

Annihilation Attacks for Multilinear Maps: Cryptanalysis of Indistinguishability Obfuscation over GGH13

Eric Miles
UCLA
enmiles@cs.ucla.edu

Amit Sahai*
UCLA
sahai@cs.ucla.edu

Mark Zhandry
MIT
mzhandry@gmail.com

Abstract

In this work, we put forward a new class of polynomial-time attacks on the original multilinear maps of Garg, Gentry, and Halevi (2013). Previous polynomial-time attacks on GGH13 were “zeroizing” attacks that generally required the availability of low-level encodings of zero. Most significantly, such zeroizing attacks were not applicable to candidate indistinguishability obfuscation (iO) schemes. iO has been the subject of intense study.

To address this gap, we introduce *annihilation attacks*, which attack multilinear maps using non-linear polynomials. Annihilation attacks can work in situations where there are no low-level encodings of zero. Using annihilation attacks, we give the first polynomial-time cryptanalysis of candidate iO schemes over GGH13. More specifically, we exhibit two simple programs that are functionally equivalent, and show how to efficiently distinguish between the obfuscations of these two programs.

Given the enormous applicability of iO, it is important to devise iO schemes that can avoid attack.

*Research supported in part from a DARPA/ARL SAFEWARE award, NSF Frontier Award 1413955, NSF grants 1228984, 1136174, 1118096, and 1065276, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the ARL under Contract W911NF-15-C-0205. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

1 Introduction

In this work, we put forward a new class of polynomial-time attacks on the original multilinear maps of Garg, Gentry, and Halevi [GGH13a]. Previous polynomial-time attacks on GGH13 were “zeroizing” attacks that either explicitly required the availability of low-level encodings of zero [GGH13a, HJ15], or required a differently represented low-level encoding of zero, in the form of an encoded matrix with a zero eigenvalue [CGH⁺15]. As a result, such zeroizing attacks were not applicable to many important applications of multilinear maps, most notably candidate indistinguishability obfuscation (iO) schemes over GGH13 [GGH⁺13b, BR14, BGK⁺14, AGIS14, MSW14, BMSZ15]. iO has been the subject of intense study. Thus, understanding the security of candidate iO schemes is of high importance. To do so, we need to develop new attacks that do not require, explicitly or implicitly, low-level encodings of zero.

Annihilation Attacks. To address this gap, we introduce *annihilation attacks*, which attack multilinear maps in a new way, using non-linear polynomials. Annihilation attacks can work in situations where there are no low-level encodings of zero. Using annihilation attacks, we give the first polynomial-time cryptanalysis of candidate iO schemes over GGH13. More specifically, we exhibit two simple programs that are functionally equivalent, and show how to efficiently distinguish between the obfuscations of these two programs. We also show how to extend our attacks to more complex candidate obfuscation schemes over GGH13, namely ones that incorporate the “dual-input” approach of [BGK⁺14]. (Note that, even without the dual-input structure, [BGK⁺14, AGIS14, MSW14, BMSZ15] were candidates for achieving iO security when implemented with [GGH13a].)

We now give an overview of our attack. The overview will introduce the main conceptual ideas and challenges in mounting our attack.

1.1 Overview of the Attack

We begin with a simplified description of the GGH13 scheme, adapted from text in [CGH⁺15].

1.1.1 The GGH13 scheme

For GGH13 [GGH13a] with k levels of multilinearity, the plaintext space is a quotient ring $R_g = R/gR$ where R is the ring of integers in a number field and $g \in R$ is a “small element” in that ring. The space of encodings is $R_q = R/qR$ where q is a “big integer”. An instance of the scheme relies on two secret elements, the generator g itself and a uniformly random denominator $z \in R_q$. A small plaintext element α is encoded “at level one” as $u = [e/z]_q$ where e is a “small element” in the coset of α , that is $e = \alpha + gr$ for some small $r \in R$.

Addition/subtraction of encodings at the same level is just addition in R_q , and it results in an encoding of the sum at the same level, so long as the numerators do not wrap around modulo q . Similarly multiplication of elements at levels i, i' is a multiplication in R_q , and as long as the numerators do not wrap around modulo q the result is an encoding of the product at level $i + i'$.

The scheme also includes a “zero-test parameter” in order to enable testing for zero at level k . Noting that a level- k encoding of zero is of the form $u = [gr/z^k]_q$, the zero-test parameter is an element of the form $\mathbf{p}_{zt} = [hz^k/g]_q$ for a “somewhat small element” $h \in R$. This lets us eliminate the z^k in the denominator and the g in the numerator by computing $[\mathbf{p}_{zt} \cdot u]_q = h \cdot r$, which is much smaller than q because both h, r are small. If u is an encoding of a non-zero α , however, then

multiplying by \mathbf{p}_{zt} leaves a term of $[h\alpha/g]_q$ which is not small. Testing for zero therefore consists of multiplying by the zero-test parameter modulo q and checking if the result is much smaller than q .

Note that above we describe the “symmetric” setting for multilinear maps where there is only one z , and its powers occur in the denominators of encodings. More generally, we will equally well be able to deal with the “asymmetric” setting where there are multiple z_i . However, we omit this generalization here as our attack is agnostic to such choices. Our attack is also agnostic to other basic parameters of the GGH13, including the specific choice of polynomial defining the ring R .

1.1.2 Setting of our attack

Recall that in our setting, we – as the attacker – will not have access to any low-level encodings of zero. Thus, in general, we are given as input a vector \vec{u} of ℓ encodings, corresponding to a vector $\vec{\alpha}$ of ℓ values being encoded, and with respect to a vector \vec{r} of ℓ random small elements. Thus, for each $i \in [\ell]$, there exists some value $j_i < k$ such that

$$u_i = \left[\frac{\alpha_i + gr_i}{z^{j_i}} \right]_q \quad \alpha_i \neq 0$$

What a distinguishing attack entails. In general, we consider a situation where there are two distributions over vectors: $\vec{\alpha}^{(0)}$ and $\vec{\alpha}^{(1)}$. Rather than thinking of these vectors as directly giving distributions over values, we can think of them as distinct vectors of multivariate polynomials over some underlying random variables. Thus, from this viewpoint, $\vec{\alpha}^{(0)}$ and $\vec{\alpha}^{(1)}$ are just two distinct vectors of polynomials, that are known to us in our role as attacker.

Then a challenger chooses a random bit $b \in \{0, 1\}$, and we set $\vec{\alpha} = \vec{\alpha}^{(b)}$. Then we are given encodings \vec{u} of the values $\vec{\alpha}$ using fresh randomness \vec{r} , and our goal in mounting an attack is to determine the challenger’s bit b .

Note that to make this question interesting, it should be the case that all efficiently computable methods of computing top-level encodings of zero using encodings of $\vec{\alpha}^{(0)}$ should also yield top-level encodings of zero using encodings of $\vec{\alpha}^{(1)}$. Otherwise, an adversary can distinguish between these encodings simply by zero testing.

1.1.3 Using annihilating polynomials

Our attack first needs to move to the polynomial ring R . In order to do so, the attack will need to build top-level encodings of zero, and then multiply by the zero-testing element \mathbf{p}_{zt} . Because we are in a setting where there are no low-level encodings of zero, top-level encodings of zero can only be created through algebraic manipulations of low-level encodings of nonzero values that lead to cancellation. Indeed, a full characterization of exactly how top-level encodings of zero can be created for candidate iO schemes over GGH13 was recently given by [BMSZ15]. In general, our attack will need to have access to a collection of valid algebraic manipulations that yield top-level encodings of zero, starting with the encodings \vec{u} .

Generally, then, a top-level encoding of zero e produced in this way would be stratified into levels corresponding to different powers of g , as follows:

$$e = \frac{g\gamma_1 + g^2\gamma_2 + \dots + g^k\gamma_k}{z^k}$$

and thus

$$f := [e \cdot \mathbf{p}_{zt}]_q = h \cdot (\gamma_1 + g\gamma_2 + \dots + g^{k-1}\gamma_k)$$

Above, each γ_i is a polynomial in the entries of $\vec{\alpha}$ and \vec{r} . As suggested by the stratification above, our main idea is to focus on just one level of the stratification. In particular, let us focus on the first level of the stratification, corresponding to the polynomial γ_1 .

A simple illustrative example. Suppose that we had three ways of generating top-level encodings of zero, e, e' , and e'' , which yield products f, f' , and f'' in the ring R . Suppose further that e, e' , and e'' contained polynomials $\gamma_1 = xr$; $\gamma'_1 = xr^2$; and $\gamma''_1 = x$, where x is a random variable underlying $\vec{\alpha}^{(0)}$. Then we observe that there is an efficiently computable *annihilating polynomial*, $Q(a, b, c) := a^2 - bc$, such that $Q(\gamma_1, \gamma'_1, \gamma''_1)$ is the zero polynomial. Further, because Q is homogeneous, $Q(h \cdot \gamma_1, h \cdot \gamma'_1, h \cdot \gamma''_1)$ is also the zero polynomial. (We will always ensure that our annihilating polynomials are homogeneous, which essentially comes for free due to the homogeneity of the γ_1 polynomials in the iO setting; see Lemma 5.3.)

Thus, if we compute $Q(f, f', f'')$, we obtain an element in the ring R that is contained in the ideal $\langle hg \rangle$.

However, consider the top-level encodings of zero e, e' , and e'' that arise from $\vec{\alpha}^{(1)}$, which is a different vector of polynomials over x than $\vec{\alpha}^{(0)}$. Suppose that in this case, the encodings e, e' , and e'' contain polynomials $\gamma_1 = x^3r$; $\gamma'_1 = xr$; and $\gamma''_1 = x$. In this scenario, the polynomial Q is no longer annihilating, and instead yields $Q(\gamma_1, \gamma'_1, \gamma''_1) = x^6r^2 - x^2r$. Thus, what we have is that if the challenge bit $b = 0$, then $Q(f, f', f'')$ is contained in the ideal $\langle hg \rangle$, but if the challenge bit $b = 1$, then $Q(f, f', f'')$ is not contained in the ideal $\langle hg \rangle$.

Obtaining this distinction in outcomes is the main new idea behind our attack.

1.1.4 Central challenge: How to compute annihilating polynomials?

While it was easy to devise an annihilating polynomial for the polynomials contained in the simple example above, in general annihilating polynomials can be hard to compute. Every set of $n + 1$ or more polynomials over n variables is algebraically dependent and hence must admit an annihilating polynomial. Indeed, therefore, if we do not worry about how to compute annihilating polynomials, our high-level attack idea as described above would apply to *every* published iO scheme that can be instantiated with GH13 maps that we are aware of, and it would work for every pair of equivalent programs that output zero sufficiently often. This is simply because every published iO candidate can be written as an algebraic expression using only a polynomial number of underlying random variables, whereas the obfuscated program can be evaluated on an exponential number of inputs.

However, unless the polynomial hierarchy collapses (specifically, unless $\mathbf{Co-NP} \subseteq \mathbf{AM}$), there are sets of (cubic) polynomials over n variables for which the annihilating polynomial cannot be represented by any polynomial-size arithmetic circuit [Kay09]. As a result, for our attack idea to be meaningful, we must show that the annihilating polynomials we seek are efficiently representable by arithmetic circuits and that such representations are efficiently computable. In particular, we seek to do this in the context of (quite complex) candidates for indistinguishability obfuscation.

We begin by looking deeper at the structure of the polynomials γ_1 that we need to annihilate. In particular, let's examine what these polynomials look like as a consequence of the stratification by powers of g . We see that by the structure of encodings in GH13, each polynomial γ_1 will be linear in the entries of \vec{r} and potentially non-linear in the entries of $\vec{\alpha}$. This is already useful,

since the \vec{r} variables are totally unstructured and unique to each encoding given out, and therefore present an obstacle to the kind of analysis that will enable us to find an annihilating polynomial.

To attack iO , we will first design two simple branching programs that are functionally equivalent but distinct as branching programs. To this end, we consider two branching programs that both compute the *always zero* functionality. The simplest such program is one where every matrix is simply the identity matrix, and this will certainly compute the constant zero functionality. To design another such program, we observe that the anti-identity matrix

$$B = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

can be useful, because it has the property that $BB = I$. Thus, to make another branching program that computes the always zero functionality, we can create a two-pass branching program, where the (two) matrices corresponding to $x_1 = 0$ are both set to B , and all other matrices are set to I .

With these branching programs in mind, we analyze the γ_1 polynomials that arise. The main method that we use to prune the search space for annihilating polynomials is to find changes of variables that can group variables together in order to minimize the number of active variables. We use a number of methods, including inclusion-exclusion formulas, to do this. By changing variables, we are able to reduce the problem to finding the annihilating polynomial for a set of polynomials over only a *constant* number of variables. When only a constant number of variables are present, exhaustive methods for finding annihilating polynomials are efficient. For further details, refer to Section 5.

Moving to more complex iO candidates. The above discussion covers the main ideas for finding annihilating polynomials, and by generalizing our methods, we show that they extend to more challenging settings. Most notably, we can extend our methods to work for the *dual-input* technique of [BGK⁺14], which has been used in several follow-up works [AGIS14, MSW14, BMSZ15]. Previously, no cryptanalysis techniques were known to apply to this setting. For further details, see Sections 4 and 5.

1.1.5 An abstract attack model

We first describe our attacks within a new abstract attack model, which is closely related to a model proposed in [CGH⁺15, App. A]. The new model is roughly the same as existing generic graded encoding models, except that a successful zero test returns an *algebraic element* rather than a bit $b \in \{0, 1\}$. These algebraic elements can then be manipulated, say, by evaluating an annihilating polynomial over them. This model captures the fact that, in the GGH13 candidate graded encoding scheme [GGH13a], the zero test actually does return an algebraic element in a polynomial ring that can be manipulated.

We describe our attacks in this abstract model to (1) highlight the main new ideas for our attack, and (2) to demonstrate the robustness of our attack to simple “fixes” for multilinear maps that have been proposed.

Theorem 1.1. *Let \mathcal{O} denote the single-input variant of the iO candidates in [BGK⁺14, AGIS14, BMSZ15] over GGH13 [GGH13a] maps. There exist two functionally-equivalent branching programs \mathbf{A}, \mathbf{A}' such that $\mathcal{O}(\mathbf{A})$ and $\mathcal{O}(\mathbf{A}')$ can be efficiently distinguished in the abstract attack model described in Section 2.*

Note that in the single input case, the [BGK⁺14, AGIS14, BMSZ15] obfuscators over GGH13 [GGH13a] maps were shown to achieve iO security in the standard generic graded encoding model. This theorem shows that such security does not extend to our more refined model.

The attack in Theorem 1.1 works by executing the obfuscated program honestly on several inputs, which produces several zero-tested top-level 0-encodings. Recall that in our model, each successful zero-test returns an algebraic element. We then give an explicit polynomial that annihilates these algebraic elements in the case of one branching program, but fails to annihilate in the other. Thus by evaluating this polynomial on the algebraic elements obtained and testing for zero, it is possible to distinguish the two cases.

1.1.6 Beyond the abstract attack model

Our abstract attack does not immediately yield an attack on actual graded encoding instances. For example, when the graded encoding is instantiated with [GGH13a], the result of an annihilating polynomial is an element in the ideal $\langle hg \rangle$, whereas if the polynomial does not annihilate, then the element is not in this ideal. However, this ideal is not explicitly known, so it is not a priori obvious how to distinguish the two cases.

We observe that by evaluating the annihilating polynomial many times on different sets of values, we get many different vectors in $\langle hg \rangle$. With enough vectors, we (heuristically) can compute a spanning set of vectors for $\langle hg \rangle$. This is the only heuristic portion of our attack analysis, and it is similar in spirit to previous heuristic analysis given in other attacks of multilinear maps (see, e.g., [CGH⁺15]). With such a spanning set, we can then test to see if another “test” vector is in this ideal or not. This is the foundation for our attack on obfuscation built from the specific [GGH13a] candidate.

Dual-Input Obfuscation and Beyond. Moving on to the dual-input setting, we do not know an explicit annihilating polynomial for the set of algebraic elements returned by our model. However, we are able to show both that such a polynomial must exist, and furthermore that it must be efficiently computable because it has constant size. Thus we demonstrate that there *exists* an efficient distinguishing adversary in the abstract attack model. As before, we can turn this into a heuristic attack on obfuscation built from [GGH13a] graded encodings. We also show that modifying the branching programs to read $d > 2$ bits at each level does not thwart the attack for constant d , because the annihilating polynomial still has constant size (albeit a larger constant).

Theorem 1.2. *Let \mathcal{O} denote the dual-input variant of the iO candidates found in [BGK⁺14, AGIS14, BMSZ15] over GGH13 [GGH13a] maps. There exist two functionally-equivalent branching programs \mathbf{A}, \mathbf{A}' such that $\mathcal{O}(\mathbf{A})$ and $\mathcal{O}(\mathbf{A}')$ can be efficiently distinguished in the abstract attack model described in Section 2.*

1.2 Defenses

The investigation of the mathematics needed to build iO remains in its infancy. In particular, our work initiates the study of annihilation attacks, and significant future study is needed to see how such attacks can be avoided.

2 Model Description

We now describe an abstract model for attacks on current multilinear map candidates. There are “hidden” variables X_1, \dots, X_n for some integer n , Z_1, \dots, Z_m for another integer m , and g . Then there are “public” variable Y_1, \dots, Y_m , which are set to $Y_i = q_i(\{X_j\}) + gZ_i$ for some polynomials q_i . All variables are defined over a field \mathbb{F} .

The adversary is allowed to make two types of queries:

- In a **Type 1** query, the adversary submits a “valid” polynomial p_k on the Y_i . Here “valid” polynomials come from some restricted set of polynomials. These restrictions are those that are enforceable using graded encodings. Next, we consider p as a polynomial of the formal variables X_j, Z_i, g . Write $p_k = p_k^{(0)}(\{X_j\}, \{Z_i\}) + gp_k^{(1)}(\{X_j\}, \{Z_i\}) + g^2 \dots$

If p_k is identically 0, then the adversary receives \perp in return. If $p_k^{(0)}$ is *not* identically 0, then the adversary receive \perp in return. If p_k is *not* 0 but $p_k^{(0)}$ is identically 0, then the adversary receives a handle to a new variable W_k , which is set to be $p_k/g = p_k^{(1)}(\{X_j\}, \{Z_i\}) + gp_k^{(3)}(\{X_j\}, \{Z_i\}) + \dots$

- In a **Type 2** query, the adversary is allowed to submit *arbitrary* polynomials r with small algebraic circuits on the W_k that it has seen so far. Consider $r(\{W_k\})$ as a polynomial of the variables X_j, Z_i, g , and write $r = r^{(0)}(\{X_j\}, \{Z_i\}) + gr^{(1)}(\{X_j\}, \{Z_i\}) + g^2 \dots$. If $r^{(0)}$ is identically zero, then the model responds with 0. Otherwise the model responds with 1.

In current graded encoding schemes, the set of “valid” polynomials is determined by the restrictions placed by the underlying set structure of the graded encoding. Here we consider a more abstract setting where the set of “valid” polynomials is arbitrary.

In the standard abstract model for graded encodings, **Type 1** queries output a bit as opposed to an algebraic element, and there are no **Type 2** queries. However, this model has been shown to improperly characterize the information received from **Type 1** queries in current candidate graded encoding schemes. The more refined model above more accurately captures the types of attacks that can be carried out on current graded encodings.

2.1 Obfuscation in the Abstract Model

We now describe an abstract obfuscation scheme that encompass the schemes of [AGIS14, BMSZ15], and can also be easily extended to incorporate the scheme of [BGK⁺14]. The obfuscator takes as input a branching program of length ℓ , input length n , and arity d . The branching program contains an input function $\text{inp} : [\ell] \rightarrow 2^{[n]}$ such that $|\text{inp}(i)| = d$ for all $i \in [\ell]$. Moreover, the branching program contains $2^d \ell + 2$ matrices $A_0, \{A_{i, S_i}\}_{i \in [\ell]}, A_{\ell+1}$ where S_i ranges over subsets of $\text{inp}(i)$, and $A_0 A_{\ell+1}$ are the “bookend” vectors. To evaluate a branching program on input x , we associate x with the set $T \subseteq [n]$ where $i \in T$ if and only if $x_i = 1$. To evaluate the branching program on input x (set T) compute the following product.

$$\mathbf{A}(T) = A_0 \times \prod_{i=1}^{\ell} A_{i, T \cap \text{inp}(i)} \times A_{\ell+1}$$

The output of the branching program is 0 if and only if $\mathbf{A}(T) = 0$.

The obfuscator first generates random matrices $\{R_i\}_{i \in [\ell+1]}$ and random scalars $\{\alpha_{i,S_i}\}_{i \in [\ell], S_i \subseteq \text{inp}(i)}$. Then it computes the randomized branching program consisting of the matrices $\widetilde{A}_{i,S_i} = \alpha_{i,S_i}(R_i \cdot A_{i,S_i} \cdot R_{i+1}^{adj})$ and bookend vectors $\widetilde{A}_0 \cdot R_1^{adj}$ and $\widetilde{A}_{\ell+1} = R_{\ell+1} \cdot A_{\ell+1}$. It is easy to see that this program computes the same function as the original branching program.

Finally, the obfuscator sets the “hidden” variables in the model to the \widetilde{A} matrices. Denote the “public” variables as $Y_{i,S} = \widetilde{A}_{i,S} + gZ_{i,S} = \alpha_{i,S}R_i \cdot A_{i,S} \cdot R_{i+1}^{adj} + gZ_{i,S}$. The set of valid **Type 1** polynomials is set up so that honest evaluations of the branching program are considered valid. That is, the polynomials

$$p_T = Y_0 \times \prod_{i=1}^{\ell} Y_{i, T \cap \text{inp}(i)} \times Y_{\ell+1}$$

are explicitly allowed. Notice that the g^0 coefficient of p_T is exactly $\widetilde{\mathbf{A}}(\mathbf{T}) \equiv \mathbf{A}(T)$, so the evaluator can run the program by querying on p_T , and checking if the result is \perp .

In the case $d = 1$, this obfuscator corresponds to the basic single-input branching program obfuscator of [BGK⁺14, AGIS14, BMSZ15]. In the more restricted model where there are no **Type 2** queries, it was shown how to set the underlying graded encodings so that the only valid **Type 1** queries are linear combinations of arbitrarily-many p_T polynomials. This is sufficient for indistinguishability obfuscation. When $d = 2$ this corresponds to the dual-input version of these obfuscators, in which it was shown how to set up the underlying graded encodings so that the linear combination has polynomial size. This is sufficient for virtual black box obfuscation (again in the more restricted model).

In any case, since for functionality the set of allowed queries *must* include honest executions of the program, we always allow queries on the p_T polynomials themselves. As such, our attacks will work by only making **Type 1** queries on honest evaluations of p_T . Thus with *any* restrictions in our abstract model that allow for such honest evaluations of p_T , we will demonstrate how to break indistinguishability security.

3 Abstract Attack

Here we describe an abstract attack on obfuscation in our generic model. For simplicity, we describe the attack for *single input* branching programs, which proves Theorem 1.1. We extend to dual-input and more generally d -input branching programs in Section 5, which proves Theorem 1.2.

3.1 The branching programs

The first branching program \mathbf{A} is defined as follows. It has $2n + 2$ layers, where the first and last layers consist of the row vector $A_0 := (0 \ 1)$ and the column vector $A_{2n+1} := (1 \ 0)^T$ respectively. The middle $2n$ layers scan through the input bits twice, once forward and once in reverse, with input selection function $\text{inp}(i) := \min(i, 2n + 1 - i)$ (so x_1 is read in layers 1 and $2n$, x_2 is read in layers 2 and $2n - 1$, etc.)¹. In each of these layers, both matrices are the identity, i.e. we have

$$A_{i,0} = A_{i,1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

¹Recall that in the single-input case, the set outputted by $\text{inp}(i)$ is just a singleton set.

for $i \in [2n]$. Here, we adopt the more standard notation for branching programs where the matrix $A_{i,b}$ is selected if $x_{\text{inp}(i)} = b$.

The branching program $\mathbf{A} = \{\text{inp}, A_0, A_{2n+1}, A_{i,b} \mid i \in [2n], b \in \{0, 1\}\}$ is evaluated in the usual way:

$$\mathbf{A}(x) := A_0 \times \prod_{i=1}^{2n} A_{i, x_{\text{inp}(i)}} \times A_{2n+1}.$$

Clearly this satisfies $\mathbf{A}(x) = 0$ for all x .

The second branching program $\mathbf{A}' = \{\text{inp}', A'_0, A'_{2n+1}, A'_{i,b} \mid i \in [2n], b \in \{0, 1\}\}$ is defined almost identically. The sole difference is that, in the layers reading bits any of the bits x_1, \dots, x_k for some integer $k \leq n$, the matrices corresponding to “ $x_i = 0$ ” are changed to be anti-diagonal. Namely, we have

$$A'_{i,0} = A'_{2n+1-i,0} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \text{ for } i \in [k]$$

and all other components remain the same (i.e. $\text{inp}' = \text{inp}$, $A'_0 = A_0$, $A'_{2n+1} = A_{2n+1}$, and $A'_{i,b} = A_{i,b}$ for all (i, b) where $b = 1$ or $i \in [k+1, 2n-k]$). We again have $\mathbf{A}'(x) = 0$ for all x , because the anti-diagonal matrix above is its own inverse and all the matrices commute.

3.2 The distinguishing attack

We now specialize the abstract obfuscation scheme from Section 2 to the single-input case. We choose invertible matrices $\{R_i \in \mathbb{Z}_p^{2 \times 2}\}_{i \in [2n+1]}$ and non-zero scalars $\{\alpha_{i,x} \in \mathbb{Z}_p\}_{i \in [2n], b \in \{0,1\}}$ uniformly at random. Next, we define

$$\tilde{A}_0 := A_0 \cdot R_1^{\text{adj}} \quad \tilde{A}_{2n+2} := R_{2n+1} \cdot A_{2n+1} \quad \tilde{A}_{i,b} := \alpha_{i,b} R_i \cdot A_{i,b} \cdot R_{i+1}^{\text{adj}}$$

for $i \in [2n]$, $b \in \{0, 1\}$, where R_i^{adj} is the adjugate matrix of R_i . Finally, each of the entries of the various \tilde{A} are what are actually encoded, meaning the “public” variables consist of

$$Y_{i,b} = \alpha_{i,b} R_i \cdot A_{i,b} \cdot R_{i+1}^{\text{adj}} + gZ_{i,b}$$

Next, by performing a change of variables on the $Z_{i,b}$, we can actually write

$$Y_{i,b} = \alpha_{i,b} R_i \cdot (A_{i,b} + gZ_{i,b}) \cdot R_{i+1}^{\text{adj}}$$

The underlying graded encodings guarantee some restrictions on the types of **Type 1** encodings allowed — however, the restrictions *must* allow evaluation of the branching program on various inputs. In particular, the query

$$p_x := Y_0 \times \prod_{i=1}^{2n} Y_{i, x_{\text{inp}(i)}} \times Y_{2n+1}$$

is allowed. Now, the coefficient of g^0 in p_x is given by

$$p_x^{(0)} := \tilde{A}_0 \times \prod_{i=1}^{2n} \widetilde{A_{i, x_{\text{inp}(i)}}} \times \widetilde{A_{2n+1}} = \rho \prod_i \alpha_{i, x_{\text{inp}(i)}} A_0 \times \prod_{i=1}^{2n} A_{i, x_{\text{inp}(i)}} \times A_{2n+1}$$

which evaluates to 0 by our choice of branching programs. Here $\rho := \prod_i \det(R_i)$ satisfies $\rho I = \prod_i R_i R_i^{adj}$, and we abuse notation by letting $Y_{0, x_{\text{inp}(0)}}$ denote Y_0 (and similarly for the other matrices).

Thus, the model, on **Type 1** query p_x , will return a handle to the variable

$$p_x^{(1)} := \rho \prod_i \alpha_{i, x_{\text{inp}(i)}} \sum_{i=1}^{2n} \left(\cdots A_{i-1, x_{\text{inp}(i-1)}} \cdot Z_{i, x_{\text{inp}(i)}} \cdot A_{i+1, x_{\text{inp}(i+1)}} \cdots \right)$$

As in Section 2, we will associate $x \in \{0, 1\}^n$ with sets $T \subset [n]$ where $i \in T$ if and only if $x_i = 1$. For $i \in [2, n]$, write $\alpha'_{i,b} = \alpha_{i,b} \alpha_{2n+1-i,b}$. Also set $\alpha'_{1,b} = \rho \alpha_{1,b} \alpha_{2n,b}$. Thus $\rho \prod_i \alpha_{i, x_{\text{inp}(i)}} = \prod_i \alpha'_{i, x_i}$. Define this quantity as $U_x = U_T$. It is straightforward to show that the U_T satisfy the following equation² when $|T| \geq 2$.

$$U_T = \prod_{S \subseteq T, |S| \leq 1} U_S^{\binom{|T|-|S|-1}{1-|S|} (-1)^{|S|+1}} = U_\emptyset^{-(|T|-1)} \cdot \prod_{j \in T} U_{\{j\}}$$

Moreover, any equation satisfied by the U_T is generated by these equations.

For the other part of $p_x^{(1)} = p_T^{(1)}$, there are two cases:

- The branching program is all-identity. Then $A_{i,0} = A_{i,1} =: A_i$. Here, we write $\beta_{i,b} = \cdots A_{i-1} \cdot Z_{i, x_{\text{inp}(i)}} \cdot A_{i+1} \cdots$. Notice that the $\beta_{i,b}$ are all independent. For $i \in [n]$, let $\beta'_{i,b} = \beta_{i,b} + \beta_{2n+1-i,b}$. Thus,

$$\sum_{i=1}^{2n} \left(\cdots A_{i-1, x_{\text{inp}(i-1)}} \cdot Z_{i, x_{\text{inp}(i)}} \cdot A_{i+1, x_{\text{inp}(i+1)}} \cdots \right) = \sum_{i=1}^n \beta'_{i, x_i}$$

Define this quantity as $V_x = V_T$. It is straightforward to show that the V_T satisfy the following equation when $|T| \geq 2$.

$$V_T = \sum_{S \subseteq T, |S| \leq 1} \left[\binom{|T|-|S|-1}{1-|S|} (-1)^{|S|+1} \right] V_S = -(|T|-1) V_\emptyset + \sum_{j \in T} V_{\{j\}}$$

Moreover, any equation satisfied by the V_T is generated by these equations.

Piecing together, we have that $p_T^{(1)} = U_T V_T$, where U_T, V_T satisfy the equations above.

- The branching program has reverse diagonals for $b = 0, i \leq k$. Consider a term $\cdots A_{i-1, x_{\text{inp}(i-1)}} \cdot Z_{i, x_{\text{inp}(i)}} \cdot A_{i+1, x_{\text{inp}(i+1)}} \cdots$. Suppose for the moment that $i \leq k+1$. Since each $A_{i,b}$ is either diagonal or anti-diagonal, we have that $\cdots A_{i-1, x_{\text{inp}(i-1)}} = \cdots A_{i-1, x_{i-1}}$ is equal to the row vector $(0 \ 1)$ if the parity of $x_{[1, i-1]}$ is zero, and is equal to $(1 \ 0)$ if the parity is 1. Similarly, $A_{i+1, x_{\text{inp}(i+1)}} \cdots$ is equal to the column vector $(1 \ 0)^T$ if the parity of $x_{[1, i-1]}$ is zero, and $(0 \ 1)^T$ otherwise³. Therefore, $\cdots A_{i-1, x_{\text{inp}(i-1)}} \cdot Z_{i, x_{\text{inp}(i)}} \cdot A_{i+1, x_{\text{inp}(i+1)}} \cdots$ is equal to $\left(Z_{i, x_{\text{inp}(i)}} \right)_{1,2}$ or $\left(Z_{i, x_{\text{inp}(i)}} \right)_{2,1}$, depending on the parity of $x_{[1, i-1]}$. Therefore, define $\gamma_{i,b,p}$ to be the result of the product when $x_i = b$ and the parity of $x_{[1, i-1]}$ is p . For $i \in [2n+k, 2n]$, the same holds, so we can absorb the product for this i into $\gamma_{i,b,p}$. For $i \in [k+2, 2n-k-1]$, the same holds

²See Section 5 for a proof of a more general case.

³The rest of the bits of x do not matter, since both matrices for each of these bits occur in the product $A_{i+1, x_{\text{inp}(i+1)}} \cdots$, and therefore cancel out.

true, except that it is only the parity of the bits $x_{[1,k]}$ that matter. Therefore, we can write the product as $\gamma_{i,b,p}$ where $x_i = b$ and the parity of $x_{[1,k]}$ is p . Notice that each of the $\gamma_{i,b,p}$ are independent.

Define

$$W_T = W_x = \sum_{i=1}^n \gamma_{i,x_i,\text{parity}(x_{[1,\min(i-1,k)]})}$$

Then we have that $p_T^{(1)} = U_T W_T$.

The W_T must satisfy *some* linear relationships, since the number of W is 2^n , but the number of γ is $4n$. We have not derived a general equation, but instead we will focus on two cases. If the bits x_1, \dots, x_k are fixed (say to 0), then the parity for these bits is always the same (0). Therefore, W_T for these T satisfy the same equations as the V_T . Thus, any equation satisfied by the $p_T^{(1)}$ for these T in the all-identity case will also be satisfied in the anti-diagonal case.

In the other case, take $T \subseteq \{1, 2, 3\}$, and suppose $k = 1$. In this simple case, it is straightforward to show that the following are the only linear relationships among these W :

$$\begin{aligned} W_{1,2,3} + W_1 &= W_{1,2} + W_{1,3} \\ W_{2,3} + W_\emptyset &= W_2 + W_3 \end{aligned}$$

These are different, and fewer, than the equations satisfied by the V_T . This will be the basis for our distinguishing attack.

To distinguish the two branching programs, it suffices to find a polynomial Q that annihilates the $p_T^{(1)}$ for $T \subseteq \{1, 2, 3\}$ in the all-identity case, but does not annihilate in the anti-identity case. Here is such a polynomial:

$$\begin{aligned} Q_{1,2,3} &= \left(p_\emptyset^{(1)} p_{1,2,3}^{(1)}\right)^2 + \left(p_1^{(1)} p_{2,3}^{(1)}\right)^2 + \left(p_2^{(1)} p_{1,3}^{(1)}\right)^2 + \left(p_3^{(1)} p_{1,2}^{(1)}\right)^2 \\ &\quad - 2 \left(p_\emptyset^{(1)} p_{1,2,3}^{(1)} p_1^{(1)} p_{2,3}^{(1)} + p_\emptyset^{(1)} p_{1,2,3}^{(1)} p_2^{(1)} p_{1,3}^{(1)} + p_\emptyset^{(1)} p_{1,2,3}^{(1)} p_3^{(1)} p_{1,2}^{(1)}\right) \\ &\quad + p_1^{(1)} p_{2,3}^{(1)} p_2^{(1)} p_{1,3}^{(1)} + p_1^{(1)} p_{2,3}^{(1)} p_3^{(1)} p_{2,3}^{(1)} + p_2^{(1)} p_{1,3}^{(1)} p_3^{(1)} p_{2,3}^{(1)} \\ &\quad + 4 \left(p_\emptyset^{(1)} p_{1,2}^{(1)} p_{1,3}^{(1)} p_{2,3}^{(1)} + p_{1,2,3}^{(1)} p_1^{(1)} p_2^{(1)} p_3^{(1)}\right) \end{aligned}$$

The fact that $Q_{1,2,3}$ annihilates in the all-identity case can be verified by tedious computation. The fact that it does *not* annihilate in the anti-diagonal case can also be verified by tedious computation as follows. Consider a generic degree 4 polynomial Q in the $p_T^{(1)}$ for $T \subseteq \{1, 2, 3\}$. The condition “ Q annihilates the $p_T^{(1)}$ ” can be expressed as a linear equation in the coefficients of Q . Since Q has degree 4 in 8 variables, the number of coefficients is bounded by a constant, so the linear constraints can be solved. The result of of this computation is that $Q = 0$ is the only solution.

By Schwartz-Zippel, if Q does not annihilate, then with overwhelming probability over the randomness of the obfuscation, the result of applying Q is non-zero.

The attack thus works as follows. First query on inputs x which are zero in every location except the first three bits. Since the branching program always evaluates to zero, the model will

return a handle to the element $p_T^{(1)}$, where $T \subseteq \{1, 2, 3\}$ is the set of bits where x is 1. Then, evaluate the polynomial $Q_{1,2,3}$ on the elements obtained. If the result is 0, then guess that we are in the all-identity case. If the result is non-zero, then guess that we are in the anti-diagonal case. As we have shown, this attack distinguishes the two cases with overwhelming probability.

We make one final observation that will be relevant for attacking the specific [GGH13a] candidate. We note that, for either branching program, the following is true. Let T_0 be some subset of $[k+1, n]$ of size 3, and write $T_0 = i_1, i_2, i_3$. Let T_1 some subset of $[1, n] \setminus T_0$. Then for any subset $T \subseteq [3]$, write $\hat{p}_T^{(1)} := p_{T'}^{(1)}$, where $T' = \{i : i \in T_1 \text{ or } i = i_j \text{ for some } j \in T\}$. If we then evaluate the above polynomial $Q_{1,2,3}$ over the $\hat{p}_T^{(1)}$, we see that it annihilates. This is because the corresponding $p_{T'}^{(1)}$ satisfy the same equations as above.

4 Attack on GGH13 Encodings

In this section we explain how the abstract attack above extends to actual obfuscation schemes [BGK⁺14, AGIS14, BMSZ15] when implemented with [GGH13a] multilinear maps. At a high level, this is done by implementing **Type 1** and **Type 2** queries ourselves, without the help of the abstract model's oracle.

Implementing **Type 1** queries is straightforward: for any honestly executed 0-output of the program, namely an encoding

$$p_x = \left[\left(p_x^{(0)}(\{X_j\}, \{Z_i\}) + gp_x^{(1)}(\{X_j\}, \{Z_i\}) + g^2 \dots \right) / z^k \right]_q$$

with $p_x^{(0)}(\{X_j\}, \{Z_i\}) = 0$, we can multiply by the zero-testing parameter $\mathbf{p}_{zt} = [hz^k/g]_q$ to obtain

$$W_x := [p_x \cdot \mathbf{p}_{zt}]_q = h \cdot \left(p_x^{(1)}(\{X_j\}, \{Z_i\}) + gp_x^{(2)}(\{X_j\}, \{Z_i\}) + g^2 \dots \right) \quad (4.1)$$

This differs from what is returned in the abstract attack because of the factor h . To handle this, we ensure that our annihilating polynomial Q is *homogeneous*, and thus $Q(\{h \cdot p_x^{(1)}(\{X_j\}, \{Z_i\})\}_x) = 0$ whenever $Q(\{p_x^{(1)}(\{X_j\}, \{Z_i\})\}_x) = 0$. (Lemma 5.3 in fact shows we can assume Q is homogeneous without loss of generality, because the $p_x^{(1)}$ are all homogeneous and of the same degree.)

To implement **Type 2** queries, we must check whether a given polynomial Q over $\{W_x\}_{x \in S}$ (for some $S \subseteq \{0, 1\}^n$) is an annihilating polynomial, i.e. whether $Q(\{h \cdot p_x^{(1)}(\{X_j\}, \{Z_i\})\}_{x \in S}) = 0$. To do this we observe that, for any such Q , $Q(\{W_x\}_{x \in S})$ produces a ring element in the ideal $\langle hg \rangle$. So, we compute many such elements $v_i = Q_i(\{W_x\}_{x \in S_i})$, where Q_i is the (homogeneous) polynomial that annihilates $\{p_x^{(1)}(\{X_j\}, \{Z_i\})\}_{x \in S_i}$ when the encodings were formed by obfuscating the all-identity branching program. More specifically, we compute enough v_i to (heuristically) form a basis of $\langle hg \rangle$. Then, we compute one more element v^* which is either in $\langle hg \rangle$ or not depending on which branching program was obfuscated, and finally we use the $\langle hg \rangle$ -basis to test this.

4.1 The attack

We use essentially the same pair of branching programs \mathbf{A}, \mathbf{A}' that were used in the abstract attack (see Section 3.1): \mathbf{A} consists of all identity matrices, while in \mathbf{A}' the two matrices corresponding to $x_1 = 0$ are changed to be anti-diagonal.

Let \mathcal{O} denote the obfuscator described in Section 2.1. This obfuscator is exactly the one from [BMSZ15], with two exceptions. First, it operates on a branching program reading only one bit per layer, while in [BMSZ15] the branching programs read two bits per layer. In Section 5, we show that our abstract attack, and thus also the concrete attack described here, extends to the dual-input setting. (In fact, we show that it extends to arity- d branching programs for any constant d .) Second, equation (4.1) (and the presence of z^k in \mathbf{p}_{zt}) assumes that all encodings output by \mathcal{O} are level-1 GGH encodings, while in [BMSZ15] a more complicated level structure is used (following [BGK⁺14, MSW14]). However, since our attack only uses these encodings to honestly execute the obfuscated program, (4.1) holds even for this level structure.

Here is our attack:

- Let $m = n^{O(1)}$ be the dimension of the underlying encodings (this is a parameter of the [GGH13a] scheme). Note that any m linearly independent elements of $\langle hg \rangle$ form a basis for $\langle hg \rangle$. Let $m' \gg m$ be an integer.
- Repeat the following for $t = 1, \dots, m'$:
 - Choose a random size-3 subset $T_0 = \{i_1, i_2, i_3\} \subseteq [n]$ that *does not* contain 1. T_0 will correspond to the set of input bits that we vary.
 - Choose a random subset $T_1 \subseteq ([n] \setminus T_0)$. T_1 will correspond to a fixing of the bits outside T_0 .
 - For each $T \subseteq [3]$,
 - * let $x_T \in \{0, 1\}^n$ be the string such that $x_i = 1$ if and only if either $i \in T_1$, or $i = i_j$ for some $j \in T$ (recall that $T_0 = \{i_1, i_2, i_3\}$).
 - * Run the obfuscated program on input x , until the zero test query. Let $p_T^{(1)}$ be the vector obtained from zero testing.
 - Evaluate the polynomial $Q_{1,2,3}$ in Section 3 on the $p_T^{(1)}$. Let the output be defined as v_t . That is, we let x_T vary over the the 8 possible values obtained by fixing all the input bits outside of T_0 , run the obfuscated program on each of the x_T , and then evaluate the polynomial $Q_{1,2,3}$ on the results to get v_t .
- Find a linearly independent subset V of the v_t .
- Choose a random size-3 subset $T_0^* = \{i_1, i_2, i_3\} \subseteq [n]$ that *does* contain 1. For each $T \subseteq [3]$, compute $p_T^{(1)}$ as above. Then evaluate the polynomial $Q_{1,2,3}$ on the $p_T^{(1)}$ to obtain a vector v^* .
- Finally, test if v^* is in the span of V . If it is, output 1. Otherwise, output 0.

Analysis of our attack. As in Section 3, let $T_0 \subseteq [n]$, and choose an arbitrary fixing of the remaining bits. Suppose we evaluate the branching program on the 8 different inputs corresponding to varying the bits in T_0 , and then run the polynomial $Q_{1,2,3}$ on the results. Then $Q_{1,2,3}$ annihilates in either of the following cases:

- T_0 does not contain 1.
- The branching program is the all-identity program, even if T_0 contains 1.

Therefore, we see that $Q_{1,2,3}$ annihilates for each $t = 1, \dots, m'$. In the case of [GGH13a], $Q_{1,2,3}$ annihilating means that the resulting vector v is an element of the ideal $\langle hg \rangle$.

Thus, each of the v_t are elements in the ideal, regardless of the branching program. We will heuristically assume that the v_t span the entire ideal. This is plausible since the number m' of v_t is much larger than the dimension of the ideal. Increasing m' relative to m should increase the likelihood of the heuristic being true.

For v^* , however, things are different. v^* is in the ideal if the branching program is the all-identity, but outside the ideal (with high probability) if the branching program has anti-diagonals, since in this case $Q_{1,2,3}$ does not annihilate. Therefore, our test for v^* being linearly independent from v will determine which branching program we were given.

5 Beyond Single-Input Branching Programs

In this section, we show an abstract attack on dual-input branching programs. More generally, we show that generalizing to d -input branching programs for any constant d will not prevent the attack.

We first recall our semantics of branching programs in the general d -ary setting. Fix integers d, ℓ and n which respectively correspond to the number of bits read by each layer of the branching program, the length of the branching program, and the input length. Let $\text{inp} : [\ell] \rightarrow 2^{[n]}$ be any function such that $|\text{inp}(i)| = d$ for all $i \in [\ell]$. A branching program of length ℓ then consists of $2^d \ell + 2$ matrices $A_0, \{A_{i,S_i}\}_{i \in [\ell]}, A_{\ell+1}$ where S_i ranges over subsets of $\text{inp}(i)$, and $A_0 A_{\ell+1}$ are the “bookend” vectors.

We associate an input x with the subset $T \in 2^{[n]}$ of indices where x is 1. To evaluate the branching program on input x (set T) compute the product

$$\mathbf{A}(T) = A_0 \times \prod_{i=1}^{\ell} A_{i, T \cap \text{inp}(i)} \times A_{\ell+1}$$

Consider the obfuscation of the branching program. Let R_i be the Kilian randomizing matrices. Let $\alpha_{i,S}$ be the extra randomization terms. Then the encoded values seen by the adversary are the matrices $Y_{i,S} = \alpha_{i,S} R_i \cdot A_{i,S} \cdot R_{i+1}^{\text{adj}} + gZ_{i,S}$

By performing a change of variables on the $Z_{i,S}$, we can actually write $Y_{i,S} = \alpha_{i,S} R_i \cdot (A_{i,S} + gZ_{i,S}) \cdot R_{i+1}^{\text{adj}}$

The encodings will guarantee some restrictions on the **Type 1** queries allowed — however they must allow evaluation of the branching program. Thus we assume that the following query is allowed for every $T \subseteq [n]$.

$$p_T = Y_0 \times \prod_{i=1}^{\ell} Y_{i, T \cap \text{inp}(i)} \times Y_{\ell+1}$$

Now we will assume a trivial branching program where (1) within each layer, all matrices are the same ($A_{i,S_i} = A_{i,S'_i}$ for any $S_i, S'_i \in \text{inp}(i)$), so in particular the program is constant, and (2), the branching program evaluates to 0 on all inputs. Therefore, the g^0 coefficient in p_T will evaluate to zero everywhere. Thus, a **Type 1** query will output a handle to the variable

$$p_T^{(1)} = \rho \left(\prod_i \alpha_{i, S \cap \text{inp}(i)} \right) \sum_i \left(\cdots A_{i, T \cap \text{inp}(i-1)} \cdot Z_{i, T \cap \text{inp}(i)} \cdot A_{i+1, T \cap \text{inp}(i+1)} \cdots \right)$$

For any sets $S' \subseteq S \subseteq [n]$ with $|S| = d$, define

$$\alpha_{S, S'} := \prod_{i: \text{inp}(i)=S} \alpha_{i, S'} \quad \beta_{S, S'} := \sum_{i: \text{inp}(i)=S} \beta_{i, S'}$$

and for any set $T \subseteq [n]$, define

$$U_T := \prod_{S: |S|=d} \alpha_{S, T \cap S} \quad V_T := \sum_{S: |S|=d} \beta_{S, T \cap S}$$

Then we have that $p_T^{(1)} = U_T V_T$.

The following theorem shows that, for $|T| > d$, U_T and V_T can each be written as rational polynomials in the variables $U_{T'}, V_{T'}$ for $|T'| \leq d$.

Theorem 5.1. *Let $T \subseteq [n]$ with $|T| > d$. Then,*

$$U_T = \prod_{T' \subseteq T: |T'| \leq d} U_{T'}^{(-1)^{d-|T'|} \binom{|T|-|T'|}{d-|T'|}}$$

and

$$V_T = \sum_{T' \subseteq T: |T'| \leq d} (-1)^{d-|T'|} \binom{|T|-|T'|}{d-|T'|} V_{T'}.$$

Proof. We prove this equation for V_T , the proof for U_T is analogous. Consider expanding the left and right sides of the equation in terms of the $\beta_{S, Z}$ and equating the coefficients of $\beta_{S, Z}$ on both sides, we see that the following claim suffices to prove the theorem:

Claim 5.2. *For any sets T, S, Z ,*

$$\sum_{T' \subseteq T: |T'| \leq d, T' \cap S = Z} \binom{|T|-|T'|}{d-|T'|} (-1)^{d-|T'|} = \begin{cases} 1 & \text{if } T \cap S = Z \\ 0 & \text{if } T \cap S \neq Z \end{cases}$$

The left hand side (resp. right hand side) of the above equation corresponds to the coefficient of $\beta_{S, Z}$ in the right hand side (resp. left hand side) of the V equation in Theorem 5.1. Hence the theorem follows from the claim.

We now prove the claim. First, suppose $Z \not\subseteq T \cap S$. Then the sum on the right is empty, so the result is zero, as desired. Next, suppose $Z \subseteq T \cap S$. Then for any T' in the sum, we can write $T' = Z \cup T''$ where $T'' \subseteq T \setminus (S \cup Z)$ and $|T''| \leq d - |Z|$. Therefore, we can think of the sum as being over T'' . The number of T'' of size i is $\binom{|T \setminus (S \cup Z)|}{i}$. Therefore, the sum on the left is equal to

$$\sum_{i=0}^{d-|Z|} \binom{|T \setminus (S \cup Z)|}{i} \binom{|T|-|Z|-i}{d-|Z|-i} (-1)^{d-i-|Z|}$$

Let $e = d - |Z|$, $t = |T| - |Z| = |T \setminus Z|$ (since $Z \subseteq (T \cap S) \subseteq T$), and $k = |T \setminus (S \cup Z)|$. Notice that $k \leq t$, and that $k = t$ if and only $Z = T \cap S$. Thus, we need to show that

$$\sum_{i=0}^e \binom{k}{i} \binom{t-i-1}{e-i} (-1)^{e-i} = \begin{cases} 1 & \text{if } k = t \\ 0 & \text{if } k < t \end{cases}$$

First, we use the identity $(-1)^s \binom{s-r-1}{s} = \binom{r}{e-i}$ with $s = e-i$ and $r = e-t$ to replace $\binom{t-i-1}{e-i} (-1)^{e-i}$ with $\binom{e-t}{e-i}$ (note that the binomial coefficients are defined for negative integers such as $e-t$).

Then we have that the left hand side becomes $\sum_{i=0}^e \binom{k}{i} \binom{e-t}{e-i}$. The Chu-Vandermonde identity shows that this is equal to $\binom{k+(e-t)}{e} = \binom{e-(t-k)}{e}$. Notice that if $t = k$, the result is 1. Moreover, if $k < t$, then the upper index of the binomial is less than the bottom index, so the result is 0. This proves the claim and hence the theorem. \square

Annihilating polynomial for $p_T^{(1)}$. We now describe our abstract attack using annihilating polynomials. The first step is to argue that it is possible to efficiently devise a non-zero polynomial Q on several of the $p_T^{(1)}$ such that Q is identically zero when the $p_T^{(1)}$ come from the obfuscation. In particular, we need Q to be identically zero as a polynomial over the α 's and β 's. Using Theorem 5.1, it suffices to find Q that is identically zero as a rational function over the U_T, V_T for $|T| \leq d$.

We will first consider the values $p_T^{(1)}$ as polynomials in the $V_T, U_T, |T| \leq d$ over the rationals. Let $k = 2d + 2$, and consider all $p_T^{(1)}$ for $T \subseteq [k]$. Then each $p_T^{(1)}$ is a rational function of the U_T, V_T for $T \subseteq [k], |T| \leq d$. There are $\sum_{i=0}^d \binom{k}{i} < 2^{2d+1}$ such T , and therefore fewer than 2^{2d+2} such U_T, V_T . Yet there are 2^{2d+2} different $p_T^{(1)}$ for $T \subseteq [k]$ of arbitrary size. Thus, there must be some algebraic dependence among the $p_T^{(1)}$. Notice moreover that the expression for $p_T^{(1)}, T \subseteq [k]$ in terms of the $U_{T'}, V_{T'}, T' \subseteq [k], |T'| \leq d$ are fixed rational functions with integer coefficients, independent of the branching program, n , or ℓ ; the only dependence is on d . Recall that we are taking d to be a constant, so the number of $p_T^{(1)}, V_{T'}, U_{T'}$ and the coefficients in the relation between them are all constants. Therefore, there is a fixed polynomial Q_d in the $p_T^{(1)}$ over the rationals such that Q_d is identically zero when the $p_T^{(1)}$ come from obfuscation.

We note that by a more tedious argument, it is actually possible to show there must be an algebraic dependence among the $p_T^{(1)}$, and hence an annihilating polynomial for them, when T varies over the subsets of $[k]$ for $k = 2d + 1$ (as opposed to $2d + 2$).

By multiplying by the LCM of the denominators of the rational coefficients, we can assume without loss of generality that Q_d has integer coefficients. Therefore, there is a fixed integer polynomial Q_d such that $Q_d(p_T^{(1)})$ is identically 0. Since the coefficients are integers, this polynomial actually also applies in any field or ring; we just need to verify that it is not identically zero in the field/ring. This will be true as long as the characteristic of the ring is larger than the largest of the coefficients. Since in our case, the ring characteristic grows (exponentially) with the security parameter, for high enough security parameter, the polynomial Q_d will be non-zero over the ring.

Computing the annihilating polynomial Q_d . In Section 3, we gave an annihilating polynomial for the case $d = 1$. For more general d , we do not know a more general expression. However, we still argue that such a Q_d can be efficiently found for any d :

- The polynomial Q_d is just a fixed polynomial over the integers; in particular it has a constant-sized description for constant d . Thus, we can assume that Q_d is simply given to the adversary.
- If we want to actually compute Q_d , this is possible using linear algebra. Using degree bounds for the annihilating polynomial due to [Kay09], we can determine an upper bound t on the degree of Q_d . Then, the statement “ Q_d annihilates the $p_T^{(1)}$ ” can be expressed as a system of linear equations in the coefficients of Q_d , where the equations themselves are determined by expressions for $p_T^{(1)}$ in terms of the U_{T^r}, V_{T^r} . By solving this system of linear equations, it is possible to obtain a polynomial Q_d . We note that, for constant d , t will be constant, the system of linear constraints will be constant, and hence it will take constant time to compute Q_d . In terms of d , the running time is necessarily exponential (since the number of variables $p_T^{(1)}$ is exponential).

The following lemma shows that we can take Q to be a homogeneous polynomial, which will be necessary for obtaining an attack over [GGH13a].

Lemma 5.3. *Let p_1, \dots, p_k be homogeneous polynomials each of the same degree d . Let Q be any polynomial that annihilates $\{p_i\}_i$, and let Q_r denote the homogeneous degree- r part of Q . Then Q_r annihilates $\{p_i\}_i$ for each $r \leq \deg(Q)$.*

Proof. If $Q_r(\{p_i\}_i) \neq 0$ for some $r \leq \deg(Q)$, then Q_r contains some degree- dr monomial m . Then because $\sum_{r=0}^{\deg(Q)} Q_r(\{p_i\}_i) = Q(\{p_i\}_i) = 0$, some $Q^{(r')}$ for $r' \neq r$ must contain the monomial $-m$. However, since $Q^{(r')}$ is homogeneous of degree $dr' \neq dr$, this is a contradiction. \square

Completing the attack. Using the annihilating polynomial above, we immediately get an attack on the abstract model of obfuscation. The attack distinguishes the trivial branching program where all matrices across each layer are the same, from a more general all-zeros branching program that always outputs zero, but has a non-trivial branching program structure.

The attack proceeds as follows: query the model on **Type 1** queries for all p_T as T ranges over the subsets of $[k]$. Since the branching program always outputs 0, the model will return a handle to the $p_T^{(1)}$ polynomials. Then evaluate the annihilating polynomial Q_d above on the obtained $p_T^{(1)}$. If the result is non-zero (as will be the case for many non-trivial branching programs), then we know the branching program was *not* the trivial branching program. In contrast, if the result *is* zero, then we can safely guess that we are in the trivial branching program case. Hence, we breach the indistinguishability security of the obfuscator.

References

- [AGIS14] Prabhanjan Ananth, Divya Gupta, Yuval Ishai, and Amit Sahai. Optimizing obfuscation: Avoiding barrington’s theorem. Cryptology ePrint Archive, Report 2014/222, 2014. <http://eprint.iacr.org/>.
- [BGK⁺14] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. In *Proc. of EuroCrypt*, 2014.
- [BMSZ15] Saikrishna Badrinarayanan, Eric Miles, Amit Sahai, and Mark Zhandry. Post-zeroizing obfuscation: The case of evasive circuits. Cryptology ePrint Archive, Report 2015/167, 2015. <http://eprint.iacr.org/>.

- [BR14] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In *TCC*, 2014.
- [CGH⁺15] Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancreède 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 *CRYPTO*, 2015.
- [GGH13a] S Garg, Craig Gentry, and S Halevi. Candidate multilinear maps from ideal lattices. *Advances in Cryptology (EuroCrypt)*, 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 *Proc. of FOCS*, 2013.
- [HJ15] Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. *IACR Cryptology ePrint Archive*, 2015:301, 2015.
- [Kay09] Neeraj Kayal. The complexity of the annihilating polynomial. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity, CCC 2009, Paris, France, 15-18 July 2009*, pages 184–193, 2009.
- [MSW14] Eric Miles, Amit Sahai, and Mor Weiss. Protecting obfuscation against arithmetic attacks. *IACR Cryptology ePrint Archive*, 2014:878, 2014.