Non-interactive key exchange in a generic multilinear group: An underwhelming lower bound

Allison Bishop   Lucas Kowalczyk   Valerio Pastro

Abstract

This paper proves a lower bound on the degree of multilinearity needed to accomplish perfectly correct non-interactive k-party key exchange in a generic multilinear group, where k is only polylogarithmic in the security parameter. We conjecture that our result can be extended to polynomial k, as well as that similar lower bounds can be proven for more complex functionalties. We were surprised at our inability to extend the result easily, and we describe here the obstruction to completing our argument for polynomial values of k.

1 Introduction

The discovery of cryptographic bilinear and more recently multilinear maps [GGH13a] has fueled an exciting explosion of candidate constructions for advancing cryptographic functionalities. First applications of bilinear maps include three party non-interactive key exchange [Jou00] and identity-based encryption [BF01]. In the years following these initial breakthroughs, bilinear maps have driven many advances in cryptographic functionality, providing the first constructions of hierarchical identity-based encryption, attribute-based encryption, inner-product encryption, and several other functionalities. Multilinear maps have similarly led to the first constructions of further functionalities, such as indistinguishability obfuscation and functional encryption for all circuits [GGH$^+$13b]. In many cases, candidate constructions have been progressively refined, starting with heuristic arguments of security in a generic bilinear or multilinear group model, and eventually converging to full security arguments from well-studied assumptions such as the decisional linear assumption in the bilinear setting or the LWE assumption in the lattice setting.

There are (informal) parallels between many bilinear and LWE constructions, such as the bilinear IBE of Boneh-Boyen [BB04] and the LWE-based IBE of Agrawal, Boneh, and Boyen [ABB10]. However, there are some functionalities that are achievable from LWE, but are not known to be achievable in the bilinear setting, regardless of the complexity assumptions one is willing to make. One very interesting example is attribute-based encryption for circuits, constructed from LWE in [GVW13]. Furthermore, there are powerful functionalities that are known only in the multilinear setting (or at least seem to be most naturally achieved in the multilinear setting), where the status of specific complexity assumptions is relatively uncertain due to recent cryptanalysis (e.g. [CGH$^+$15, CLLT16, CLLT17]), and efficiency degrades very substantially as the required degree of multilinearity grows.

Thus we arrive at a point in the development of bilinear (and multilinear) cryptography where our understanding of what is achievable is tantalizing but disjointed. If we work towards a more systematic understanding, we can hope to answer questions like, "is circuit ABE from bilinear maps possible?" and more generally:

"what is the minimal level of multilinearity required to instantiate a particular cryptographic primitive?"

Generic Group Models Generic group models (e.g. [Sho97, BBG05]) present a natural starting ground for such systematic study, as they idealize the possible security properties of bilinear and multilinear groups and free us from the burden of reducing to specific complexity assumptions. For constructive results, relying on generic group models means that our security analyses are weaker and hold only heuristically, but for impossibility results, working in generic group models makes our results stronger, as ruling out schemes that are generically secure automatically rules out schemes that derive their security from any specific assumption as well. Of course, there are likely many interesting lower bounds/impossibility results that such models cannot capture. For instance, it would be nice to have a framework capable of proving bounds on tradeoffs between scheme parameter sizes and assumption sizes under certain classes of reductions. For now, we begin by focusing on questions of what is and is not achievable by generic groups at particular levels of multilinearity and ignore the complications of reductions to concrete complexity assumptions.

In this work, we take an initial step in building new machinery for proving impossibility results in generic bilinear and multilinear group models, though we get stuck at an unexpectedly early stage. Specifically, we consider the minimal level of multilinearity required to construct generically secure, perfectly correct non-interactive key exchange for k parties. Intuitively, we expect the answer here to be $k-1$. We consider a basic, symmetric model of a generic $(k-1)$-linear group G of prime order p generated by an element g such that there is a $(k-1)$-linear map $e_{k-1} : G^{k-1} \to G_T$. Informally, the generic group model considers an artificial setting where parties do not have direct access to the group elements, but rather there is an oracle that provides parties with random "handles" corresponding to group elements in G and G_T. Parties can perform the group operations in G and G_T, exponentiation in G and G_T, and the multilinear map e_{k-1} by submitting queries to the oracle that respectively contain two previously obtained handles (both for G or both for G_T), a single handle and an exponent in $\mathbb{Z}_p$, or $k-1$ handles for G. The oracle performs the requested operation internally and returns the handle for the resulting element of G or G_T. We will further stipulate that handles are unique, which implicitly allows parties to test equality of any two elements. (Equivalently, this can be thought of as "zero test" that allows parties to recognize the identity element of either group. Such a zero test can be applied to the difference of any two elements to yield an equality test.) Essentially, this captures a scenario where these requested operations plus zero/equality tests are the *only* efficiently computable relations among the elements of G and G_T, an idealized scenario for security.

In this model, there is a simple and generically secure non-interactive k-party key exchange protocol. Each party P_i (for i from 1 to k) chooses a uniformly random exponent x_i from $\mathbb{Z}_p$ and publishes g^{x_i}. The secret key is $e_{k-1}(g, \ldots, g)^{\prod_{i=1}^{k} x_i}$. A party P_i can compute this shared secret key by applying the $(k-1)$-linear map to the published group elements g^{x_j} for all $j \neq i$ and then raising the result to the power x_i in G_T. However, an attacker who only knows the published elements $g^{x_1}, \ldots, g^{x_k}$ cannot compute this shared key, and in the generic group model the scheme is thus secure (in fact, its security follows from the rather simple, falsifiable decisional assumption that given $g, g^{x_1}, \ldots, g^{x_k}$, it is hard to distinguish $e_{k-1}(g, \ldots, g)^{\prod_{i=1}^{k} x_i}$ from a uniformly random element of G_T). When we set the parameter $k = 3$, this corresponds to the familiar 3-party non-interactive key exchange protocol in bilinear groups whose security follows from the decisional bilinear Diffie-Hellman assumption [Jou00].

Our Result This begs the question: "can 3-party non-interactive key exchange be achieved without a bilinear map, say from just regular DDH?" Or "could we possibly achieve 4-party non-interactive key exchange from a bilinear map?" More generally,

"can a k-party non-interactive key exchange be achieved from a $(k-2)$-linear map?"

One might conjecture the answer to this questions to be a rather easy "no," but formalizing this is quite subtle. First, we must formalize what a key exchange protocol in a generic $(k-2)$-linear group can and cannot do. In particular, we must rule out schemes that use group elements merely as a means of encoding an alternate scheme that is secure based on some other assumption. For example, suppose we allow each party not only to publish group elements like g^x, but also to publish scalar values in $\mathbb{Z}_p$. These values could correspond to the public values of any arbitrary key exchange scheme outside of the generic group model, and the key reconstruction procedure run by each party could simply ignore the published group elements and just run this unrelated scheme on the scalar values.

We might conjecture that such behavior could be ruled out by stipulating that parties publish only group elements and not scalars. However, some distributions of exponents for these published values can still be used to encode arbitrary bits and hence allow the scheme to escape the bonds of the generic group and accomplish key exchange through other means. As an example, suppose there exists a non-interactive, k-party key exchange protocol where each party samples some ℓ secret bits and then publishes a bit string of length L. A party P_i could simulate this scheme in the generic group model as follows. It would first choose a secret bit string s_i of length ℓ, and then compute the corresponding public bit string S_i of length L. by choosing L random exponents $a_1, \ldots, a_L$ and then publishing $2L$ elements of G, comprised of L pairs. The j^{th} pair would have its first element equal to g^{a_j} and its second element equal to g^{a_j} again if the j^{th} bit of S_i equals 0 and equal to the identity element if the j^{th} bit of S_i equals 1. Now, every party can compute the strings $S_1, \ldots, S_k$ from the published group elements, as equality tests of group elements are enabled in the generic group model. Thus, each party P_i can run the key reconstruction procedure from the pre-existing scheme on its private value s_i and the collection of public values $S_1, \ldots, S_k$.

To cleanly rule out this kind of bypass of the generic group, we will consider only key agreement protocols in which an individual party P_i chooses independently, uniformly random exponents $x_{i,1}, \ldots, x_{i,\ell}$ (for some ℓ) and publishes $g^{x_{i,1}}, \ldots, g^{x_{i,\ell}}$. The key that such a P_i derives must be a handle for an element of G or G_T that can be derived via the generic group oracle from knowledge of the exponents $x_{i,1}, \ldots, x_{i,\ell}$ as well as the handles published by the other parties.

Our original goal was to prove that a perfectly correct and (generically) secure scheme in this setting for k parties cannot exist in a multilinear group with linearity lower than $k-1$. However, we were only able to prove this for values of k that are poly-logarithmic in the security parameter. We suspect the result holds for any k that is polynomial in the security parameter, but somewhat to our surprise, our proof technique broke down as k increased, and we were unable to extend it. We intended to submit this result to TCC 2015 when we thought we could prove it for general polynomial values of k, but we decided to shelve it temporarily when we realized the limitation of our technique. Our "temporary" hold lasted several years as we focused on other efforts, and so here we are! In this paper, we explain our original idea and the proof for the case that k is poly-logarithmic, and discuss how the analysis fails to extend to general polynomial k.

Our Techniques To see how one might analyze a hypothetical scheme of linearity $\leq k-2$ under our restrictions, we can start by considering what kind of keys a single party, say P_1, can possibly compute. Using the generic group oracle, P_1 can produce the handle of a group element in G or G_T whose exponent is a multivariate polynomial over the variables $\{X_{i,1}, \ldots, X_{i,\ell}\}_{i \in [k]}$. (We'll use the notational convention that $X_{i,j}$ denotes a random variable, whereas $x_{i,j}$ denotes a sample of that random variable.) The monomials in this polynomial must be of total degree less or equal to $k-2$ in the variables $\{X_{i,1}, \ldots, X_{i,\ell}\}_{i \neq 1}$, but can be arbitrary degree in $X_{1,1}, \ldots, X_{1,\ell}$. (The restriction of total degree $\leq k-2$ on the unknown exponents of other parties follows from the fact that we are assuming a maximum $(k-2)$-linear map.) Intuitively, we expect that monomials of high degree in his own variables $X_{1,1}, \ldots, X_{1,\ell}$ will not be useful, as other parties will not be able to compute them and hence their inclusion in the polynomial will cause disagreement among the reconstructed values. However, it is somewhat delicate to turn this intuition into a proof, as polynomial representations as linear combinations of monomials over $\mathbb{Z}_p$ are not unique when some variables can have arbitrarily high degree. In addition, we cannot naïvely apply typical tools such as the Schwartz-Zippel lemma out of the box, as every individual party can produce polynomials of high total degree due to knowledge of their own secret exponents.

To address these challenges, we first prove that we can reduce to a unique representation over a basis of monomials where all individual monomial degrees are $\leq k-2$, and then argue that any polynomial in the exponent that can be computed by all the parties can also be computed by an attacker from the published values. Hence any scheme that is perfectly correct is also insecure. But this conclusion is only fully justified when the total number of monomials is bounded by a polynomial in the security parameter. Otherwise, the fact that an attacker could compute each individual term in the monomial from the published group elements using the group oracles does not imply an overall polynomial attack time on the shared secret key. This is the crux of why our technique does not suffice to prove the result for values of k that are beyond poly-logarithmic.

Future Directions For the sake of the reader's amusement, here is the optimistic language we wrote in 2015 before we realized even our seemingly innocuous target of a proof for general polynomial values of k was not yet reached:

"Since we are merely at the starting point of an ambitious program to systematically understand the capabilities of generic bilinear and multilinear groups, we have narrowly transcribed a palatable and instructional base result and intend to pursue many extensions and further functionalities. Firstly, many of the restrictions we have placed on schemes can likely be relaxed. For example, we suspect our result can be extended to rule out statistically correct (vs. perfectly correct) schemes. We also suspect many kinds of non-independent distributions of published exponents could be allowed (though as discussed above this must be delicate and avoid the kind of counterexamples of encoding alternate schemes). Particular types of dependence, such as setting $x_{1_2} = (x_{1,1})^5$, in a generic bilinear group seem innocuous, as to the other parties this will be indistinguishable from a uniformly random, independent value for $x_{1,2}$. More generally, correlations between exponents that are higher degree than the multilinearity and hence untestable seem reasonable to handle. An additional way to cleanly rule out counterexamples while allowing more variation in the distributions of published values may be to consider a setting augmented by a powerful oracle that, say, inverts all supposedly

one-way functions but still cannot see "inside" the generic group. This could be an approach for ruling out reliance on hardness outside of the generic group without placing as stringent conditions on the form of allowable schemes.

Looking beyond key exchange protocols, we would next like to study primitives like broadcast encryption, where higher degrees of multilinearity appear to allow more compact ciphertexts and parameters (e.g. comparing [BW06] to [BWZ14]). It would be interesting to prove lower bounds, for example, on the number of group elements needed in public parameters and ciphertexts as a function of the linearity of the employed group. Going further, we would also like to study the class of access policies achievable for attribute-based encryption in a generic bilinear group. This class is known to include monotone span programs (of polynomial size), but not known to include polynomially-sized circuits."

Oh what it would feel like to be young and optimistic again.

2 Preliminaries

Lemma 2.1. *Let $m \in \mathbb{Z}_{\geq 1}$. Let $f \in \mathbb{F}_p[X_1, \dots, X_m]$ such that $\deg_{X_i}(f) \leq p-2$ for all $i \in [m]$. If $f(x_1, \dots, x_m) = 0$ for all $(x_1, \dots, x_m) \in \mathbb{F}_p^m$, then f is the zero polynomial $0 \in \mathbb{F}_p[X_1, \dots, X_m]$.*

Proof. By induction on m.

We start with the base case $m = 1$. If $f \in \mathbb{F}_p[X_1]$ evaluates to zero at every point in $\mathbb{F}_p$, then it must be either the zero polynomial or a polynomial of degree at least p in X_1 (since f has p roots). Since we assume that $\deg_{X_1}(f) \leq p-2$, then f must be the zero polynomial.

For the induction step, let $m \in \mathbb{Z}_{\geq 2}$. Assume that any $f' \in \mathbb{F}_p[X_1, \dots, X_{m-1}]$ such that $\deg_{X_i}(f') \leq p-2$ for all $i \in [m]$ and such that $f'(x_1, \dots, x_{m-1}) = 0$ for all $(x_1, \dots, x_{m-1}) \in \mathbb{F}_p^{m-1}$ is the zero polynomial $0 \in \mathbb{F}_p[X_1, \dots, X_{m-1}]$. Let $f \in \mathbb{F}_p[X_1, \dots, X_m]$ be such that $\deg_{X_i}(f) \leq p-2$ and $f(x_1, \dots, x_m) = 0$ for all $(x_1, \dots, x_m) \in \mathbb{F}_p^m$. Notice that for all $x_m \in \mathbb{F}_p$ we have that $f(X_1, \dots, X_{m-1}, x_m)$ is a polynomial in $\mathbb{F}_p[X_1, \dots, X_{m-1}]$ of degree less or equal to $p-2$ in each of its variables and it evaluates to zero at every point. This means that $f(X_1, \dots, X_{m-1}, x_m)$ is the zero polynomial $0 \in \mathbb{F}_p[X_1, \dots, X_{m-1}]$, by inductive hypothesis. Therefore, if we write

$$f(X_1, \dots, X_m) = \sum_{j=0}^{p-2} g_j(X_1, \dots, X_{m-1}) \cdot X_m^j \tag{2.1}$$

for some $g_j(X_1, \dots, X_{m-1}) \in \mathbb{F}_p[X_1, \dots, X_{m-1}]$ we have

$$\underbrace{\begin{pmatrix} 1 & 1 & \dots & 1 \\ 1 & 2 & \dots & 2^{p-2} \\ \vdots & \vdots & & \vdots \\ 1 & p-1 & \dots & (p-1)^{p-2} \end{pmatrix}}_{V \in \mathbb{F}_p^{p-1 \times p-1}} \underbrace{\begin{pmatrix} g_0(X_1, \dots, X_{m-1}) \\ \vdots \\ g_{p-2}(X_1, \dots, X_{m-1}) \end{pmatrix}}_{\vec{g} \in \mathbb{F}_p[X_1, \dots, X_{m-1}]^{p-1}} = \begin{pmatrix} f(X_1, \dots, X_{m-1}, 1) \\ \vdots \\ f(X_1, \dots, X_{m-1}, p-1) \end{pmatrix} = \underbrace{\begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}}_{\vec{0} \in \mathbb{F}_p[X_1, \dots, X_{m-1}]^{p-1}}$$

Since V is a square Vandermonde matrix, it is invertible, and by multiplying the above equation by V^{-1} we get

$$\vec{g} = V^{-1} V \vec{g} = V^{-1} \vec{0} = \vec{0}$$

which means that each $g_j(X_1, \dots, X_{m-1})$ is the zero polynomial $0 \in \mathbb{F}_p[X_1, \dots, X_{m-1}]$. This, combined with equation 2.1 implies that f is the zero polynomial $0 \in \mathbb{F}_p[X_1, \dots, X_m]$. □

3 Schematic for Scheme

Here we describe the format we are prescribing for a k-party non-interactive key exchange in a generic multilinear group:

Scheme Generation: Given as input the security parameter λ and the number of players k, the scheme generation outputs a $(k-2)$-linear group G, G_T of prime order p sufficiently large with respect to λ.

Publishing: For $j = 1, \ldots, \ell$, party P_i does the following:

- sample uniform $x_{i,j} \leftarrow \mathbb{F}_p$
- compute $h_{i,j} \leftarrow g^{x_{i,j}}$
- publish $h_{i,j} \leftarrow g^{x_{i,j}}$

Key Reconstruction: P_i runs $K_i \leftarrow \mathsf{Rec}_i((h_{1,j}, \ldots, h_{k,j}, x_{i,j})_{j \in [m]})$ where Rec_i is a function composed by a sequence of the following operations:

- group operation: on input two group elements $g^{a_1}, g^{a_2} \in G$ (respectively $g_T^{a_1}, g_T^{a_2} \in G_T$) for some $a_1, a_2 \in \mathbb{F}_p$, output $g^{a_1+a_2}$ (respectively $g_T^{a_1+a_2}$)
- multilinear operation: on input $k-2$ group elements $g^{a_1}, \ldots, g^{a_{k-2}} \in G$ for $a_i \in \mathbb{F}_p$, output $e_{k-2}(g^{a_1}, \ldots, g^{a_{k-2}})$
- exponentiation operation: on input a group element $g^a \in G$ (respectively $g_T^a \in G_T$), an ℓ-variate polynomial f, and the values $x_{i,1}, \ldots, x_{i,\ell}$, output $(g^a)^{f(x_{i,1},\ldots,x_{i,\ell})}$ (respectively $(g_T^a)^{f(x_{i,1},\ldots,x_{i,\ell})}$).

Theorem 3.1. *When k is poly-logarithmic in the security parameter, a perfectly correct key exchange following the schematic above cannot be implemented securely via an $(k-2)$-multilinear map.*

Proof. This proof is divided into separate steps. We show the following:

1. the reconstruction outputs a handle for a specific class of polynomials;
2. a handle to any polynomial in this class can be computed via a polynomial number of invocations of the group operation oracles and the use of the $(k-2)$-multilinear map oracle on published handles.

Without loss of generality, we can assume that any uniform value sampled by the players has a corresponding handle that is published. Let

$$V_i = \{X_{i,1}, \ldots, X_{i,m_i}\}$$

be the set of formal random variables belonging to P_i. Let

$$V = \bigcup_{i=1}^{k} V_i$$

be the set of all formal random variables in the scheme.

Without loss of generality, the handle reconstructed by P_i corresponds to a polynomial l_i of the following form:

$$l_i = \sum_{\substack{I \subseteq V \setminus V_i \\ |I| \leq k-2}} \sum_{\substack{\vec{b}_I \in [p]^{|I|} \\ \|\vec{b}_I\|_1 \leq k-2}} m_{I,\vec{b}_I} \cdot f_{I,\vec{b}_I}$$

where $\vec{b}_I = (b_{j,n})_{X_{j,n} \in I}$ is a vector of positive exponents (whose sum doesn't exceed $k-2$), $m_{I,\vec{b}_I} = \prod_{X_{j,n} \in I} X_{j,n}^{b_{j,n}}$, is a monomial in the variables that do not belong to P_i (at most $k-2$ of them), and $f_{I,\vec{b}_I}$ is a polynomial in variables belonging to P_i. This captures the fact that P_i has freedom to exponentiate (by any polynomial in his own variables) any handle corresponding to a monomial in other players' variables, of degree at most $k-2$ (this monomial is a by-product of the $(k-2)$-multilinear map).

We'll consider a particular monomial $m_{I,\vec{b}_I}$ in l_1. Since its total degree is $\leq k-2$, there must be some party P_j such that the variables V_j associated with party P_j do not appear in the monomial $m_{I,\vec{b}_I}$. If we sample the variables in V_1 and V_j, we can replace l_1 and l_j by the resulting polynomials over the remaining common variables $V - (V_1 \cup V_j)$, and we'll denote these by $\tilde{l}_1$ and $\tilde{l}_j$. Since $\tilde{l}_1 - \tilde{l}_j$ is a polynomial of total degree $\leq k-2 < p-2$, we can apply Lemma 2.1 to conclude that this is the zero polynomial.

Now let's examine the coefficients of $m_{I,\vec{b}_I}$ in $\tilde{l}_1$ and in $\tilde{l}_j$. The coefficient of this term in $\tilde{l}_1$ is:

$$f_{I,\vec{b}_I}(x_{1,1}, \ldots x_{1,n}),$$

where $x_{1,1}, \ldots, x_{1,n}$ are the sampled values of the variables in V_1. In $\tilde{l}_j$, the coefficient of this term is a polynomial of degree $\leq k-2$ in $x_{1,1}, \ldots, x_{1,n}$. Since we are assuming perfect correctness of our key exchange reconstruction, we must have that for all values of $x_{1,1}, \ldots, x_{1,n}$, this polynomial is equal to $f_{I,\vec{b}_I}$. Hence $f_{I,\vec{b}_I}$ can be considered to have degree $\leq k-2$. More specifically, it can be assume to have degree $\leq k-2$ minus the degree of $m_{I,\vec{b}_I}$.

Since this argument can be made for every monomial $m_{I,\vec{b}_I}$, we can conclude that all of the arbitrary polynomials $f_{I,\vec{b}_I}$ behave identically to polynomials of bounded degrees so that ultimately we can write $l_1 = l$ for a polynomial l over all the variables in V with total degree $\leq k-2$.

As a consequence, an attacker who is not a legitimate party to the key exchange can reconstruct a handle for each monomial in l using the group operation oracles as well as the $k-2$-linear map oracle, starting from the published handles.

When k is poly-logarithmic in the security parameter, the number of monomials in l is a polynomial in the security parameter, and hence this implies a polynomial time attack on the security of the key exchange scheme. □

Interestingly, it is *not clear* that this leads to a polynomial time attack in the case that k is polynomial in the security parameter, as the number of monomial terms in l may be exponential in that case. One might suspect that this would make it difficult for the legitimate parties to reconstruct the key in polynomial time, and we suspect that this is indeed true, but alas, we have not been able to prove it.

The challenge is to rule out the possibility that the higher degrees possible for the coefficient polynomials f employed by the legitimate parties may yield more efficient reconstruction algorithms in terms of the underlying group operations, even if they are ultimately equivalent to lower degree coefficients. We think that resolving this may provide a core component of the more powerful lower bound proof toolkit that should exist to approach more sophisticated questions about inherent capabilities and limitations of generic groups. We must admit we have not kept up with the latest lower bound techniques since our failed attempt to prove this in 2015, so it is quite possible that techniques developed in more recent works (e.g. [PS16, MMN16]) may provide some insight here.

References

[ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, pages 98–115, 2010.

[BB04] Dan Boneh and Xavier Boyen. Efficient selective-id secure identity-based encryption without random oracles. In *Advances in Cryptology - EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings*, pages 223–238, 2004.

[BBG05] Dan Boneh, Xavier Boyen, and Eu-Jin Goh. Hierarchical identity based encryption with constant size ciphertext. In *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings*, pages 440–456, 2005.

[BF01] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the weil pairing. In *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, pages 213–229, 2001.

[BW06] Dan Boneh and Brent Waters. A fully collusion resistant broadcast, trace, and revoke system. In *Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, Alexandria, VA, USA, Ioctober 30 - November 3, 2006*, pages 211–220, 2006.

[BWZ14] Dan Boneh, Brent Waters, and Mark Zhandry. Low overhead broadcast encryption from multilinear maps. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, pages 206–223, 2014.

[CGH$^+$15] Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancrède Lepoint, Hemanta K. Maji, Eric Miles, Mariana Raykova, Amit Sahai, and Mehdi Tibouchi. Zeroizing without low-level zeroes: New MMAP attacks and their limitations. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, pages 247–266, 2015.

[CLLT16] Jean-Sébastien Coron, Moon Sung Lee, Tancrède Lepoint, and Mehdi Tibouchi. Cryptanalysis of GGH15 multilinear maps. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, pages 607–628, 2016.

[CLLT17] Jean-Sébastien Coron, Moon Sung Lee, Tancrède Lepoint, and Mehdi Tibouchi. Zeroizing attacks on indistinguishability obfuscation over CLT13. In *Public-Key Cryptography - PKC 2017 - 20th IACR International Conference on Practice and Theory in Public-Key Cryptography, Amsterdam, The Netherlands, March 28-31, 2017, Proceedings, Part I*, pages 41–58, 2017.

[GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, pages 1–17, 2013.

[GGH$^+$13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 40–49, 2013.

[GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 545–554, 2013.

[Jou00] Antoine Joux. A one round protocol for tripartite diffie-hellman. In *Algorithmic Number Theory, 4th International Symposium, ANTS-IV, Leiden, The Netherlands, July 2-7, 2000, Proceedings*, pages 385–394, 2000.

[MMN16] Mohammad Mahmoody, Ameer Mohammed, and Soheil Nematihaji. On the impossibility of virtual black-box obfuscation in idealized models. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, pages 18–48, 2016.

[PS16] Rafael Pass and Abhi Shelat. Impossibility of VBB obfuscation with ideal constant-degree graded encodings. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, pages 3–17, 2016.

[Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In *Advances in Cryptology - EUROCRYPT '97, International Conference on the Theory and Application of Cryptographic Techniques, Konstanz, Germany, May 11-15, 1997, Proceeding*, pages 256–266, 1997.