

The Kernel Matrix Diffie-Hellman Assumption^{*}

Paz Morillo¹, Carla Ràfols², and Jorge L. Villar¹

¹ Universitat Politècnica de Catalunya, Spain
{paz.morillo,jorge.villar}@upc.edu

² Universitat Pompeu Fabra, Spain
carla.rafols@upf.edu

Abstract. We put forward a new family of computational assumptions, the Kernel Matrix Diffie-Hellman Assumption. Given some matrix \mathbf{A} sampled from some distribution \mathcal{D} , the kernel assumption says that it is hard to find “in the exponent” a nonzero vector in the kernel of \mathbf{A}^\top . This family is a natural computational analogue of the Matrix Decisional Diffie-Hellman Assumption (MDDH), proposed by Escala *et al.* As such it allows to extend the advantages of their algebraic framework to computational assumptions.

The k -Decisional Linear Assumption is an example of a family of decisional assumptions of strictly increasing hardness when k grows. We show that for any such family of MDDH assumptions, the corresponding Kernel assumptions are also strictly increasingly weaker. This requires ruling out the existence of some black-box reductions between flexible problems (*i.e.*, computational problems with a non unique solution).

Keywords: Matrix Assumptions, Computational Problems, Black-Box Reductions, Structure Preserving Cryptography

1 Introduction

It is commonly understood that cryptographic assumptions play a crucial role in the development of secure, efficient protocols with strong functionalities. For instance, upon referring to the rapid development of pairing-based cryptography, X. Boyen [8] says that “it has been supported, in no small part, by a dizzying array of tailor-made cryptographic assumptions”. Although this may be a reasonable price to pay for constructing new primitives or improve their efficiency, one should not lose sight of the ideal of using standard and simple assumptions. This is an important aspect of provable security. Indeed, Goldreich [16], for instance, cites “having clear definitions of one’s assumptions” as one of the three main ingredients of good cryptographic practice.

There are many aspects to this goal. Not only it is important to use clearly defined assumptions, but also to understand the relations between them: to

^{*} Work supported by the Spanish research project MTM2013-41426-R and by a Sofja Kovalevskaja Award of the Alexander von Humboldt Foundation and the German Federal Ministry for Education and Research.

see, for example, if two assumptions are equivalent or one is weaker than the other. Additionally, the definitions should allow to make accurate security claims. For instance, although technically it is correct to say that unforgeability of the Waters’ signature scheme [42] is implied by the DDH Assumption, defining the CDH Assumption allows to make a much more precise security claim.

A notable effort in reducing the “dizzying array” of cryptographic assumptions is the work of Escala *et al.* [11]. They put forward a new family of decisional assumptions in a prime order group \mathbb{G} , the *Matrix Diffie-Hellman* Assumption ($\mathcal{D}_{\ell,k}$ -MDDH). It says that, given some matrix $\mathbf{A} \in \mathbb{Z}_q^{\ell \times k}$ sampled from some distribution $\mathcal{D}_{\ell,k}$, it is hard to decide membership in $\text{Im } \mathbf{A}$, the subspace spanned by the columns of \mathbf{A} , in the exponent. Rather than as new assumption, it should be seen as an algebraic framework for decisional assumptions which includes as a special case the widely used k -Lin family.

This framework has some obvious conceptual advantages. For instance, it allows to explain all the members of the k -Lin assumption family (and also others, like the uniform assumption, appeared previously in [13, 14, 41]) as a single assumption and unify different constructions of the same primitive in the literature (e.g., the Naor-Reingold PRF [36] and the Lewko-Waters PRF [29] are special cases of the same construction instantiated with the 1-Lin and the 2-Lin Assumption, respectively). Another of its advantages is that it avoids arbitrary choices and instead points out to a trade-off between efficiency and security (a scheme based on any $\mathcal{D}_{\ell,k}$ -MDDH Assumption can be instantiated with many different assumptions, some leading to stronger security guarantees and others leading to more efficient schemes). But follow-up work has also illustrated other possibly less obvious advantages. For instance, Herold *et al.* [21] have used the Matrix Diffie-Hellman abstraction to extend the model of composite-order to prime-order transformation of Freeman [13] and to derive efficiency improvements which were proven to be impossible in the original model.³ We believe this illustrates that the benefits of conceptual clarity can translate into concrete improvements as well.

The security notions for cryptographic protocols can be classified mainly in hiding and unforgeability ones. The former typically appear in encryption schemes and commitments and the latter in signature schemes and soundness in zero-knowledge proofs. Although it is theoretically possible to base the hiding property on computational problems, most of the practical schemes achieve this notion either information theoretically or based on decisional assumptions, at least in the standard model. Likewise, unforgeability naturally comes from computational assumptions (typically implied by stronger, decisional assumptions). Thus, a natural question is if one can find a computational analogue of their MDDH Assumption which can be used in “unforgeability type” of security notions.

³ More specifically, we are referring to the lower bounds on the image size of a projecting bilinear map of [39] which were obtained in Freeman model [13]. The results of [21] by-passed this lower bounds allowing to save on pairing operations for projecting maps in prime order groups.

Most computational problems considered in the literature are search problems with a unique solution like the discrete logarithm or CDH. But unforgeability actually means the inability to produce one among many solutions to a given problem (*e.g.*, in many signature schemes or zero knowledge proofs). Thus, unforgeability is more naturally captured by a *flexible computational problem*, namely, a problem which admits several solutions⁴. This maybe explains why several new flexible assumptions have appeared recently when considering “unforgeability-type” security notions in structure-preserving cryptography [2]. Thus a useful computational analogue of the MDDH Assumption should not only consider problems with a unique solution but also flexible problems which can naturally capture this type of security notions.

1.1 Our Results

In the following $\mathcal{G} = (\mathbb{G}, q, \mathcal{P})$, being \mathbb{G} some group in additive notation of prime order q generated by \mathcal{P} , that is, the elements of \mathbb{G} are $\mathcal{Q} = a\mathcal{P}$ where $a \in \mathbb{Z}_q$. They will be denoted as $[a] := a\mathcal{P}$. This notation naturally extends to vectors and matrices as $[\mathbf{v}] = (v_1\mathcal{P}, \dots, v_n\mathcal{P})$ and $[\mathbf{A}] = (A_{ij}\mathcal{P})$.

Computational Matrix Assumptions. In our first attempt to design a computational analogue of the MDDH Assumption, we introduce the *Matrix Computational DH Assumption*, (MCDH) which says that, given a uniform vector $[\mathbf{v}] \in \mathbb{G}^k$ and some matrix $[\mathbf{A}]$, $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$ for $\ell > k$, it is hard to extend $[\mathbf{v}]$ to a vector in \mathbb{G}^ℓ in the image of $[\mathbf{A}]$, $\text{Im}[\mathbf{A}]$. Although this assumption is natural and is weaker than the MDDH one, we argue that it is equivalent to CDH.

We then propose the *Kernel Matrix DH Assumption* ($\mathcal{D}_{\ell,k}$ -KerMDH). This new flexible assumption states that, given some matrix $[\mathbf{A}]$, $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$ for some $\ell > k$, it is hard to find a vector $[\mathbf{v}] \in \mathbb{G}^\ell$ in the kernel of \mathbf{A}^\top . We observe that for some special instances of $\mathcal{D}_{\ell,k}$, this assumption has appeared in the literature in [2, 18, 19, 27, 32] under different names, like *Simultaneous Pairing*, *Simultaneous Double Pairing (SDP in the following)*, *Simultaneous Triple Pairing*, *1-Flexible CDH*, *1-Flexible Square CDH*. Thus, the new KerMDH Assumption allows us to organize and give a unified view on several useful assumptions. This suggests that the KerMDH Assumption (and not the MCDH one) is the right computational analogue of the MDDH framework. Indeed, for any matrix distribution the $\mathcal{D}_{\ell,k}$ -MDDH Assumption implies the corresponding $\mathcal{D}_{\ell,k}$ -KerMDH Assumption. As a unifying algebraic framework, it offers the advantages mentioned above: it highlights the algebraic structure of any construction based on it, and it allows writing many instantiations of a given scheme in a compact way.

The power of Kernel Assumptions. At Eurocrypt 2015, our KerMDH Assumptions were applied to design simpler QA-NIZK proofs of membership in

⁴ In the cryptographic literature we sometimes find the term “strong” as an alternative to “flexible”, like the Strong RSA or the Strong DDH.

linear spaces [26]. They have also been used to give more efficient constructions of structure preserving signatures [25], to generalize and simplify the results on quasi-adaptive aggregation of Groth-Sahai proofs [17] (given originally in [24]) and to construct a tightly secure QA-NIZK argument for linear subspaces with unbounded simulation soundness in [15]. The power of a KerMDH Assumption is that it allows to guarantee uniqueness. This has been used by Kiltz and Wee [26], for instance, to compile some secret key primitives to the public key setting. Indeed, Kiltz and Wee [26] modify a hash proof system (which is only designated verifier) to allow public verification (a QA-NIZK proof of membership). In a hash proof system for membership in some linear subspace of \mathbb{G}^n spanned by the columns of some matrix $[\mathbf{M}]$, the public information is $[\mathbf{M}^\top \mathbf{K}]$, for some secret matrix \mathbf{K} , and given the proof $[\boldsymbol{\pi}]$ that $[\mathbf{y}]$ is in the subspace, verification tests if $[\boldsymbol{\pi}] \stackrel{?}{=} [\mathbf{y}^\top \mathbf{K}]$.

The core argument to compile this to a public key primitive is that given $([\mathbf{A}], [\mathbf{KA}])$, $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$ and any pair $[\mathbf{y}], [\boldsymbol{\pi}]$, the previous test is equivalent to $e([\boldsymbol{\pi}^\top], [\mathbf{A}]) = e([\mathbf{y}^\top], [\mathbf{KA}])$, under the $\mathcal{D}_{\ell,k}$ -KerMDH Assumption. Indeed,

$$\begin{aligned} e([\boldsymbol{\pi}^\top], [\mathbf{A}]) = e([\mathbf{y}^\top], [\mathbf{KA}]) &\iff e([\boldsymbol{\pi}^\top - \mathbf{y}^\top \mathbf{K}], [\mathbf{A}]) = [\mathbf{0}] \stackrel{\mathcal{D}_{\ell,k}\text{-KerMDH}}{\implies} \\ &\implies [\boldsymbol{\pi}] = [\mathbf{y}^\top \mathbf{K}]. \end{aligned} \tag{1}$$

That is, although potentially there are many possible proofs which satisfy the public verification equation (left hand side of Equation (1)), the $\mathcal{D}_{\ell,k}$ -KerMDH Assumption guarantees that only one of them is efficiently computable, so verification gives the same guarantees as in the private key setting (right hand side of Equation (1)). This property is also used in a very similar way in [15] and also in the context of structure preserving signatures in [25]. In Section 5 we use it to argue that, of all the possible openings of a commitment, only one is efficiently computable, *i.e.* to prove computational soundness of a commitment scheme. Moreover, some previous works, notably in the design of structure preserving cryptographic primitives [1–3, 31], implicitly used this property for one specific KerMDH Assumption: the Simultaneous (Double) Pairing Assumption.

On the other hand, we have already discussed the importance of having a precise and clear language when talking about cryptographic assumptions. This justifies the introduction of a framework specific to computational assumptions, because one should properly refer to the assumption on which security is actually based, rather than just saying “security is based on an assumption weaker than $\mathcal{D}_{\ell,k}$ -MDDH”. A part from being imprecise, a problem with such a statement is that might lead to arbitrary, not optimal choices. For example, the signature scheme of [30] is based on the SDP Assumption but a slight modification of it can be based on the \mathcal{L}_2 -KerMDH Assumption. If the security guarantee is “the assumption is weaker than 2-Lin” then the modified scheme achieves shorter public key and more efficient verification with no loss in security. Further, the claim that security is based on the MDDH decisional assumptions when only computational ones are necessary might give the impression that a certain tradeoff is in place when this is not known to be the case. For instance, Jutla and Roy [24]

construct constant-size QA-NIZK arguments of membership in linear spaces under what they call the “Switching Lemma”, which is proven under a certain $\mathcal{D}_{k+1,k}$ -MDDH Assumption. However, a close look at the proof reveals that in fact it is based on the corresponding $\mathcal{D}_{k+1,k}$ -KerMDH Assumption⁵. For these assumptions, prior to our work, it was unclear whether the choice of larger k gives any additional guarantees.

Strictly Increasing Families of Kernel Assumptions. An important problem is that it is not clear whether there are increasingly weaker families of KerMDH Assumptions. That is, some decisional assumptions families parameterized by k like the k -Lin Assumption are known to be strictly increasingly weaker. The proof of increasing hardness is more or less immediate and the term *strictly* follows from the fact that every two $\mathcal{D}_{\ell,k}$ -MDDH and $\mathcal{D}_{\tilde{\ell},\tilde{k}}$ -MDDH problems with $\tilde{k} < k$ are separated by an oracle computing a k -linear map. For the computational case, increasing hardness is also not too difficult, but nothing is known about *strictly* increasing hardness (see Fig. 1). This means that, as opposed to the decisional case, prior to our work, for protocols based on KerMDH Assumptions there was no-known tradeoff between larger k (less efficiency) and security.

In this paper, we prove that the families of matrix distributions in [11], $\mathcal{U}_{\ell,k}$, \mathcal{L}_k , \mathcal{SC}_k , \mathcal{C}_k and \mathcal{RL}_k , as well as a new distribution we propose in Section 6, the *circulant* family $\mathcal{CI}_{k,d}$, define families of kernel problems with increasing hardness. For this we show a tight reduction from the smaller to the larger problems in each family. Our main result (Theorem 2) is to prove that the hardness of these problems is *strictly* increasing. For this, we prove that there is no black-box reduction from the larger to the smaller problems in the multilinear generic group model. These new results correspond to the dotted arrows in Fig. 1.




Fig. 1. Implication and separation results between Matrix Assumptions (dotted arrows correspond to the new results).

Having in mind that the computational problems we study in the paper are defined in a generic way, that is without specifying any particular group, the

⁵ To see this, note that in the proof of their “Switching Lemma” on which soundness is based, they use the output of the adversary to decide if $\mathbf{f} \stackrel{?}{\in} \text{Im } \mathbf{A}$, $\mathbf{A} \leftarrow \mathcal{RL}_k$, by checking whether $[\mathbf{f}]$ is orthogonal to the adversary’s output (equation (1), proof of Lemma 1, [24], full version), and where \mathcal{RL}_k is the matrix distribution of Section 2.3.

generic group approach arises naturally as the setting for the analysis of their hardness and reducibility relations. Otherwise, we would have to rely on specific properties of the representation of the elements of particular group families, not captured by the generic model.

The proof of Theorem 2 requires dealing with the notion of black-box reduction between flexible problems. A black-box reduction must work for any possible behavior of the oracle, but, contrary to the normal (unique answer) black-box reductions, here the oracle has to choose among the set of valid answers in every call. Ruling out the existence of a reduction implies that for any reduction there is an oracle behavior for which the reduction fails. This is specially subtle when dealing with multiple oracle calls. We think that the proof technique we introduce to deal with these issues can be considered as a contribution in itself and can potentially be used in future work.

Combining the black-box techniques and the generic group model is not new in the literature. For instance Dodis et al. [10] combine the black-box reductions and a generic model for the group \mathbb{Z}_n^* to show some uninstantiability results for FDH-RSA signatures.

Theorem 2 supports the intuition that there is a tradeoff between the size of the matrix — which typically results in less efficiency — and the hardness of the KerMDH Problems, and justifies the generalization of several protocols to different choices of k given in [17, 24–26].

Applications. The discussion of our results given so far should already highlight some of the advantages of using the new Kernel family of assumptions and the power of these new assumptions, which have already been used in compelling applications in follow-up work in [17, 25, 26]. To further illustrate the usefulness of the new framework, we apply it to the study of trapdoor commitments. First, we revisit the Pedersen commitment [38] to vectors of scalars and its extension to vectors of group elements of Abe *et al.* [2] in bilinear groups. We unify these two constructions and we generalize to commit vectors of elements at each level \mathbb{G}_r , for any $0 \leq r \leq m$ under the extension of KerMDH Assumptions to the ideal m -graded encodings setting. In particular, when $m = 2$ we recover in a single construction as a special case both the original Pedersen and Abe *et al.* commitments.

The (generalized) Pedersen commitment maps vectors in \mathbb{G}_r to vectors in \mathbb{G}_{r+1} , is perfectly hiding and computationally binding under any Kernel Assumption. In Section 5.2 we use it as a building block to construct a “group-to-group” commitment, which maps vectors in \mathbb{G}_r to vectors in the same group \mathbb{G}_r . These commitments were defined in [3] because they are a good match to Groth-Sahai proofs. In [3], two constructions were given, one in asymmetric and the other in symmetric bilinear groups. Both are optimal in terms of commitment size and number of verification equations. Rather surprisingly, we show that both constructions in [3] are special instances of our group-to-group commitment for some specific matrix distributions.

A New Family of MDDH Assumptions of Optimal Representation Size.

We also propose a new interesting family of Matrix distributions, the circulant matrix distribution, $\mathcal{CI}_{k,d}$, which defines new MDDH and KerMDH assumptions. This family generalizes the Symmetric Cascade Distribution (\mathcal{SC}_k) defined in [11] to matrices of size $\ell \times k$, $\ell = k+d > k+1$. We prove that it has optimal representation size d independent of k among all matrix distributions of the same size. The case $\ell > k+1$ typically arises when one considers commitments/encryption in which the message is a vector of group elements instead of a single group element and the representation size typically affects the size of the public parameters.

We prove the hardness of the $\mathcal{CI}_{k,d}$ -KerMDH Problem, by proving that the $\mathcal{CI}_{k,d}$ -MDDH Problem is generically hard in k -linear groups. Analyzing the hardness of a family of decisional problems (depending on a parameter k) can be rather involved, specially when an efficient k -linear map is supposed to exist. This is why in [11], the authors gave a practical criterion for generic hardness when $\ell = k+1$ in terms of irreducibility of some polynomials involved in the description of the problem. This criterion was used then to prove the generic hardness of several families of MDDH Problems. To analyze the generic hardness of the $\mathcal{CI}_{k,d}$ -MDDH Problem for any d , the techniques in [11] are not practical enough, and we need some extensions of these techniques for the case $\ell > k+1$, recently introduced in [20]. However, we could not avoid the explicit computation of a large (but well-structured) Gröbner basis of an ideal associated to the matrix distribution. The new assumption can be used to instantiate the commitment schemes of Section 5 with shorter public parameters and improved efficiency.

2 Preliminaries

For $\lambda \in \mathbb{N}$, we write 1^λ for the string of λ ones. For a set S , $s \leftarrow S$ denotes the process of sampling an element s from S uniformly at random. For an algorithm \mathcal{A} , we write $z \leftarrow \mathcal{A}(x, y, \dots)$ to indicate that \mathcal{A} is a (probabilistic) algorithm that outputs z on input (x, y, \dots) . For any two computational problems \mathbb{P}_1 and \mathbb{P}_2 we recall that $\mathbb{P}_1 \Rightarrow \mathbb{P}_2$ denotes the fact that \mathbb{P}_1 reduces to \mathbb{P}_2 , and then ‘ \mathbb{P}_1 is hard’ \Rightarrow ‘ \mathbb{P}_2 is hard’. Thus, we will use ‘ \Rightarrow ’ both for computational problems and for the corresponding hardness assumptions.

Let Gen denote a cyclic group instance generator, that is a probabilistic polynomial time (PPT) algorithm that on input 1^λ returns a description $\mathcal{G} = (\mathbb{G}, q, \mathcal{P})$ of a cyclic group \mathbb{G} of order q for a λ -bit prime q and a generator \mathcal{P} of \mathbb{G} . We use additive notation for \mathbb{G} and its elements are $a\mathcal{P}$, for $a \in \mathbb{Z}_q$ and will be denoted as $[a] := a\mathcal{P}$. The notation extends to vectors and matrices in the natural way as $[\mathbf{v}] = (v_1\mathcal{P}, \dots, v_n\mathcal{P})$ and $[\mathbf{A}] = (A_{ij}\mathcal{P})$. For a matrix $\mathbf{A} \in \mathbb{Z}_q^{\ell \times k}$, $\text{Im } \mathbf{A}$ denotes the subspace of \mathbb{Z}_q^ℓ spanned by the columns of \mathbf{A} . Thus, $\text{Im}[\mathbf{A}]$ is the corresponding subspace of \mathbb{G}^ℓ .

2.1 Multilinear Maps

In the case of groups with a bilinear map, or more generally with a k -linear map for $k \geq 2$, we consider a generator producing the tuple $(e_k, \mathbb{G}_1, \mathbb{G}_k, q, \mathcal{P}_1, \mathcal{P}_k)$,

where $\mathbb{G}_1, \mathbb{G}_k$ are cyclic groups of prime-order q , \mathcal{P}_i is a generator of \mathbb{G}_i and e_k is a non-degenerate efficiently computable k -linear map $e_k : \mathbb{G}_1^k \rightarrow \mathbb{G}_k$, such that $e_k(\mathcal{P}_1, \dots, \mathcal{P}_1) = \mathcal{P}_k$. We actually consider graded encodings which offer a richer structure. For any fixed $k \geq 1$, let MGen_k be a PPT algorithm that on input 1^λ returns a description of a graded encoding $\mathcal{MG}_k = (e, \mathbb{G}_1, \dots, \mathbb{G}_k, q, \mathcal{P}_1, \dots, \mathcal{P}_k)$, where $\mathbb{G}_1, \dots, \mathbb{G}_k$ are cyclic groups of prime-order q , \mathcal{P}_i is a generator of \mathbb{G}_i and e is a collection of non-degenerate efficiently computable bilinear maps $e_{i,j} : \mathbb{G}_i \times \mathbb{G}_j \rightarrow \mathbb{G}_{i+j}$, for $i+j \leq k$, such that $e(\mathcal{P}_i, \mathcal{P}_j) = \mathcal{P}_{i+j}$. For simplicity we will omit the subindexes of e when they become clear from the context. Sometimes \mathbb{G}_0 is used to refer to \mathbb{Z}_q . For group elements we use the following implicit notation: for all $i = 1, \dots, k$, $[a]_i := a\mathcal{P}_i$. The notation extends in a natural way to vectors and matrices and to linear algebra operations. We sometimes drop the index when referring to elements in \mathbb{G}_1 , *i.e.*, $[a] := [a]_1 = a\mathcal{P}_1$. In particular, it holds that $e([a]_i, [b]_j) = [ab]_{i+j}$.

Additionally, for the asymmetric case, let AGen_2 be a PPT algorithm that on input 1^λ returns a description of an asymmetric bilinear group $\mathcal{AG}_2 = (e, \mathbb{G}, \mathbb{H}, \mathbb{T}, q, \mathcal{P}, \mathcal{Q})$, where $\mathbb{G}, \mathbb{H}, \mathbb{T}$ are cyclic groups of prime-order q , \mathcal{P} is a generator of \mathbb{G} , \mathcal{Q} is a generator of \mathbb{H} and $e : \mathbb{G} \times \mathbb{H} \rightarrow \mathbb{T}$ is a non-degenerate, efficiently computable bilinear map. In this case we refer to group elements as: $[a]_G := a\mathcal{P}$, $[a]_H := a\mathcal{Q}$ and $[a]_T := ae(\mathcal{P}, \mathcal{Q})$.

2.2 A Generic Model For Groups With Graded Encodings

In this section we describe a (purely algebraic) generic model for the graded encodings functionality, in order to obtain meaningful results about the hardness and separations of computational problems. The model is an adaptation of Maurer's generic group model [33, 34] including the k -graded encodings, but in a completely algebraic formulation that follows the ideas in [5, 12, 20]. Since the k -graded encodings functionality implies the k -linear group functionality, the former gives more power to the adversaries or reductions working within the corresponding generic model. This in particular means that non-existential results proven in the richer k -graded encodings generic model also imply the same results in the k -linear group generic model. Therefore, in this paper we consider the former model. Due to the space limitations, we can only give a very succinct description of the model. See the full version of the paper [35] for a detailed and more formal description.

In a first approach we consider Maurer's model adapted to the graded encodings functionality, but still not phrased in a purely algebraic language. In this model, an algorithm \mathcal{A} does not deal with proper group elements in $[y]_a \in \mathbb{G}_a$, but only with labels (Y, a) , and it has access to an additional oracle internally performing the group operations, so that \mathcal{A} cannot benefit from the particular way the group elements are represented. Namely, on start all the group elements $[x_1]_{a_1}, \dots, [x_\alpha]_{a_\alpha}$ in the input intended for \mathcal{A} are replaced by the labels $(X_1, a_1), \dots, (X_\alpha, a_\alpha)$. Then, \mathcal{A} actually receives as input the set of labels, and possibly some other non-group elements (*i.e.*, that do not belong to any of the groups $\mathbb{G}_1, \dots, \mathbb{G}_k$), denoted as \tilde{x} , and considered as a bit string. For each group

\mathbb{G}_a two additional labels $(0, a), (1, a)$, corresponding to the neutral element and the generator, are implicitly given to \mathcal{A} . Then \mathcal{A} can adaptively make the following queries to an oracle implementing the k -graded encodings:

- **GroupOp** $((Y_1, a), (Y_2, a))$: group operation in \mathbb{G}_a for two previously issued labels in \mathbb{G}_a resulting in a new label (Y_3, a) in \mathbb{G}_a .
- **GroupInv** $((Y, a))$: similarly for group inversion in \mathbb{G}_a .
- **GroupPair** $((Y_1, a), (Y_2, b))$: bilinear map for two previously issued labels in \mathbb{G}_a and \mathbb{G}_b , $a + b \leq k$, resulting in a new label $(Y_3, a + b)$ in \mathbb{G}_{a+b} .
- **GroupEqTest** $((Y_1, a), (Y_2, a))$: test two previously issued labels in \mathbb{G}_a for equality of the corresponding group elements, resulting in a bit ($1 = \text{equality}$).

In addition, the oracle performs the actual computations with the group elements, and it uses them to answer the **GroupEqTest** queries. Every badly formed query (for instance, containing a label not previously issued by the oracle or as an input to \mathcal{A}) is answered with a special rejection symbol \perp . Following the usual step in generic group model proofs (see for instance [5, 11, 20]), we use polynomials as labels to group elements. Namely, labels in \mathbb{G}_a are polynomials in $\mathbb{Z}_q[\mathbf{X}]$, where the algebraic variables $\mathbf{X} = (X_1, \dots, X_\alpha)$ are just formal representations of the group elements in the input of \mathcal{A} . Now the oracle computes the new labels using the natural polynomial operations: **GroupOp** $((Y_1, a), (Y_2, a)) = (Y_1 + Y_2, a)$, **GroupInv** $((Y, a)) = (-Y, a)$ and **GroupPair** $((Y_1, a), (Y_2, b)) = (Y_1 Y_2, a + b)$. It is easy to see that for any valid label (Y, a) , $\deg Y \leq a$.⁶

The output of \mathcal{A} consists only of some labels $(Y_1, b_1), \dots, (Y_\beta, b_\beta)$ (given at some time by the oracle) corresponding to group elements $[y_1]_{b_1}, \dots, [y_\beta]_{b_\beta}$, along with some non-group elements, denoted as \tilde{y} . Therefore, for any fixed random tape of \mathcal{A} and any choice of the non-group elements \tilde{x} , there exist polynomials $Y_1, \dots, Y_\beta \in \mathbb{Z}_q[\mathbf{X}]$ of degrees upper bounded by b_1, \dots, b_β respectively, with coefficients known to \mathcal{A} . Notice that \mathcal{A} itself can predict all answers given by the oracle except for some **GroupEqTest** queries. In particular, some **GroupEqTest** queries trivially result in 1, due to the group structure (*e.g.*, **GroupOp** $((Y, a), \text{GroupInv}((Y, a)))$ is the same as $(0, a)$), or due to the (known) *a priori* constraints in the input group elements (*i.e.*, the definition of the problem instance given to \mathcal{A}). The answers to nontrivial **GroupEqTest** queries (*i.e.*, queries that cannot be trivially predicted by \mathcal{A}) are the only effective information \mathcal{A} can receive from the generic group oracle.

We now introduce a “purely algebraic” version of the generic model. For that, we need to assume that the distribution of \mathbf{x} can be sampled by evaluating a polynomial map f of constant degree at a random point.⁷ This is not an actual restriction in our context since all Matrix Diffie-Hellman problems fulfil this requirement. In the “purely algebraic” model we redefine the oracle **GroupEqTest** to answer 1 if and only if \mathcal{A} can itself predict the positive answer. Namely **GroupEqTest** $((Y_1, a), (Y_2, a)) = 1$ if and only if $Y_1 \circ f = Y_2 \circ f$ as polynomials

⁶ It clearly holds for the input group elements (since $\deg Y = 1$), and the inequality is preserved by **GroupOp**, **GroupInv** and **GroupPair**.

⁷ A formal definition of this notion is given in the full version of the paper.

over \mathbb{Z}_q . With this change the behavior of \mathcal{A} can only differ negligibly from the original,⁸ meaning that generic algorithms perform almost equally in Maurer’s model and its purely algebraic version. But now, any generic algorithm is just modelled by a set of polynomials. As we need to handle elements in different groups, we will use the shorter vector notation $[\mathbf{x}]_{\mathbf{a}} = ([x_1]_{a_1}, \dots, [x_\alpha]_{a_\alpha}) = (x_1\mathcal{P}_{a_1}, \dots, x_\alpha\mathcal{P}_{a_\alpha}) \in \mathbb{G}_{a_1} \times \dots \times \mathbb{G}_{a_\alpha}$. Note that the length of a vector of indices \mathbf{a} is denoted by a corresponding Greek letter α . We will also use a tilde to denote variables containing only non-group elements (*i.e.*, elements not in any of $\mathbb{G}_1, \dots, \mathbb{G}_k$).

Lemma 1. *Let \mathcal{A} be an algorithm in the (purely algebraic) generic multilinear group model. Let $([\mathbf{x}]_{\mathbf{a}}, \tilde{x})$ and $([\mathbf{y}]_{\mathbf{b}}, \tilde{y})$ respectively be the input and output of \mathcal{A} . Then, for every choice of \tilde{x} and any choice of the random tape of \mathcal{A} , there exist polynomials $Y_1, \dots, Y_\beta \in \mathbb{Z}_q[\mathbf{X}]$ of degree upper bounded by b_1, \dots, b_β such that $\mathbf{y} = \mathbf{Y}(\mathbf{x})$, for all possible $\mathbf{x} \in \mathbb{Z}_q^n$, where $\mathbf{Y} = (Y_1, \dots, Y_\beta)$. Moreover, \tilde{y} does not depend on \mathbf{x} .*

The proof of the lemma comes from the above discussion.

As usually, the proposed generic model reduces the analysis of the hardness of some problems to solving a merely algebraic problem related to polynomials. As an example, consider a computational problem \mathcal{P} which instances are entirely described by some group elements in the base group \mathbb{G}_1 , $[\mathbf{x}] \leftarrow \mathcal{P}.\text{InstGen}(1^\lambda)$, and its solutions are also described by some group elements $[\mathbf{y}]_{\mathbf{b}} \in \mathcal{P}.\text{Sol}([\mathbf{x}])$. We also assume that $\mathcal{P}.\text{InstGen}$ just samples \mathbf{x} by evaluating polynomial functions of constant degree at a random point. Then, \mathcal{P} is hard in the purely algebraic generic multilinear group model if and only if for all (randomized) polynomials $Y_1, \dots, Y_\beta \in \mathbb{Z}_q[\mathbf{X}]$ of degrees upper bounded by b_1, \dots, b_β respectively,

$$\Pr([\mathbf{y}]_{\mathbf{b}} \in \mathcal{P}.\text{Sol}([\mathbf{x}]) : [\mathbf{x}] \leftarrow \mathcal{P}.\text{InstGen}(1^\lambda), \mathbf{y} = \mathbf{Y}(\mathbf{x}) \in \text{negl}(\lambda))$$

where $\mathbf{Y} = (Y_1, \dots, Y_m)$ and the probability is computed with respect the random coins of the instance generator and the randomized polynomials.⁹ In a few words, this means that the set $\mathcal{P}.\text{Sol}([\mathbf{x}])$ cannot be hit by polynomials of the given degree evaluated at \mathbf{x} .

This model extends naturally to algorithms with oracle access (*e.g.*, black-box reductions) but only when the oracles fit well into the generic model. Let us consider the algorithm $\mathcal{A}^\mathcal{O}$, with oracle access to \mathcal{O} . A completely arbitrary oracle (specified in the plain model) could have access to the internal representation of the group elements, and then it could leak some information about the group

⁸ As a standard argument used in proofs in the generic group model, the difference between the original model and its purely algebraic reformulation amounts to a negligible probability, which is typically upper-bounded by using Schwartz-Zippel Lemma and the union bound, as shown for instance in [5, 12, 20].

⁹ We can similarly deal with problems with non-group elements both in the instance description and the solution, but this would require a more sophisticated formalization, in which both the polynomials and the non-group elements in the solution could depend on the non-group elements in the instance, but in an efficient way.

elements that is outside the generic group model. Thus, we will impose the very limiting constraint that the oracles are also “algebraic”, meaning that the oracle’s input/output behavior respects the one-wayness of the graded encodings, and it only performs polynomial operations on the input labels.

Definition 1. Let $([\mathbf{u}]_d, \tilde{u})$ and $([\mathbf{v}]_e, \tilde{v})$ respectively be a query to an oracle \mathcal{O} and its corresponding answer, where \tilde{u} and \tilde{v} contain the respective non-group elements. The oracle \mathcal{O} is called algebraic if for any choice of \tilde{u} there exist polynomials $V_1, \dots, V_\epsilon \in \mathbb{Z}_q[\mathbf{U}, \mathbf{R}]$, $\mathbf{R} = (R_1, \dots, R_\rho)$, of constant degree (in the security parameter) such that

- for the specific choice of \tilde{u} , $v_i = V_i(\mathbf{u}, \mathbf{r})$, $i = 1, \dots, \epsilon$, for all $\mathbf{u} \in \mathbb{Z}_q^\epsilon$ and $\mathbf{r} \in \mathbb{Z}_q^\rho$, where $\mathbf{r} = (r_1, \dots, r_\rho)$ are random parameters defined and uniformly sampled by the oracle,
- \tilde{v} does not depend on \mathbf{u}, \mathbf{r} (thus, \mathbf{r} can only have influence in the group elements in the answer),
- V_j does not depend on any U_i such that $e_j < d_i$ (in order to preserve the one-wayness of the graded encodings),

The parameters \mathbf{r} capture the behavior of an oracle solving a problem with many solutions (called here a “flexible” problem). They could be independent or not across different oracle calls, depending on whether the oracle is stateless or stateful. For technical reasons we consider only the stateless case with uniform sampling. Observe that the first two requirements in the definition mean that \mathbf{v} depends algebraically on \mathbf{u}, \mathbf{r} and no extra information about \mathbf{u}, \mathbf{r} can be leaked through \tilde{v} . Removing any of these requirements from the definition results in that a generic algorithm using such an oracle will no longer be algebraically generic. Also notice that after a call to an algebraic oracle, there is no guarantee that labels (Y, a) fulfil the bound $\deg Y \leq a$.

Although the notion of algebraic oracle looks very limiting (*e.g.*, it excludes a Discrete Logarithm oracle, as it destroys the one-wayness property of the graded encodings, but oracles solving CDH or the Bilinear Computational Diffie-Hellman problem fit well in the definition), it is general enough for our purposes. We will need the following generalization of Lemma 1:

Lemma 2. Let $\mathcal{A}^\mathcal{O}$ be an oracle algorithm in the (purely algebraic) generic multilinear group model, making a constant number of calls Q to an algebraic oracle \mathcal{O} . Let $([\mathbf{x}]_a, \tilde{x})$ and $([\mathbf{y}]_b, \tilde{y})$ respectively be the input and output of \mathcal{A} . Then, for every choice of \tilde{x} and the random tape, there exist polynomials of constant degree $Y_1, \dots, Y_\beta \in \mathbb{Z}_q[\mathbf{X}, \mathbf{R}_1, \dots, \mathbf{R}_Q]$, such that $\mathbf{y} = \mathbf{Y}(\mathbf{x}, \mathbf{r}_1, \dots, \mathbf{r}_Q)$, for all possible inputs, where $\mathbf{Y} = (Y_1, \dots, Y_\beta)$, and $\mathbf{r}_1, \dots, \mathbf{r}_Q$ are the parameters introduced in Definition 1 for the Q queries. Moreover, \tilde{y} does not depend on \mathbf{x} or $\mathbf{r}_1, \dots, \mathbf{r}_Q$.

The proof of this lemma is given in Appendix A.

2.3 The Matrix Decisional Diffie-Hellman Assumption

We recall here the definition of the decisional assumptions introduced in [11], which are the starting point of our flexible computational matrix problems.

Definition 2. [11], Let $\ell, k \in \mathbb{N}$ with $\ell > k$. We call $\mathcal{D}_{\ell,k}$ a matrix distribution if it outputs (in polynomial time, with overwhelming probability) matrices in $\mathbb{Z}_q^{\ell \times k}$ of full rank k . We denote $\mathcal{D}_k := \mathcal{D}_{k+1,k}$.

Definition 3 ($\mathcal{D}_{\ell,k}$ -MDDH Assumption). [11] Let $\mathcal{D}_{\ell,k}$ be a matrix distribution. The $\mathcal{D}_{\ell,k}$ -Matrix Diffie-Hellman ($\mathcal{D}_{\ell,k}$ -MDDH) Problem is telling apart the two probability distributions $(\mathbb{G}, q, \mathcal{P}, [\mathbf{A}], [\mathbf{A}\mathbf{w}])$ and $(\mathbb{G}, q, \mathcal{P}, [\mathbf{A}], [\mathbf{z}])$, where $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}, \mathbf{w} \leftarrow \mathbb{Z}_q^k, \mathbf{z} \leftarrow \mathbb{Z}_q^\ell$.

We say that the $\mathcal{D}_{\ell,k}$ -Matrix Diffie-Hellman ($\mathcal{D}_{\ell,k}$ -MDDH) Assumption holds relative to Gen if the corresponding problem is hard, that is, if for all PPT adversaries \mathcal{A} , the advantage

$$\text{Adv}_{\mathcal{D}_{\ell,k}, \text{Gen}}(\mathcal{A}) = \Pr[\mathcal{A}(\mathcal{G}, [\mathbf{A}], [\mathbf{A}\mathbf{w}]) = 1] - \Pr[\mathcal{A}(\mathcal{G}, [\mathbf{A}], [\mathbf{z}]) = 1] \in \text{negl}(\lambda),$$

where the probability is taken over $\mathcal{G} = (\mathbb{G}, q, \mathcal{P}) \leftarrow \text{Gen}(1^\lambda)$, $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}, \mathbf{w} \leftarrow \mathbb{Z}_q^k, \mathbf{z} \leftarrow \mathbb{Z}_q^\ell$ and the coin tosses of adversary \mathcal{A} .

In the case of asymmetric bilinear groups or symmetric k -linear groups, we similarly say that the $\mathcal{D}_{\ell,k}$ -MDDH Assumption holds relative to AGen_2 or MGen_k , respectively. In the former we specify if the assumption holds in the left (\mathcal{A} receives $[\mathbf{A}]_G, [\mathbf{A}\mathbf{w}]_G$ or $[\mathbf{z}]_G$), or in the right (\mathcal{A} receives $[\mathbf{A}]_H, [\mathbf{A}\mathbf{w}]_H$ or $[\mathbf{z}]_H$).

Definition 4. A matrix distribution $\mathcal{D}_{\ell,k}$ is hard if the corresponding $\mathcal{D}_{\ell,k}$ -MDDH problem is hard in the generic k -linear group model.

Many different matrix distributions appear in the literature. Namely, the cascade \mathcal{C}_k and symmetric cascade \mathcal{SC}_k distributions were presented in [11], while the uniform $\mathcal{U}_{\ell,k}$, the linear \mathcal{L}_k , the randomized linear \mathcal{RL}_k and the square polynomial $\mathcal{P}_{\ell,2}$ distributions were implicitly used in some previous works. We give their explicit definitions in Appendix B.

3 The Matrix Diffie-Hellman Computational Problems

In this section we introduce two families of search problems naturally related to the Matrix Decisional Diffie-Hellman problems. In the first family, given a matrix $[\mathbf{A}]$, where $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$, and the first k components of a vector $[\mathbf{z}]$, the problem is completing it so that $\mathbf{z} \in \text{Im } \mathbf{A}$.

Definition 5 ($\mathcal{D}_{\ell,k}$ -MCDH). Given a matrix distribution $\mathcal{D}_{\ell,k}$, such that the upper $k \times k$ submatrix of $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$ has full rank with overwhelming probability, the computational matrix Diffie-Hellman Problem is given $([\mathbf{A}], [\mathbf{z}_0])$, with $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}, \mathbf{z}_0 \leftarrow \mathbb{Z}_q^k$, compute $[\mathbf{z}_1] \in \mathbb{G}^{\ell-k}$ such that $(\mathbf{z}_0 \| \mathbf{z}_1) \in \text{Im } \mathbf{A}$.

The full-rank condition ensures the existence of solutions to the $\mathcal{D}_{\ell,k}$ -MCDH problem instance. Thus, we tolerate the existence of a negligible fraction of unsolvable problem instances. Indeed, all known interesting matrix distributions fulfil this requirement. Notice that CDH and the computational k -Lin problems

are particular examples of MCDH problems. Namely, CDH is exactly \mathcal{L}_1 -MCDH and the computational k -Lin problem is \mathcal{L}_k -MCDH. Indeed, the \mathcal{L}_1 -MCDH problem is given $[1], [a], [z_1]$, compute $[z_2]$ such that (z_1, z_2) is collinear with $(1, a)$, or equivalently, $z_2 = z_1 a$, which is solving the CDH problem. All MCDH problems have a unique solution and they appear naturally in some scenarios using MDDH problems. For instance, the one-wayness of the encryption scheme in [11] is equivalent to the corresponding MCDH assumption.

There is an immediate relation between any MCDH problem and its decisional counterpart. Not surprisingly, for any matrix distribution $\mathcal{D}_{\ell,k}$, $\mathcal{D}_{\ell,k}$ -MDDH \Rightarrow $\mathcal{D}_{\ell,k}$ -MCDH.

We are not going to study the possible reductions between MCDH problems, due to the fact that, essentially, any MCDH problem amounts to computing some polynomial on the elements of \mathbf{A} , and it is then equivalent to CDH ([4, 23]), although the tightness of the reduction depends on the degree of the polynomial.

In the second family of computational problems, given a matrix $[\mathbf{A}]$, where $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$, the problem is finding $[\mathbf{x}]$ such that $\mathbf{x} \in \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}$. It is notable that some computational problems in the literature are particular cases of this second family.

Definition 6 ($\mathcal{D}_{\ell,k}$ -KerMDH). *Given a matrix distribution $\mathcal{D}_{\ell,k}$, the Kernel Diffie-Hellman Problem is given $[\mathbf{A}]$, with $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$, find a nonzero vector $[\mathbf{x}] \in \mathbb{G}^\ell$ such that \mathbf{x} is orthogonal to $\text{Im } \mathbf{A}$, that is, $\mathbf{x} \in \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}$.*

Definition 6 naturally extends to asymmetric bilinear groups. There, given $[\mathbf{A}]_H$, the problem is to find $[\mathbf{x}]_G$ such that $\mathbf{x} \in \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}$. A solution can be obviously verified by checking if $e([\mathbf{x}]_G, [\mathbf{A}]_H) = [\mathbf{0}]_T$. We can also consider an extension of this problem in which the goal is to solve the same problem but giving the solution in a different group \mathbb{G}_r , in some ideal graded encoding \mathcal{MG}_m , for some $0 \leq r \leq \min(m, k-1)$. The case $r = 1$ corresponds to the previous problem defined in a m -linear group.

Definition 7 ($(r, m, \mathcal{D}_{\ell,k})$ -KerMDH). *Given a matrix distribution $\mathcal{D}_{\ell,k}$ over a m -linear group \mathcal{MG}_m and r an integer $0 \leq r \leq \min(m, k-1)$, the $(r, m, \mathcal{D}_{\ell,k})$ -KerMDH Problem is to find $[\mathbf{x}]_r \in \mathbb{G}_r^\ell$ such that $\mathbf{x} \in \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}$.*

When the precise degree of multilinearity m is not an issue, we will write $(r, \mathcal{D}_{\ell,k})$ -KerMDH instead of $(r, m, \mathcal{D}_{\ell,k})$ -KerMDH, for any $m \geq r$. We excluded the case $r \geq k$ because the problem is easy.

Lemma 3. *For all integers $k \leq r \leq m$ and for all matrix distributions $\mathcal{D}_{\ell,k}$, the $(r, m, \mathcal{D}_{\ell,k})$ -KerMDH Problem is easy.*

The kernel problem is also harder than the corresponding decisional problem, in multilinear groups.

Lemma 4. *In a m -linear group, $\mathcal{D}_{\ell,k}$ -MDDH \Rightarrow $(r, m, \mathcal{D}_{\ell,k})$ -KerMDH for any matrix distribution $\mathcal{D}_{\ell,k}$ and for any $0 \leq r \leq m-1$. In particular, for $m \geq 2$, $\mathcal{D}_{\ell,k}$ -MDDH \Rightarrow $\mathcal{D}_{\ell,k}$ -KerMDH.*

The proofs of Lemmas 3, and 4 can be found in the full version of this paper [35].

3.1 The Kernel DH Assumptions in the Multilinear Maps Candidates

We have shown that for any hard matrix distribution $\mathcal{D}_{\ell,k}$ the $\mathcal{D}_{\ell,k}$ -KerMDH problem is generically hard in m -linear groups. We emphasize that all our results refer to generic, ideal multilinear maps (in fact, to graded encodings, which have more functionality). Our aim is only to give necessary condition for the assumptions to hold in candidate multilinear maps. The status of current candidate multilinear maps is rather uncertain, e.g. it is described in [28] as “break-and-repair mode”. Thus, it is hard to argue if our assumptions hold in any concrete instantiation and we leave this as an open question for further investigation.

3.2 A Unifying View on Computational Matrix Problems

In this section we show how some computational problems in the cryptographic literature are unified as particular instances of KerMDH problems. Their explicit definitions are given in Appendix C. It is straightforward to see that Find-Rep [9] Assumption is just $(0, \mathcal{U}_{\ell,1})$ -KerMDH, the Simultaneous Double Pairing Assumption (SDP) [2] is \mathcal{RL}_2 -KerMDH, the Simultaneous Triple Pairing [18] Assumption is \mathcal{U}_2 -KerMDH, the Simultaneous Pairing [19] Assumption is $\mathcal{P}_{\ell,2}$ -KerMDH. The Double Pairing (DP) [18] Assumption corresponds to \mathcal{U}_1 -KerMDH in an asymmetric bilinear setting. On the other hand, the 1-Flexible Diffie-Hellman (1-FlexDH) [32] Assumption is \mathcal{C}_2 -KerMDH, the 1-Flexible Square Diffie-Hellman (1-FlexSDH) [27] Assumption is \mathcal{SC}_2 -KerMDH, and the ℓ -Flexible Diffie-Hellman (ℓ -FlexDH) [32] Assumption for $\ell > 1$ is the only one which is not in the KerMDH family. However, ℓ -FlexDH \Rightarrow $\mathcal{C}_{\ell+1}$ -KerMDH. Getting the last three results requires a bit more work, and they are proven in the full version [35].

4 Reduction and Separation of Kernel Diffie-Hellman Problems

In this section we prove that the most important matrix distribution families $\mathcal{U}_{\ell,k}$, \mathcal{L}_k , \mathcal{SC}_k , \mathcal{C}_k and \mathcal{RL}_k (see Appendix B) define families of KerMDH problems with **strictly** increasing hardness, as we precisely state in Theorem 2, at the end of the section. By ‘strictly increasing’ we mean that (1) there are known reductions of the smaller problems to the larger problems (in terms of k) within each family, and (2) there are no black-box reductions in the other way in the multilinear generic group model. This result shows the necessity of using $\mathcal{D}_{\ell,k}$ -KerMDH Assumptions for $k > 2$. A similar result is known for the corresponding $\mathcal{D}_{\ell,k}$ -MDDH problems. Indeed, one can easily prove a separation between large and small decisional problems. Observe that any efficient m -linear map can efficiently solve any $\mathcal{D}_{\ell,k}$ -MDDH problem with $k \leq m - 1$, and therefore every two $\mathcal{D}_{\ell,k}$ -MDDH and $\mathcal{D}_{\ell,\tilde{k}}$ -MDDH problems with $\tilde{k} < k$ are separated by an oracle computing a k -linear map. However, when dealing with the computational $\mathcal{D}_{\ell,k}$ -KerMDH family, no such a trivial argument is known to exist. Actually, a

m -linear map does not seem to help to solve any $\mathcal{D}_{\ell,k}$ -KerMDH problem with $k > 1$. Furthermore, the m -linear map seems to be useless for any (reasonable) reduction between KerMDH problems defined over the same group. Indeed, all group elements involved in the problem instances and their solutions belong to the base group \mathbb{G} , and the result of computing any m -linear map is an element in \mathbb{G}_m , where no efficient map from \mathbb{G}_m back to \mathbb{G} is supposed to exist.

4.1 Separation

In this section we firstly show the non-existential part of Theorem 2. Namely, we show that there is no black-box reduction in the generic group model (described in Section 2.2) from $\mathcal{D}_{\ell,k}$ -KerMDH to $\mathcal{D}_{\tilde{\ell},\tilde{k}}$ -KerMDH for $k > \tilde{k}$, assuming that the two matrix distributions $\mathcal{D}_{\ell,k}$ and $\mathcal{D}_{\tilde{\ell},\tilde{k}}$ are hard (see Definition 4). Before proving the main result we need some technical lemmas and also a new geometrical notion defined on a family of subspaces of a vector space, named *t-Elusiveness*.

In the first lemma we show that the natural (black-box, algebraic) reductions between KerMDH problems have a very special form. Observe that a black-box reduction to a flexible problem must work for any adversary solving it. In particular, the reduction should work for **any** solution given by this adversary, or for **any** probability distribution of the solutions given by it. Informally, the lemma states that the output of a successful reduction can always be computed in essentially two ways: (1) By just applying a (randomized) linear map to the answer given by the adversary in the last call. Therefore, all possibly existing previous calls to the adversary are just used to prepare the last one. (2) By just ignoring the last call to the adversary and using only the information gathered in the previous ones.




Fig. 2. Splitting of the black-box reduction.

Let $\mathcal{R}^{\mathcal{O}}$ be a black-box reduction of $\mathcal{D}_{\ell,k}$ -KerMDH to $\mathcal{D}_{\tilde{\ell},\tilde{k}}$ -KerMDH, in the purely algebraic generic multilinear group model, discussed in Section 2.2, for some matrix distributions $\mathcal{D}_{\ell,k}$ and $\mathcal{D}_{\tilde{\ell},\tilde{k}}$. Namely, $\mathcal{R}^{\mathcal{O}}$ solves $\mathcal{D}_{\ell,k}$ -KerMDH with

a non-negligible probability by making $Q \geq 1$ queries to an oracle \mathcal{O} solving $\mathcal{D}_{\tilde{\ell}, \tilde{k}}$ -KerMDH with probability one. As we aim at ruling out the existence of some reductions, we just consider the best possible case any black-box reduction must be able to handle. Now we split the reduction as $\mathcal{R}^{\mathcal{O}} = (\mathcal{R}_0^{\mathcal{O}}, \mathcal{R}_1)$, where the splitting point is the last oracle call, as shown in Figure 2. We actually use the same splitting in the proof of Lemma 2 in Appendix D. More formally, on the input of $[\mathbf{A}]$, for $\mathbf{A} \leftarrow \mathcal{D}_{\ell, k}$, and after making $Q - 1$ oracle calls, $\mathcal{R}_0^{\mathcal{O}}$ stops by outputting the last query to \mathcal{O} , that is a matrix $[\tilde{\mathbf{A}}]$, where $\tilde{\mathbf{A}} \in \mathcal{D}_{\tilde{\ell}, \tilde{k}}$, together with some state information s for \mathcal{R}_1 . Next, \mathcal{R}_1 resumes the execution from s and the answer $[\mathbf{w}] \in \mathbb{G}^{\tilde{\ell}}$ given by the oracle, and finally outputs $[\mathbf{v}] \in \mathbb{G}^{\ell}$. Without loss of generality, we assume that both stages $\mathcal{R}_0^{\mathcal{O}}$ and \mathcal{R}_1 receive the same random tape, $\$$ (\mathcal{R}_1 can redo the computations performed by $\mathcal{R}_0^{\mathcal{O}}$).

Lemma 5. *There exists an algebraic oracle \mathcal{O} (in the sense of Definition 1), that solves the $\mathcal{D}_{\ell, k}$ -KerMDH Problem with probability one.*

All the proofs in Section 4 are given in Appendix D.

Lemma 2 applied to $\mathcal{R}_0^{\mathcal{O}}$ (and using also Lemma 5) implies that only the group elements in s can depend on \mathbf{A} . Indeed, the non-group elements in s can only depend on the random tape $\$$. Now, from Lemma 1 applied to \mathcal{R}_1 , we know that its output $[\mathbf{v}]$ is determined by a polynomial map of total degree at most one in the input group elements (*i.e.*, $\tilde{\mathbf{A}}$ and the group elements in s), and the coefficients of this polynomial can only depend on $\$$, and the non-group elements in s , which in turn only depend on $\$$. Therefore, splitting the polynomial map into two parts, for every fixed $\$$ and every fixed oracle behavior in the first $Q - 1$ oracle calls there exists a vector $\mathbf{u} \in \mathbb{Z}_q^{\ell}$ and a linear map $\eta : \mathbb{Z}_q^{\tilde{\ell}} \rightarrow \mathbb{Z}_q^{\ell}$ such that we can write $\mathbf{v} = \mathbf{u} + \eta(\mathbf{w})$, where \mathbf{u} actually depends on the group elements in s . The important fact here is that η can only depend on $\$$, but not on \mathbf{A} .

Lemma 6. *Let $\mathcal{R}^{\mathcal{O}} = (\mathcal{R}_0^{\mathcal{O}}, \mathcal{R}_1)$ be a black-box reduction from $\mathcal{D}_{\ell, k}$ -KerMDH to $\mathcal{D}_{\tilde{\ell}, \tilde{k}}$ -KerMDH, in the purely algebraic generic multilinear group model, making $Q \geq 1$ calls to an oracle \mathcal{O} solving the latter with probability one. If $\mathcal{R}^{\mathcal{O}}$ succeeds with a non negligible probability ε then, for every possible behavior of the oracle, either $\Pr(\eta(\mathbf{w}) \in S') > \text{negl}$ or $\Pr(\mathbf{u} \in S') > \text{negl}$, where $S' = \ker \mathbf{A}^{\top} \setminus \{\mathbf{0}\}$, $[\mathbf{A}]$ is the input of $\mathcal{R}^{\mathcal{O}}$, and its output is written as $[\mathbf{u} + \eta(\mathbf{w})]$, for some \mathbf{u} only depending on the state output by $\mathcal{R}_0^{\mathcal{O}}$, $[\mathbf{w}]$ is the answer to the Q -th oracle query, and $\eta : \mathbb{Z}_q^{\tilde{\ell}} \rightarrow \mathbb{Z}_q^{\ell}$ is a (randomized) linear map that only depends on the random tape of $\mathcal{R}^{\mathcal{O}}$.*

The following property of the hard matrix distributions allows us to prove that indeed in the last lemma $\Pr(\eta(\mathbf{w}) \in S \setminus \{\mathbf{0}\}) \in \text{negl}$.

Definition 8 (t -Elusiveness). *A family of subspaces \mathcal{S} of a vector space X over the finite field \mathbb{Z}_q is called t -elusive for some $t < \dim X$ if for all t -dimensional subspaces $F \subset X$, $\Pr(F \cap S \neq \{\mathbf{0}\}) \in \text{negl}$, where the probability is computed with respect to the choice of $S \in \mathcal{S}$. A matrix distribution $\mathcal{D}_{\ell, k}$ is called t -elusive if the family $\{\ker \mathbf{A}^{\top}\}_{\mathbf{A} \in \mathcal{D}_{\ell, k}}$ is t -elusive.*

Lemma 7. *If a matrix distribution $\mathcal{D}_{\ell,k}$ is hard (as given in Definition 4) then $\mathcal{D}_{\ell,k}$ is k -elusive.*

In the next theorem we use the k -elusiveness to prove that $\Pr(\mathbf{u} \in \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}) > \text{negl}$ for all possible behaviors of the oracle in the first $Q - 1$ calls. This actually implies that the reduction can directly output \mathbf{u} , and only $Q - 1$ oracle calls are actually needed. Therefore, by the descent method we show that no successful reduction exists unless $\mathcal{D}_{\ell,k}$ -KerMDH is easy.

Theorem 1. *Let $\mathcal{D}_{\ell,k}$ be k -elusive. If there exists a black-box reduction in the purely algebraic generic multilinear group model from $\mathcal{D}_{\ell,k}$ -KerMDH to another problem $\mathcal{D}_{\tilde{\ell},\tilde{k}}$ -KerMDH with $\tilde{k} < k$, then $\mathcal{D}_{\ell,k}$ -KerMDH is easy.*

Now we consider the contrapositive statement, that directly applies to the known families of hard matrix distributions.

Corollary 1. *If a matrix distribution family $\{\mathcal{D}_{\ell,k}\}$ is hard then for any $\mathcal{D}_{\ell,k}$ and $\mathcal{D}_{\tilde{\ell},\tilde{k}}$ in the family with $k > \tilde{k}$ there is no black-box reduction in the generic group model from $\mathcal{D}_{\ell,k}$ -KerMDH to $\mathcal{D}_{\tilde{\ell},\tilde{k}}$ -KerMDH.*

Proof. Since all $\mathcal{D}_{\ell,k}$ -MDDH problems in the family are generically hard on a k -linear group, we know that $\mathcal{D}_{\ell,k}$ is k -elusive by Lemma 7, and also $\mathcal{D}_{\ell,k}$ -KerMDH is hard in that group (otherwise, any solution to $\mathcal{D}_{\ell,k}$ -KerMDH can be used to solve $\mathcal{D}_{\ell,k}$ -MDDH). By the above theorem, no black-box reduction in the generic group model from $\mathcal{D}_{\ell,k}$ -KerMDH to $\mathcal{D}_{\tilde{\ell},\tilde{k}}$ -KerMDH can exist for $k > \tilde{k}$.

4.2 Increasing Families of KerMDH Problems

Most matrix distributions, like $\mathcal{U}_{\ell,k}$, \mathcal{L}_k , \mathcal{SC}_k , \mathcal{C}_k and \mathcal{RL}_k , are indeed families parameterized by their size k . The negative results in Corollary 1 prevent us from finding reductions from larger to smaller KerMDH problems. Nevertheless, we provide here some examples of (tight) reductions going in the other way, within each of the previous families.

Lemma 8. $\mathcal{U}_{\tilde{\ell},\tilde{k}}$ -KerMDH \Rightarrow $\mathcal{U}_{\ell,k}$ -KerMDH for any $\tilde{k} \leq k$, $\tilde{\ell} > \tilde{k}$ and $\ell > k$.

Proof. We divide the proof into two steps: Firstly, assume that $\tilde{\ell} = \tilde{k} + 1$, $k \geq \tilde{k}$, $\ell \geq k + 1$. Given an instance $[\tilde{\mathbf{A}}]$, with $\tilde{\mathbf{A}} \leftarrow \mathcal{U}_{\tilde{k}+1,\tilde{k}}$, we choose a full-rank matrix $\mathbf{L} \in \mathbb{Z}_q^{\ell \times (k+1)}$ and compute $[\mathbf{A}] = \mathbf{L}([\tilde{\mathbf{A}}] \oplus [\mathbf{I}])$, where \mathbf{I} is the identity matrix of size $(k - \tilde{k}) \times (k - \tilde{k})$ and \oplus operation denotes diagonal block matrix concatenation. That is

$$U \oplus V = \begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix}.$$

Clearly, the probability distribution of the new matrix is statistically close to the uniform distribution in $\mathbb{Z}_q^{\ell \times k}$. Any vector $[\mathbf{x}]$, obtained from a solver of $\mathcal{U}_{\ell,k}$ -KerMDH, such that $\mathbf{x} \in \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}$ can be transformed into $[\tilde{\mathbf{x}}]$ such

that $\tilde{\mathbf{x}} \in \ker \tilde{\mathbf{A}}^\top \setminus \{\mathbf{0}\}$ with overwhelming probability,¹⁰ by just letting $[\tilde{\mathbf{x}}]$ to be the first $\tilde{k} + 1$ components of $\mathbf{L}^\top [\mathbf{x}]$. Thus, we have built a tight reduction $\mathcal{U}_{\tilde{k}+1, \tilde{k}}\text{-KerMDH} \Rightarrow \mathcal{U}_{\ell, k}\text{-KerMDH}$.

The second step, $k = \tilde{k}$, $\tilde{\ell} > \ell = \tilde{k} + 1$, is simpler. Given an instance $[\tilde{\mathbf{A}}]$, with $\tilde{\mathbf{A}} \leftarrow \mathcal{U}_{\tilde{\ell}, \tilde{k}}$, define the matrix $[\mathbf{A}]$ to be the upper $\tilde{k} + 1$ rows of $[\tilde{\mathbf{A}}]$. Clearly \mathbf{A} follows the uniform distribution in $\mathbb{Z}_q^{(\tilde{k}+1) \times \tilde{k}}$. Now, any vector $[\mathbf{x}]$ such that $\mathbf{x} \in \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}$ can be transformed into $[\tilde{\mathbf{x}}]$ such that $\tilde{\mathbf{x}} \in \ker \tilde{\mathbf{A}}^\top \setminus \{\mathbf{0}\}$, by just padding \mathbf{x} with $\tilde{\ell} - \tilde{k} - 1$ zeros. Thus, $\mathcal{U}_{\tilde{\ell}, \tilde{k}}\text{-KerMDH} \Rightarrow \mathcal{U}_{\tilde{k}+1, \tilde{k}}\text{-KerMDH}$. By concatenating the two tight reductions we obtain the general case.

Lemma 9. For $\mathcal{D}_k = \mathcal{L}_k, \mathcal{SC}_k, \mathcal{C}_k$ and \mathcal{RL}_k , $\mathcal{D}_k\text{-KerMDH} \Rightarrow \mathcal{D}_{k+1}\text{-KerMDH}$.

Proof. We start with the case $\mathcal{D}_k = \mathcal{L}_k$. Observe that given a matrix $\tilde{\mathbf{A}} \leftarrow \mathcal{L}_k$, with parameters a_1, \dots, a_k , we can build a matrix \mathbf{A} following the distribution \mathcal{L}_{k+1} , by adding an extra row and column to $\tilde{\mathbf{A}}$ corresponding to new random parameter $a_{k+1} \in \mathbb{Z}_q$. Moreover, given $\mathbf{x} = (x_1, \dots, x_{k+2}) \in \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}$, the vector $\tilde{\mathbf{x}} = (x_1, \dots, x_k, x_{k+2})$ is in $\ker \tilde{\mathbf{A}}^\top \setminus \{\mathbf{0}\}$ (except for a negligible probability due to the possibility that $a_{k+1} = 0$ and $\tilde{\mathbf{x}} = \mathbf{0}$, while $\mathbf{x} \neq \mathbf{0}$). The reduction consists of choosing a random a_{k+1} , then building $[\mathbf{A}]$ from $[\tilde{\mathbf{A}}]$ as above, and finally obtaining $[\tilde{\mathbf{x}}]$ from $[\mathbf{x}]$ by deleting the $(k + 1)$ -th coordinate.

Similarly, from a matrix $\tilde{\mathbf{A}} \leftarrow \mathcal{SC}_k$, with parameter a , we can obtain a matrix \mathbf{A} following \mathcal{SC}_{k+1} by adding a new row and column to $\tilde{\mathbf{A}}$. Now given $\mathbf{x} = (x_1, \dots, x_{k+2}) \in \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}$, it is easy to see that the vector $\tilde{\mathbf{x}} = (x_1, \dots, x_{k+1})$ is always in $\ker \tilde{\mathbf{A}}^\top \setminus \{\mathbf{0}\}$.

$\mathcal{C}_k\text{-KerMDH} \Rightarrow \mathcal{C}_{k+1}\text{-KerMDH}$ and $\mathcal{RL}_k\text{-KerMDH} \Rightarrow \mathcal{RL}_{k+1}\text{-KerMDH}$ are proven using the same ideas.

By combining Corollary 1 with the explicit reductions given above, we can now state our main result in this section.

Theorem 2. The matrix distribution families $\{\mathcal{U}_{\ell, k}\}$, $\{\mathcal{L}_k\}$, $\{\mathcal{SC}_k\}$, $\{\mathcal{C}_k\}$ and $\{\mathcal{RL}_k\}$ define families of KerMDH problems with **strictly** increasing hardness. Namely, for any $\mathcal{D}_{\ell, k}$ and $\mathcal{D}_{\tilde{\ell}, \tilde{k}}$ belonging to one of the previous families, such that $\tilde{k} < k$,

1. there exists a tight reduction, $\mathcal{D}_{\tilde{\ell}, \tilde{k}}\text{-KerMDH} \Rightarrow \mathcal{D}_{\ell, k}\text{-KerMDH}$,
2. there is no black-box reduction in the generic group model in the opposite direction.

5 Applications

We have already mentioned that the Kernel Matrix Diffie-Hellman Assumptions have already found applications in follow-up work, more concretely: a) to generalize and improve previous constructions of QA-NIZK proofs for linear spaces

¹⁰ Actually, $\tilde{\mathbf{x}} = \mathbf{0}$ depends on the $(\tilde{k} + 1)$ -th column of \mathbf{L} which is independent of \mathbf{A} .

[26], b) to construct more efficient structure preserving signatures starting from affine algebraic MACS [25], c) to improve and generalize aggregation of Groth-Sahai proofs [17] or d) to construct a tightly secure QA-NIZK argument for linear subspaces with unbounded simulation soundness [15].

As a new application, we use our new framework to abstract two constructions of trapdoor commitments. See for instance [3] for the formal definition of a trapdoor commitment scheme $C = (\mathbf{K}, \text{Comm}, \text{Vrfy}, \text{TrapdoorEquiv})$ and Section 6 for a discussion on the advantages of instantiating these commitments with the new circulant matrix distribution.

5.1 Generalized Pedersen Comments in Multilinear Groups

In a group $(\mathbb{G}, q, \mathcal{P})$ where the discrete logarithm is hard, the Pedersen commitment is a statistically hiding and computationally binding commitment to a scalar. It can be naturally generalized to several scalars. Abe *et al.* [2] show how to do similar Pedersen type commitments to vectors of group elements. With our new assumption family we can write both the Pedersen commitment and the commitment of [2] as a single construction and generalize it to (ideal) graded encodings.

- $\mathbf{K}(1^\lambda, d, m)$: Let $\mathcal{MG}_m = (e, \mathbb{G}_1, \mathbb{G}_2, \dots, \mathbb{G}_m, q, \mathcal{P}_1, \dots, \mathcal{P}_m) \leftarrow \text{MGen}_m(1^\lambda)$. Sample $\mathbf{A} \leftarrow \mathcal{D}_{k+d,k}$. Let $\overline{\mathbf{A}}$ be the first k rows of \mathbf{A} and $\underline{\mathbf{A}}$ the remaining d rows and $\mathbf{T} := \underline{\mathbf{A}}\overline{\mathbf{A}}^{-1}$ (w.l.o.g. we can assume $\overline{\mathbf{A}}$ is invertible). Output $ck := (\mathcal{MG}_m, [\mathbf{A}]_1)$, $tk := (\mathbf{T})$.
- $\text{Comm}(ck, [\mathbf{v}]_r)$: To commit to a vector $[\mathbf{v}]_r \in \mathbb{G}_r^d$, for any $r < m$, pick $\mathbf{s} \leftarrow \mathbb{Z}_q^k$, and output $[\mathbf{c}]_{r+1} := e([\mathbf{s}^\top \parallel \mathbf{v}^\top])_r, [\mathbf{A}]_1 = [(\mathbf{s}^\top \parallel \mathbf{v}^\top) \mathbf{A}]_{r+1} \in \mathbb{G}_{r+1}^k$, and the opening $Op = ([\mathbf{s}]_r)$.
- $\text{Vrfy}(ck, [\mathbf{v}]_r, Op)$: Given a message $[\mathbf{v}]_r$ and opening $Op = ([\mathbf{s}]_r)$, this algorithm outputs 1 if $[\mathbf{c}]_{r+1} = e([\mathbf{s}^\top \parallel \mathbf{v}^\top])_r, [\mathbf{A}]_1$.
- $\text{TrapdoorEquiv}(ck, tk, [\mathbf{c}]_{r+1}, [\mathbf{v}]_r, Op, [\mathbf{v}']_r)$: On a commitment $[\mathbf{c}]_{r+1} \in \mathbb{G}_{r+1}^k$ to message $[\mathbf{v}]_r$ with opening $Op = ([\mathbf{s}]_r)$, compute: $[\mathbf{s}']_r := [\mathbf{s}]_r + \mathbf{T}^\top([\mathbf{v} - \mathbf{v}']_r) \in \mathbb{G}_r^k$. Output $Op' = ([\mathbf{s}']_r)$ as the opening of $[\mathbf{c}]_{r+1}$ to $[\mathbf{v}']_r$.

The analysis is almost identical to [2]. The correctness of the trapdoor opening is straightforward. The hiding property of the commitment is unconditional, while the soundness (at level r) is based on the $(r, m, \mathcal{D}_{\ell,k})$ -KerMDH Assumption. Indeed, given two messages $[\mathbf{v}]_r, [\mathbf{v}']_r$ with respective openings $[\mathbf{s}]_r, [\mathbf{s}']_r$, it obviously follows that $[\mathbf{w}] := [((\mathbf{s} - \mathbf{s}')^\top \parallel (\mathbf{v} - \mathbf{v}')^\top)]_r$ is a nonzero element in the kernel (in \mathbb{G}_r) of \mathbf{A}^\top , *i.e.* $e([\mathbf{w}^\top]_r, [\mathbf{A}]_1) = [\mathbf{0}]_{r+1}$.

Notice that the Pedersen commitment (to multiple elements) is for messages in \mathbb{G}_0 and $\mathbf{A} \leftarrow \mathcal{U}_{d+1,1}$ and soundness is based on the $(0, m, \mathcal{U}_{d+1,1})$ -KerMDH. The construction proposed in [2] is for an asymmetric bilinear group \mathcal{AG}_2 , and in this case messages are vectors in the group \mathbb{H} and the commitment key consists of elements in \mathbb{G} , *i.e.* $ck = (\mathcal{AG}_2, [\mathbf{A}]_G)$, $\mathbf{A} \leftarrow \mathcal{U}_{d+1,1}$. Further, a previous version of the commitment scheme of [2] in symmetric bilinear groups (in [18]) corresponds to our construction with $\mathbf{A} \leftarrow \mathcal{U}_{2+d,2}$.

5.2 Group-to-Group Commitments

The commitments of the previous section are “shrinking” because they map a vector of length d in the group \mathbb{G}_r to a vector of length k , for some k independent of and typically smaller than d . Abe *et al.* [3] noted that in some applications it is useful to have “group-to-group” commitments, *i.e.* commitments which are defined in the same group as the vector message. The motivation for doing so in the bilinear case is that these commitments are better compatible with Groth-Sahai proofs.

There is a natural construction of group-to-group commitments which uses the generalized Pedersen commitment of section 5.1, which is denoted as $\text{Ped.C} = (\widetilde{\text{K}}, \widetilde{\text{Comm}}, \widetilde{\text{Vrfy}}, \widetilde{\text{TrapdoorEquiv}})$ in the following.

- $\text{K}(1^\lambda, d, m)$: Run $(\widetilde{ck}, \widetilde{tk}) \leftarrow \widetilde{\text{K}}(1^\lambda, m, d)$, output $ck = \widetilde{ck}$ and $tk = \widetilde{tk}$.
- $\text{Comm}(ck, [\mathbf{v}]_r)$: To commit to a vector $[\mathbf{v}]_r \in \mathbb{G}_r^d$, for any $0 < r < m$, pick $[\mathbf{t}]_{r-1} \leftarrow [\mathbb{G}]_{r-1}^k$. Let $([\widetilde{\mathbf{c}}]_r, \widetilde{Op} = ([\mathbf{s}]_{r-1})) \leftarrow \widetilde{\text{Comm}}(ck, [\mathbf{t}]_{r-1})$ and output $c := ([\mathbf{t} + \mathbf{v}]_r, [\widetilde{\mathbf{c}}]_r)$ and the opening $Op = ([\mathbf{s}]_r)$.
- $\text{Vrfy}(ck, c, [\mathbf{v}]_r, Op)$: On input $c = ([\mathbf{y}]_r, [\widetilde{\mathbf{c}}]_r)$, this algorithm computes $[\widetilde{\mathbf{c}}]_{r+1}$ and outputs 1 if $[\mathbf{t}]_r := [\mathbf{y} - \mathbf{v}]_r$ satisfies that $1 \leftarrow \widetilde{\text{Vrfy}}(ck, [\widetilde{\mathbf{c}}]_{r+1}, [\mathbf{t}]_r, [\mathbf{s}]_r)$, else it outputs 0.
- $\text{TrapdoorEquiv}(ck, tk, c, [\mathbf{v}]_r, Op, [\mathbf{v}']_r)$: On a commitment $c = ([\mathbf{y}]_r, [\widetilde{\mathbf{c}}]_r)$ with opening $Op = ([\mathbf{s}]_r)$, if $[\mathbf{t}]_r := [\mathbf{y} - \mathbf{v}]_r$ and $[\mathbf{t}']_r := [\mathbf{y} - \mathbf{v}']_r$, this algorithm computes $[\widetilde{\mathbf{c}}]_{r+1}$ and runs the algorithm $\widetilde{Op} \leftarrow \widetilde{\text{TrapdoorEquiv}}(ck, tk, [\widetilde{\mathbf{c}}]_{r+1}, [\mathbf{t}]_r, [\mathbf{s}]_r, [\mathbf{t}']_r)$, and outputs \widetilde{Op} .

A commitment is a vector of size $k + d$ and an opening is of size k . The required security properties follow easily from the properties of the generalized Pedersen commitment.

Theorem 3. *C is a perfectly hiding, computationally binding commitment.*

Proof. Since the generalized Pedersen commitment is perfectly hiding, then $([\mathbf{t} + \mathbf{v}]_r, \widetilde{\text{Comm}}(\widetilde{ck}, [\mathbf{t}]_{r-1}))$ perfectly hides $[\mathbf{v}]_r$ because $[\mathbf{t}]_r$ acts as a one-time pad. Similarly, it is straightforward to see that the computationally binding property of C follows from the computationally binding property of the generalized Pedersen commitment.

Interestingly, this construction explains the two instantiations of “group-to-group” commitments given in [3] (see the full version [35] for more details).

6 A New Matrix Distribution and Its Applications

Both of our commitment schemes of Section 5 base security on some $\mathcal{D}_{k+d,k}$ -KerMDH assumptions, where d is the length of the committed vector. When $d > 1$, the only example of $\mathcal{D}_{k+d,k}$ -MDDH Assumption considered in [11] is the one corresponding to the uniform matrix distribution $\mathcal{U}_{k+d,k}$, which is the

weakest MDDH Assumption of size $(k + d) \times k$. Another natural assumption for $d > 1$ is the one associated to the matrix distribution resulting from sampling from an arbitrary hard distribution $\mathcal{D}_{k+1,k}$ (e.g., \mathcal{L}_k) and adding $d - 1$ new random rows. Following the same ideas in the proof of Lemma 8, it is easy to see that the resulting $\mathcal{D}_{k+d,k}$ -MDDH assumption is equivalent to the original $\mathcal{D}_{k+1,k}$ -MDDH assumption. However, for efficiency reasons, we would like to have a matrix distributions with an even smaller representation size. This motivates us to introduce a new family of matrix distributions, the $\mathcal{CI}_{k,d}$ family.

Definition 9 (Circulant Matrix Distribution). We define $\mathcal{CI}_{k,d}$ as

$$\mathbf{A} = \begin{pmatrix} a_1 & & & 0 \\ \vdots & a_1 & & \\ a_d & \vdots & \ddots & \\ 1 & a_d & & a_1 \\ & 1 & \ddots & \vdots \\ & & \ddots & a_d \\ 0 & & & 1 \end{pmatrix} \in \mathbb{Z}_q^{(k+d) \times k}, \quad \text{where } a_i \leftarrow \mathbb{Z}_q$$

Matrix \mathbf{A} is such that each column can be obtained by rotating one position the previous column, which explains the name. Notice that when $d = 1$, $\mathcal{CI}_{k,d}$ is exactly the symmetric cascade distribution \mathcal{SC}_k , introduced in [11]. It can be shown that the representation size of $\mathcal{CI}_{k,d}$, which is the number of parameters d , is the optimal among all hard matrix distributions $\mathcal{D}_{k+d,k}$ defined by linear polynomials in the parameters. A similar argument shows that the circulant assumption is also optimal in the sense that it has a minimal number of nonzero entries among all hard matrix distributions $\mathcal{D}_{k+d,k}$. It can also be proven that $\mathcal{CI}_{k,d}$ -MDDH holds generically in k -linear groups, which implies the hardness of the corresponding KerMDH problem. To prove the generic hardness of the assumption, we turn to a result of Herold [20, Thm. 5.15 and corollaries]. It states that if all matrices produced by the matrix distribution are full-rank, $\mathcal{CI}_{k,d}$ is a hard matrix distribution. Indeed, an algorithm solving the $\mathcal{CI}_{k,d}$ -MDDH problem in the generic k -linear group model must be able to compute a polynomial in the ideal $\mathfrak{H} \subset \mathbb{Z}_q[a_1, \dots, a_d, z_1, \dots, z_{k+d}]$ generated by all the $(k + 1)$ -minors of $\mathbf{A}\|\mathbf{z}$ as polynomials in $a_1, \dots, a_d, z_1, \dots, z_{k+d}$. Although this ideal can actually be generated using only a few of the minors, we need to build a Gröbner basis of \mathfrak{H} to reason about the minimum degree a nonzero polynomial in \mathfrak{H} can have. We show that, carefully selecting a monomial order, the set of all $(k + 1)$ -minors of $\mathbf{A}\|\mathbf{z}$ form a Gröbner basis, and all these minors have total degree exactly $k + 1$. Therefore, all nonzero polynomials in \mathfrak{H} have degree at least $k + 1$, and then they cannot be evaluated by any algorithm in the generic k -linear group model. The full proof of both properties of $\mathcal{CI}_{k,d}$ can be found in the full version [35].

As for other matrix distribution families, we can combine Corollary 1 and the techniques used in Lemma 9 to show that for any fixed $d \geq 1$ the $\mathcal{CI}_{k,d}$ -KerMDH problem family has strictly increasing hardness.

Theorem 4. For any $d \geq 1$ and for any k, \tilde{k} such that $\tilde{k} < k$

1. there exists a tight reduction, $\mathcal{CI}_{\tilde{k},d}\text{-KerMDH} \Rightarrow \mathcal{CI}_{k,d}\text{-KerMDH}$,
2. there is no black-box reduction in the generic group model in the opposite direction.

The new assumption gives new instantiations of the commitment schemes of Section 5 with public parameters of size d , independent of k . Further, because the matrix $\mathbf{A} \leftarrow \mathcal{CI}_{k,d}$ has a many zero entries, the number of exponentiations computed by the Commit algorithm, and the number of pairings of the verification algorithm is kd — as opposed to $k(k+d)$ for the uniform assumption. This seems to be optimal — but we do not prove this formally.

7 Acknowledgements

The authors thank E. Kiltz and G. Herold for improving this work through very fruitful discussions. Also G. Herold gave us the insight and guidelines to prove the hardness of the circulant matrix distribution.

A Deferred Proofs from Section 2.2

Lemma 2. Let $\mathcal{A}^\mathcal{O}$ be an oracle algorithm in the (purely algebraic) generic multilinear group model, making a constant number of calls Q to an algebraic oracle \mathcal{O} . Let $([\mathbf{x}]_{\mathbf{a}}, \tilde{x})$ and $([\mathbf{y}]_{\mathbf{b}}, \tilde{y})$ respectively be the input and output of \mathcal{A} . Then, for every choice of \tilde{x} and the random tape, there exist polynomials of constant degree $Y_1, \dots, Y_\beta \in \mathbb{Z}_q[\mathbf{X}, \mathbf{R}_1, \dots, \mathbf{R}_Q]$, such that $\mathbf{y} = \mathbf{Y}(\mathbf{x}, \mathbf{r}_1, \dots, \mathbf{r}_Q)$, for all possible inputs, where $\mathbf{Y} = (Y_1, \dots, Y_\beta)$, and $\mathbf{r}_1, \dots, \mathbf{r}_Q$ are the parameters introduced in Definition 1 for the Q queries. Moreover, \tilde{y} does not depend on \mathbf{x} or $\mathbf{r}_1, \dots, \mathbf{r}_Q$.

Proof. We proceed by induction in Q . The first step, $Q = 0$, follows immediately from Lemma 1, because $\mathcal{A}^\mathcal{O}$ is just an algorithm (without oracle access). For $Q \geq 1$, we split $\mathcal{A}^\mathcal{O}$ into two sections $\mathcal{A}_0^\mathcal{O}$ and \mathcal{A}_1 , separated exactly at the last query point (see Figure 3). Let $([\mathbf{z}]_{\mathbf{c}}, \tilde{z})$ be the state information (group and non-group elements) that $\mathcal{A}_0^\mathcal{O}$ passes to \mathcal{A}_1 , $([\mathbf{u}]_{\mathbf{d}}, \tilde{u})$ be the Q -th query to \mathcal{O} , and $([\mathbf{v}]_{\mathbf{d}}, \tilde{v})$ be its corresponding answer. We assume that $\mathcal{A}_0^\mathcal{O}$ and \mathcal{A}_1 receive the same random tape, $\$,$ (perhaps introducing some redundant computations in \mathcal{A}_1). Observe that the output of $\mathcal{A}_0^\mathcal{O}$ consists of $([\mathbf{z}]_{\mathbf{c}}, \tilde{z})$ and $([\mathbf{u}]_{\mathbf{d}}, \tilde{u})$.

By the induction assumption, for any choice of \tilde{x} and $\$,$ there exist some polynomials of constant degree $Z_1, \dots, Z_\gamma \in \mathbb{Z}_q[\mathbf{X}, \mathbf{R}_1, \dots, \mathbf{R}_{Q-1}]$ and $U_1, \dots, U_\delta \in \mathbb{Z}_q[\mathbf{X}, \mathbf{R}_1, \dots, \mathbf{R}_{Q-1}]$ such that $\mathbf{z} = \mathbf{Z}(\mathbf{x}, \mathbf{r}_1, \dots, \mathbf{r}_{Q-1})$, where $\mathbf{Z} = (Z_1, \dots, Z_\gamma)$, and $\mathbf{u} = \mathbf{U}(\mathbf{x}, \mathbf{r}_1, \dots, \mathbf{r}_{Q-1})$, where $\mathbf{U} = (U_1, \dots, U_\delta)$, for all possible $\mathbf{x} \in \mathbb{Z}_q^\alpha$ and $\mathbf{r}_1, \dots, \mathbf{r}_{Q-1} \in \mathbb{Z}_q^\rho$. Moreover, \tilde{z} and \tilde{u} only depend on \tilde{x} and $\$$.

Now, the algorithm \mathcal{A}_1 receives as input $([\mathbf{z}]_{\mathbf{c}}, \tilde{z})$ and $([\mathbf{v}]_{\mathbf{e}}, \tilde{v})$. By Definition 1, \mathbf{v} also depend polynomially on \mathbf{u} and \mathbf{r}_Q . Namely, for every choice of \tilde{u} , there exist polynomials of constant degree $V_1, \dots, V_\epsilon \in \mathbb{Z}_q[\mathbf{U}, \mathbf{R}_Q]$ such that $\mathbf{v} = \mathbf{V}(\mathbf{u}, \mathbf{r}_Q)$, where $\mathbf{V} = (V_1, \dots, V_\epsilon)$, while \tilde{v} only depends on \tilde{u} .




Fig. 3. Splitting of the oracle algorithm in Lemma 2.

Since \mathcal{A}_1 is just an algorithm without oracle access, by Lemma 1, for any choice of \tilde{v} , \tilde{z} and $\$$, there exist polynomials of constant degree $Y_1, \dots, Y_\beta \in \mathbb{Z}_q[\mathbf{V}, \mathbf{Z}]$ such that $\mathbf{y} = \mathbf{Y}(\mathbf{v}, \mathbf{z})$, where $\mathbf{Y} = (Y_1, \dots, Y_\beta)$, for all $\mathbf{v} \in \mathbb{Z}_q^\epsilon$ and $\mathbf{z} \in \mathbb{Z}_q^\gamma$, while \tilde{y} only depends on \tilde{v} , \tilde{z} and $\$$. By composition of all the previous polynomials, we show that \mathbf{y} depend polynomially on \mathbf{x} and $\mathbf{r}_1, \dots, \mathbf{r}_Q$, where the polynomials depend only on $\$$ and \tilde{x} . Indeed

$$\mathbf{y} = \mathbf{Y}(\mathbf{V}(\mathbf{U}(\mathbf{x}, \mathbf{r}_1, \dots, \mathbf{r}_{Q-1}), \mathbf{r}_Q), \mathbf{Z}(\mathbf{x}, \mathbf{r}_1, \dots, \mathbf{r}_{Q-1}))$$

and all the polynomials involved depend only on \tilde{x} , \tilde{z} , \tilde{u} , \tilde{v} and $\$$, but all in turn only depend on \tilde{x} and $\$$. In addition, for the same reason, \tilde{y} only can depend on \tilde{x} and $\$$, which concludes the proof.

B Examples of Matrix Distributions

Some particular families of matrix distributions were presented in [11]. Namely,

$$\mathcal{SC}_k : \mathbf{A} = \begin{pmatrix} a & & 0 \\ 1 & \ddots & \\ & \ddots & a \\ 0 & & 1 \end{pmatrix} \quad \mathcal{C}_k : \mathbf{A} = \begin{pmatrix} a_1 & & 0 \\ 1 & \ddots & \\ & \ddots & a_k \\ 0 & & 1 \end{pmatrix} \quad \mathcal{L}_k : \mathbf{A} = \begin{pmatrix} a_1 & & 0 \\ & \ddots & \\ 0 & & a_k \\ 1 & \dots & 1 \end{pmatrix},$$

where $a, a_i \leftarrow \mathbb{Z}_p$, and $\mathcal{U}_{\ell,k}$ which is simply the uniform distribution in $\mathbb{Z}_p^{\ell \times k}$. The \mathcal{SC}_k -MDDH Assumption is the Symmetric Cascade Assumption, the \mathcal{C}_k -MDDH Assumption is the Cascade Assumption, which were proposed for the first time. $\mathcal{U}_{\ell,k}$ -MDDH is the Uniform Assumption, which appeared under other names in [7, 37]. \mathcal{L}_k -MDDH is the Decisional Linear Assumption [6, 22, 40]. For instance, we can consider the case $k = 2$, in which the \mathcal{L}_2 -MDDH problem is given $([1], [a_1], [a_2])$, tell apart the two distributions $([1], [a_1], [a_2], [w_1 a_1], [w_2 a_2], [w_1 +$

$w_2]$) and $([1], [a_1], [a_2], [z_1], [z_2], [z_3])$, where $a_1, a_2, w_1, w_2, z_1, z_2, z_3$ are random. This is exactly the 2-Lin Problem, since we can always set $z_1 = w_1 a_1$ and $z_2 = w_2 a_2$. We also give examples of matrix distributions which did not appear in [11] but that are implicitly used in the problems 2 and 4 in Appendix C. The Randomized Linear and the Square Polynomial distributions are respectively given by the matrices

$$\mathcal{RL}_k : \mathbf{A} = \begin{pmatrix} a_1 & & 0 \\ & \ddots & \\ 0 & & a_k \\ b_1 & \cdots & b_k \end{pmatrix} \quad \mathcal{P}_{\ell,2} : \mathbf{A} = \begin{pmatrix} a_1 & a_1^2 \\ a_2 & a_2^2 \\ \vdots & \vdots \\ a_\ell & a_\ell^2 \end{pmatrix}$$

where $a_i \leftarrow \mathbb{Z}_q$ and $b_i \leftarrow \mathbb{Z}_q^\times$. Jutla and Roy [24] referred to \mathcal{RL}_k -MDDH Assumption as the k -lifted Assumption.

C Flexible Problems That Fit into the New Framework

In this section we recall some computational problems in the cryptographic literature that we unify as particular instances of KerMDH problems. These problems are listed below, as they appear in the cited references. In the following, all parameters a_i and b_i are assumed to be randomly chosen in \mathbb{Z}_q .

1. Find-Rep [9]: Given $([a_1], \dots, [a_\ell])$, find a nonzero tuple (x_1, \dots, x_ℓ) such that $x_1 a_1 + \dots + a_\ell x_\ell = 0$.
2. Simultaneous Double Pairing (SDP) [2]: Given the two tuples, $([a_1], [b_1])$ and $([a_2], [b_2])$, find a nonzero tuple $([x_1], [x_2], [x_3])$ such that $x_1 b_1 + x_2 a_1 = 0$, $x_1 b_2 + x_3 a_2 = 0$.
3. Simultaneous Triple Pairing [18]: Given the two tuples, $([a_1], [a_2], [a_3])$ and $([b_1], [b_2], [b_3])$, find a nonzero tuple $([x_1], [x_2], [x_3])$ such that $x_1 a_1 + x_2 a_2 + x_3 a_3 = 0$, $x_1 b_1 + x_2 b_2 + x_3 b_3 = 0$.
4. Simultaneous Pairing [19]: Given $([a_1], [a_2], \dots, [a_\ell])$ and $([a_1^2], [a_2^2], \dots, [a_\ell^2])$, find a nonzero tuple $([x_1], \dots, [x_\ell])$ such that $\sum_{i=1}^{\ell} x_i a_i = 0$, $\sum_{i=1}^{\ell} x_i a_i^2 = 0$.
5. 1-Flexible Diffie-Hellman (1-FlexDH) [32]: Given $([1], [a], [b])$, find a triple $([r], [ra], [rab])$ with $r \neq 0$.
6. 1-Flexible Square Diffie-Hellman (1-FlexSDH) [27]: Given $([1], [a])$, find a triple $([r], [ra], [ra^2])$ with $r \neq 0$.
7. ℓ -Flexible Diffie-Hellman (ℓ -FlexDH) [32]: Given $([1], [a], [b])$, find a $(2\ell + 1)$ -tuple $([r_1], \dots, [r_\ell], [r_1 a], [r_1 r_2 a], \dots, [(\prod_{i=1}^{\ell} r_i) a], [(\prod_{i=1}^{\ell} r_i) ab])$ such that $r_j \neq 0$ for all $j = 1, \dots, \ell$.
8. Double Pairing (DP) [18]: In an asymmetric group $(\mathbb{G}, \mathbb{H}, \mathbb{T})$, given a pair of random elements $([a_1]_H, [a_2]_H) \in \mathbb{H}^2$, find a nonzero tuple $([x_1]_G, [x_2]_G)$ such that $[x_1 a_1 + x_2 a_2]_T = [0]_T$.

D Deferred Proofs from Section 4

Lemma 5. *There exists an algebraic oracle \mathcal{O} (in the sense of Definition 1), that solves the $\mathcal{D}_{\ell,k}$ -KerMDH Problem with probability one.*

Proof. Observe that $\mathcal{D}_{\ell,k}$ only uses group elements both in the instance description and in the solution to the problem. In addition, the problem (input/output relation) can be described by a polynomial map. Indeed, one can use the k -minors of \mathbf{A} , which are just polynomials of degree k , to obtain a basis of $\ker \mathbf{A}^\top$. Then the oracle can use parameters $r_1, \dots, r_{\ell-k}$ as the coefficients of an arbitrary linear combination of the basis vectors. Sampling these parameters uniformly results in an oracle answer uniformly distributed in $\ker \mathbf{A}^\top$.

Lemma 6. *Let $\mathcal{R}^\mathcal{O} = (\mathcal{R}_0^\mathcal{O}, \mathcal{R}_1)$ be a black-box reduction from $\mathcal{D}_{\ell,k}$ -KerMDH to $\mathcal{D}_{\ell,k}$ -KerMDH, in the purely algebraic generic multilinear group model, making $Q \geq 1$ calls to an oracle \mathcal{O} solving the latter with probability one. If $\mathcal{R}^\mathcal{O}$ succeeds with a non negligible probability ε then, for every possible behavior of the oracle, either $\Pr(\eta(\mathbf{w}) \in S') > \text{negl}$ or $\Pr(\mathbf{u} \in S') > \text{negl}$, where $S' = \ker \mathbf{A}^\top \setminus \{\mathbf{0}\}$, $[\mathbf{A}]$ is the input of $\mathcal{R}^\mathcal{O}$, and its output is written as $[\mathbf{u} + \eta(\mathbf{w})]$, for some \mathbf{u} only depending on the state output by $\mathcal{R}_0^\mathcal{O}$, $[\mathbf{w}]$ is the answer to the Q -th oracle query, and $\eta : \mathbb{Z}_q^\ell \rightarrow \mathbb{Z}_q^\ell$ is a (randomized) linear map that only depends on the random tape of $\mathcal{R}^\mathcal{O}$.*

Proof. Let us denote $S = \ker \mathbf{A}^\top$, where $[\mathbf{A}]$ is the input to $\mathcal{R}^\mathcal{O}$, and $S' = S \setminus \{\mathbf{0}\}$. Analogously, $\tilde{S} = \ker \tilde{\mathbf{A}}^\top$, where $[\tilde{\mathbf{A}}]$ is the Q -th oracle query, and $\tilde{S}' = \tilde{S} \setminus \{\mathbf{0}\}$. From the discussion preceding the lemma, we know that \mathbf{u} and η are well-defined and fulfil the required properties. In particular, η depends only on the random tape, $\$,$ of $\mathcal{R}^\mathcal{O}$. As a black-box reduction, $\mathcal{R}^\mathcal{O}$ is successful means that it is successful for every possible behavior of the oracle in its Q queries, with a success probability at least ε . We arbitrarily fix its behavior in the first $Q - 1$ queries. Concerning the last one, for all $\mathbf{w} \in \tilde{S}'$, $\Pr(\mathbf{u} + \eta(\mathbf{w}) \in S') > \varepsilon$, where the probability is computed with respect to $\$$ and the randomness of $[\mathbf{A}]$. Now, defining

$$\begin{aligned} p_{\mathbf{w}} &= \Pr(\mathbf{u} \in S \wedge \mathbf{u} + \eta(\mathbf{w}) \in S') \\ r_{\mathbf{w}} &= \Pr(\mathbf{u} \notin S \wedge \mathbf{u} + \eta(\mathbf{w}) \in S') \end{aligned}$$

we have $p_{\mathbf{w}} + r_{\mathbf{w}} > \varepsilon$. But not all $r_{\mathbf{w}}$ can be non-negligible since the corresponding events are disjoint. Indeed, for any vector $\mathbf{w} \neq \mathbf{0}$ and any different $\alpha_1, \alpha_2 \in \mathbb{Z}_q^\times$,

$$\mathbf{u} + \eta(\alpha_1 \mathbf{w}) \in S, \mathbf{u} + \eta(\alpha_2 \mathbf{w}) \in S \quad \Rightarrow \quad (\alpha_2 - \alpha_1) \mathbf{u} \in S \quad \Rightarrow \quad \mathbf{u} \in S$$

and then $\sum_{\alpha \in \mathbb{Z}_q^\times} r_{\alpha \mathbf{w}} \leq 1$. Thus, there exists α_m such that $r_{\alpha_m \mathbf{w}} \leq \frac{1}{q-1}$, which implies $p_{\alpha_m \mathbf{w}} > \varepsilon - \frac{1}{q-1}$. Now, we split $p_{\alpha_m \mathbf{w}}$, depending on whether $\mathbf{u} \in S'$ or $\mathbf{u} = \mathbf{0}$,

$$\begin{aligned} p_{\alpha_m \mathbf{w}} &= \Pr(\mathbf{u} = \mathbf{0} \wedge \eta(\mathbf{w}) \in S') + \Pr(\mathbf{u} \in S' \wedge \mathbf{u} + \eta(\alpha_m \mathbf{w}) \in S') \leq \\ &\leq \Pr(\eta(\mathbf{w}) \in S') + \Pr(\mathbf{u} \in S') \end{aligned}$$

and conclude that either $\Pr(\mathbf{u} \in S') > \text{negl}$ or for all nonzero $\mathbf{w} \in \tilde{S}'$, $\Pr(\eta(\mathbf{w}) \in S') > \text{negl}$. However, which one is true could depend on the particular behavior of the oracle in the first $Q - 1$ calls.

The next lemma is needed in other subsequent proofs.

Lemma 10. *Consider integers $l = k + d$, $\tilde{l} = \tilde{k} + \tilde{d}$ such that $k, d, \tilde{k}, \tilde{d} > 0$ and $k > \tilde{k}$. Let $\eta : \mathbb{Z}_q^{\tilde{l}} \rightarrow \mathbb{Z}_q^l$ be a linear map. Then, there exists a subspace F of $\text{Im} \eta$ of dimension at most k such that for all \tilde{d} -dimensional subspaces \tilde{S} of $\mathbb{Z}_q^{\tilde{l}}$, either $\tilde{S} \subset \ker \eta$ or $\dim F \cap \eta(\tilde{S}) \geq 1$.*

Proof. If $\text{rank} \eta \leq k$ it suffices to take $F = \text{Im} \eta$. Indeed, if $\tilde{S} \not\subset \ker \eta$, i.e., $\eta(\tilde{S}) \neq \{\mathbf{0}\}$, then $\dim F \cap \eta(\tilde{S}) = \dim \eta(\tilde{S}) \geq 1$. Otherwise, $\text{rank} \eta > k$, let F a subspace of $\text{Im} \eta$ of dimension k , using the Grassman's formula,

$$\begin{aligned} \dim F \cap \eta(\tilde{S}) &= \dim F + \dim \eta(\tilde{S}) - \dim(F + \eta(\tilde{S})) \geq k + \dim \eta(\tilde{S}) - \text{rank} \eta \geq \\ &\geq k + \dim \tilde{S} - \dim \ker \eta - \text{rank} \eta = k + \tilde{d} - \tilde{l} = k - \tilde{k} \geq 1 \end{aligned}$$

Lemma 7. *If a matrix distribution $\mathcal{D}_{\ell,k}$ is hard (as given in Definition 4) then $\mathcal{D}_{\ell,k}$ is k -elusive.*

Proof. By definition, given a non- k -elusive matrix distribution $\mathcal{D}_{\ell,k}$, there exists a k -dimensional vector subspace $F \subset \mathbb{Z}_q^\ell$ such that $\Pr_{\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}}(F \cap \ker \mathbf{A}^\top \neq \{\mathbf{0}\}) = \varepsilon > \text{negl}$. F can be efficiently computed from the description of $\mathcal{D}_{\ell,k}$ with standard tools from linear algebra.

Let $\mathbf{M} \in \mathbb{Z}_q^{k \times \ell}$ be a maximal rank matrix such that $\text{Im} \mathbf{M}^\top = F$. Then, $\dim(F \cap \ker \mathbf{A}^\top) = \dim(\text{Im} \mathbf{M}^\top \cap \ker \mathbf{A}^\top) \leq \dim \ker(\mathbf{A}^\top \mathbf{M}^\top) = \dim \ker(\mathbf{M}\mathbf{A})^\top = \dim \ker(\mathbf{M}\mathbf{A})$, as $\mathbf{M}\mathbf{A}$ is a $k \times k$ square matrix. Thus, we know that

$$\Pr_{\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}}(\text{rank}(\mathbf{M}\mathbf{A}) < k) \geq \varepsilon$$

Now we show how to solve the $\mathcal{D}_{\ell,k}$ -MDDH problem with advantage almost ε on some k -linear group \mathbb{G} , by means of a k -linear map. Let $[(\mathbf{A} \parallel \mathbf{z})]$ be an instance of the $\mathcal{D}_{\ell,k}$ -MDDH problem. In a 'real' instance $\mathbf{z} = \mathbf{A}\mathbf{x}$ for a uniformly distributed vector $\mathbf{x} \in \mathbb{Z}_q^k$, while in a 'random' instance, \mathbf{z} is uniformly distributed \mathbb{Z}_q^ℓ . A distinguisher can efficiently compute $[\mathbf{M}\mathbf{A}]$ and $[\mathbf{M}\mathbf{z}]$. Observe that in a 'real' instance $\text{rank}(\mathbf{M}\mathbf{A} \parallel \mathbf{M}\mathbf{z}) = \text{rank}(\mathbf{M}\mathbf{A} \parallel \mathbf{M}\mathbf{A}\mathbf{x}) = \text{rank}(\mathbf{M}\mathbf{A})$, while in a 'random' instance $\mathbf{M}\mathbf{z}$ is uniformly distributed in \mathbb{Z}_q^k . Therefore, for a 'random' instance there is a non-negligible probability, greater than $\varepsilon - \frac{1}{q}$, that $\text{rank}(\mathbf{M}\mathbf{A}) < k$ and $\text{rank}(\mathbf{M}\mathbf{A} \parallel \mathbf{M}\mathbf{z}) = \text{rank}(\mathbf{M}\mathbf{A}) + 1$, because $\mathbf{M}\mathbf{z} \in \text{Im}(\mathbf{M}\mathbf{A})$ occurs only with a negligible probability $< \frac{1}{q}$. Then, the distinguisher can efficiently tell apart the two cases because with a k -linear map at hand computing the rank of a $k \times k$ or a $k \times k + 1$ matrix can be done efficiently.

Theorem 1. *Let $\mathcal{D}_{\ell,k}$ be k -elusive. If there exists a black-box reduction in the purely algebraic generic multilinear group model from $\mathcal{D}_{\ell,k}$ -KerMDH to another problem $\mathcal{D}_{\tilde{\ell},\tilde{k}}$ -KerMDH with $\tilde{k} < k$, then $\mathcal{D}_{\ell,k}$ -KerMDH is easy.*

Proof. Let us assume the existence of the claimed reduction, $\mathcal{R}^\mathcal{O} = (\mathcal{R}_0^\mathcal{O}, \mathcal{R}_1)$, making $Q \geq 1$ oracle queries, where Q is minimal, and with a success probability ε . Then, by Lemma 6, its output can be written as $[\mathbf{u} + \eta(\mathbf{w})]$, where $\eta : \mathbb{Z}_q^l \rightarrow \mathbb{Z}_q^l$ is a (randomized) linear map that does not depend on the particular choice of the matrix \mathbf{A} in the $\mathcal{D}_{\ell,k}$ -KerMDH input instance, but only on the random tape of the reduction. Let us denote as above $S = \ker \mathbf{A}^\top$, and $S' = S \setminus \{\mathbf{0}\}$. Analogously, $\tilde{S} = \ker \tilde{\mathbf{A}}^\top$, where $\tilde{\mathbf{A}} \leftarrow \mathcal{D}_{\tilde{\ell}, \tilde{k}}$ and $\tilde{S}' = \tilde{S} \setminus \{\mathbf{0}\}$.

We now prove that in Lemma 6, for any possible behavior of the oracle in the first $Q - 1$ calls, there exists a particular behavior in the last call such that $\Pr(\eta(\mathbf{w}) \in S')$ is negligible. Namely, the Q -th query is answered by \mathcal{O} by choosing a uniformly distributed $\mathbf{w} \in \tilde{S}'$ (as required to be algebraic, according to Definition 1). Indeed, $\Pr(\eta(\mathbf{w}) \in S') = \Pr(\eta(\mathbf{w}) \in S) - \Pr(\eta(\mathbf{w}) = \mathbf{0})$. Now, developing the second term,

$$\begin{aligned} \Pr(\eta(\mathbf{w}) = \mathbf{0}) &= \Pr(\eta(\mathbf{w}) = \mathbf{0} \mid \tilde{S} \subset \ker \eta) \Pr(\tilde{S} \subset \ker \eta) + \\ &\quad + \Pr(\eta(\mathbf{w}) = \mathbf{0} \mid \tilde{S} \not\subset \ker \eta) \Pr(\tilde{S} \not\subset \ker \eta) = \\ &= \Pr(\tilde{S} \subset \ker \eta) + \Pr(\mathbf{w} \in \tilde{S} \cap \ker \eta \mid \tilde{S} \not\subset \ker \eta) \Pr(\tilde{S} \not\subset \ker \eta) = \\ &= \Pr(\tilde{S} \subset \ker \eta) + \text{negl} \end{aligned}$$

where the last equality uses that the probability that a vector uniformly distributed in \tilde{S}' belongs to a proper subspace of \tilde{S}' is negligible. Analogously,

$$\begin{aligned} \Pr(\eta(\mathbf{w}) \in S) &= \Pr(\eta(\mathbf{w}) \in S \mid \eta(\tilde{S}) \subset S) \Pr(\eta(\tilde{S}) \subset S) + \\ &\quad + \Pr(\eta(\mathbf{w}) \in S \mid \eta(\tilde{S}) \not\subset S) \Pr(\eta(\tilde{S}) \not\subset S) = \\ &= \Pr(\eta(\tilde{S}) \subset S) + \Pr(\mathbf{w} \in \tilde{S} \cap \eta^{-1}(S) \mid \eta(\tilde{S}) \not\subset S) \Pr(\eta(\tilde{S}) \not\subset S) = \\ &= \Pr(\eta(\tilde{S}) \subset S) + \text{negl} \end{aligned}$$

Thus, $\Pr(\eta(\mathbf{w}) \in S') = \Pr(\eta(\tilde{S}) \subset S) - \Pr(\tilde{S} \subset \ker \eta) + \text{negl}$. Now, using Lemma 10, we know that there exists a subspace F of dimension at most k such that if $\tilde{S} \not\subset \ker \eta$, then $\dim F \cap \eta(\tilde{S}) \geq 1$. Therefore $\Pr(\eta(\tilde{S}) \subset S) - \Pr(\tilde{S} \subset \ker \eta) \leq \Pr(\eta(\tilde{S}) \subset S \wedge \dim F \cap \eta(\tilde{S}) \geq 1) \leq \Pr(\dim F \cap S \geq 1)$. Due to the k -elusiveness of $\mathcal{D}_{\ell,k}$, from Lemma 7, the last probability is negligible. Namely, it is upper bounded by $\mathbf{Adv}_{\mathcal{D}_{\ell,k}\text{-MDDH}} + \frac{1}{q}$, where $\mathbf{Adv}_{\mathcal{D}_{\ell,k}\text{-MDDH}}$ denotes the advantage of a distinguisher for the $\mathcal{D}_{\ell,k}$ -MDDH problem. By Lemma 6,

$$\Pr(\mathbf{u} \in S \setminus \{\mathbf{0}\}) > \varepsilon - \frac{1}{q-1} - \mathbf{Adv}_{\mathcal{D}_{\ell,k}\text{-MDDH}} - \frac{1}{q},$$

for any possible behavior of the oracle in the first $Q - 1$ calls. Therefore, we can modify the reduction \mathcal{R} to output \mathbf{u} , without making the Q -th oracle call. The modified reduction is also successful, essentially with the same probability ε , with only $Q - 1$ oracle calls, which contradicts the assumption that Q is minimal. In summary, if the claimed reduction exists then there also exists an algorithm (a “reduction with $Q = 0$ ”) directly solving $\mathcal{D}_{\ell,k}$ -KerMDH without the help of any oracle and with the same success probability.

References

1. M. Abe, M. Chase, B. David, M. Kohlweiss, R. Nishimaki, and M. Ohkubo. Constant-size structure-preserving signatures: Generic constructions and simple assumptions. In X. Wang and K. Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 4–24, Beijing, China, Dec. 2–6, 2012. Springer, Heidelberg, Germany.
2. M. Abe, G. Fuchsbauer, J. Groth, K. Haralambiev, and M. Ohkubo. Structure-preserving signatures and commitments to group elements. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 209–236, Santa Barbara, CA, USA, Aug. 15–19, 2010. Springer, Heidelberg, Germany.
3. M. Abe, K. Haralambiev, and M. Ohkubo. Group to group commitments do not shrink. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 301–317, Cambridge, UK, Apr. 15–19, 2012. Springer, Heidelberg, Germany.
4. F. Bao, R. H. Deng, and H. Zhu. Variations of Diffie-Hellman problem. In S. Qing, D. Gollmann, and J. Zhou, editors, *ICICS 03*, volume 2836 of *LNCS*, pages 301–312, Huhahaote, China, Oct. 10–13, 2003. Springer, Heidelberg, Germany.
5. G. Barthe, E. Fagerholm, D. Fiore, J. C. Mitchell, A. Scedrov, and B. Schmidt. Automated analysis of cryptographic assumptions in generic group models. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 95–112, Santa Barbara, CA, USA, Aug. 17–21, 2014. Springer, Heidelberg, Germany.
6. D. Boneh, X. Boyen, and H. Shacham. Short group signatures. In M. Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 41–55, Santa Barbara, CA, USA, Aug. 15–19, 2004. Springer, Heidelberg, Germany.
7. D. Boneh, S. Halevi, M. Hamburg, and R. Ostrovsky. Circular-secure encryption from decision Diffie-Hellman. In D. Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 108–125, Santa Barbara, CA, USA, Aug. 17–21, 2008. Springer, Heidelberg, Germany.
8. X. Boyen. The uber-assumption family (invited talk). In S. D. Galbraith and K. G. Paterson, editors, *PAIRING 2008*, volume 5209 of *LNCS*, pages 39–56, Egham, UK, Sept. 1–3, 2008. Springer, Heidelberg, Germany.
9. S. Brands. Untraceable off-line cash in wallets with observers (extended abstract). In D. R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 302–318, Santa Barbara, CA, USA, Aug. 22–26, 1994. Springer, Heidelberg, Germany.
10. Y. Dodis, I. Haitner, and A. Tentes. On the instantiability of hash-and-sign RSA signatures. In R. Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 112–132, Taormina, Sicily, Italy, Mar. 19–21, 2012. Springer, Heidelberg, Germany.
11. A. Escala, G. Herold, E. Kiltz, C. Ràfols, and J. Villar. An algebraic framework for Diffie-Hellman assumptions. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147, Santa Barbara, CA, USA, Aug. 18–22, 2013. Springer, Heidelberg, Germany.
12. A. Escala, G. Herold, E. Kiltz, C. Ràfols, and J. Villar. An Algebraic Framework for Diffie-Hellman Assumptions. *Journal of Cryptology*, pages 1–47, 2015.
13. D. M. Freeman. Converting pairing-based cryptosystems from composite-order groups to prime-order groups. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 44–61, French Riviera, May 30 – June 3, 2010. Springer, Heidelberg, Germany.

14. D. Galindo, J. Herranz, and J. L. Villar. Identity-based encryption with master key-dependent message security and leakage-resilience. In S. Foresti, M. Yung, and F. Martinelli, editors, *ESORICS 2012*, volume 7459 of *LNCS*, pages 627–642, Pisa, Italy, Sept. 10–12, 2012. Springer, Heidelberg, Germany.
15. R. Gay, D. Hofheinz, E. Kiltz, and H. Wee. Tightly CCA-Secure Encryption Without Pairings. In M. Fischlin and J. Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part I*, volume 9665 of *Lecture Notes in Computer Science*, pages 1–27. Springer, 2016.
16. O. Goldreich. On post-modern cryptography. Cryptology ePrint Archive, Report 2006/461, 2006. <http://eprint.iacr.org/2006/461>.
17. A. González, A. Hevia, and C. Ràfols. QA-NIZK arguments in asymmetric groups: New tools and new constructions. In T. Iwata and J. H. Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 605–629, Auckland, New Zealand, Nov. 30 – Dec. 3, 2015. Springer, Heidelberg, Germany.
18. J. Groth. Homomorphic trapdoor commitments to group elements. Cryptology ePrint Archive, Report 2009/007, 2009. <http://eprint.iacr.org/2009/007>.
19. J. Groth and S. Lu. A non-interactive shuffle with pairing based verifiability. In K. Kurosawa, editor, *ASIACRYPT 2007*, volume 4833 of *LNCS*, pages 51–67, Kuching, Malaysia, Dec. 2–6, 2007. Springer, Heidelberg, Germany.
20. G. Herold. Applications of classical algebraic geometry to cryptography. *PhD Thesis, Ruhr-Universität Bochum*, 2014.
21. G. Herold, J. Hesse, D. Hofheinz, C. Ràfols, and A. Rupp. Polynomial spaces: A new framework for composite-to-prime-order transformations. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 261–279, Santa Barbara, CA, USA, Aug. 17–21, 2014. Springer, Heidelberg, Germany.
22. D. Hofheinz and E. Kiltz. Secure hybrid encryption from weakened key encapsulation. In A. Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 553–571, Santa Barbara, CA, USA, Aug. 19–23, 2007. Springer, Heidelberg, Germany.
23. A. Joux and A. Rojatz. Security ranking among assumptions within the uber assumption framework. Cryptology ePrint Archive, Report 2013/291, 2013. <http://eprint.iacr.org/2013/291>.
24. C. S. Jutla and A. Roy. Switching lemma for bilinear tests and constant-size NIZK proofs for linear subspaces. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 295–312, Santa Barbara, CA, USA, Aug. 17–21, 2014. Springer, Heidelberg, Germany.
25. E. Kiltz, J. Pan, and H. Wee. Structure-preserving signatures from standard assumptions, revisited. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 275–295, Santa Barbara, CA, USA, Aug. 16–20, 2015. Springer, Heidelberg, Germany.
26. E. Kiltz and H. Wee. Quasi-adaptive NIZK for linear subspaces revisited. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 101–128, Sofia, Bulgaria, Apr. 26–30, 2015. Springer, Heidelberg, Germany.
27. F. Laguillaumie, P. Paillier, and D. Vergnaud. Universally convertible directed signatures. In B. K. Roy, editor, *ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 682–701, Chennai, India, Dec. 4–8, 2005. Springer, Heidelberg, Germany.
28. T. Lepoint. Zeroizing attacks on multilinear maps. ECRYPT-CSA Workshop on Tools for Asymmetric Cryptanalysis, 2015. Slides available at <http://cryptool.hgi.rub.de/program.html>.

29. A. B. Lewko and B. Waters. Efficient pseudorandom functions from the decisional linear assumption and weaker variants. In E. Al-Shaer, S. Jha, and A. D. Keromytis, editors, *ACM CCS 09*, pages 112–120, Chicago, Illinois, USA, Nov. 9–13, 2009. ACM Press.
30. B. Libert, T. Peters, M. Joye, and M. Yung. Linearly homomorphic structure-preserving signatures and their applications. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 289–307, Santa Barbara, CA, USA, Aug. 18–22, 2013. Springer, Heidelberg, Germany.
31. B. Libert, T. Peters, M. Joye, and M. Yung. Non-malleability from malleability: Simulation-sound quasi-adaptive NIZK proofs and CCA2-secure encryption from homomorphic signatures. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 514–532, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
32. B. Libert and D. Vergnaud. Multi-use unidirectional proxy re-signatures. In P. Ning, P. F. Syverson, and S. Jha, editors, *ACM CCS 08*, pages 511–520, Alexandria, Virginia, USA, Oct. 27–31, 2008. ACM Press.
33. U. M. Maurer. Towards the equivalence of breaking the Diffie-Hellman protocol and computing discrete algorithms. In Y. Desmedt, editor, *CRYPTO'94*, volume 839 of *LNCS*, pages 271–281, Santa Barbara, CA, USA, Aug. 21–25, 1994. Springer, Heidelberg, Germany.
34. U. M. Maurer. Abstract models of computation in cryptography (invited paper). In N. P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12, Cirencester, UK, Dec. 19–21, 2005. Springer, Heidelberg, Germany.
35. P. Morillo, C. Ràfols, and J. L. Villar. Matrix computational assumptions in multilinear groups. Cryptology ePrint Archive, Report 2015/353, 2015. <http://eprint.iacr.org/2015/353>.
36. M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudorandom functions. In *38th FOCS*, pages 458–467, Miami Beach, Florida, Oct. 19–22, 1997. IEEE Computer Society Press.
37. M. Naor and G. Segev. Public-key cryptosystems resilient to key leakage. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 18–35, Santa Barbara, CA, USA, Aug. 16–20, 2009. Springer, Heidelberg, Germany.
38. T. P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In J. Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 129–140, Santa Barbara, CA, USA, Aug. 11–15, 1992. Springer, Heidelberg, Germany.
39. J. H. Seo. On the (im)possibility of projecting property in prime-order setting. In X. Wang and K. Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 61–79, Beijing, China, Dec. 2–6, 2012. Springer, Heidelberg, Germany.
40. H. Shacham. A cramer-shoup encryption scheme from the linear assumption and from progressively weaker linear variants. Cryptology ePrint Archive, Report 2007/074, 2007. <http://eprint.iacr.org/2007/074>.
41. J. L. Villar. Optimal reductions of some decisional problems to the rank problem. In X. Wang and K. Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 80–97, Beijing, China, Dec. 2–6, 2012. Springer, Heidelberg, Germany.
42. B. R. Waters. Efficient identity-based encryption without random oracles. In R. Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 114–127, Aarhus, Denmark, May 22–26, 2005. Springer, Heidelberg, Germany.