 $B := \beta + \sum_{i=0}^m a_i v_i(\tau) + s\delta$</p> <p>08 $C := \frac{1}{\delta} \left(\sum_{i=\ell+1}^m a_i (\beta u_i(\tau) + \alpha v_i(\tau) + w_i(\tau)) + h(\tau)t(\tau) \right) + As + rB - rs\delta$</p>
<p>Test(R, ϕ)</p> <p>09 Return $T \in \mathbb{F}^{(m+2n+7) \times (m+2n+7)}$ corresponding to the test</p> <p style="text-align: center;">$A \cdot B = \alpha \cdot \beta + \sum_{i=0}^{\ell} a_i \frac{1}{\gamma} (\beta u_i(\tau) + \alpha v_i(\tau) + w_i(\tau)) \cdot \gamma + C \cdot \delta$</p>

Fig. 16. Groth’s NILP (LinSetup, PrfMtrx, Test).

Theorem 7. *Let SNK denote Groth’s [Gro16] SNARK for degree- n QAPs defined over (symmetric) bilinear group \mathcal{G} with $n^2 \leq (p-1)/8$. Then we have $q\text{-dlog} \xrightarrow{(3,1)}_{\text{alg}} \mathbf{knw}\text{-snd}_{SNK}$ with $q := 2n - 1$.*

Let us start with a proof overview. Consider an algebraic adversary A_{alg} against knowledge soundness that returns $(\phi, [\vec{\Pi}]_P)$ on input $(R, \langle \vec{\sigma} \rangle)$. Since A_{alg} is algebraic and its group-element inputs are $\langle \vec{\sigma} \rangle$, we have $\vec{\Pi} = \langle P\vec{\sigma} \rangle$ with $P \in \mathbb{F}^{3 \times \mu}$. By the definition of Vfy , the proof $\vec{\Pi}$ is valid iff $P\vec{\sigma}$ satisfies LinVfy . By Groth’s theorem (Theorem 6) there exists an extractor X , which on input P such that $P\vec{\sigma}$ satisfies LinVfy extracts a witness (see game $\mathbf{k}\text{-snd}\text{-aff}_{\text{NILP},R}$ in Fig. 14).

So it seems this extractor X should also work for A_{alg} (which returns P as required). However, X is only guaranteed to succeed if $P\vec{\sigma}$ verifies for a *randomly* sampled $\vec{\sigma}$, whereas for A_{alg} in $\mathbf{knw}\text{-snd}_{SNK,R}$ it suffices to return P so that $P\vec{\sigma}$ verifies for the specific $\vec{\sigma}$ for which it received $\langle \vec{\sigma} \rangle$. To prove knowledge soundness, it now suffices to show that an adversary can only output P which works for *all* choices of $\vec{\sigma}$ (from which X will then extract a witness).

In the generic group model this follows rather straight-forwardly, since the adversary has no information about the concrete $\vec{\sigma}$. In the AGM however, A is given $\langle \vec{\sigma} \rangle$. Examining the structure of a NILP CRS $\vec{\sigma}$ (Fig. 16), we see that its components are defined as multivariate (Laurent) polynomials evaluated at a random point $\vec{x} = (\alpha, \beta, \gamma, \delta, \tau)$.

Now what does it mean for A_{alg} to output a valid P ? By the definition of LinVfy via Test (cf. Eq. (4) with $\vec{\pi} := P\vec{\sigma}$), it means that A_{alg} found P such that

$$(\vec{\sigma}^\top | (P\vec{\sigma})^\top)^\top T (\vec{\sigma}^\top | (P\vec{\sigma})^\top)^\top = 0. \quad (7)$$

If we interpret the components of $\vec{\sigma}$ as polynomials over X_1, \dots, X_5 (corresponding to $\vec{x} = (\alpha, \beta, \gamma, \delta, \tau)$) then the left-hand side of (7) defines a polynomial $Q_P(\vec{X})$.

On the other hand, what does it mean that $P\vec{\sigma}$ verifies only for the specific $\vec{\sigma}$ from A_{alg} ’s input? It means that $Q_P \neq 0$, that is, Q_P is not the zero polynomial, since otherwise (7) would hold for *any* choice of \vec{x} , that is $P\vec{\sigma}'$ would verify for any $\vec{\sigma}'$.

We now bound the probability that A_{alg} behaves “badly”, that is, it returns a proof that only holds with respect to its specific CRS. To do so, we bound the probability that given $\langle \vec{\sigma} \rangle$, A_{alg} returns a nonzero polynomial Q_P which vanishes at \vec{x} , that is, the point that defines $\vec{\sigma}$. By factoring Q_P , we can then extract information about \vec{x} , which was only given as group elements $\langle \vec{\sigma} \rangle$. In particular, we embed a q -DLog instance into the CRS $\langle \vec{\sigma} \rangle$, for which we need q to be at least the maximum of the total degrees of the polynomials defining σ , which for Groth’s NILP is $2n - 1$. We then factor the polynomial to obtain its roots, which yields the challenge discrete logarithm.

Proof (of Theorem 7). Let R be a QAP of degree n (cf. (5)). Let $\text{NILP} = (\text{LinSetup}, \text{PrfMtrx}, \text{Test})$ denote Groth's NILP (Fig. 16). By Theorem 6 there exists an extractor X , which on input R , statement $\phi \in L_R$, and $P \in \mathbb{F}^{\mu \times \kappa}$ such that $\text{LinVfy}(R, \vec{\sigma}, \phi, P\vec{\sigma}) = 1$ for $\vec{\sigma} \stackrel{\$}{\leftarrow} \text{LinSetup}(R)$ returns a witness w with probability $\text{Adv}_{\text{NILP}, R, X, F}^{\text{k-snd-aff}}$ for any affine F .

Let SNK denote Groth's SNARK obtained from NILP via the transformation in Fig. 15 and let A_{alg} be an algebraic adversary in the game $\text{knw-snd}_{\text{SNK}, R}$. From X we construct an extractor X_A for A_{alg} in Fig. 17. Note that since A_{alg} is algebraic, we have $\vec{\Pi} = \langle P\vec{\sigma} \rangle$. We thus have

$$\text{Vfy}(R, \vec{\Sigma}, \phi, \vec{\Pi}) = \text{Vfy}(R, \vec{\Sigma}, \phi, \langle P\vec{\sigma} \rangle) = \text{LinVfy}(R, \vec{\sigma}, \phi, P\vec{\sigma}), \quad (8)$$

where the last equality follows from the definition of Vfy (Fig. 15). Game $\text{knw-snd}_{\text{SNK}, R}^{A_{\text{alg}}, X_A}$ is written out in Fig. 17 and our goal is to upperbound $\text{Adv}_{\text{SNK}, R, A_{\text{alg}}, X_A}^{\text{knw-snd}}$.

Consider the affine prover A' in Fig. 18 and $\text{k-snd-aff}_{\text{NILP}, R}^{X, A'}$, with the code of A' written out, also in Fig. 18.

$X_A(R, \langle \vec{\sigma} \rangle)$	$\text{knw-snd}_{\text{SNK}, R}^{A_{\text{alg}}, X_A}$
00 $(\phi, [\vec{\Pi}]_P) \stackrel{\$}{\leftarrow} A_{\text{alg}}(R, \langle \vec{\sigma} \rangle)$	03 $\vec{\sigma} \stackrel{\$}{\leftarrow} \text{LinSetup}(R)$
01 $w \stackrel{\$}{\leftarrow} X(R, \phi, P)$	04 $(\phi, [\vec{\Pi}]_P) \stackrel{\$}{\leftarrow} A_{\text{alg}}(R, \langle \vec{\sigma} \rangle)$
02 Return w .	05 $w \stackrel{\$}{\leftarrow} X(R, \phi, P)$
	06 Return $((\phi, w) \notin R \wedge \text{LinVfy}(R, \vec{\sigma}, \phi, P\vec{\sigma}) = 1)$

Fig. 17. Extractor X_A defined from X and A_{alg} (left) and knowledge soundness game knw-snd for a SNARK built from $\text{NILP} = (\text{LinSetup}, \text{PrfMtrx}, \text{Test})$, algebraic adversary A_{alg} and X_A (right).

$A'(R)$	$\text{k-snd-aff}_{\text{NILP}, R}^{X, A'}$
00 $\vec{\sigma} \stackrel{\$}{\leftarrow} \text{LinSetup}(R)$	03 $\vec{\rho} \stackrel{\$}{\leftarrow} \text{LinSetup}(R);$
01 $(\phi, [\vec{\Pi}]_P) \stackrel{\$}{\leftarrow} A_{\text{alg}}(R, \langle \vec{\sigma} \rangle)$	04 $\vec{\sigma} \stackrel{\$}{\leftarrow} \text{LinSetup}(R)$
02 Return (ϕ, P) .	05 $(\phi, [\vec{\Pi}]_P) \stackrel{\$}{\leftarrow} A_{\text{alg}}(R, \langle \vec{\sigma} \rangle)$
	06 $w \stackrel{\$}{\leftarrow} X(R, \phi, P)$
	07 If $((\phi, w) \notin R \wedge \text{LinVfy}(R, \vec{\sigma}, \phi, P\vec{\sigma}) = 1$
	$\wedge \text{LinVfy}(R, \vec{\rho}, \phi, P\vec{\rho}) = 0)$
	08 Then $\text{bad} := 1$
	09 Return $((\phi, w) \notin R \wedge \text{LinVfy}(R, \vec{\rho}, \phi, P\vec{\rho}) = 1)$

Fig. 18. Affine prover A' defined from A_{alg} (left) and game k-snd-aff for NILP, extractor X and A' (right).

Comparing the right-hand sides of Figs. 17 and 18, we see that **knw-snd** returns 1 whereas **k-snd-aff** returns 0 if LinVfy returns 0 for $P\vec{\rho}$ w.r.t. $\vec{\rho}$, but it returns 1 for $P\vec{\sigma}$ w.r.t. $\vec{\sigma}$. Let **bad** denote the event when this happens; formally defined as a flag in game **k-snd-aff** in Fig. 18. By definition, we have

$$\text{Adv}_{\text{SNK}, R, \text{Alg}, X_A}^{\text{knw-snd}} \leq \text{Adv}_{\text{NILP}, R, X, A'}^{\text{k-snd-aff}} + \Pr[\text{bad} = 1]. \quad (9)$$

In order to simplify our analysis, we first make a syntactical change to NILP by multiplying out all denominators, that is, we let LinSetup (cf. Fig. 16) return

$$\vec{\sigma} := \left(\delta\gamma, \alpha\delta\gamma, \beta\delta\gamma, \delta\gamma^2, \delta^2\gamma, \{\delta\gamma\tau^i\}_{i=0}^{n-1}, \{\delta(\beta u_i(\tau) + \alpha v_i(\tau) + w_i(\tau))\}_{i=0}^\ell, \right. \\ \left. \{\gamma(\beta u_i(\tau) + \alpha v_i(\tau) + w_i(\tau))\}_{i=\ell+1}^m, \{\gamma\tau^i t(\tau)\}_{i=0}^{n-2} \right). \quad (10)$$

Note that this does not affect the distribution of the SNARK CRS as running the modified LinSetup amounts to the same as choosing $g' \stackrel{\$}{\leftarrow} \mathbb{G}$ and running the original setup with $g := (g')^{\delta\gamma}$, which again is a uniformly random generator.

Observe that the components of LinSetup defined in (10) can be described via multivariate polynomials $S_i(\vec{x})$ of total degree at most $2n - 1$ with $\vec{x} := (\alpha, \beta, \gamma, \delta, \tau)$, and LinSetup can be defined as picking a random point $\vec{x} \stackrel{\$}{\leftarrow} (\mathbb{F}^*)^5$ and returning the evaluations $\sigma_i := S_i(\vec{x})$ of these polynomials.

Let T be as defined by Test in Fig. 16. By (4) we have

$$\text{LinVfy}(R, \vec{\sigma}, \phi, P\vec{\sigma}) = 1 \iff \vec{\sigma}^\top ((\text{Id} | P^\top) \cdot T \cdot (\text{Id} | P^\top)^\top) \vec{\sigma} = 0.$$

Let \vec{S} be the vector of polynomials defined by LinSetup . For a matrix P define the following multivariate polynomial

$$Q_P(\vec{X}) := (\vec{S}(\vec{X}))^\top ((\text{Id} | P^\top) \cdot T \cdot (\text{Id} | P^\top)^\top) \vec{S}(\vec{X}) \quad (11)$$

of degree at most $(2n - 1)^2$. Then for any $\vec{x} \in (\mathbb{F}^*)^5$ and $\vec{\sigma} := \vec{S}(\vec{x})$ we have

$$\text{LinVfy}(R, \vec{\sigma}, \phi, P\vec{\sigma}) = 1 \iff Q_P(\vec{x}) = 0. \quad (12)$$

Groth [Gro16] proves Theorem 6 by arguing that an affine prover without knowledge of $\vec{\sigma}$ can only succeed in the game **k-snd-aff**_{NILP,R} by making LinVfy return 1 on *every* $\vec{\sigma}$, or stated differently using (12), by returning P with $Q_P \equiv 0$. He then shows that from such P one can efficiently extract a witness.

The adversary's probability of succeeding despite $Q_P \not\equiv 0$ is bounded via the Schwartz-Zippel lemma: the total degree of Q_P is at most $d = (2n - 1)^2$ (using the modified $\vec{\sigma}$ from (10)). The probability that $Q_P(\vec{x}) = 0$ for a random $\vec{x} \stackrel{\$}{\leftarrow} (F^*)^5$ is thus bounded by $\frac{d}{p-1}$. This yields

$$\text{Adv}_{\text{NILP}, R, X, A'}^{\text{k-snd-aff}} \leq \frac{(2n-1)^2}{p-1}. \quad (13)$$

In order to bound $\mathbf{Adv}_{\text{SNK},R,\mathbf{A}_{\text{alg}},X_A}^{\text{knw-snd}}$ in (9) we will construct an adversary \mathbf{B}_{alg} such that

$$\Pr [\mathbf{bad} = 1] \leq \left(1 - \frac{(2n-1)^2}{p-1}\right) \cdot \mathbf{Adv}_{\mathcal{G},\mathbf{B}_{\text{alg}}}^{q\text{-dlog}} \quad \text{with } q = 2n - 1. \quad (14)$$

For \mathbf{bad} to be set to 1, \mathbf{A}_{alg} 's output P must be such that $Q_P \not\equiv 0$: otherwise, LinVfy returns 1 for *any* \vec{x} and in particular $\text{LinVfy}(R, \vec{\rho}, \phi, P\vec{\rho}) = 1$.

$\mathbf{bad} = 1$ implies thus that \mathbf{A}_{alg} on input $\langle \vec{\sigma} \rangle = \langle \vec{S}(\vec{x}) \rangle$ returns P such that

$$Q_P \not\equiv 0 \quad \text{and} \quad Q_P(\vec{x}) = 0. \quad (15)$$

We now use such \mathbf{A}_{alg} to construct an adversary \mathbf{B}_{alg} against q -DLog with $q := 2n - 1$.

Adversary $\mathbf{B}_{\text{alg}}(\langle z \rangle, \langle z^2 \rangle, \dots, \langle z^q \rangle)$: On input a q -DLog instance, \mathbf{B}_{alg} simulates $\mathbf{k-snd-aff}_{\text{NILP},R}^{X,A'}$ for \mathbf{A}_{alg} . It first picks a random value $\vec{y} \leftarrow (\mathbb{F}^*)^5$, (implicitly) sets $x_i := y_i z$, that is,

$$\alpha := y_1 z \quad \beta := y_2 z \quad \gamma := y_3 z \quad \delta := y_4 z \quad \tau := y_5 z$$

and generates a CRS $\langle \vec{\sigma} \rangle := \langle \vec{S}(\vec{x}) \rangle = \langle \vec{S}(\alpha, \beta, \gamma, \delta, \tau) \rangle$ as defined in (10). Since the total degree of the polynomials S_i defining $\vec{\sigma}$ is bounded by $2n - 1$ (the degree of the last component of $\vec{\sigma}$), \mathbf{B}_{alg} can compute $\vec{\sigma}$ from its q -DLog instance.

Next, \mathbf{B}_{alg} runs $(\phi, [\vec{\Pi}]_P) \stackrel{s}{\leftarrow} \mathbf{A}_{\text{alg}}(R, \langle \vec{\sigma} \rangle)$ and from P computes the multivariate polynomial $Q_P(\vec{X})$ as defined in (11). If $Q_P \equiv 0$ or $Q_P(\vec{x}) \neq 0$ (by (8) and (12) the latter is equivalent to $\text{Vfy}(R, \vec{\Sigma}, \phi, \vec{\Pi}) = 0$), then \mathbf{B}_{alg} aborts. (*) Otherwise \mathbf{B}_{alg} defines the univariate polynomial

$$Q'_P(X) := Q_P(y_1 X, \dots, y_5 X).$$

If $Q'_P \equiv 0$ then \mathbf{B}_{alg} aborts. (**)

Otherwise \mathbf{B}_{alg} factors Q'_P to obtain its roots (of which by (11) there are at most $(2n - 1)^2$), checks them against its DLog instance to determine whether z is among them, and if so, returns z .

First note that \mathbf{B}_{alg} perfectly simulates $\mathbf{k-snd-aff}_{\text{NILP},R}^{X,A'}$ for \mathbf{A}_{alg} . Let us analyze the probability that \mathbf{B}_{alg} finds the target z provided that $\mathbf{bad} = 1$. In this case \mathbf{B}_{alg} will not abort at (*).

Since $Q'_P(z) = Q_P(y_1 z, \dots, y_5 z) = Q_P(\vec{x})$, by (15) we have $Q'_P(z) = 0$. Thus if $Q'_P \not\equiv 0$ then \mathbf{B}_{alg} finds z by factoring Q'_P . It remains to argue that $Q'_P \not\equiv 0$. By (15) we have $Q_P \not\equiv 0$. By the Schwartz-Zippel lemma, the probability that for a random $\vec{y} \stackrel{s}{\leftarrow} (\mathbb{F}^*)^5$, we have $Q_P(\vec{y}) = 0$ is bounded by $\frac{d}{p-1}$ where d is the total degree of Q_P , which is bounded by $(2n - 1)^2$. Since $Q_P(\vec{y}) = Q'_P(1)$, we have $Q'_P \not\equiv 0$ with probability at least $1 - \frac{(2n-1)^2}{p-1}$. Since the choice of \vec{y} is perfectly hidden from the adversary's view this shows

$$\mathbf{Adv}_{\mathcal{G},\mathbf{B}_{\text{alg}}}^{q\text{-dlog}} \geq \left(1 - \frac{(2n-1)^2}{p-1}\right) \cdot \Pr [\mathbf{bad} = 1] \geq \frac{1}{2} \cdot \Pr [\mathbf{bad} = 1],$$

where the last inequality comes from $n^2 \leq (p-1)/8$. Putting this together with (13), we have shown that

$$\mathbf{Adv}_{\text{SNK}, R, \mathcal{A}_{\text{alg}}, \mathcal{X}_A}^{\text{knw-snd}} \leq \frac{q^2}{p-1} + 2 \cdot \mathbf{Adv}_{\mathcal{G}, \mathcal{B}_{\text{alg}}}^{q\text{-dlog}}.$$

Following the generic bound for Boneh and Boyen's SDH assumption [BB08], we may assume that $\mathbf{Adv}_{\mathcal{G}, \mathcal{B}_{\text{alg}}}^{q\text{-dlog}} \geq \frac{q^2}{p-1}$. The above equation thus implies

$$\mathbf{Adv}_{\text{SNK}, R, \mathcal{A}_{\text{alg}}, \mathcal{X}_A}^{\text{knw-snd}} \leq 3 \cdot \mathbf{Adv}_{\mathcal{G}, \mathcal{B}_{\text{alg}}}^{q\text{-dlog}},$$

which concludes the proof. \square

Corollary 6. *It is $(\frac{3t^2q+3q^3}{p}, t)$ -hard to break knowledge soundness of Groth's SNARK [Gro16] in the generic group model.*

The corollary follows from the generic $(\frac{t^2q+q^3}{p}, t)$ -hardness of $q\text{-dlog}$, which is derived analogously to the bound for Boneh and Boyen's SDH assumption [BB08].

We remark that the above result is not specific to Groth's SNARK; it applies to any SNARK built from a NILP whose setup evaluates multivariate polynomials on a random position. The maximal total degree of these polynomials determines the parameter q in the $q\text{-DLog}$ instance.

Acknowledgments. We thank Dan Brown for valuable comments and Pooya Farshim for discussions on polynomials. We also thank Helger Lipmaa for sharing with us his independent security proof for Groth's SNARK. The first author is supported by the French ANR EFTREC project (ANR-16-CE39-0002). The second author was supported in part by ERC Project ERCC (FP7/615074) and by DFG SPP 1736 Big Data. The third author was supported by ERC Project ERCC (FP7/615074).

References

- [ABM15] Abdalla, M., Benhamouda, F., MacKenzie, P.: Security of the J-PAKE password-authenticated key exchange protocol. In: 2015 IEEE Symposium on Security and Privacy, pp. 571–587. IEEE Computer Society Press, May 2015. 35
- [ABR01] Abdalla, M., Bellare, M., Rogaway, P.: The oracle Diffie-Hellman assumptions and an analysis of DHIES. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 143–158. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45353-9_12. 35, 45
- [ABS16] Ambrona, M., Barthe, G., Schmidt, B.: Automated unbounded analysis of cryptographic constructions in the generic group model. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 822–851. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_29. 39
- [ACdM05] Ateniese, G., Camenisch, J., de Medeiros, B.: Untraceable RFID tags via insubvertible encryption. In: Atluri, V., Meadows, C., Juels, A. (eds.) ACM CCS 2005, pp. 92–101. ACM Press, November 2005. 36

- [ACHdM05] Ateniese, G., Camenisch, J., Hohenberger, S., de Medeiros, B.: Practical group signatures without random oracles. Cryptology ePrint Archive, Report 2005/385 (2005). <http://eprint.iacr.org/2005/385>. 36
- [AGO11] Abe, M., Groth, J., Ohkubo, M.: Separating short structure-preserving signatures from non-interactive assumptions. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 628–646. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_34. 35
- [AM09] 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). https://doi.org/10.1007/978-3-642-01001-9_2. 34, 37
- [BB08] Boneh, D., Boyen, X.: Short signatures without random oracles and the SDH assumption in bilinear groups. *J. Cryptol.* **21**(2), 149–177 (2008). 58
- [BCI+13] Bitansky, N., Chiesa, A., Ishai, Y., Paneth, O., Ostrovsky, R.: Succinct non-interactive arguments via linear interactive proofs. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 315–333. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36594-2_18. 52
- [BCL04] Bangerter, E., Camenisch, J., Lysyanskaya, A.: A cryptographic framework for the controlled release of certified data. In: Security Protocols Workshop, pp. 20–24 (2004). 36
- [BCPR16] Bitansky, N., Canetti, R., Paneth, O., Rosen, A.: On the existence of extractable one-way functions. *SIAM J. Comput.* **45**(5), 1910–1952 (2016). 38
- [BCS05] Backes, M., Camenisch, J., Sommer, D.: Anonymous yet accountable access control. In: WPES, pp. 40–46 (2005). 36
- [BDZ03] Bao, F., Deng, R.H., Zhu, H.F.: Variations of Diffie-Hellman problem. In: Qing, S., Gollmann, D., Zhou, J. (eds.) ICICS 2003. LNCS, vol. 2836, pp. 301–312. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-39927-8_28. 43
- [BFF+14] Barthe, G., Fagerholm, E., Fiore, D., Mitchell, J.C., Scedrov, A., Schmidt, B.: Automated analysis of cryptographic assumptions in generic group models. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 95–112. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44371-2_6. 37, 39
- [BFW16] Bernhard, D., Fischlin, M., Warinschi, B.: On the hardness of proving CCA-security of signed ElGamal. In: Cheng, C.-M., Chung, K.-M., Persiano, G., Yang, B.-Y. (eds.) PKC 2016, Part I. LNCS, vol. 9614, pp. 47–69. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49384-7_3. 35
- [BG04] Brown, D.R.L., Gallant, R.P.: The static Diffie-Hellman problem. Cryptology ePrint Archive, Report 2004/306 (2004). <http://eprint.iacr.org/2004/306>. 37
- [BL96] Boneh, D., Lipton, R.J.: Algorithms for black-box fields and their application to cryptography (extended abstract). In: Kobitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 283–297. Springer, Heidelberg (1996). https://doi.org/10.1007/3-540-68697-5_22. 33
- [BLS04] Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil pairing. *J. Cryptol.* **17**(4), 297–319 (2004). 36, 48
- [BMV08] Bresson, E., Monnerat, J., Vergnaud, D.: Separation results on the “one-more” computational problems. In: Malkin, T. (ed.) CT-RSA 2008. LNCS, vol. 4964, pp. 71–87. Springer, Heidelberg (2008). 35

- [Boy08] Boyen, X.: The uber-assumption family (invited talk). In: Galbraith, S.D., Paterson, K.G. (eds.) Pairing 2008. LNCS, vol. 5209, pp. 39–56. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85538-5_3. 38, 39
- [BR93] Bellare, M., Rogaway, P.: Random oracles are practical: a paradigm for designing efficient protocols. In: Ashby, V. (ed.) ACM CCS 1993, pp. 62–73. ACM Press, November 1993. 41
- [BR04] Bellare, M., Rogaway, P.: Code-based game-playing proofs and the security of triple encryption. Cryptology ePrint Archive, Report 2004/331 (2004). <http://eprint.iacr.org/2004/331>. 39
- [BV98] Boneh, D., Venkatesan, R.: Breaking RSA may not be equivalent to factoring. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 59–71. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0054117>. 34, 35, 40
- [CGH98] Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited (preliminary version). In: 30th ACM STOC, pp. 209–218. ACM Press, May 1998. 38
- [Che06] Cheon, J.H.: Security analysis of the strong Diffie-Hellman problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 1–11. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_1. 37, 47
- [CHK+06] Camenisch, J., Hohenberger, S., Kohlweiss, M., Lysyanskaya, A., Meyerovich, M.: How to win the clonewars: efficient periodic n-times anonymous authentication. In: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) ACM CCS 2006, pp. 201–210. ACM Press, October/November 2006. 36
- [CHL05] Camenisch, J., Hohenberger, S., Lysyanskaya, A.: Compact e-cash. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 302–321. Springer, Heidelberg (2005). https://doi.org/10.1007/11426639_18. 36
- [CHP07] Camenisch, J., Hohenberger, S., Pedersen, M.Ø.: Batch verification of short signatures. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 246–263. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-72540-4_14. 36
- [CL04] Camenisch, J., Lysyanskaya, A.: Signature schemes and anonymous credentials from bilinear maps. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 56–72. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-28628-8_4. 35, 45
- [CM14] Chase, M., Meiklejohn, S.: Déjà Q: using dual systems to revisit q-type assumptions. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 622–639. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_34. 38
- [Cor02] Coron, J.-S.: Optimal security proofs for PSS and other signature schemes. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 272–287. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_18. 35, 36
- [Dam92] Damgård, I.: Towards practical public key systems secure against chosen ciphertext attacks. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 445–456. Springer, Heidelberg (1992). https://doi.org/10.1007/3-540-46766-1_36. 35, 38
- [Den02] 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). https://doi.org/10.1007/3-540-36178-2_6. 38

- [DH76] Diffie, W., Hellman, M.E.: New directions in cryptography. *IEEE Trans. Inf. Theory* **22**(6), 644–654 (1976). 35, 43
- [GBL08] Garg, S., Bhaskar, R., Lokam, S.V.: Improved bounds on security reductions for discrete log based signatures. In: Wagner, D. (ed.) *CRYPTO 2008*. LNCS, vol. 5157, pp. 93–107. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_6. 35
- [GG17] Ghadafi, E., Groth, J.: Towards a classification of non-interactive computational assumptions in cyclic groups. In: Takagi, T., Peyrin, T. (eds.) *ASIACRYPT 2017*. LNCS, vol. 10625, pp. 66–96. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70697-9_3. 38
- [GGPR13] Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct NIZKs without PCPs. In: Johansson, T., Nguyen, P.Q. (eds.) *EUROCRYPT 2013*. LNCS, vol. 7881, pp. 626–645. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_37. 51
- [Gro16] Groth, J.: On the size of pairing-based non-interactive arguments. In: Fischlin, M., Coron, J.-S. (eds.) *EUROCRYPT 2016, Part II*. LNCS, vol. 9666, pp. 305–326. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_11. 36, 50, 52, 53, 54, 56, 58
- [HP78] Hellman, M.E., Pohlig, S.C.: An improved algorithm for computing logarithms over $GF(p)$ and its cryptographic significance. *IEEE Trans. Inf. Theory* **24**(1), 106–110 (1978). 34
- [JR10] Jager, T., Rupp, A.: The semi-generic group model and applications to pairing-based cryptography. In: Abe, M. (ed.) *ASIACRYPT 2010*. LNCS, vol. 6477, pp. 539–556. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_31. 34, 38
- [JR15] Joux, A., Rojat, A.: Security ranking among assumptions within the *Uber Assumption* framework. In: Desmedt, Y. (ed.) *ISC 2013*. LNCS, vol. 7807, pp. 391–406. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-27659-5_28. 38
- [JS09] Jager, T., Schwenk, J.: On the analysis of cryptographic assumptions in the generic ring model. In: Matsui, M. (ed.) *ASIACRYPT 2009*. LNCS, vol. 5912, pp. 399–416. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10366-7_24. 37
- [Kil01] Kiltz, E.: A tool box of cryptographic functions related to the Diffie-Hellman function. In: Rangan, C.P., Ding, C. (eds.) *INDOCRYPT 2001*. LNCS, vol. 2247, pp. 339–349. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45311-3_32. 38
- [KK12] Kakvi, S.A., Kiltz, E.: Optimal security proofs for full domain hash, revisited. In: Pointcheval, D., Johansson, T. (eds.) *EUROCRYPT 2012*. LNCS, vol. 7237, pp. 537–553. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_32. 36
- [KMP16] Kiltz, E., Masny, D., Pan, J.: Optimal security proofs for signatures from identification schemes. In: Robshaw, M., Katz, J. (eds.) *CRYPTO 2016, Part II*. LNCS, vol. 9815, pp. 33–61. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53008-5_2. 35
- [KSW08] Katz, J., Sahai, A., Waters, B.: Predicate encryption supporting disjunctions, polynomial equations, and inner products. In: Smart, N.P. (ed.) *EUROCRYPT 2008*. LNCS, vol. 4965, pp. 146–162. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_9. 39

- [LR06] Leander, G., Rupp, A.: On the equivalence of RSA and factoring regarding generic ring algorithms. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 241–251. Springer, Heidelberg (2006). https://doi.org/10.1007/11935230_16. 34
- [LRSW99] Lysyanskaya, A., Rivest, R.L., Sahai, A., Wolf, S.: Pseudonym systems. In: Heys, H.M., Adams, C.M. (eds.) SAC 1999. LNCS, vol. 1758, pp. 184–199. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-46513-8_14. 35, 37, 45
- [Mau05] Maurer, U.M.: Abstract models of computation in cryptography (invited paper). In: Smart, N.P. (ed.) Cryptography and Coding 2005. LNCS, vol. 3796, pp. 1–12. Springer, Heidelberg (2005). https://doi.org/10.1007/11586821_1. 33, 37, 39, 42
- [MRV16] Morillo, P., Ràfols, C., Villar, J.L.: The kernel matrix Diffie-Hellman assumption. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016, Part I. LNCS, vol. 10031, pp. 729–758. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53887-6_27. 38
- [MW98] Maurer, U.M., Wolf, S.: Lower bounds on generic algorithms in groups. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 72–84. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0054118>. 33, 39
- [MW99] Maurer, U., Wolf, S.: The relationship between breaking the Diffie-Hellman protocol and computing discrete logarithms. SIAM J. Comput. **28**(5), 1689–1721 (1999). 43
- [Nec94] Nechaev, V.I.: Complexity of a determinate algorithm for the discrete logarithm. Math. Notes **55**(2), 165–172 (1994). 33
- [Pol78] Pollard, J.M.: Monte Carlo methods for index computation mod p . Math. Comput. **32**, 918–924 (1978). 34
- [PV05] Paillier, P., Vergnaud, D.: Discrete-log-based signatures may not be equivalent to discrete log. In: Roy, B.K. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 1–20. Springer, Heidelberg (2005). 34, 35, 40
- [Riv04] Rivest, R.L.: On the notion of pseudo-free groups. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 505–521. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24638-1_28. 34
- [RLB+08] Rupp, A., Leander, G., Bangerter, E., Dent, A.W., Sadeghi, A.-R.: Sufficient conditions for intractability over black-box groups: generic lower bounds for generalized DL and DH problems. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 489–505. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-89255-7_30. 39
- [Sho97] 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). https://doi.org/10.1007/3-540-69053-0_18. 33, 37, 38, 39, 42
- [Sho04] Shoup, V.: Sequences of games: a tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332 (2004). <http://eprint.iacr.org/2004/332>
- [SS01] Sadeghi, A.-R., Steiner, M.: Assumptions related to discrete logarithms: why subtleties make a real difference. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 244–261. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44987-6_16. 38
- [FKL17] Fuchsbauer, G., Kiltz, E., Loss, J.: The algebraic group model and its applications. Cryptology ePrint Archive, Report 2017/620 (2017). <http://eprint.iacr.org/2004/332>. 45, 47