

# Post-Zeroizing Obfuscation: new mathematical tools, and the case of evasive circuits\*

Saikrishna Badrinarayanan  
UCLA  
saikrishna@cs.ucla.edu

Eric Miles  
UCLA  
enmiles@cs.ucla.edu

Amit Sahai<sup>†</sup>  
UCLA  
sahai@cs.ucla.edu

Mark Zhandry  
Stanford University and MIT  
mzhandry@gmail.com

## Abstract

Recent devastating attacks by Cheon et al. [Eurocrypt'15] and others have highlighted significant gaps in our intuition about security in candidate multilinear map schemes, and in candidate obfuscators that use them. The new attacks, and some that were previously known, are typically called “zeroizing” attacks because they all crucially rely on the ability of the adversary to create encodings of 0.

In this work, we initiate the study of *post-zeroizing obfuscation*, and we obtain a key new mathematical tool to understand security in a post-zeroizing world. Our new mathematical tool allows for analyzing polynomials constructed by the adversary when given encodings of randomized matrices arising from a general matrix branching program. This technique shows that the types of encodings an adversary can create are much more restricted than was previously known, and is a crucial step toward achieving post-zeroizing security. We also believe the technique is of independent interest, as it yields efficiency improvements for existing schemes – efficiency improvements that have already found application in other settings.

Finally, we show how to apply our new mathematical tool to the special case of evasive functions. We show that our obfuscator survives *all known attacks* on the underlying multilinear maps, by proving that no encodings of 0 can be created by a generic-model adversary. Previous obfuscators (for both evasive and general functions) were either analyzed in a less-conservative “pre-zeroizing” model that *does not* capture recent attacks, or were proved secure relative to assumptions that no longer have any plausible instantiation due to zeroizing attacks.

---

\*This paper subsumes a previous work of Sahai and Zhandry [SZ14].

<sup>†</sup>Supported in part by a DARPA/ONR PROCEED award, NSF grants 1228984, 1136174, 1118096, 1065276, 0916574 and 0830803, 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 U.S. Office of Naval Research under Contract N00014-11-1-0389. 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

Over the past three years, *all* candidate constructions [GGH13a, CLT13, GGH15, CLT15, Hal15] of multilinear maps, also called graded encoding schemes, have been shown to suffer from “zeroizing” attacks [GGH13a, CHL<sup>+</sup>15, BWZ14, CGH<sup>+</sup>15, HJ15, BGH<sup>+</sup>15, Hal15, CLR15, MF15] — and these attacks have in many cases been devastating.

Given this state of affairs, one would expect that the most-studied application of graded encodings schemes – indistinguishability obfuscation [BGI<sup>+</sup>01, GGH<sup>+</sup>13b] – would be similarly devastated. However, quite surprisingly, *none* of the zeroizing attacks discovered so far have placed current obfuscation schemes over prime-order<sup>1</sup> graded encodings in jeopardy. In this paper, we ask: *Why is this the case?* Given the profound level of interest in obfuscation over the past two years [GGH<sup>+</sup>13b, BR14, BGK<sup>+</sup>14, PST14a, GLSW14, AGIS14, Zim14, MSW14, AB15, SW14, GGG<sup>+</sup>14, GGHR14, BZ14], and given that so far all proposed obfuscation schemes rely on graded encoding schemes, we believe this question is of paramount importance. And indeed, before our work, no security analysis for obfuscation used a model or assumption that took into account the impact of zeroizing attacks.

**Our Approach.** While the primary goal of this paper is to answer the question, “Why are there no attacks on obfuscation schemes?”, there is a greater purpose as well. This paper seeks to initiate a research program whose aim is to build fully secure obfuscation schemes out of *weakened* graded encodings schemes – graded encoding schemes that are subject to zeroizing attacks. As the research cycle of construct-and-attack over the past 3 years has shown, building fully secure graded encoding schemes is a challenging task. Thus, our approach is to take a pessimistic view and see if, in fact, even weakened forms of graded encoding schemes suffice for constructing fully secure obfuscation. We note that even if future constructions of graded encoding schemes are successful in avoiding zeroizing attacks, the research program initiated by our work would still be valuable, because it will help to identify the minimal security properties actually needed by graded encoding schemes to achieve secure obfuscation. This could lead to greater efficiency.

While the future may be brighter, the current situation is dire. Because zeroizing attacks on composite-order graded encoding schemes [CLT13, CHL<sup>+</sup>15, BWZ14, CGH<sup>+</sup>15, CLT15, CLR15, MF15] have generally been more devastating than zeroizing attacks on prime-order graded encoding schemes [GGH13a, CGH<sup>+</sup>15, HJ15], in this work we will focus on understanding the security of obfuscation schemes in the prime-order setting. In particular, we will focus on variants of the original family of obfuscators from [GGH<sup>+</sup>13b, BR14, BGK<sup>+</sup>14, AGIS14].

The central contribution of our work is a new mathematical tool that characterizes when an adversary can set up the most basic requirement for a zeroizing attack, namely a top-level encoding of zero. Furthermore, we present this mathematical tool in a very general form, which even has consequences for efficiency of prime-order obfuscation constructions. We believe that our characterization lemma will prove valuable in the long-term study of building secure obfuscation from

---

<sup>1</sup> We note that certain *simplified* versions of obfuscation schemes over composite-order graded encoding schemes [Zim14] have been broken by zeroizing attacks [CGH<sup>+</sup>15], but no such attacks exist over prime-order graded encoding schemes, which are the focus of this paper. There are no published methods for converting the composite-order obfuscation schemes of [Zim14, AB15] to the prime-order setting. Furthermore, zeroizing attacks over prime-order graded encoding schemes typically apply when the multiplicity of zero achieved at the top level is greater than zero, and if a prime-order conversion was attempted, care would need to be taken to ensure that such higher multiplicities do not occur.

weak graded encodings. We demonstrate this by applying our lemma to the case of evasive circuits, for which we can show security even using only extremely weak graded encodings whose security completely breaks down when a top-level encoding of zero is found.

**Background - Obfuscation.** Obfuscation is a cryptographic tool that offers a powerful capability: software that can keep a secret. That is, consider a piece of software that makes use of a secret to perform its computation. Obfuscation allows us to transform this software so that it can be run *publicly*: anyone can obtain the full code of the program, run it, and see its outputs, but no one can learn anything about the embedded secret, beyond what can be learned by examining the outputs of the program.

The first candidate construction for a general-purpose obfuscator was given by Garg, Gentry, Halevi, Raykova, Sahai, and Waters [GGH<sup>+</sup>13b]. This construction, and all subsequent works constructing candidate obfuscators [BR14, BGK<sup>+</sup>14, PST14a, GLSW14, AGIS14, Zim14, MSW14, AB15], are built on top of another cryptographic primitive called a graded encoding scheme. In a graded encoding scheme, plaintext elements are encoded at various levels, and can be added and multiplied subject to algebraic restrictions relating to these levels. Further, there is a “top” level at which one can test whether an element encodes 0.

**Background - Zeroizing Attacks.** The many zeroizing attacks differ somewhat in their details, but each attack obeys the algebraic restrictions imposed by the graded encodings schemes, and critically they all share the need to create top-level 0-encodings. Indeed, many such encodings are needed for each attack, and the attacks require that these encodings have further structure.

Several of the works constructing candidate obfuscators prove security in an idealized “generic multilinear model” that seeks to capture the algebraic restrictions imposed by the graded encoding scheme candidates. However, the known zeroizing attacks use extra information that is provided by the zero-testing procedure, which is not captured in the standard generic model. Thus, a proof of security in the generic multilinear model by itself is no longer a persuasive argument of security. In particular, it is now crucial to gain a better understanding of exactly what types of top-level 0-encodings can be constructed in prime-order graded encoding schemes.

**Our contribution.** We introduce a new mathematical tool for analyzing how an adversary can create 0-encodings given a set of randomized matrices. This tool both shows that the types of 0-encodings an adversary can create are much more restricted than was previously known, and that the adversary’s behavior can be so controlled in a much richer set of circumstances than was previously known. We stress that this new tool, Theorem 4.2, was not present in any previous work, and allows for a much more fine-grained analysis of the adversary’s behavior in prime-order settings than was previously available.

Briefly, we first consider an obfuscator  $\mathcal{O}$  that can use much wider class of matrix branching programs than was previously known, most notably this class includes matrix branching programs that involve low-rank matrices. Theorem 4.2 shows that any polynomial  $p$  over (the encodings in) the obfuscation  $\mathcal{O}(f)$  can be efficiently mapped to a poly-size set of inputs  $X$  such that  $p$  evaluates to an encoding of 0 if and only if *every*  $x \in X$  satisfies  $f(x) = 0$ . For context, previous works both could not handle the case of low-rank matrices, and only gave a map that allowed the evaluation of  $p$  to be *simulated* given the set  $\{f(x) \mid x \in X\}$ , but did not show the stronger precise characterization

of 0-encodings that we obtain. We stress that we do not know of any simpler way of obtaining such a characterization.

We now elaborate on Theorem 4.2, how it is proved, and how previous works that did not consider zeroizing attacks did not need and did not achieve such a theorem. Following that, we mention two applications of this new theorem, namely improving the efficiency of obfuscation, and obfuscating *evasive* functions in a model that captures all known attacks on graded encoding schemes.

## 1.1 Our Techniques

As stated above, the main technical challenge in our paper is to show that any polynomial  $p$  over the obfuscation  $\mathcal{O}(f)$  can be efficiently mapped to a set of inputs  $X$  such that  $p$  evaluates to an encoding of 0 if and only if every  $x \in X$  satisfies  $f(x) = 0$ .

One ingredient in our paper is the notion of strong straddling sets from [MSW14], as this tool allows us to show that low-level encodings of 0 can be efficiently transformed into top-level encodings of 0. Thus, the only obstacle that remains is to prove Theorem 4.2 for top-level encodings.

**The technical barrier – Kilian’s statistical simulation.** Before we proceed to provide intuition about our proof, let us consider the technical roots of how security was shown in previous works. In every paper constructing secure obfuscation for matrix branching programs so far [GGH<sup>+</sup>13b, BR14, BGK<sup>+</sup>14, PST14a, GLSW14, AGIS14, MSW14] and in every different model that has been considered, one theorem has played a starring role in all security analyses: Kilian’s statistical simulation theorem [Kil88]. As relevant here, Kilian’s theorem considers the setting where we randomize each matrix in a sequence of matrices as follows:

$$\widehat{\mathbf{B}}_i = \mathbf{R}_{i-1}^{-1} \mathbf{B}_i \mathbf{R}_i$$

where  $\mathbf{R}_i$  are random invertible matrices for  $i \in [\ell - 1]$ , and identity otherwise. Note that this randomization does not affect the iterated product. Then, for any particular input  $x$ , if the iterated product is  $M$ , Kilian’s theorem states that we can statistically simulate the collection of matrices  $\{\widehat{\mathbf{B}}_i\}_{i \in [\ell]}$  knowing only  $M$  but with no knowledge of the original matrices  $\{\mathbf{B}_i\}$ .

Kilian’s statistical simulation theorem has been a keystone in all previous analyses of obfuscation: in one way or another, all previous security analyses for obfuscation methods have found some way to isolate the adversary’s view of the obfuscation to a single input. Once this isolation was accomplished, Kilian’s theorem provided the assurance that the adversary’s view of the obfuscation, as it related to this single input, only encoded information about the output of the computation within  $M$ , and nothing more.

However, Kilian’s statistical simulation theorem only allows for simulation. It does not rule out the possibility that an encoding of 0 may result no matter what the function outputs on the input in question. Indeed, it is not hard to construct an obfuscator that is secure in the generic model but allows for encodings of 0 even when the function being obfuscated always outputs 1. Moreover, Kilian’s theorem only applies when the branching program matrices are full-rank. Indeed, if the matrices are allowed to be arbitrary rank, then it is *impossible* to simulate each of the matrices just given the product  $M$ , as there is no way to determine what the rank of each matrix should be, nor the ranks of various subproducts of matrices. (In the next subsection, we discuss the efficiency benefits of allowing low-rank matrices.) Because of this impossibility, we know of no way to generalize Kilian’s theorem or its proof to obtain our theorem.

**Our Approach.** To obtain our result, we directly analyze what kinds of polynomials an adversary can generate using multilinear operations. We model the multilinear setting as follows. There is a universe set  $[\ell]$ , and for every subset  $S \subseteq [\ell]$ , we have a copy of  $\mathbb{Z}_q$  that we name  $G_S$ . Then, the adversary has access to the following operations:

- **ADD:**  $G_S \times G_S \rightarrow G_S$ , for every subset  $S \subseteq [\ell]$ .
- **MULT:**  $G_S \times G_T \rightarrow G_{S \cup T}$ , for every pair  $S, T \subseteq [\ell] : S \cap T = \emptyset$ .
- **ZEROTEST:**  $G_{[\ell]} \rightarrow \{\text{TRUE}, \text{FALSE}\}$ .

This is sometimes called the “asymmetric” multilinear setting, is natively supported by known instantiations of prime-order graded encoding schemes [GGH13a], and was used in previous works. Observe that in this setting, if the adversary is given a matrix entirely encoded in  $G_{\{1\}}$ , then for example it is not possible for it to compute the rank of this matrix. This is because no two entries within this matrix can be multiplied together, since they both reside in the same group  $G_{\{1\}}$ , and multiplication is only possible across elements of groups corresponding to disjoint index sets.

Even though we do not rely on Kilian’s simulation theorem, our obfuscator uses a matrix randomization scheme that is essentially<sup>2</sup> identical to the one used when applying Kilian’s randomization. Our analysis then proceeds by considering the most general polynomial that the adversary can construct in  $G_{[\ell]}$ . More precisely, we consider every possible monomial  $m$  that can exist over the matrix entries that are given to the adversary, and we associate each such monomial  $m$  with a coefficient  $\alpha_m$  that the adversary could potentially choose arbitrarily. Thus, the adversary’s polynomial is a giant sum  $p = \sum_m \alpha_m m$  over all these potential monomials.

Observe that the adversary can only extract useful information from this polynomial by passing it to **ZEROTEST**, thereby determining if it is zero or not. However, recall that the randomizing matrices  $\{\mathbf{R}_i\}$  are chosen uniformly during obfuscation. Therefore, by the Schwartz-Zippel lemma, we know that unless the adversary’s polynomial  $p$  is the zero polynomial over the entries of the  $\mathbf{R}_i$  matrices, **ZEROTEST** will declare the polynomial to be nonzero with overwhelming probability. So, we restrict ourselves to analyzing adversarial polynomials that are identically zero over the entries of the  $\mathbf{R}_i$  matrices.

Our new analysis differs at a fundamental level from Kilian’s analysis. At the heart of the analysis is an argument based on the structure of random square matrices  $R$  and their inverses  $R^{-1}$  that allows us to argue about how terms that arise in  $R^{-1}$  can be cancelled using terms from  $R$ . In particular, we use the fact that:

$$R_{i,\ell}^{-1} = \frac{1}{\det(R)} \sum_{\sigma:\sigma(i)=\ell} \text{sign}(\sigma) \left( \prod_{t \neq i} R_{\sigma(t),t} \right)$$

Our analysis is obtained by carefully considering different types of permutations  $\sigma$  that arise in the expression above, and how different permutations interfere with each other. (This exemplifies our conceptual departure from the proof of Kilian’s theorem.) Our analysis shows that multilinear polynomials that allow for cancellation of  $R$  and  $R^{-1}$  terms are extremely constrained.

---

<sup>2</sup>Because we consider rectangular matrices in general, we do need to modify this slightly. Also, for technical simplification, we consider the adjugate matrix rather than the inverse. However, for the purposes of this technical overview, these variations can be ignored.

From this analysis, we conclude that any adversarial polynomial that is identically zero over the entries of the  $\{\mathbf{R}_i\}$  matrices must in fact be the result of an honest iterated matrix multiplication (or a constant multiple thereof), which corresponds to evaluating  $f(x)$  for some input  $x$ . In other words, such an adversarial polynomial will result in an encoding of 0 if and only if  $f(x) = 0$ , as desired. Even though the analysis as presented here is not efficient, we are still able to use it to yield an efficient simulator in our generic model. At a high level, this is done by using the Schwartz-Zippel lemma to “weed out” most adversarial polynomials without needing to examine their structure at all.

## 1.2 Applications

We now discuss two applications of our new analysis tool.

**Efficiency of obfuscation.** Current techniques, while being asymptotically polynomial-time, lead to incredibly inefficient implementations of obfuscation. For example, the recent implementation of Apon et al. [AHKM14] for obfuscating (only) 16-bit point functions resulted in a 31GB obfuscated program, which took over 6 hours to generate and about 11 minutes to run on each input.

A major source of inefficiency is that the direct application of current obfuscators to circuits requires overhead that grows exponentially with the depth. This occurs because the level of multilinearity required grows exponentially with the depth, while current multilinear map candidates have complexity that grows polynomially with the level of multilinearity.

The work of Garg et al. [GGH<sup>+</sup>13b] shows that, nevertheless, such a “core” obfuscator can be used to obfuscate general (high depth) circuits with a polynomial overhead through a “bootstrapping” procedure (see also [GIS<sup>+</sup>10, BR14, App14]). However, bootstrapping based on existing core obfuscators entails overheads that are asymptotically polynomial but easily reach above  $2^{100}$ . Such large overheads primarily arise due to the depth of the circuit processed by the core obfuscator (though, asymptotically, this circuit has depth logarithmic in the security parameter). Indeed, similarly large overheads arise when attempting to apply the core obfuscator to other programs represented in circuit form, since few interesting and non-learnable families of circuits have depth below, say, 50.

This suggests that practical implementations of obfuscation will only be able to handle functionalities that require a polynomial level of multilinearity, and not exponential. One such class of functionalities are those computable by small matrix branching programs, where evaluation corresponds to evaluating an iterated matrix product. This class of functionalities includes, among others, finite automata.

Unfortunately, natural representations of finite automata and other simple programs as branching programs require *low-rank* matrices. Though these representations can be made full rank by using much larger matrices, this results in substantial efficiency loss. The reason for this, intuitively, is that branching programs with invertible matrices model *reversible* computation, whereas general computation allows for previous states to be forgotten. While it is possible to convert an irreversible computation into a reversible one, the cost is a significant loss in efficiency. The ability to handle low-rank matrices is thus crucial to obtaining efficient obfuscators even for simple functionalities. As detailed in the preceding subsection, all previous constructions critically relied on full-rank matrices.

Armed with our new tool (Theorem 4.2), our construction *no longer requires full-rank matrices*, and even non-square matrices are allowed.<sup>3</sup> That is, we show for the first time how to obfuscate matrix branching programs that are represented with low-rank, rectangular matrices. This leads to more efficient obfuscators, even beyond previous works that lack a post-zeroizing proof of security; for details, see Appendix B. Our analysis also extends to other settings besides obfuscation: for example, Boneh et al. [BLR<sup>+</sup>14] rely on our analysis to obtain implementable constructions of order-revealing encryption.

**Obfuscating evasive circuits in a post-zeroizing model.** We view Theorem 4.2 as the first step on a path towards achieving obfuscation in a post-zeroizing world. As a “proof of concept” for this goal, we construct an obfuscator for a natural class of functions that, for the first time, is provably secure in a model that captures *all* known attacks on graded encoding schemes.

In previous works that prove security in a generic model, the graded encoding scheme’s zero-test procedure is modeled as a Boolean function (i.e. one that returns a yes/no answer). In candidate constructions however, a successful zero-test actually returns an algebraic element in the ring of encodings, and this fact is crucially exploited in the zeroizing attacks. By contrast, our new model considers *any* encoding of 0 to be a complete break, thereby capturing these attacks.

We show how to obfuscate *evasive* functions [BBC<sup>+</sup>14] in this model, namely functions for which it is hard to find an input that evaluates to 0. (Typically one defines evasive functions as having hidden 1-outputs, but in terms of their functionality this is only a semantic difference.) A natural example of an evasive function is the “password check” function (typically called a point function), which evaluates to 0 on only a single, secret input. Obfuscating general evasive functions has many applications, including most notably obfuscating important classes of software patches that check for rare inputs on which the unpatched software is known to misbehave (see [BBC<sup>+</sup>14] for further discussion).

Prior to our work, except as a special case of general obfuscation, the only work that considered obfuscating general classes of evasive functions is that of [BBC<sup>+</sup>14]. However, the positive results in [BBC<sup>+</sup>14] were based on assumptions over approximate multilinear maps that are now known to be false when instantiated with current multilinear map candidates. Furthermore, the positive results in [BBC<sup>+</sup>14] did not consider completely arbitrary distributions of evasive circuits, as we do here.

Using our new analysis techniques, we prove the following.

**Theorem 1.1** (informal). *There exists a PPT obfuscator  $\mathcal{O}$  such that, for any evasive function family  $\mathcal{C}$  on  $n$ -bit inputs and any efficient generic-model adversary  $\mathcal{A}$ ,*

$$\Pr [\mathcal{A}(\mathcal{O}(C)) \text{ constructs an encoding of } 0] < \text{negl}(n)$$

where the probability is over the choice of  $C \leftarrow \mathcal{C}$  and the coins of  $\mathcal{A}$  and  $\mathcal{O}$ .

Theorem 1.1 in particular implies the first *witness encryption* scheme [GGSW13] with a generic model proof that captures zeroizing attacks<sup>4</sup>. Indeed, in the original witness encryption protocol of [GGSW13] the attacker *can* produce top-level encodings of zero, and therefore the protocol is

---

<sup>3</sup>We do require a mild natural technical condition, called *non-shortcutting*, on the branching program. Non-shortcutting can be achieved generically on any branching program with minimal overhead.

<sup>4</sup>When building witness encryption from obfuscation, witness encryption security only requires the obfuscator to be secure when obfuscating functions that always evaluate to 0, which are in particular evasive.

not secure in the post-zeroizing model. Subsequent witness encryption protocols [GLW14, Zha14] also allow top-level encodings of zero to be constructed.

In proving Theorem 1.1, we show that the “bootstrapping” theorem of [GGH<sup>+</sup>13b] extends to the setting of evasive functions. (As mentioned above, this theorem transforms a core obfuscator for a “small” class of functions into an obfuscator for all efficient functions.) We observe that the proof of this theorem only uses the core obfuscator on evasive functions, and we show that it holds only assuming the core obfuscator’s security on such functions. In particular, we show that Theorem 1.1 applies to all evasive functions and not only those on which the core obfuscator operates. Interestingly, the more recent bootstrapping technique of Applebaum [App14] cannot be used for our purposes, because it inherently produces encodings of 0 regardless of the function being obfuscated.

**Directions for future work.** An obvious next step is to consider obfuscating non-evasive functions. To do so, we will need to look precisely at the kinds of post-zero-test information that can be obtained using zeroizing attacks during zero testing for general (non-evasive) functions. We note that our paper answers a critical first question toward this goal: we show that in our scheme, the *only* way that the adversary can create top-level encodings of zero are the prescribed ways of evaluating the function at a particular input. This is a necessary first step in understanding what kinds of information can arise in the general case.

**Organization.** In Section 2 we give some preliminary definitions and background information. In Section 3 we define our obfuscator for matrix branching programs. In Section 4 we prove the key technical theorem that analyzes adversarially-constructed polynomials over the obfuscation. The proof of VBB security is outlined in Section 5, and the full details are given in Section 7. In Section 6 we prove that, when obfuscating evasive functions, no encodings of zero can be created.

## 2 Preliminaries

### 2.1 Evasive circuits

We define evasive circuit collections as in Barak et al. [BBC<sup>+</sup>14], except that in our definition it is hard to find a 0-output (typically one says that it is hard to find a 1-output).

**Definition 2.1.** A function family  $\{\mathcal{C}_\ell\}_{\ell \in \mathbb{N}}$  is *evasive* if for every oracle-aided adversary  $\mathcal{A}^{(\cdot)}$  that makes at most  $\text{poly}(\ell)$  queries on input  $1^\ell$ , and every  $\ell \in \mathbb{N}$ :

$$\Pr_{C \leftarrow \mathcal{C}_\ell} \left[ C \left( \mathcal{A}^C \left( 1^\ell \right) \right) = 0 \right] = \text{negl}(\ell).$$

$\{\mathcal{C}_\ell\}_{\ell \in \mathbb{N}}$  is *evasive with auxiliary input*  $\text{Aux}$  for a (possibly randomized) function  $\text{Aux} : \mathcal{C}_\ell \rightarrow \{0, 1\}^*$  if  $\mathcal{A}$  additionally receives  $\text{Aux}(C)$  when its oracle is  $C$ .

### 2.2 Obfuscation

We now give the definition of virtual black-box obfuscation in an idealized model, identical to the model studied in Barak et al. [BGK<sup>+</sup>14] and Ananth et al. [AGIS14], with one exception: we also consider giving both the adversary and simulator an auxiliary input determined by the program.



**Definition 2.2** (Virtual Black-Box Obfuscation in an  $\mathcal{M}$ -idealized model). For a (possibly randomized) oracle  $\mathcal{M}$ , a circuit class  $\{\mathcal{C}_\ell\}_{\ell \in \mathbb{N}}$ , and an efficiently computable deterministic function  $\text{Aux}_\ell : \mathcal{C}_\ell \rightarrow \{0, 1\}^{t_\ell}$ , we say that a uniform PPT oracle machine  $\mathcal{O}$  is a “Virtual Black-Box” Obfuscator for  $\{\mathcal{C}_\ell\}_{\ell \in \mathbb{N}}$  in the  $\mathcal{M}$ -idealized model with respect to auxiliary information  $\text{Aux}_\ell$ , if the following conditions are satisfied:

- Functionality: For every  $\ell \in \mathbb{N}$ , every  $C \in \mathcal{C}_\ell$ , every input  $x$  to  $C$ , and for every possible coins for  $\mathcal{M}$ :

$$\Pr[(\mathcal{O}^{\mathcal{M}}(C))(x) \neq C(x)] \leq \text{negl}(|C|) ,$$

where the probability is over the coins of  $\mathcal{C}$ .

- Polynomial Slowdown: there exist a polynomial  $p$  such that for every  $\ell \in \mathbb{N}$  and every  $C \in \mathcal{C}_\ell$ , we have that  $|\mathcal{O}^{\mathcal{M}}(C)| \leq p(|C|)$ .
- Virtual Black-Box: for every PPT adversary  $\mathcal{A}$  there exist a PPT simulator  $\text{Sim}$ , and a negligible function  $\mu$  such that for all PPT distinguishers  $D$ , for every  $\ell \in \mathbb{N}$  and every  $C \in \mathcal{C}_\ell$ :

$$\left| \Pr \left[ D \left( \mathcal{A}^{\mathcal{M}} \left( \mathcal{O}^{\mathcal{M}}(C), \text{Aux}_\ell(C) \right) \right) = 1, \right] - \Pr \left[ D \left( \text{Sim}^C \left( 1^{|C|}, \text{Aux}_\ell(C) \right) \right) = 1 \right] \right| \leq \mu(|C|) ,$$

where the probabilities are over the coins of  $D, \mathcal{A}, \text{Sim}, \mathcal{O}$  and  $\mathcal{M}$ .

Note that in this model, both the obfuscator and the evaluator have access to the oracle  $\mathcal{M}$  but the function family that is being obfuscated does not have access to  $\mathcal{M}$ .

We also define the average-case version of VBB obfuscation, which is the correct security notion when obfuscating evasive circuit collections.

**Definition 2.3** (Average-case Virtual Black-Box Obfuscation in an  $\mathcal{M}$ -idealized model). Let  $\mathcal{M}$ ,  $\{\mathcal{C}_\ell\}_{\ell \in \mathbb{N}}$ , and  $\text{Aux}_\ell$  be as in Def. 2.2. We say that a uniform PPT oracle machine  $\mathcal{O}$  is an average-case Virtual Black-Box Obfuscator for  $\{\mathcal{C}_\ell\}_{\ell \in \mathbb{N}}$  in the  $\mathcal{M}$ -idealized model with respect to auxiliary information  $\text{Aux}_\ell$ , if it satisfies all properties in Def. 2.2 except that in the Virtual Black-Box property the probabilities are over a uniform choice of  $C \leftarrow \mathcal{C}_\ell$  (as opposed to  $\forall C \in \mathcal{C}_\ell$ ).

**Definition 2.4** (Average-case Indistinguishability Obfuscation in an  $\mathcal{M}$ -idealized model). For a (possibly randomized) oracle  $\mathcal{M}$ , a circuit class  $\{\mathcal{C}_\ell\}_{\ell \in \mathbb{N}}$ , we say that a uniform PPT oracle machine  $\mathcal{O}$  is an Average-case Indistinguishability Obfuscator for  $\{\mathcal{C}_\ell\}_{\ell \in \mathbb{N}}$  in the  $\mathcal{M}$ -idealized model if the following conditions are satisfied:

- Functionality: Same as in the definition of VBB.
- Polynomial Slowdown: Same as in the definition of VBB.
- Indistinguishability: For every PPT Distinguisher  $D$ , there exists a negligible function  $\mu$  such that the following holds : for every  $\ell \in \mathbb{N}$ , for a uniform choice of circuit  $C \in \mathcal{C}_\ell$  and for every pair of circuits  $C_0, C_1 \in \mathcal{C}_\ell$  that compute the same function as  $C$ , we have :

$$\left| \Pr \left[ D(\mathcal{O}^{\mathcal{M}}(C_0)) = 1 \right] - \Pr \left[ D(\mathcal{O}^{\mathcal{M}}(C_1)) = 1 \right] \right| \leq \mu(|C|) ,$$

where the probabilities are over the coins of  $D, \mathcal{O}, \mathcal{M}$  and the choice of  $C$ .

Note that in this model, both the obfuscator and the evaluator have access to the oracle  $\mathcal{M}$  but the function family that is being obfuscated does not have access to  $\mathcal{M}$ .

## 2.3 Branching Programs

Here we define the main type of branching program we consider. A detailed description of other types of branching programs, and how to build these branching programs from other computational models, can be found in Appendix A.

**Definition 2.5.** A dual-input generalized matrix branching program of length  $\ell$  and shape  $(d_0, d_1, \dots, d_\ell) \in (\mathbb{Z}^+)^{\ell+1}$  for  $n$ -bit inputs is given by a sequence

$$BP = \left( \text{inp}_0, \text{inp}_1, \{\mathbf{B}_{i,b_0,b_1}\}_{i \in [\ell], b_0, b_1 \in \{0,1\}} \right)$$

where  $\mathbf{B}_{i,b_0,b_1} \in \mathbb{Z}^{d_{i-1} \times d_i}$  are  $d_{i-1} \times d_i$  matrices, and  $\text{inp} : [\ell] \rightarrow [n]$  is the evaluation function of  $BP$ .  $BP$  defines the following three functions:

- $BP_{arith} : \{0, 1\}^n \rightarrow \mathbb{Z}^{d_0 \times d_\ell}$  computed as  $BP_{arith}(x) = \prod_{i=1}^n \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}}$
- $BP_{bool} : \{0, 1\}^n \rightarrow \{0, 1\}^{d_0 \times d_\ell}$  computed as  $BP_{bool}(x)_{j,k} = \begin{cases} 0 & \text{if } BP_{arith}(x)_{j,k} = 0 \\ 1 & \text{if } BP_{arith}(x)_{j,k} \neq 0 \end{cases}$
- $BP_{bool(q)} : \{0, 1\}^n \rightarrow \{0, 1\}^{d_0 \times d_\ell}$  computed as  $BP_{bool(q)}(x)_{j,k} = \begin{cases} 0 & \text{if } BP_{arith}(x)_{j,k} = 0 \pmod q \\ 1 & \text{if } BP_{arith}(x)_{j,k} \neq 0 \pmod q \end{cases}$

A matrix branching program is  $t$ -bounded if  $|BP_{arith}(x)_{j,k}| \leq t$  for all  $x, j, k$ .

Next, we define a notion of non-shortcutting for matrix branching programs, which roughly states that it is not possible to determine any of the output components of  $BP_{arith/bool}$  without carrying out the entire matrix product. In the case  $d_0 = d_\ell = 1$  (so that the branching program outputs just a single element), this translates to requiring that no strict sub-product  $\prod_{i=i_0}^{i_1} \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}}$  for  $(i_0, i_1) \neq (1, n)$  of the overall matrix product evaluates to an all-zero matrix. Clearly, if some sub-product evaluates to zero, the entire product would evaluate to zero, and so the evaluation could stop after computing just the sub-product. We call this a short-cut, and non-shortcutting is the requirement that there are no shortcuts for any inputs. In the more general case of arbitrary  $d_0, d_\ell$ , the condition becomes slightly more technical, and is given below:

**Definition 2.6.** A dual-input generalized matrix branching program is *non-shortcutting* if, for any input  $x$ , and any  $j \in [d_0]$  and any  $k \in [d_\ell]$ , the following holds:

$$\mathbf{e}_j^T \cdot \left( \prod_{i=1}^{\ell-1} \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}} \right) \neq 0^{d_\ell-1} \quad \text{and} \quad \left( \prod_{i=2}^{\ell} \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}} \right) \cdot \mathbf{e}_k \neq 0^{d_1}$$

where  $\mathbf{e}_j$  and  $\mathbf{e}_k$  are the  $j$ th and  $k$ th standard basis vectors of the correct dimension. Equivalently, each row of the product  $\prod_{i=1}^{\ell-1} \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}}$  and each column of the product  $\prod_{i=2}^{\ell} \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}}$  has at least one non-zero entry.

**Matrix Branching Program Samplers.** We now define a matrix branching program sampler (MBPS). Roughly, an MBPS is a procedure that takes as input a modulus  $q$ , and outputs a matrix branching program  $BP$ . However, we will be interested mainly in the function  $BP_{bool(q)}$ .

**Definition 2.7.** A matrix branching program sampler (MBPS) is a possibly randomized procedure  $BP^S$  that takes as input a modulus  $q$  satisfying  $q > t$  for some bound  $t$ . It outputs a matrix branching program.

**Fact 2.8.** Any matrix branching program  $BP$  with bound  $t$  can trivially be converted into a matrix branching program sampler  $BP^S$  with the same bound  $t$ , such that if  $BP' \leftarrow BP^S(q)$ , then  $BP'_{bool(q)}(x) = BP_{bool}(x)$ .

## 2.4 The Ideal Graded Encoding Model

In this section, we describe the ideal graded encoding model. This section has been taken almost verbatim from [BGK<sup>+</sup>14] and [AGIS14]. All parties have access to an oracle  $\mathcal{M}$ , implementing an ideal graded encoding. The oracle  $\mathcal{M}$  implements an idealized and simplified version of the graded encoding schemes from [GGH13a]. The parties are provided with encodings of various elements at different levels. They are allowed to perform arithmetic operations of addition/multiplication and testing equality to zero as long as they respect the constraints of the multilinear setting. We start by defining an algebra over the elements.

**Definition 2.9.** Given a ring  $R$  and a universe set  $\mathbb{U}$ , an element is a pair  $(\alpha, S)$  where  $\alpha \in R$  is the *value* of the element and  $S \subseteq \mathbb{U}$  is the *index* of the element. Given an element  $e$  we denote by  $\alpha(e)$  the value of the element, and we denote by  $S(e)$  the index of the element. We also define the following binary operations over elements:

- For two elements  $e_1, e_2$  such that  $S(e_1) = S(e_2)$ , we define  $e_1 + e_2$  to be the element  $(\alpha(e_1) + \alpha(e_2), S(e_1))$ , and  $e_1 - e_2$  to be the element  $(\alpha(e_1) - \alpha(e_2), S(e_1))$ .
- For two elements  $e_1, e_2$  such that  $S(e_1) \cap S(e_2) = \emptyset$ , we define  $e_1 \cdot e_2$  to be the element  $(\alpha(e_1) \cdot \alpha(e_2), S(e_1) \cup S(e_2))$ .

We will often use the notation  $[\alpha]_S$  to denote the element  $(\alpha, S)$ . Next, we describe the oracle  $\mathcal{M}$ .  $\mathcal{M}$  is a stateful oracle mapping elements to “generic” representations called *handles*. Given handles to elements,  $\mathcal{M}$  allows the user to perform operations on the elements.  $\mathcal{M}$  will implement the following interfaces:

**Initialization.**  $\mathcal{M}$  will be initialized with a ring  $R$ , a universe set  $\mathbb{U}$ , and a list  $L$  of initial elements. For every element  $e \in L$ ,  $\mathcal{M}$  generates a handle. We do not specify how the handles are generated, but only require that the value of the handles are independent of the elements being encoded, and that the handles are distinct (even if  $L$  contains the same element twice).  $\mathcal{M}$  maintains a handle table where it saves the mapping from elements to handles.  $\mathcal{M}$  outputs the handles generated for all the elements in  $L$ . After  $\mathcal{M}$  has been initialized, all subsequent calls to the initialization interface fail.

**Algebraic operations.** Given two input handles  $h_1, h_2$  and an operation  $\circ \in \{+, -, \cdot\}$ ,  $\mathcal{M}$  first locates the relevant elements  $e_1, e_2$  in the handle table. If any of the input handles does not appear in the handle table (that is, if the handle was not previously generated by  $\mathcal{M}$ ) the call to  $\mathcal{M}$  fails. If the expression  $e_1 \circ e_2$  is undefined (i.e.,  $S(e_1) \neq S(e_2)$  for  $\circ \in \{+, -\}$ , or  $S(e_1) \cap S(e_2) \neq \emptyset$  for  $\circ \in \{\cdot\}$ ) the call fails. Otherwise,  $\mathcal{M}$  generates a new handle for  $e_1 \circ e_2$ , saves this element and the new handle in the handle table, and returns the new handle.

**Zero testing.** Given an input handle  $h$ ,  $\mathcal{M}$  first locates the relevant element  $e$  in the handle table. If  $h$  does not appear in the handle table (that is, if  $h$  was not previously generated by  $\mathcal{M}$ ) the call to  $\mathcal{M}$  fails. If  $S(e) \neq \mathbb{U}$ , the call fails. Otherwise,  $\mathcal{M}$  returns 1 if  $\alpha(e) = 0$ , and returns 0 if  $\alpha(e) \neq 0$ .

## 2.5 Straddling Set Systems

We use the *strong* straddling set system of [MSW14], which modifies the straddling set system of [BGK<sup>+</sup>14] to obtain a denser intersection graph between the subsets. This extra power is used in Section 6 when showing that the adversary cannot create low-level encodings of 0.

**Definition 2.10** (Strong straddling set system). A *strong straddling set system* with  $n$  entries is a collection of sets  $\mathbb{S} = \{S_{i,b} : i \in [n], b \in \{0, 1\}\}$  over a universe  $\mathbb{U}$ , such that  $\cup_{i \in [n]} S_{i,0} = \mathbb{U} = \cup_{i \in [n]} S_{i,1}$ , and the following holds.

- (Collision at universe.) If  $C, D \subseteq \mathbb{S}$  are distinct non-empty collections of disjoint sets such that  $\cup_{S \in C} S = \cup_{S \in D} S$ , then  $\exists b \in \{0, 1\}$  such that  $C = \{S_{i,b}\}_{i \in [n]}$  and  $D = \{S_{i,1-b}\}_{i \in [n]}$ .
- (Strong intersection.) For every  $i, j \in [n]$ ,  $S_{i,0} \cap S_{j,1} \neq \emptyset$ .

We will need the following simple lemma.

**Lemma 2.11.** *Let  $\mathbb{S} = \{S_{i,b} : i \in [n], b \in \{0, 1\}\}$  be a strong straddling set system over a universe  $\mathbb{U}$ . Then for any  $T \subsetneq \mathbb{U}$  that can be written as a disjoint union of sets from  $\mathbb{S}$ , there is a unique  $b \in \{0, 1\}$  such that  $T = \cup_{i \in I} S_{b,i}$  for some  $I \subseteq [n]$ .*

**Proof.** By the second property of Def. 2.10, any pairwise disjoint collection of sets from  $\mathbb{S}$  must be either all of the form  $S_{i,0}$  or all of the form  $S_{i,1}$ . If there are two sets  $I_0, I_1 \subseteq [n]$  such that  $\cup_{i \in I_0} S_{i,0} = T = \cup_{i \in I_1} S_{i,1}$ , then by the first property of Def. 2.10 we must have  $T = \mathbb{U}$  which contradicts our assumption.  $\square$

We use the following construction from [MSW14].

**Construction 2.12** (Strong straddling set system). Define  $\mathbb{S} = \{S_{i,b} : i \in [n], b \in \{0, 1\}\}$  over a universe  $\mathbb{U} = \{1, 2, \dots, n^2\}$  as follows for all  $1 \leq i \leq n$ .

$$S_{i,0} = \{n(i-1) + 1, n(i-1) + 2, \dots, ni\} \quad S_{i,1} = \{i, n+i, 2n+i, \dots, n(n-1)+i\}$$

### 3 Obfuscator for Low-Rank Branching Programs

We now describe our obfuscator for generalized matrix branching programs. Our obfuscator is essentially the same as the obfuscator of Ananth et al. [AGIS14]. The differences are as follows:

- We view branching programs as including the bookends. While the bookends of previous works did not depend on the input, they can in our obfuscator. However, for [AGIS14], this distinction is superficial: the bookends of [AGIS14] can be “absorbed” into the branching program by merging them with the left-most and right-most matrices of the branching program. This does not change functionality, since this merging always happens during evaluation, and it does not change security, since the adversary can perform the merging himself.
- We allow our branching program to have singular and rectangular matrices. We do, however, require the branching program to be non-shortcutting. Note that a branching program with square invertible internal matrices and non-zero bookend vectors, such as in [AGIS14], necessarily is non-shortcutting.
- We allow branching programs to output multiple bits — that is, the function computed by our obfuscated program will be  $BP_{bool}$ , which is a matrix of 0/1 entries. In order to prove security, we will have to perform additional randomization. However, in the case of single-bit outputs, this additional randomization is redundant.

**Input.** The input to our obfuscator is a dual-input matrix branching program sampler  $BP^S$  of length  $\ell$ , shape  $(d_0, d_1, \dots, d_\ell)$ , and bound  $t$ . The first step is to choose a large prime  $q$  for the graded encodings. Then sample  $BP \leftarrow BP^S(q)$ . Write

$$BP = (\text{inp}_0, \text{inp}_1, \{\mathbf{B}_{i,b_0,b_1}\})$$

We require  $BP^S$  to output  $BP$  satisfying the following properties:

- $BP$  is non-shortcutting.
- For each  $i$ ,  $\text{inp}_0(i) \neq \text{inp}_1(i)$
- For each pair  $(j, k) \in [n]^2$ , there exists an  $i \in [\ell]$  such that  $(\text{inp}_0(i), \text{inp}_1(i)) = (j, k)$  or  $(\text{inp}_1(i), \text{inp}_0(i)) = (j, k)$

For ease of notation in our security proof, we will also assume that each input bit is used exactly  $m$  times, for some integer  $m$ . In other words, for each  $i \in [n]$ , the sets  $\text{ind}(i) = \{j : \text{inp}_b(j) = i \text{ for some } b \in \{0, 1\}\}$  have the same size. This requirement, however, is not necessary for security.

**Step 1: Randomize  $BP$ .** First, similar to previous works, we use Kilian [Kil88] to randomize  $BP$ , obtaining a randomized branching program  $BP'$ . This is done as follows.

- Let  $q$  be a sufficiently large prime of  $\Omega(\lambda)$  bits.
- For each  $i \in [\ell - 1]$ , choose a random matrix  $\mathbf{R}_i \in \mathbb{Z}_q^{d_i \times d_i}$ . Set  $\mathbf{R}_0, \mathbf{R}_\ell$  to be identity matrices of the appropriate size. Define

$$\widehat{\mathbf{B}}_{i,b_0,b_1} = \mathbf{R}_{i-1}^{adj} \cdot \mathbf{B}_{i,b_0,b_1} \cdot \mathbf{R}_i$$

- For each  $s \in [d_0]$ , choose a random  $\beta_s$  and set  $\mathbf{S}$  to be the  $d_0 \times d_0$  diagonal matrix with the  $\beta_s$  along the diagonal. For each  $t \in [d_\ell]$ , choose a random  $\gamma_t$  and set  $\mathbf{T}$  to be the  $d_\ell \times d_\ell$  diagonal matrix with  $\gamma_t$  along the diagonal. Set

$$\mathbf{C}_{1,b_0,b_1} = \mathbf{S} \cdot \widehat{\mathbf{B}}_{1,b_0,b_1} \quad \mathbf{C}_{\ell,b_0,b_1} = \widehat{\mathbf{B}}_{1,b_0,b_1} \cdot \mathbf{T} \quad \mathbf{C}_{i,b_0,b_1} = \widehat{\mathbf{B}}_{i,b_0,b_1} \text{ for each } i \in [2, \ell - 1]$$

We note that this additional randomization step is not present in previous works, but is required to handle multi-bit outputs

- For each  $i \in [\ell]$ ,  $b_0, b_1 \in \{0, 1\}$ , choose a random  $\alpha_{i,b_0,b_1} \in \mathbb{Z}_p$ , and define

$$\mathbf{D}_{i,b_0,b_1} = \alpha_{i,b_0,b_1} \mathbf{C}_{i,b_0,b_1}$$

Then define  $BP' = (\text{inp}_0, \text{inp}_1, \{\mathbf{D}_{i,b_0,b_1}\})$ . Observe that  $BP'_{\text{bool}(q)}(x) = BP_{\text{bool}(q)}(x)$  for all  $x$ .

**Step 2: Create set systems.** Consider a universe  $\mathbb{U}$ , and a partition  $\mathbb{U}_1, \dots, \mathbb{U}_\ell$  of  $\mathbb{U}$  into equal sized disjoint sets:  $|\mathbb{U}_i| = 2m - 1$ . Let  $\mathbb{S}^j$  be a straddling set system over the elements of  $\mathbb{U}_j$ . Note that  $\mathbb{S}^j$  will have  $m$  entries, corresponding to the number of times each input bit is used. We now associate the elements of  $\mathbb{S}_j$  to the indices of  $BP$  that depend on  $x_j$ :

$$\mathbb{S}^j = \{S_{k,b}^j : k \in \text{ind}(j), b \in \{0, 1\}\}$$

Next, we associate a set to each element output by the randomization step. Recall that in a dual-input relaxed matrix branching program, each step depends on two fixed bits in the input defined by the evaluation functions  $\text{inp}_0$  and  $\text{inp}_1$ . For each step  $i \in [n]$ ,  $b_0, b_1 \in \{0, 1\}$ , we define the set  $S(i, b_0, b_1)$  using the straddling sets for input bits  $\text{inp}_1(i)$  and  $\text{inp}_2(i)$  as follows:

$$S_{i,b_0,b_1} = S_{i,b_0}^{\text{inp}_0(i)} \cup S_{i,b_1}^{\text{inp}_1(i)}$$

**Step 3: Initialization.**  $\mathcal{O}$  initializes the oracle  $\mathcal{M}$  with the ring  $\mathbb{Z}_p$  and the universe  $\mathbb{U}$ . Then it asks for the encodings of the following elements:

$$\{(\mathbf{D}_{i,b_0,b_1}[j, k], S_{i,b_0,b_1})\}_{i \in [\ell], b_0, b_1 \in \{0, 1\}, j \in [d_{i-1}], k \in [d_i]}$$

$\mathcal{O}$  receives a list of handles back from  $\mathcal{M}$ . Let  $[\beta]_S$  denote the handle for  $(\beta, S)$ , and for a matrix  $M$ , let  $[M]_S$  denote the matrix of handles ( $[M]_S[j, k] = [M[j, k]]_S$ ). Thus,  $\mathcal{O}$  receives the handles:

$$\{[\mathbf{D}_{i,b_0,b_1}]_{S_{i,b_0,b_1}}\}_{i \in [\ell], b_0, b_1 \in \{0, 1\}}$$

**Output.**  $\mathcal{O}(BP^S)$  outputs these handles, along with the length  $\ell$ , shape  $d_0, \dots, d_\ell$ , and input functions  $\text{inp}_0, \text{inp}_1$ , as the obfuscated program. Denote the resulting obfuscated branching program as  $BP^\mathcal{O}$

**Evaluation.** To evaluate  $BP^\mathcal{O}$  on input  $x$ , use the oracle  $\mathcal{M}$  to add and multiply encodings in order to compute the product

$$h = \left[ \prod_{i \in [\ell]} \mathbf{D}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}} \right]_{\mathbb{U}} = \prod_{i \in [\ell]} \left[ \mathbf{D}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}} \right]_{S_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}}}$$

$h$  is a  $d_0 \times d_\ell$  matrix of encodings relative to  $\mathbb{U}$ . Next, use  $\mathcal{M}$  to test each of the components of  $h$  for zero, obtaining a matrix  $h_{\text{bool}} \in \{0, 1\}^{d_0 \times d_\ell}$ . That is, if the zero test on returns a 1 on  $h[s, t]$ ,  $h_{\text{bool}}[s, t]$  is 0, and if the zero test returns a 0,  $h_{\text{bool}}[s, t]$  is 1.

**Correctness of evaluation.** The following shows that all calls to the oracle  $\mathcal{M}$  succeed:

**Lemma 3.1** (Adapted from [AGIS14]). *All calls made to the oracle  $\mathcal{M}$  during obfuscation and evaluation succeed.*

It remains to show that the obfuscated program computes the correct function. Fix an input  $x$ , and define  $b_c^i = x_{\text{inp}_c(i)}$  for  $i \in [\ell], c \in \{0, 1\}$ . From the description above,  $BP^\mathcal{O}$  outputs 0 at position  $[s, t]$  if and only if

$$\begin{aligned} 0 &= \left( \prod_{i \in [\ell]} \mathbf{D}_{i, b_0^i, b_1^i} \right) [s, t] = \beta_s \gamma_t \left( \prod_{i \in [\ell]} \alpha_{i, b_0^i, b_1^i} \mathbf{R}_{i-1}^{\text{adj}} \cdot \mathbf{B}_{i, b_0^i, b_1^i} \cdot \mathbf{R}_i \right) [s, t] \\ &= \beta_s \gamma_t \left( \left( \prod_{i \in [\ell]} \alpha_{i, b_0^i, b_1^i} \right) \left( \prod_{i \in [\ell]} \mathbf{B}_{i, b_0^i, b_1^i} \right) \right) [s, t] = \left( \beta_s \gamma_t \prod_{i \in [\ell]} \alpha_{i, b_0^i, b_1^i} \right) (BP_{\text{arith}}(x)[s, t]) \end{aligned}$$

With high probability  $\beta_s, \gamma_t, \alpha_{i, b_0, b_1} \neq 0$ , meaning  $BP_{\text{arith}}(x)[s, t] = 0 \pmod q$  if and only if the zero test procedure on position  $[s, t]$  gives 0. Therefore,  $BP^\mathcal{O}(x) = BP_{\text{bool}(q)}(x)$  for the branching program  $BP$  sampled from  $BP^S$ .

## 4 Polynomials on Kilian-Randomized Matrices

In this section, we prove a theorem about polynomials on the Kilian-randomized matrices from the previous section. Our high level goal is to show polynomials the adversary tries to construct other than the correct matrix products will be useless to the adversary. In this section, we focus on a simpler case where the polynomial is only over matrices corresponding to a single input. In the following section, we use the results of this section to prove the general case.

Previous works showed the single-input case using Kilian simulation [BR14, BGK<sup>+</sup>14], or a variant of it [PST14b, AGIS14]. Namely, these works queried the function oracle to determine what the result of the matrix product  $P(x)$  should be. Then, they tested the polynomial on random matrices, subject to the requirement that the product equaled  $P(x)$ , to see what the result was. Unfortunately, this step of the analysis does indicate what the outputs of the polynomial may be, only that they can be simulated. If the polynomial were to output zero, this would correspond to the adversary obtaining a zero encoding, which would violate security in our post-zeroizing model.

Moreover, previous works crucially relied on the fact that the matrices the polynomial is tested on come from the same distribution as the matrices would in the branching program. This requires

the branching program to consist of square invertible matrices. However, we need to be able to handle generalized matrix branching programs with rectangular and low-rank matrices.

In light of the two issues above, we need to replace the Kilian randomization theorem with a new theorem suitable in our setting.

Let  $d_1, \dots, d_{n-1}$  be positive integers and  $d_0 = d_n = 1$ . Let  $\widehat{\mathbf{A}}_k$  for  $k \in [n]$  be  $d_{k-1} \times d_k$  matrices of variables.

**Definition 4.1.** Let  $d_k, \widehat{\mathbf{A}}_k$  be as above. Consider a multilinear polynomial  $p$  on the variables in  $\{\widehat{\mathbf{A}}_k\}_{k \in [n]}$ . We call  $p$  *allowable* if each monomial in the expansion of  $p$  contains at most one variable from each of the  $\widehat{\mathbf{A}}_k$ .

As an example of an allowable polynomial, consider the the *matrix product polynomial*  $\widehat{\mathbf{A}}_1 \cdot \widehat{\mathbf{A}}_2 \cdot \dots \cdot \widehat{\mathbf{A}}_n$ .

Now fix a field  $\mathbb{F}$ , and let  $\mathbf{A}_k \in \mathbb{F}^{d_{k-1} \times d_k}$  for  $k = 1, \dots, n$  be a collection of matrices over  $\mathbb{F}$ . Let  $\mathbf{R}_k$  be  $d_k \times d_k$  matrices of variables for  $k \in [n]$ , and let  $\mathbf{R}_k^{adj}$  be the adjugate matrix of  $\mathbf{R}_k$ . Let  $\mathbf{R}_0 = \mathbf{R}_{n+1} = 1$ . Now suppose we set

$$\widehat{\mathbf{A}}_k = \mathbf{R}_{k-1}^{adj} \cdot \mathbf{A}_k \cdot \mathbf{R}_k$$

**Theorem 4.2.** Let  $\mathbb{F}, d_k, \mathbf{A}_k, \mathbf{R}_k, \widehat{\mathbf{A}}_k$  be as above. Consider an allowable polynomial  $p$  in the  $\widehat{\mathbf{A}}_k$ , and suppose  $p$ , after making the substitution  $\widehat{\mathbf{A}}_k = \mathbf{R}_{k-1}^{adj} \cdot \mathbf{A}_k \cdot \mathbf{R}_k$ , is identically 0 as a polynomial over the  $\mathbf{R}_k$ . Then the following is true:

- If  $\mathbf{A}_1 \cdot \mathbf{A}_2 \cdot \dots \cdot \mathbf{A}_n \neq 0$ , then  $p$  is identically zero as a polynomial over its formal variables, namely the  $\widehat{\mathbf{A}}_k$ .
- If  $\mathbf{A}_1 \cdot \mathbf{A}_2 \cdot \dots \cdot \mathbf{A}_n = 0$  but

$$\begin{aligned} \mathbf{A}_1 \cdot \mathbf{A}_2 \cdot \dots \cdot \mathbf{A}_{n-1} &\neq 0^{1 \times d_n} \\ \mathbf{A}_2 \cdot \dots \cdot \mathbf{A}_{n-1} \cdot \mathbf{A}_n &\neq 0^{d_2 \times 1} \end{aligned}$$

then  $p$ , as a polynomial over the  $\widehat{\mathbf{A}}_k$ , is a constant multiple of the matrix product polynomial  $\widehat{\mathbf{A}}_1 \cdot \widehat{\mathbf{A}}_2 \cdot \dots \cdot \widehat{\mathbf{A}}_n$ .

**Proof.** If  $n = 1$ , there are no  $\mathbf{R}_k$  matrices, a single  $\mathbf{A}_1$  matrix of dimension  $1 \times 1$ , with entry  $a$ . Then  $p = p(a) = ca$  for some constant  $c$ . As a polynomial over the (non-existent)  $\mathbf{R}_i$  matrices,  $p$  is just a constant polynomial, so  $p = 0$  means  $ca = 0$ . In the first case above,  $a \neq 0$ , so  $c = 0$ , meaning  $p$  is identically 0. The second case above is trivially satisfied since the matrix product polynomial is also a constant.

We will assume that  $\mathbf{A}_1$  is non-zero in every coordinate. At the end of the proof, we will show this is without loss of generality.

Now we proceed by induction on  $n$ . Assume Theorem 4.2 is proved for  $n - 1$ . Consider an arbitrary allowable polynomial  $p$ . We can write  $p$  as

$$p = \sum_{j_1, i_2, j_2, \dots, j_n, i_{n+1}} \alpha_{j_1, i_2, \dots, j_{n-1}, i_n} \widehat{A}_{1,1,j_1} \widehat{A}_{2,i_2,j_2} \dots \widehat{A}_{n-1,i_{n-1},j_{n-1}} \widehat{A}_{n,i_n,1}$$

Where  $i_{k+1}, j_k \in [d_k]$ , and  $\widehat{A}_{k,i,j}$  is the  $(i, j)$  entry of the matrix  $\widehat{\mathbf{A}}_k$ . From this point forward, for convenience, we will no longer explicitly refer to the bounds  $d_k$  on the  $i_{k+1}, j_k$ .



Now we can expand  $p$  in terms of the  $R_1$  matrix:

$$\begin{aligned} p &= \sum_{j_1, i_2, j_2, \dots, j_n, i_{n+1}, m, \ell} \alpha_{j_1, i_2, \dots, j_{n-1}, i_n} A_{1,1,m} R_{1,m,j_1} R_{1,i_2,\ell}^{adj} (\mathbf{A}_2 \cdot \mathbf{R}_2)_{\ell, j_2} \widehat{A}_{3, i_3, j_3} \cdots \widehat{A}_{n, i_n, 1} \\ &= \sum_{j, i, \ell, m} \alpha'_{j, i, \ell} A_{1,1,m} R_{1,m,j} R_{1,i,\ell}^{adj} \end{aligned}$$

where

$$\alpha'_{j, i, \ell} = \sum_{j_2, \dots, j_n, i_{n+1}} \alpha_{j, i, \dots, j_{n-1}, i_n} (\mathbf{A}_2 \cdot \mathbf{R}_2)_{\ell, j_2} \widehat{A}_{3, i_3, j_3} \cdots \widehat{A}_{n, i_n, 1}$$

Recall that

$$R_{1,i,\ell}^{adj} = \sum_{\sigma: \sigma(i)=\ell} \text{sign}(\sigma) \left( \prod_{t \neq i} R_{1, \sigma(t), t} \right)$$

where the sum is over all permutations satisfying  $\sigma(i) = \ell$ . Thus we can write  $p$  as

$$p = \sum_{j, i, \sigma, m} \text{sign}(\sigma) \alpha'_{j, i, \sigma(i)} A_{1,1,m} R_{1,m,j} \left( \prod_{t \neq i} R_{1, \sigma(t), t} \right)$$

Now, since  $p$  is identically zero as a polynomial over the  $\mathbf{R}_k$  matrices, it must be that for each product  $R_{1,m,j} \left( \prod_{t \neq i} R_{1, \sigma(t), t} \right)$ , the coefficient of the product (which is a polynomial over the  $\mathbf{R}_k : k \geq 2$  matrices) must be identically 0. We now determine the coefficients.

First, we examine the types of products of entries in  $\mathbf{R}_1$  that are possible. Products can be thought of as arising from the following process. Choose a permutation  $\sigma$ , which corresponds to selecting  $d_1$  entries of  $\mathbf{R}_1$  such that each row and column of  $\mathbf{R}_1$  contain exactly one selected entry. Then, for some  $i$ , un-select the selected entry from column  $i$  and instead select any entry from  $\mathbf{R}_1$  (possibly selecting the same entry twice). We observe that the following products are possible:

- $\prod_t R_{1, \sigma(t), t}$  for a permutation  $\sigma$ . This corresponds to re-selecting the un-selected entry from column  $i$ . The resulting list of entries determines the permutation  $\sigma$  used to select the original entries (since it is identical to the original list), but allows the column  $i$  of the un-selected/re-selected entry to vary. Thus in the summation above, this fixes  $\sigma$ ,  $j = i$  and  $m = \sigma(i)$ , but allows  $i$  to vary over all values, corresponding to the fact that if we remove any entry and replace it with itself, the result is independent of which entry we removed. Call such products *well-formed*. Well-formed products give the following equation:

$$\sum_i \alpha'_{i, i, \sigma(i)} A_{1,1, \sigma(i)} = 0 \text{ for all } \sigma \quad (4.1)$$

- $R_{1,m,j} \prod_{t \neq i} R_{1, \sigma(t), t}$  where  $j \neq i$  and  $m \neq \sigma(i)$ . This corresponds to, after un-selecting the entry in column  $i$ , selecting a another entry that is in both a different row and a different column. Note that, given final list of selected entries, it is possible to determine the newly selected entry as the unique selected entry that shares both a column with another selected entry and a row with another selected entry. It is also possible to determine the un-selected entry as the only entry that shares no column nor row with another entry. Therefore, the original entry selection is determined as well. Thus, in the summation above, the selected

entries fix  $\sigma$ ,  $i$ ,  $j$ , and  $m$ . In other words, there is no other selection process that gives the same list of entries from  $\mathbf{R}_1$ .

We call such products *malformed type 1*. Malformed type 1 products have the coefficient

$$\alpha'_{j,i,\sigma(i)} A_{1,1,m}$$

Given any  $i, j \neq i, m, \ell \neq m$ , pick  $\sigma$  so that  $\sigma(i) = \ell$ . Since  $A_{1,1,m} \neq 0$  for all  $m$ , this gives

$$\alpha'_{j,i,\ell} = 0 \text{ for all } i, j \neq i, \ell \quad (4.2)$$

- $R_{1,m,i} \prod_{t \neq i} R_{1,\sigma(t),t}$  where  $m \neq \sigma(i)$ . This corresponds to, after un-selecting the entry  $R_{1,\sigma(i),i}$ , selecting a different entry  $R_{1,m,i}$  in the same column. Let  $i', m', \sigma'$  be some other selection process that leads to the same product.

Given the final selection of entries, it is possible to determine  $m' = m$  as the only row with two selected entries. It is also possible to determine  $\sigma'(i') = \sigma(i)$  as the only row with no selected entries (though  $i'$  has not been determined yet). Moreover,  $i'$  must be one of the two columns selected in row  $m$ , call the other  $i''$ . All entries outside of these two rows must have come from the original selection of entries, so this determines  $\sigma'(t) = \sigma(t)$  on all inputs outside of  $i, i''$ . Notice that if  $i = i'$ , then  $\sigma'$  agrees with  $\sigma$  on  $d_1 - 1$  entries, and since they are both permutations, this sets  $\sigma' = \sigma$ . In this case,  $(i', m', \sigma') = (i, m, \sigma)$ .

Otherwise  $i' \neq i$ , so  $i'' = i$ , which leaves  $\sigma'(i) = \sigma(i') = m$ . At this point,  $\sigma'$  is fully determined as  $\sigma \circ (i \ i')$  where  $(i \ i')$  is the transposition swapping  $i$  and  $i'$ . Therefore, there are two possibilities leading to this product, one corresponding to  $i$  and the other corresponding to  $i'$ .

We call these products *malformed type 2*. Notice that  $\sigma'$  and  $\sigma$  only differ by a transposition swapping  $i$  and  $i'$ , and so they have opposite parity, meaning the corresponding terms in  $p$  have the opposite sign. Given  $i, i' \neq i, m, \ell \neq m$ , choose  $\sigma$  so that  $\sigma(i) = \ell$ . This gives us  $(\alpha'_{i,i,\ell} - \alpha'_{i',i',\ell}) A_{1,1,m} = 0$ . Since  $A_{1,1,m} \neq 0$  for all  $m$ , we therefore have that  $\alpha'_{i,i,\ell} = \alpha'_{i',i',\ell}$  for all  $i, i'$ . We can thus choose  $\beta_\ell$  such that:

$$\alpha'_{i,i,\ell} = \beta_\ell \text{ for all } i, \ell \quad (4.3)$$

- $R_{1,\sigma(i),j} \prod_{t \neq i} R_{1,\sigma(t),t}$  where  $j \neq i$ . We call such products *malformed type 3*. The coefficients of these products are linear combinations of the  $\alpha'_{i,j,\ell}$  for  $i \neq j$ , which we already know to be 0. Therefore, these equations are redundant, and we will not need to consider them.

Setting  $\sigma(i) = i$  in Equation 4.1 and combining with Equation 4.3, we have that

$$\sum_{\ell} \beta_{\ell} A_{1,1,\ell} = 0 \quad (4.4)$$

Now we can expand  $\alpha'_{j,i,\ell}$  and  $\beta_i$  in Equations 4.2 and 4.4, obtaining:

$$0 = \alpha'_{i,j,\ell} = \sum_{j_2, i_3, \dots, j_{n-1}, i_n} \alpha_{j,i,j_2, i_3, \dots, j_{n-1}, i_n} (\mathbf{A}_2 \cdot \mathbf{R}_2)_{\ell, j_2} \hat{A}_{3, i_3, j_3} \dots \hat{A}_{n, i_n, 1} \text{ for all } \ell, i, j \neq i \quad (4.5)$$

$$\begin{aligned} 0 = \sum_{\ell} \beta_{\ell} A_{1,1,\ell} &= \sum_{\ell, j_2, i_3, \dots, j_{n-1}, i_n} \alpha_{i,i,j_2, i_3, \dots, j_{n-1}, i_n} A_{1,1,\ell} (\mathbf{A}_2 \cdot \mathbf{R}_2)_{\ell, j_2} \hat{A}_{3, i_3, j_3} \dots \hat{A}_{n, i_n, 1} \\ &= \sum_{j_2, i_3, \dots, j_{n-1}, i_n} \alpha_{i,i,j_2, i_3, \dots, j_{n-1}, i_n} (\mathbf{A}_1 \cdot \mathbf{A}_2 \cdot \mathbf{R}_2)_{1, j_2} \hat{A}_{3, i_3, j_3} \dots \hat{A}_{n, i_n, 1} \text{ for all } i \end{aligned} \quad (4.6)$$

Now we invoke the inductive step multiple times. Let  $\mathbf{A}_{2,\ell}$  be the  $\ell$ th row of  $\mathbf{A}_2$ , and let  $\widehat{\mathbf{A}}_{2,\ell} = \mathbf{A}_{2,\ell} \cdot \mathbf{R}_2$ . Since  $\mathbf{A}_2 \cdot \mathbf{A}_3 \dots \mathbf{A}_n \neq 0$ , there is some  $\ell$  such that  $\mathbf{A}_{2,\ell} \cdot \mathbf{A}_3 \dots \mathbf{A}_n \neq 0$ . Then the matrices  $\mathbf{A}_{2,\ell}, \mathbf{A}_3, \dots, \mathbf{A}_n$  satisfy the first set of requirements of Theorem 4.2 for  $n - 1$ . Moreover, the right side of Equation 4.5 gives an allowable polynomial that is identically zero as a polynomial over the  $\mathbf{R}_k, k \geq 2$ , and therefore, by induction, it is identically 0 as a polynomial over  $\widehat{\mathbf{A}}_{2,\ell}, \widehat{\mathbf{A}}_3, \dots, \widehat{\mathbf{A}}_n$ . This shows us that

$$\alpha_{j,i,j_2,i_3,\dots,j_{n-1},i_n} = 0 \text{ for all } j \neq i \quad (4.7)$$

Next, Let  $\mathbf{A}'_2 = \mathbf{A}_1 \cdot \mathbf{A}_2$ , and let  $\widehat{\mathbf{A}}'_2 = \mathbf{A}'_2 \cdot \mathbf{R}_2$ . There are two cases:

- $\mathbf{A}_1 \cdot \mathbf{A}_2 \dots \mathbf{A}_n \neq 0$ . Then  $\mathbf{A}'_2 \cdot \mathbf{A}_3 \dots \mathbf{A}_n \neq 0$ . Therefore,  $\mathbf{A}'_2, \mathbf{A}_3, \dots, \mathbf{A}_n$  satisfy the first set of requirements in Theorem 4.2. Moreover, for each  $i$ , Equation 4.6 gives an allowable polynomial that is identically zero as a polynomial over the  $\mathbf{R}_k, k \geq 2$ . Therefore, by induction, the polynomial is identically zero as a polynomial over  $\widehat{\mathbf{A}}'_2, \widehat{\mathbf{A}}_3, \dots, \widehat{\mathbf{A}}_n$ . This means

$$\alpha_{i,i,j_2,i_3,\dots,j_{n-1},i_n} = 0 \text{ for all } i$$

Combining with Equation 4.7, we have that all the  $\alpha$  values are 0. Therefore  $p$  is identically zero as a polynomial over the  $\widehat{\mathbf{A}}_1, \widehat{\mathbf{A}}_2, \dots, \widehat{\mathbf{A}}_n$ .

- $\mathbf{A}_1 \cdot \mathbf{A}_2 \dots \mathbf{A}_n = 0$ . Then  $\mathbf{A}'_2 \cdot \mathbf{A}_3 \dots \mathbf{A}_n = 0$ . However,  $\mathbf{A}'_2 \cdot \mathbf{A}_3 \dots \mathbf{A}_{n-1} = \mathbf{A}_1 \cdot \mathbf{A}_2 \dots \mathbf{A}_{n-1} \neq 0$  and  $\mathbf{A}_3 \dots \mathbf{A}_4 \dots \mathbf{A}_n \neq 0$  (since otherwise  $\mathbf{A}_2 \dots \mathbf{A}_3 \dots \mathbf{A}_n = 0$ , contradicting the assumptions of Theorem 4.2). Therefore,  $\mathbf{A}'_2, \mathbf{A}_3, \dots, \mathbf{A}_n$  satisfy the second set of requirements in Theorem 4.2. By induction, for each  $i$ , the polynomial in Equation 4.6 must therefore be a multiple  $\gamma_i \widehat{\mathbf{A}}'_2 \cdot \widehat{\mathbf{A}}_3 \dots \widehat{\mathbf{A}}_n$  of the matrix product polynomial. This is equivalent to

$$\begin{aligned} \alpha_{i,i,j_2,i_3,\dots,j_{n-1},i_n} &= 0 \text{ if } j_k \neq i_{k+1} \text{ for any } k \\ \alpha_{i,i,i_3,i_3,\dots,i_n,i_n} &= \gamma_i \end{aligned}$$

This means we can write

$$\begin{aligned} \alpha'_{j,i,\ell} &= 0 \text{ for all } j \neq i \text{ (by Equation 4.7 and the definition of } \alpha'_{i,j,\ell}) \\ \alpha'_{i,i,\ell} &= \gamma_i \sum_{i_3,\dots,i_n} (\mathbf{A}_2 \cdot \mathbf{R}_2)_{\ell,i_3} \widehat{A}_{3,i_3,i_4} \dots \widehat{A}_{n,i_n,1} = \gamma_i (\mathbf{A}_2 \cdot \mathbf{A}_3 \dots \mathbf{A}_n)_{\ell,1} \end{aligned}$$

Since  $\alpha'_{i,i,\ell} = \beta_\ell$  for all  $i$  and the product  $\mathbf{A}_2 \cdot \mathbf{A}_3 \dots \mathbf{A}_n$  is non-zero, we have that  $\gamma_i = \gamma$  is the same for all  $i$ . Therefore,

$$\alpha_{i,i,i_3,i_3,\dots,i_n,i_n} = \gamma \text{ for all } i, i_3, \dots, i_n$$

meaning  $p$  is a multiple of the matrix product polynomial, as desired.

It remains to show the case where  $\mathbf{A}_1$  has zero entries. Since  $\mathbf{A}$  is non-zero (as a consequence of our assumptions), and  $\mathbf{A}$  is a single row vector, it is straightforward to build an invertible matrix  $\mathbf{B}$  such that  $\mathbf{A}'_1 = \mathbf{A}_1 \cdot \mathbf{B}$  is non-zero in every coordinate.

Let  $\mathbf{A}'_2 = \mathbf{B}^{-1} \mathbf{A}_2$ . Let  $\mathbf{R}'_1 = \mathbf{B}^{-1} \cdot \mathbf{R}_1$ ,  $\widehat{\mathbf{A}}'_1 = \mathbf{A}'_1 \cdot \mathbf{R}'_1 = \widehat{\mathbf{A}}_1$ , and  $\widehat{\mathbf{A}}'_2 = (\mathbf{R}'_1)^{adj} \cdot \mathbf{A}'_2 \cdot \mathbf{R}_2 = \widehat{\mathbf{A}}_2$ . Now  $\mathbf{A}'_1, \mathbf{A}'_2, \mathbf{A}_3, \dots, \mathbf{A}_n$  satisfy the same conditions of Theorem 4.2 as the original  $\mathbf{A}_k$ . Moreover,  $p$

is still allowable as a polynomial over  $\widehat{\mathbf{A}}'_1, \widehat{\mathbf{A}}'_2, \widehat{\mathbf{A}}'_3, \dots, \widehat{\mathbf{A}}'_n$ . Moreover, we can relate  $p$  as a polynomial over  $\mathbf{R}_k$  to  $p$  as a polynomial over  $\mathbf{R}'_1, \mathbf{R}_2, \dots, \mathbf{R}_{n-1}$  by a linear transformation on the  $\mathbf{R}_1$  variables. Therefore,  $p$  is identically zero as a polynomial over the  $\mathbf{R}_k$  if and only if it is identically zero as a polynomial over  $\mathbf{R}'_1, \mathbf{R}_2, \dots, \mathbf{R}_n$ . Thus we can invoke Theorem 4.2 on  $\mathbf{A}'_1, \mathbf{A}'_2, \dots, \mathbf{A}_n$  using the same polynomial  $p$ , and arrive at the desired conclusion. This completes the proof.  $\square$

## 5 Sketch of VBB Security Proof

We now explain how to use Theorem 4.2 to prove the VBB security of our obfuscator. The full security proof appears in Section 7. In this sketch, we pay special attention to the steps in our proof that deviate from previous works [BGK<sup>+</sup>14, AGIS14]. We also state a definition and lemma that will be used in Section 6 to prove that encodings of zero cannot be created when the function being obfuscated is evasive.

The adversary is given an obfuscation of a branching program  $BP$ , which consists of a list of handles corresponding to elements in the graded encoding. The adversary can operate on these handles using the graded encoding interface, which allows performing algebraic operations and zero testing. Our goal is to build a simulator that has oracle access only to the output of  $BP$ , and is yet able to simulate all of the handles and interfaces seen by the adversary. Formally, we prove the following theorem.

**Theorem 5.1.** *If  $BP^S$  outputs non-shortcutting branching programs, then for any PPT adversary  $\mathcal{A}$ , there is a PPT simulator  $\text{Sim}$  such that*

$$\left| \Pr[\mathcal{A}^{\mathcal{M}}(\mathcal{O}^{\mathcal{M}}(BP^S)) = 1] - \Pr_{BP \leftarrow BP^S}[\text{Sim}^{BP}(\ell, d_0, \dots, d_\ell, \text{inp}_0, \text{inp}_1) = 1] \right| < \text{negl}.$$

The simulator will choose random handles for all of the encodings in the obfuscation, leaving the actual entries of the  $\mathbf{D}_{i,b_0,b_1}$  as formal variables<sup>5</sup>. Simulating the algebraic operations is straightforward; the bulk of the security analysis goes in to answering zero-test queries. Any handle the adversary queries the zero test oracle on corresponds to some polynomial  $p$  on the variables  $\mathbf{D}_{i,b_0,b_1}$ , which the adversary can determine by inspecting the queries made by the adversary so far.

The simulator's goal is to decide if  $p$  evaluates to zero, when the formal variables in the  $\mathbf{D}_{i,b_0,b_1}$  are set to the values in the randomized matrix branching program  $BP'$ . However, the simulator does not know  $BP'$ , and must instead determine if  $p$  gives zero knowing only the outputs of  $BP$ .

The analysis of [BGK<sup>+</sup>14] and [AGIS14] first simplifies the problem of determining if  $p$  evaluates to zero, using Lemma 5.3 below.

**Definition 5.2.** A *single-input element* for an input  $x$  is a polynomial  $p_x$  whose variables are the  $\mathbf{C}_{i,x_{\text{inp}_0(i)},x_{\text{inp}_1(i)}}$  matrices, and  $p_x$  is allowable in the sense of Definition 4.1: each monomial in the expansion of  $p_x$  contains exactly one variable from each of the  $\mathbf{C}_{i,x_{\text{inp}_0(i)},x_{\text{inp}_1(i)}}$  matrices.

**Lemma 5.3** (Adapted from [BGK<sup>+</sup>14, AGIS14]). *Any polynomial  $p$  over the obfuscation  $\mathcal{O}^{\mathcal{M}}(BP^S)$  can be efficiently decomposed into a sum  $p = \sum_{x \in D} \alpha_x p_x$ , where  $\alpha_x = \prod_{i \in [\ell]} \alpha_{i,x_{\text{inp}_0(i)},x_{\text{inp}_1(i)}}$ , each  $p_x$  is a single-input element for input  $x$ , and  $|D|$  is polynomial in the circuit size of  $p$ .*

<sup>5</sup>The simulator does not know the branching program, and so it has no way of actually sampling the  $\mathbf{D}_{i,b_0,b_1}$ .

Due to the independence of the  $\alpha_x$  variables, it can be shown that  $p$  evaluates to zero iff each of the polynomials  $p_x$  do. Thus Lemma 5.3, along with some extra analysis of our own to handle multi-bit outputs, reduces the general problem to the following simpler problem. There is an unknown sequence of matrices  $\mathbf{A}_i \in \mathbb{Z}_q^{d_{i-1} \times d_i}$  for  $i \in [\ell]$ , where  $d_0 = d_\ell = 1$  (the shapes of the  $\mathbf{A}_i$  ensure that the product  $\prod_{i \in [\ell]} \mathbf{A}_i$  is valid and results in a scalar). We are also given an allowable polynomial  $p'$  on matrices of random variables  $\widehat{\mathbf{A}}_i$ . Our goal is to determine, if the  $\widehat{\mathbf{A}}_i$  are set to the Kilian-randomized matrices  $\widehat{\mathbf{A}}_i = \mathbf{R}_{i-1} \cdot \mathbf{A} \cdot \mathbf{R}_i^{adj}$ , whether or not  $p'$  evaluates to zero. We note that by applying the Schwartz-Zippel lemma, it suffices to decide if  $p'$  is *identically* zero, when considered a polynomial over the formal variables  $\mathbf{R}_i$ .

It is not hard to see that this simpler problem is impossible in general:  $p'$  could be the polynomial computing the iterated matrix product  $\prod_{k \in [\ell]} \widehat{\mathbf{A}}_k$ , which is equal to  $\prod_{i \in [\ell]} \mathbf{A}_i$ . Therefore, to decide if  $p'$  is identically zero in this case, we at a minimum need to know if  $\prod_{i \in [\ell]} \mathbf{A}_i$  evaluates to 0.

The analysis shows that the  $\mathbf{A}_i$  are actually equal to  $\mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}}$  for some (known) input  $x$ , where  $\mathbf{B}_{i, b_0, b_1}$  are the matrices in the branching program  $BP$ . Therefore, we can determine if  $\prod_{i \in [\ell]} \mathbf{A}_i = 0$  by querying the  $BP$  oracle on  $x$ . In the case where  $p'$  is the iterated matrix product, this allows us to determine if  $p'$  is identically 0. What about other, more general, polynomials  $p'$ ?

In previous works,  $\mathbf{A}_1$  and  $\mathbf{A}_\ell$  are bookend vectors, and the  $\mathbf{A}_i$  for  $k \in [2, \ell - 1]$  are square invertible matrices. In this setting, Kilian's statistical simulation theorem allows us to sample from the distribution of  $\widehat{\mathbf{A}}_i$  knowing only the product of the  $\mathbf{A}_i$ , but not the individual values. Then we can apply  $p'$  to the sample, and the Schwartz-Zippel lemma shows that  $p'$  will evaluate to zero, with high probability, if and only if it is identically zero. This allows deciding if  $p'$  is identically zero.

In our case, we cannot sample from the correct distribution of  $\widehat{\mathbf{A}}_i$ . Instead, we observe that our branching program is non-shortcutting, which means the  $\mathbf{A}_i$  and  $p'$  satisfy the requirements of Theorem 4.2. Theorem 4.2 implies something remarkably strong: if  $p'$  is not (a multiple of) the iterated matrix product, it *cannot possibly* be identically zero as a polynomial over the formal variables  $\mathbf{R}_k$ . Thus, we first decide if  $p'$  is a multiple of the iterated matrix product, which is possible using the Schwartz-Zippel lemma. If  $p'$  is a multiple, then we know it is identically zero if and only if the product  $\prod_{i \in [\ell]} \mathbf{A}_i$  is zero, and we know whether this product is zero by using our  $BP$  oracle.

## 6 Obfuscating Evasive Functions with No Zero Encodings

In this section we show that when the obfuscator of Section 3 is applied to an *evasive* function, any poly-time adversary will have only negligible probability in constructing an encoding of 0.

**Definition 6.1.** We say that an adversary  $\mathcal{A}$  *constructs an encoding of 0* if it ever receives a handle  $h$  from  $\mathcal{M}$  such that (a)  $h$  maps to an encoding of 0 in  $\mathcal{M}$ 's table, and (b) the polynomial that produced the encoding is not identically zero as a polynomial over its formal variables.

**Theorem 6.2.** *Let  $\mathcal{O}$  be the obfuscator from Section 3, and let  $BP^S$  sample an evasive function family. Then for any PPT adversary  $\mathcal{A}$ :*

$$\Pr \left[ \mathcal{A}^{\mathcal{M}}(\mathcal{O}^{\mathcal{M}}(BP^S)) \text{ constructs an encoding of } 0 \right] < \text{negl}(\ell).$$

One can never prevent an adversary from constructing a trivial encoding of 0 by computing  $e - e$  for some encoding  $e$  that it has. (More generally, any identically zero polynomial will produce a

trivial encoding of 0.) However in all candidate constructions of graded encoding schemes, such an operation always produces *the integer* 0, which contains no information. Indeed, it seems unlikely that a plausible candidate would not have this property.

To prove Theorem 6.2, we first show that any element that is not at the top level  $\mathbb{U}$  can be “completed” to the top level by multiplying with other basic elements output by the obfuscator. This is a consequence of our use of strong straddling sets.

**Definition 6.3.** For  $i \in [\ell]$  and  $b \in \{0, 1\}$ , an element encoded at level  $S_{j, b_0, b_1}$  implies  $x_i = b$  if either  $\text{inp}_0(j) = i$  and  $b_0 = b$  or  $\text{inp}_1(j) = i$  and  $b_1 = b$ .

**Lemma 6.4.** Let  $R := \{\mathbf{D}_{i, b_0, b_1}]_{S_{i, b_0, b_1}}\}$  be the basic elements output by the obfuscator  $\mathcal{O}$ , and let  $[r]_S$  be any valid element created by a polynomial  $p$  over  $R$ .

Then there exists a set of elements  $R' \subseteq R$  such that  $[r]_S \times \prod_{z \in R'} z$  is a valid element at level  $\mathbb{U}$ , and further  $R'$  can be efficiently found.

**Proof.** We say that  $p$  touches layer  $j \in [n]$  if any leaf of  $p$  is a basic element from layer  $j$  (cf. [MSW14, Def. 4.2]).  $S$  uniquely determines the layers touched by  $p$  and vice versa (though not necessarily the specific matrices touched in each layer); in particular,  $p$  touches every layer iff  $S = \mathbb{U}$ . Thus we construct  $R'$  to contain one basic element from each layer that is not touched by  $p$ . If  $S = \mathbb{U}$  then the lemma holds trivially with  $R' := \emptyset$ , so assume  $S \neq \mathbb{U}$  and let  $J \subseteq [n]$  be the set of layers not touched by  $p$ . Let  $I := \{\text{inp}_0(j), \text{inp}_1(j) \mid j \in J\} \subseteq [\ell]$  be the set of all indices that are read in some untouched layer.

We claim that there is a sequence  $(b_i)_{i \in I} \in \{0, 1\}^{|I|}$  such that for every  $i \in I$ ,  $p$ 's leaves do not contain any basic element that implies  $x_i = 1 - b_i$ . Fix any  $i \in I$ . Recall that  $\mathbb{U}_i \subset \mathbb{U}$  is the universe set for index  $i$ , and note that we must have  $\mathbb{U}_i \not\subseteq S$  because some layer that reads index  $i$  is untouched. If  $\mathbb{U}_i \cap S = \emptyset$ , then  $p$ 's leaves do not contain a basic element that implies  $x_i = 0$  nor one that implies  $x_i = 1$ ; in this case we can take  $b_i = 0$ . If instead  $\mathbb{U}_i \cap S \not\subseteq \{\emptyset, \mathbb{U}_i\}$ , then by Lemma 2.11 there is a unique  $b_i \in \{0, 1\}$  for which there exists  $J' \subset [n]$  such that

$$\mathbb{U}_i \cap S = \bigcup_{j' \in J'} S_{j', b_i}^i.$$

(Recall that each  $S_{j', b_i}^i$  comes from the strong straddling set system over  $\mathbb{U}_i$ .) Thus  $p$ 's leaves do not contain any basic element that implies  $x_i = 1 - b_i$ .

Finally let  $R'$  contain, for each  $j \in J$ , an arbitrary entry from the  $(b_{\text{inp}_0(j)}, b_{\text{inp}_1(j)})$ th matrix in layer  $j$ . Formally,  $R' := \{\mathbf{D}_{j, b_{\text{inp}_0(j)}, b_{\text{inp}_1(j)}}[0, 0] \mid j \in J\}$  which can be efficiently computed given  $e$ . Then  $[r]_S \times \prod_{z \in R'} z$  is valid by construction, and it is at level  $\mathbb{U}$  because it touches every layer.  $\square$

We now prove the main theorem of this section. The proof uses the simulator  $\text{Sim}$  of Theorem 5.1 in a non-black-box way, and specifically relies on properties of the decomposition  $p = \sum_x \alpha_x p_x$  given by Lemma 5.3. (We remark that evasive functions are inherently single-bit output, and thus we will not need the further decomposition given by Lemma 7.2 in Section 7.)

*Proof of Theorem 6.2.* For any PPT adversary  $\mathcal{A}$ , denote

$$\mathcal{P}'(\mathcal{A}) := \Pr \left[ \mathcal{A}^{\mathcal{M}}(\mathcal{O}^{\mathcal{M}}(BP^S)) \text{ constructs a level-}\mathbb{U} \text{ encoding of } 0 \right].$$

We first show that if  $\mathcal{P}'(\mathcal{A})$  is a noticeable function of  $\ell$  for some PPT  $\mathcal{A}$ , then  $BP^S$  cannot be evasive, in contradiction to our assumption. Next we use Lemma 6.4 to remove the assumption that  $\mathcal{A}$ 's encoding of 0 is at level  $\mathbb{U}$ .

Let  $f \leftarrow BP^S$  denote the function being obfuscated. Let  $\mathcal{A}$  be any PPT, and let  $\text{Sim}$  denote the corresponding simulator given by Theorem 5.1. We construct a new adversary  $\mathcal{B}$ , with oracle access to  $f$ , that finds an input  $x$  such that  $f(x) = 0$ .

$\mathcal{B}^f(1^\ell)$ :

1. Run  $\text{Sim}^f$ , which itself is running  $\mathcal{A}$ , up until the point where  $\mathcal{A}$  constructs a level- $\mathbb{U}$  encoding.
2. Decompose  $p = \sum_{x \in D} \alpha_x p_x$  as in Lemma 5.3. Check if  $f(x) = 0$  for any  $x \in D$ . If so, stop and output  $x$ ; otherwise, continue running  $\text{Sim}$  until  $\mathcal{A}$ 's next level- $\mathbb{U}$  encoding, and repeat.
3. If  $\text{Sim}$  halts, then output a random  $x \in \{0, 1\}^\ell$ .

Note that  $\mathcal{B}$ 's simulation of  $\mathcal{A}$ 's view is correct up to statistical distance  $\text{negl}(\ell)$ , because  $\text{Sim}$ 's is. The proof of Theorem 5.1 establishes that for any level- $\mathbb{U}$   $p$  constructed by  $\mathcal{A}$ ,

$$\Pr[p \text{ is an encoding of 0 but some } p_x \text{ is not}] < \text{negl}(\ell).$$

Further, Theorem 4.2 establishes that if  $p_x$  is not identically zero (and some  $p_x$  must not be since  $p$  is not), then  $p_x$  is a multiple of the honest matrix product polynomial corresponding to input  $x$ . Thus  $p_x$  is an encoding of 0 iff  $f(x) = 0$ , and we have established  $\forall$  PPT  $\mathcal{A} \exists$  PPT  $\mathcal{B}$ :

$$\Pr[f(\mathcal{B}^f(1^\ell)) = 0] \geq \mathcal{P}'(\mathcal{A}) - \text{negl}(\ell). \quad (6.1)$$

Finally, let

$$\mathcal{P}(\mathcal{A}) := \Pr[\mathcal{A}^{\mathcal{M}}(\mathcal{O}^{\mathcal{M}}(BP^S)) \text{ constructs an encoding of 0}]$$

be the probability that we want to bound. We claim that  $\forall$  PPT  $\mathcal{A} \exists$  PPT  $\mathcal{A}'$ :  $\mathcal{P}'(\mathcal{A}') \geq \mathcal{P}(\mathcal{A})$ . Namely  $\mathcal{A}'$  runs  $\mathcal{A}$ , and for every encoding  $[r]_S$  with  $S \neq \mathbb{U}$  created by  $\mathcal{A}$ ,  $\mathcal{A}'$  also creates the level- $\mathbb{U}$  encoding  $[r']_{\mathbb{U}} := [r]_S \times \prod_{z \in R'} z$  guaranteed by Lemma 6.4. Note that if  $[r]_S$  encodes 0 then  $[r']_{\mathbb{U}}$  must encode 0 as well, so we have  $\mathcal{P}'(\mathcal{A}') \geq \mathcal{P}(\mathcal{A})$ . Combining this with (6.1), we complete the proof: if  $\exists$  PPT  $\mathcal{A}$  such that  $\mathcal{P}(\mathcal{A})$  is a noticeable function of  $\ell$ , then  $BP^S$  does not sample an evasive function family.  $\square$

## 6.1 Bootstrapping to P/poly

In this subsection, we show that if there exists a VBB (resp. iO) obfuscator  $\mathcal{O}_{\text{NC}^1}$  for evasive functions computed by log-depth circuits (or more generally poly-size BPs), and a semantically-secure fully homomorphic encryption (FHE) scheme, then there exists a VBB (resp. iO) obfuscator  $\mathcal{O}_{\text{P/poly}}$  for evasive functions computed by poly-size circuits (i.e. P/poly). These proofs follow previous ‘‘bootstrapping’’ proofs, namely from [GGH<sup>+</sup>13b] in the case of iO and from [BR14] in the case of VBB. However, we take care to ensure that the proofs go through *even under the assumption that  $\mathcal{O}_{\text{NC}^1}$  is only an obfuscator for evasive functions*; indeed, this is the main reason that we re-prove these theorems here.

We refer the reader to [GGH<sup>+</sup>13b, App. B] for background material on FHE, statistical simulation-sound non-interactive zero-knowledge proofs, and low-depth proofs. Briefly, a non-interactive proof is low-depth if the verifier can be implemented in  $\text{NC}^1$ .

### 6.1.1 Bootstrapping for VBB obfuscation

Let  $\mathcal{C}$  be an evasive circuit collection on  $n$ -bit inputs, where each  $C \in \mathcal{C}$  has size  $|C| \leq p(n) = \text{poly}(n)$ . Let  $U$  be a universal circuit evaluating circuits of size  $p(n)$  on  $n$ -bit inputs; in particular,  $U(C, x) = C(x)$  for every  $C \in \mathcal{C}$  and  $x \in \{0, 1\}^n$ . Let  $\mathcal{O}_{\text{NC}^1}$  be an VBB obfuscator for all evasive circuit collections in  $\text{NC}^1$ . We denote an FHE scheme by the tuple of PPT algorithms  $(\text{FHE.KeyGen}, \text{FHE.Enc}, \text{FHE.Dec}, \text{FHE.Eval})$ , and we assume that  $\text{FHE.Dec}$  can be implemented in  $\text{NC}^1$ . The construction has two algorithms: Obfuscate and Evaluate.

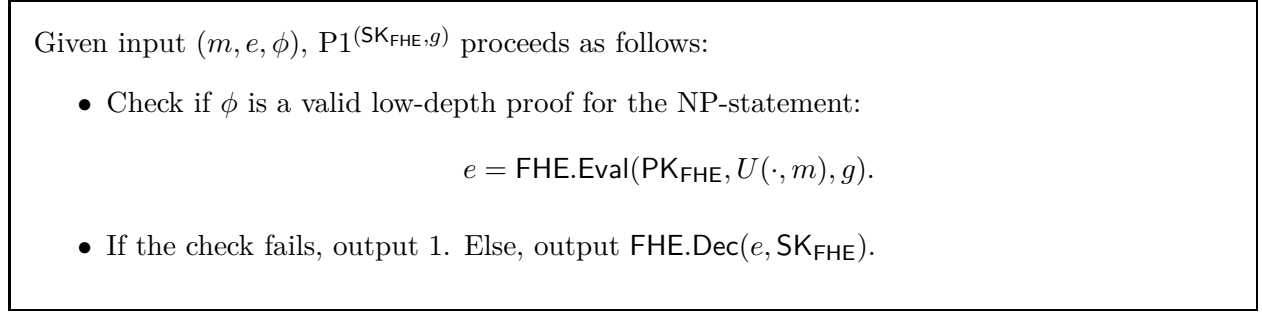


Figure 1

**Obfuscate** $(1^n, C \in \mathcal{C})$ :

- Generate  $(\text{PK}_{\text{FHE}}, \text{SK}_{\text{FHE}}) \leftarrow \text{FHE.KeyGen}(1^n)$ . If we are using a leveled FHE scheme, the number of levels should be set to the depth of  $U$ .
- Encrypt  $g = \text{FHE.Enc}(\text{PK}_{\text{FHE}}, C)$ .
- Generate  $O := \mathcal{O}_{\text{NC}^1}(\text{P1}^{(\text{SK}_{\text{FHE}}, g)})$ , where  $\text{P1}$  is the  $\text{NC}^1$  circuit described in Figure 1.
- Output the obfuscation  $\sigma := (O, \text{PK}_{\text{FHE}}, g)$ .

**Evaluate** $(\sigma = (O, \text{PK}_{\text{FHE}}, g), m)$ : The Evaluate algorithm takes in the obfuscation output  $\sigma$  along with an input  $m$ , and computes  $C(m)$  as follows.

- Compute  $e = \text{FHE.Eval}(\text{PK}_{\text{FHE}}, U(\cdot, m), g)$ .
- Compute a low depth proof  $\phi$  that  $e$  was computed correctly.
- Run  $O(m, e, \phi)$  and output the result.

**Correctness.** First, we verify that the size of the circuit evaluating  $\text{P1}$  is in  $\text{NC}^1$ . The first step of the program is in  $\text{NC}^1$  because we apply a low-depth proof and so the the verifier can be implemented in  $\text{NC}^1$ . The next step of the program is also in  $\text{NC}^1$  because we use an FHE scheme in which decryption can be done in  $\text{NC}^1$ . Hence, the whole circuit describing  $\text{P1}$  is in  $\text{NC}^1$ .

The correctness of FHE encryption means that  $g$  will be an encryption of the circuit  $C$ . And the correctness of FHE evaluation means that  $e$  will be an encryption of  $C(m)$ . The correctness



of  $\mathcal{O}_{\text{NC}^1}$  guarantees that the obfuscation of the program  $P1(SK_{\text{FHE}}, g)$  will be executed faithfully. Thus, on a correct execution, step 1 of the program  $P1$  will pass and the output will be  $C(m)$ .

**Security Proof.** Before proving security for the scheme described above, we observe that the function family  $\{P1^{(SK_{\text{FHE}}, g)}\}$  defined above is evasive when  $\mathcal{C}$  is, even in the presence of  $(PK_{\text{FHE}}, g)$ .

**Lemma 6.5.** *Let  $\mathcal{P}1$  be the collection of circuits  $\{P1^{(SK_{\text{FHE}}, g)}\}$  which is indexed by  $C \leftarrow \mathcal{C}$  and  $(PK_{\text{FHE}}, SK_{\text{FHE}}) \leftarrow \text{FHE.KeyGen}$ , where  $g := \text{FHE.Enc}(PK_{\text{FHE}}, C)$ .*

*If  $\mathcal{C}$  is evasive, then  $\mathcal{P}1$  is evasive with auxiliary input  $\text{Aux}(P1^{(SK_{\text{FHE}}, g)}) := (PK_{\text{FHE}}, g)$ .*

**Proof.** We first show that  $\mathcal{C}$  is evasive with auxiliary input  $\text{Aux}$ . Let  $C'$  be any fixed circuit of the same size as those in  $\mathcal{C}$ . By the semantic security of the FHE scheme, for every  $C \in \mathcal{C}$ , no PPT adversary can distinguish between  $(PK_{\text{FHE}}, g := \text{FHE.Enc}(PK_{\text{FHE}}, C))$  and  $(PK_{\text{FHE}}, g' := \text{FHE.Enc}(PK_{\text{FHE}}, C'))$  except with negligible probability. This holds even for PPTs that know  $C$  and  $C'$ , and in particular for PPTs with oracle access to  $C$ .

Let  $\mathcal{B}^{(\cdot)}$  be any PPT attacking the evasiveness of  $\mathcal{C}$ , and let  $\mathcal{D}^C$  be the PPT that, on input  $x$ , runs  $\mathcal{B}^C(x)$  and outputs  $C(\mathcal{B}^C(x))$ . We have

$$|\Pr[\mathcal{D}^C(PK_{\text{FHE}}, g) = 0] - \Pr[\mathcal{D}^C(PK_{\text{FHE}}, g') = 0]| < \text{negl}(n)$$

by semantic security. Because the distribution  $(PK_{\text{FHE}}, g')$  is independent of the distribution on  $\mathcal{D}$ 's oracle, we have  $\Pr[\mathcal{D}^C(PK_{\text{FHE}}, g') = 0] < \text{negl}(n)$  by  $\mathcal{C}$ 's evasiveness. Thus

$$\Pr[C(\mathcal{B}^C(PK_{\text{FHE}}, g)) = 0] = \Pr[\mathcal{D}^C(PK_{\text{FHE}}, g) = 0] < \text{negl}(n). \quad (6.2)$$

We now prove the lemma. Let  $P := P1^{(SK_{\text{FHE}}, g)}$  over a uniform choice of  $C \leftarrow \mathcal{C}$  and  $(PK_{\text{FHE}}, SK_{\text{FHE}}) \leftarrow \text{FHE.KeyGen}$ . Assume there exists a PPT  $\mathcal{A}^{(\cdot)}$  such that

$$\Pr\left[P\left(\mathcal{A}^P(PK_{\text{FHE}}, g)\right) = 0\right] \geq \epsilon(n)$$

for a noticeable function  $\epsilon(\cdot)$ . We will show there exists another PPT  $\mathcal{B}$  that contradicts (6.2).

First, observe that if  $\mathcal{A}$  outputs  $(m, e, \phi)$  such that  $P(m, e, \phi) = 0$ , then the input must pass the check step (first step inside  $P1^{(SK_{\text{FHE}}, g)}$ ), because the output of the program is 1 if the check step fails. This is a consequence of the perfect soundness of the interactive proof system we use (cf. [GGH<sup>+</sup>13b, App. B]). Further, because our FHE scheme has perfect correctness and  $g$  is an encryption of  $C$ , if  $(m, e, \phi)$  passes the check step and  $P(m, e, \phi) = 0$  then  $C(m) = 0$ .

Now we describe the PPT  $\mathcal{B}^C$ . On input  $(PK_{\text{FHE}}, g)$ ,  $\mathcal{B}^C$  runs  $\mathcal{A}^{(\cdot)}(PK_{\text{FHE}}, g)$ . When  $\mathcal{A}$  makes an oracle query  $(m, e, \phi)$ ,  $\mathcal{B}$  checks if  $C(m) = 0$ . If so  $\mathcal{B}$  halts and outputs  $m$ , otherwise it returns 1 to  $\mathcal{A}$  and continues the simulation. If  $\mathcal{A}$  halts with output  $(m, e, \phi)$ ,  $\mathcal{B}$  outputs  $m$ .

By the previous observations,  $\mathcal{B}$ 's simulation of  $\mathcal{A}$  is perfect until  $\mathcal{A}$  makes a query  $(m, e, \phi)$  such that  $P(m, e, \phi) = 0$ . As we have shown, this query must satisfy  $C(m) = 0$ . Thus we have

$$\Pr\left[C\left(\mathcal{B}^C(PK_{\text{FHE}}, g)\right) = 0\right] \geq \Pr\left[P\left(\mathcal{A}^P(PK_{\text{FHE}}, g)\right) = 0\right] \geq \epsilon(n)$$

which contradicts (6.2). □

**Theorem 6.6.** *Let  $\mathcal{O}_{\text{NC}^1}$  be an VBB obfuscator for all  $\text{NC}^1$  evasive circuit collections, and let FHE be a perfectly correct FHE scheme whose decryption circuit is in  $\text{NC}^1$ . Then the obfuscation scheme described above is a VBB obfuscator for all poly-size evasive circuit collections.*

**Proof.** The proof follows [BR14, Lemma 4.3]. Let  $\mathcal{O}_{P/\text{poly}}$  denote the **Obfuscate** algorithm described above. Let  $\mathcal{C}$  be any poly-size evasive circuit collection, and let  $C \leftarrow \mathcal{C}$  be chosen uniformly. For any adversary  $\mathcal{A}$  taking inputs of the form  $(O, \text{PK}_{\text{FHE}}, g) = \mathcal{O}_P(C)$ , we define the simulator  $\text{Sim}_P^{\mathcal{C}}(1^n)$  with oracle access to  $C$  as follows:

- Generate  $(\text{PK}_{\text{FHE}}, \text{SK}_{\text{FHE}}) \leftarrow \text{FHE.KeyGen}(1^n)$ .
- Generate  $g'$  as an encryption under  $\text{PK}_{\text{FHE}}$  of a size- $p(n)$  circuit that always outputs 1.
- Define the program  $\text{P1}'$  as follows:

Given input  $(m, e, \phi)$ ,  $\text{P1}'$  proceeds as follows:

- Check if  $\phi$  is a valid low-depth proof for the NP-statement:
$$e = \text{FHE.Eval}(\text{PK}_{\text{FHE}}, U(\cdot, m), g').$$
- If the check fails, output 1. Else, output  $C(m)$ .

Note that the output of  $\text{P1}'$  can be computed by  $\text{Sim}_P^{\mathcal{C}}$  as it has oracle access to  $C$ .

- Let  $\mathcal{A}'$  be  $\mathcal{A}$  but with its second and third inputs hardwired to  $\text{PK}_{\text{FHE}}$  and  $g'$ . Let  $\text{Sim}_{\text{NC}^1}^{(\cdot)}$  be the corresponding simulator for  $\mathcal{O}_{\text{NC}^1}$  when considered against adversary  $\mathcal{A}'$ . Execute  $\text{Sim}_{\text{NC}^1}^{\text{P1}'}(1^n)$  (answering oracle queries by running  $\text{P1}'$ ), and return its output.

To prove that the above simulation succeeds (i.e. that  $\text{Sim}_P^{\mathcal{C}}(1^n)$  simulates the view of  $\mathcal{A}(\mathcal{O}_P(C))$ ), we use a hybrid argument as described below.

- **Hybrid 1 :**  
In the first hybrid, we just run  $\text{Sim}_P^{\mathcal{C}}(1^n)$  as described above and give its output.
- **Hybrid 2 :**  
We change  $g'$  to be the actual encryption of  $C$  rather than the encryption of a dummy circuit. The rest of the simulator works in the same manner. This hybrid is indistinguishable from the previous one because  $\text{Sim}_P^{\mathcal{C}}$  does not use  $\text{SK}_{\text{FHE}}$ , and therefore distinguishing this hybrid from the previous would contradict the semantic security of the homomorphic encryption scheme.
- **Hybrid 3 :**  
We now replace the execution of  $\text{Sim}_{\text{NC}^1}^{(\cdot)}$  by answering its oracle queries with  $\text{P1}$  rather than  $\text{P1}'$ . The verification process done by both programs, along with the correctness of  $\text{FHE.Eval}$  and  $\text{FHE.Dec}$ , guarantee that the output of the simulator in this hybrid is statistically close to that in the previous hybrid.
- **Hybrid 4 :**  
We now replace the execution of  $\text{Sim}_{\text{NC}^1}^{(\cdot)}$  with the actual execution of  $\mathcal{A}'(\mathcal{O}_{\text{NC}^1}(\text{P1})) = \mathcal{A}(\mathcal{O}_{\text{NC}^1}(\text{P1}), \text{PK}_{\text{FHE}}, g) = \mathcal{A}(\mathcal{O}_P(C))$ . Thus this hybrid is indistinguishable from the previous by the security of  $\mathcal{O}_{\text{NC}^1}$ . Note that here, it is crucial that  $\text{P1}$  is an evasive function even for an adversary that knows  $\text{PK}_{\text{FHE}}$  and  $g$ , which is proved in Lemma 6.5.

We conclude that  $\text{Sim}_P^{\mathcal{C}}(1^n) \equiv \text{Hybrid 1}$  is indistinguishable from  $\mathcal{A}(\mathcal{O}_P(C)) \equiv \text{Hybrid 4}$ .  $\square$

### 6.1.2 Bootstrapping for iO

Let  $\mathcal{C}$  be an evasive circuit collection on  $n$ -bit inputs, where each  $C \in \mathcal{C}$  has size  $|C| \leq p(n) = \text{poly}(n)$ . Let  $U$  be a universal circuit evaluating circuits of size  $p(n)$  on  $n$ -bit inputs; in particular,  $U(C, x) = C(x)$  for every  $C \in \mathcal{C}$  and  $x \in \{0, 1\}^n$ . Let  $\mathcal{O}_{\text{NC}^1}$  be an Indistinguishability obfuscator for all evasive circuit collections in  $\text{NC}^1$ . We denote an FHE scheme by the tuple of PPT algorithms  $(\text{FHE.KeyGen}, \text{FHE.Enc}, \text{FHE.Dec}, \text{FHE.Eval})$ , and we assume that  $\text{FHE.Dec}$  can be implemented in  $\text{NC}^1$ . The construction has two algorithms: Obfuscate and Evaluate.

**Obfuscate** $(1^n, C \in \mathcal{C})$ :

- Generate  $(\text{PK}_{\text{FHE}}^1, \text{SK}_{\text{FHE}}^1) \leftarrow \text{FHE.KeyGen}(1^n)$  and  $(\text{PK}_{\text{FHE}}^2, \text{SK}_{\text{FHE}}^2) \leftarrow \text{FHE.KeyGen}(1^n)$ . If we are using a leveled FHE scheme, the number of levels should be set to the depth of  $U$ .
- Encrypt  $g_1 = \text{FHE.Enc}(\text{PK}_{\text{FHE}}^1, C)$  and  $g_2 = \text{FHE.Enc}(\text{PK}_{\text{FHE}}^2, C)$ .
- Generate  $O := \mathcal{O}_{\text{NC}^1}(\text{P1}^{(\text{SK}_{\text{FHE}}^1, g_1, g_2)})$ , where P1 is the  $\text{NC}^1$  circuit described in Figure 2.
- Output the obfuscation  $\sigma := (O, \text{PK}_{\text{FHE}}^1, \text{PK}_{\text{FHE}}^2, g_1, g_2)$ .

**Evaluate** $(\sigma = (O, \text{PK}_{\text{FHE}}^1, \text{PK}_{\text{FHE}}^2, g_1, g_2), m)$ : The Evaluate algorithm takes in the obfuscation output  $\sigma$  along with an input  $m$ , and computes  $C(m)$  as follows.

- Compute  $e_1 = \text{FHE.Eval}(\text{PK}_{\text{FHE}}^1, U(\cdot, m), g_1)$  and  $e_2 = \text{FHE.Eval}(\text{PK}_{\text{FHE}}^2, U(\cdot, m), g_2)$ .
- Compute a low depth proof  $\phi$  that  $e_1$  and  $e_2$  were computed correctly.
- Run  $O(m, e_1, e_2, \phi)$  and output the result.

**Correctness:** First, we verify that the size of the circuit evaluating P1 is in  $\text{NC}^1$ . The first step of the program is in  $\text{NC}^1$  because we apply a low-depth proof and so the the verifier can be implemented in  $\text{NC}^1$ . The next step of the program is also in  $\text{NC}^1$  because we use an FHE scheme in which decryption can be done in  $\text{NC}^1$ . Hence, the whole circuit describing P1 is in  $\text{NC}^1$ .

The correctness of FHE encryption means that  $g_1$  and  $g_2$  will be encryptions of  $C$ . The correctness of FHE evaluation means that  $e_1$  and  $e_2$  will be encryptions of  $C(m)$ . The correctness of  $\mathcal{O}_{\text{NC}^1}$  guarantees that the obfuscation of the program  $\text{P1}^{(\text{SK}_{\text{FHE}}^1, g_1, g_2)}$  will be executed faithfully. Thus, on a correct execution, step 1 of the program P1 will pass and the output will be  $C(m)$ .

**Security proof.** We prove that for all  $C_0, C_1 \in \mathcal{C}$  there can be no polynomial time attacker  $\mathcal{A}$  that wins the indistinguishability obfuscation security game (Def. 2.4) with noticeable advantage.

**Theorem 6.7.** *Let  $\mathcal{O}_{\text{NC}^1}$  be an iO obfuscator for all  $\text{NC}^1$  evasive circuit collections, and let FHE be a perfectly correct FHE scheme whose decryption circuit is in  $\text{NC}^1$ . Then the obfuscation scheme described above is an iO obfuscator for all poly-size evasive circuit collections.*

Given input  $(m, e_1, e_2, \phi)$ ,  $P1^{(SK_{FHE}^1, g_1, g_2)}$  proceeds as follows:

- Check if  $\phi$  is a valid low-depth proof for the NP-statement:

$$e_1 = \text{FHE.Eval}(\text{PK}_{FHE}^1, U(\cdot, m), g_1) \wedge e_2 = \text{FHE.Eval}(\text{PK}_{FHE}^2, U(\cdot, m), g_2).$$

- If the check fails, output 1. Else, output  $\text{FHE.Dec}(e_1, \text{SK}_{FHE}^1)$ .

Figure 2

Given input  $(m, e_1, e_2, \phi)$ ,  $P2^{(SK_{FHE}^2, g_1, g_2)}$  proceeds as follows:

- Check if  $\phi$  is a valid low-depth proof for the NP-statement:

$$e_1 = \text{FHE.Eval}(\text{PK}_{FHE}^1, U(\cdot, m), g_1) \wedge e_2 = \text{FHE.Eval}(\text{PK}_{FHE}^2, U(\cdot, m), g_2)$$

- If the check fails, output 1. Else, output  $\text{FHE.Dec}(e_2, \text{SK}_{FHE}^2)$ .

Figure 3

We organize our proof into a sequence of hybrids. In the first hybrid, the challenger obfuscates  $C_0$ . We then gradually change the obfuscation in multiple hybrids into an obfuscation of  $C_1$ . We show that each successive hybrid is indistinguishable from the previous one, thus showing our obfuscator to have indistinguishability security.

- Hybrid 0:  
This hybrid corresponds to an honest execution of the indistinguishability obfuscation game where  $C_0$  is obfuscated.
- Hybrid 1:  
Same as hybrid 0 except we now generate  $g_1 = \text{FHE.Enc}(\text{PK}_{FHE}^1, C_0)$  and  $g_2 = \text{FHE.Enc}(\text{PK}_{FHE}^2, C_1)$ . Now  $g_1$  and  $g_2$  encrypt different circuits.
- Hybrid 2:  
We still generate  $g_1 = \text{FHE.Enc}(\text{PK}_{FHE}^1, C_0)$  and  $g_2 = \text{FHE.Enc}(\text{PK}_{FHE}^2, C_1)$  as in the previous hybrid. Now,  $O$  is generated as  $\mathcal{O}_{\text{NC}^1}(P2^{(SK_{FHE}^2, g_1, g_2)})$ , where P2 is defined in Figure 3.
- Hybrid 3:  
We now generate  $g_1 = \text{FHE.Enc}(\text{PK}_{FHE}^1, C_1)$  and  $g_2 = \text{FHE.Enc}(\text{PK}_{FHE}^2, C_1)$ .  $O$  is still generated as  $\mathcal{O}_{\text{NC}^1}(P2^{(SK_{FHE}^2, g_1, g_2)})$ .
- Hybrid 4:  
We still generate  $g_1 = \text{FHE.Enc}(\text{PK}_{FHE}^1, C_1)$  and  $g_2 = \text{FHE.Enc}(\text{PK}_{FHE}^2, C_1)$  as in the previous

hybrid. Now,  $O$  is generated as  $\mathcal{O}_{\text{NC}^1}(\text{P1}^{\text{SK}_{\text{FHE}}^1, g_1, g_2})$ . Thus this hybrid corresponds to an honest execution of the indistinguishability obfuscation game where  $C_1$  is obfuscated.

We now show that each pair of adjacent hybrids are indistinguishable, which proves Theorem 6.7.

**Claim 6.8.** *If the FHE scheme is IND-CPA secure, then no polynomial time attacker can distinguish between Hybrid 0 and Hybrid 1 with non-negligible probability.*

**Proof.** We show that if there is a polynomial time attacker A that has a non-negligible difference in advantage between Hybrid 0 and Hybrid 1, then there is a polynomial time algorithm B that breaks the IND-CPA security of our FHE scheme. B begins by running A and receiving  $C_0, C_1$ .

B first gets  $\text{PK}_{\text{FHE}}^2$  from the IND-CPA challenger. It generates  $(\text{PK}_{\text{FHE}}^1, \text{SK}_{\text{FHE}}^1) \leftarrow \text{FHE.KeyGen}(1^n)$ . Next, it gives the IND-CPA challenger  $C_0, C_1$  and gets back a ciphertext  $g$ . It sets  $g_2 = g$ . B then sets  $g_1 = \text{FHE.Enc}(\text{PK}_{\text{FHE}}^1, C_0)$  and  $O$  is created as  $\mathcal{O}_{\text{NC}^1}(\text{P1}^{\text{SK}_{\text{FHE}}^1, g_1, g_2})$ . Note that only the first secret key  $\text{SK}_{\text{FHE}}^1$  is needed to create  $O$  and B already has this.

If the IND-CPA challenger used the first message  $C_0$ , then we are exactly in hybrid 0. If it chose the second message  $C_1$ , then we are in Hybrid 1. Therefore, if an attacker A can distinguish between the two hybrids with non-negligible advantage, it will break the IND-CPA property of the FHE scheme.  $\square$

**Claim 6.9.** *If the indistinguishability property holds for  $\mathcal{O}_{\text{NC}^1}$ , then no polynomial time attacker can distinguish between Hybrid 1 and Hybrid 2.*

**Proof.** We show that if there is a polynomial time attacker A that has a non-negligible difference in advantage between Hybrid 1 and Hybrid 2, then there is a polynomial time algorithm B that breaks the  $\text{NC}^1$  indistinguishability obfuscator. B begins by running A and receiving  $C_0, C_1$ .

B first generates the two FHE private keys itself, keeping both secret keys. That is, it generates  $(\text{PK}_{\text{FHE}}^1, \text{SK}_{\text{FHE}}^1) \leftarrow \text{FHE.KeyGen}(1^n)$  and  $(\text{PK}_{\text{FHE}}^2, \text{SK}_{\text{FHE}}^2) \leftarrow \text{FHE.KeyGen}(1^n)$ . It then creates  $g_1 = \text{FHE.Enc}(\text{PK}_{\text{FHE}}^1, C_0)$  and  $g_2 = \text{FHE.Enc}(\text{PK}_{\text{FHE}}^2, C_1)$ . Next, it submits the circuits  $\text{P1}^{\text{SK}_{\text{FHE}}^1, g_1, g_2}$  and  $\text{P2}^{\text{SK}_{\text{FHE}}^2, g_1, g_2}$  to the indistinguishability obfuscator challenger. It receives back a program  $O'$  and sets  $O = O'$ .

Suppose that output of  $C_0$  is equivalent to the output of  $C_1$  on all inputs. Then both programs  $\text{P1}^{\text{SK}_{\text{FHE}}^1, g_1, g_2}$  and  $\text{P2}^{\text{SK}_{\text{FHE}}^2, g_1, g_2}$  will have the same output on all inputs. Both programs will halt and output 1 if the check does not pass. If the check does pass, this means that  $e_1$  and  $e_2$  are encryptions of the same message. This is due to the perfect correctness of FHE encryption and evaluation together with the fact that  $U(C_0, m) = U(C_1, m)$  for all  $m$ . Since both of the programs give the same output on all inputs, we are in a valid instance of the assumption.

If the indistinguishability obfuscator used the first circuit  $\text{P1}^{\text{SK}_{\text{FHE}}^1, g_1, g_2}$ , then we are exactly in Hybrid 1. If it chose the second circuit  $\text{P2}^{\text{SK}_{\text{FHE}}^2, g_1, g_2}$ , then we are in Hybrid 2. Therefore, if an attacker can distinguish between the two hybrids with non-negligible advantage, it will break the indistinguishability security of  $\mathcal{O}_{\text{NC}^1}$ . To apply  $\mathcal{O}_{\text{NC}^1}$  here, we note that both P1 and P2 are evasive even given  $\text{PK}_{\text{FHE}}^1, \text{PK}_{\text{FHE}}^2, g_1, g_2$  by the same argument as in Lemma 6.5.  $\square$

**Claim 6.10.** *If the FHE scheme is IND-CPA secure, then no polynomial time attacker can distinguish between Hybrid 2 and Hybrid 3 with non-negligible probability.*

**Proof.** We show that if there is a polynomial time attacker A that has a non-negligible difference in advantage between Hybrid 2 and Hybrid 3, then there is a polynomial time algorithm B that breaks the IND-CPA security of our FHE scheme. B begins by running A and receiving  $C_0, C_1$ .

B first gets  $\text{PK}_{\text{FHE}}^1$  from the IND-CPA challenger. It generates  $(\text{PK}_{\text{FHE}}^2, \text{SK}_{\text{FHE}}^2) \leftarrow \text{FHE.KeyGen}(1^n)$ . Next, it gives the IND-CPA challenger  $C_0, C_1$  and gets back a ciphertext  $g$ . It sets  $g_1 = g$ . B then sets  $g_2 = \text{FHE.Enc}(\text{PK}_{\text{FHE}}^2, C_1)$  and  $O$  is created as  $\mathcal{O}_{\text{NC}^1}(\text{P2}^{\text{SK}_{\text{FHE}}^2, g_1, g_2})$ . Note that only the second secret key  $\text{SK}_{\text{FHE}}^2$  is needed to create  $O$  and B already has this.

If the IND-CPA challenger used the first message  $C_0$ , then we are exactly in hybrid 2. If it chose the second message  $C_1$ , then we are in Hybrid 3. Therefore, if an attacker A can distinguish between the two hybrids with non-negligible advantage, it will break the IND-CPA property of the FHE scheme.  $\square$

**Claim 6.11.** *If the indistinguishability property holds for  $\mathcal{O}_{\text{NC}^1}$ , then no polynomial time attacker can distinguish between Hybrid 3 and Hybrid 4.*

**Proof.** We show that if there is a polynomial time attacker A that has a non-negligible difference in advantage between Hybrid 3 and Hybrid 4, then there is a polynomial time algorithm B that breaks the  $\text{NC}^1$  indistinguishability obfuscator. B begins by running A and receiving  $C_0, C_1$ .

B first generates the two FHE private keys itself, keeping both secret keys. That is, it generates  $(\text{PK}_{\text{FHE}}^1, \text{SK}_{\text{FHE}}^1) \leftarrow \text{FHE.KeyGen}(1^n)$  and  $(\text{PK}_{\text{FHE}}^2, \text{SK}_{\text{FHE}}^2) \leftarrow \text{FHE.KeyGen}(1^n)$ . It then creates  $g_1 = \text{FHE.Enc}(\text{PK}_{\text{FHE}}^1, C_1)$  and  $g_2 = \text{FHE.Enc}(\text{PK}_{\text{FHE}}^2, C_1)$ . Next, it submits the circuits  $\text{P1}^{\text{SK}_{\text{FHE}}^1, g_1, g_2}$  and  $\text{P2}^{\text{SK}_{\text{FHE}}^2, g_1, g_2}$  to the indistinguishability obfuscator challenger. It receives back a program  $O'$  and sets  $O = O'$ .

Both programs  $\text{P1}^{\text{SK}_{\text{FHE}}^1, g_1, g_2}$  and  $\text{P2}^{\text{SK}_{\text{FHE}}^2, g_1, g_2}$  will have the same output on all inputs. This is because both programs will halt and output 1 if the check does not pass and if the check does pass, since both  $e_1$  and  $e_2$  are encryptions of the same message, and due to perfect correctness of FHE encryption and evaluation, both programs will output  $C_1(m)$  for all  $m$ . Since both of the programs give the same output on all inputs, we are in a valid instance of the assumption.

If the indistinguishability obfuscator used the circuit  $\text{P2}^{\text{SK}_{\text{FHE}}^2, g_1, g_2}$ , then we are exactly in Hybrid 3. If it chose the circuit  $\text{P1}^{\text{SK}_{\text{FHE}}^1, g_1, g_2}$ , then we are in Hybrid 4. Therefore, if an attacker can distinguish between the two hybrids with non-negligible advantage, it will break the indistinguishability security of the  $\mathcal{O}_{\text{NC}^1}$ . To apply  $\mathcal{O}_{\text{NC}^1}$  here, we note that both P1 and P2 are evasive even given  $\text{PK}_{\text{FHE}}^1, \text{PK}_{\text{FHE}}^2, g_1, g_2$  by the same argument as in Lemma 6.5.  $\square$

## 7 VBB Security of our Construction

We now prove the virtual black box security of our construction; for an overview of the proof, see Section 5. The main result of this section is the following theorem, restated for convenience.

**Theorem 5.1.** *If  $BP^S$  outputs non-shortcutting branching programs, then for any PPT adversary  $\mathcal{A}$ , there is a PPT simulator  $\text{Sim}$  such that*

$$\left| \Pr[\mathcal{A}^{\mathcal{M}}(\mathcal{O}^{\mathcal{M}}(BP^S)) = 1] - \Pr_{BP \leftarrow BP^S}[\text{Sim}^{BP}(\ell, d_0, \dots, d_\ell, \text{inp}_0, \text{inp}_1) = 1] \right| < \text{negl}.$$

**Proof.** We construct a simulator  $\text{Sim}$  that, on input a description of an adversary  $\mathcal{A}$ , simulates the view of  $\mathcal{A}$  on input  $BP^{\mathcal{O}} = \mathcal{O}(BP^S)$ , given only oracle access to  $BP$ .  $\text{Sim}$  is also given  $\ell, d_0, \dots, d_\ell, \text{inp}_0, \text{inp}_1$ .

Most steps in the simulator are identical to [AGIS14], which the exception being the simulation of zero-test queries. First, the simulator emulates the obfuscator  $\mathcal{O}$  on  $BP$ . Since  $\text{Sim}$  only has

oracle access to  $BP$  and thus has no way to determine the matrices  $B_{i,b_0,b_1}$ ,  $\text{Sim}$  instead initializes  $\mathcal{M}$  with formal variables. More precisely,  $\text{Sim}$  will maintain a table of handles and corresponding level of encodings that have been initialized so far.  $\text{Sim}$  initially creates the table with random handles corresponding to the randomized matrices  $C_{i,b_0,b_1}$ .  $\text{Sim}$  then easily emulates all of the interfaces of  $\mathcal{M}$  except for zero testing. The simulator also computes the set system used for the encodings from  $\text{inp}_0, \text{inp}_1$ .

**Simulating Zero-test queries.** We now describe how to simulate zero-test queries by the adversary, given only oracle access to  $BP$ . As in [AGIS14], when the adversary submits a handle  $h$  for zero testing,  $\text{Sim}$  looks up the corresponding polynomial  $p$  in its table. As a first step, we decompose  $p$  into *single-input elements*; the following definition and lemma are restated from Section 5.

**Definition 5.2.** A *single-input element* for an input  $x$  is a polynomial  $p_x$  whose variables are the  $\widehat{C}_{i,x_{\text{inp}_0(i)},x_{\text{inp}_1(i)}}$  matrices, and  $p_x$  is allowable in the sense of Definition 4.1: each monomial in the expansion of  $p_x$  contains exactly one variable from each of the  $\widehat{C}_{i,x_{\text{inp}_0(i)},x_{\text{inp}_1(i)}}$  matrices.

**Lemma 5.3** (Adapted from [BGK<sup>+</sup>14, AGIS14]). *Any polynomial  $p$  over the obfuscation  $\mathcal{O}^{\mathcal{M}}(BP^S)$  can be efficiently decomposed into a sum  $p = \sum_{x \in D} \alpha_x p_x$ , where  $\alpha_x = \prod_{i \in [\ell]} \alpha_{i,x_{\text{inp}_0(i)},x_{\text{inp}_1(i)}}$ , each  $p_x$  is a single-input element for input  $x$ , and  $|D|$  is polynomial in the circuit size of  $p$ .*

The first part of Lemma 5.3 follows from the decomposition in previous works. The absence of bookends, the multi-bit outputs, and the singular and rectangular matrices does not affect this part of the simulation. Further, the strong straddling set systems we use satisfy all the properties required of the (standard) straddling set systems in [BGK<sup>+</sup>14, AGIS14]. The fact that  $p_x$  are allowable is not mentioned or proved in previous works, but follows easily from the graded encoding structure.

Because we have multi-bit outputs, we will actually need to decompose the polynomials even further.

**Definition 7.1.** A *single-input/single-output element* for an input  $x$  and output position  $(s, t) \in [d_0] \times [d_\ell]$  is a polynomial  $p_{x,s,t}$  whose variables are the  $\widehat{B}_{i,x_{\text{inp}_0(i)},x_{\text{inp}_1(i)}}$  matrices, and  $p_x$  is allowable in the sense, and  $p_{x,s,t}$  is allowable in the sense of Definition 4.1: each monomial in the expansion of  $p_x$  contains exactly one variable from each of the  $\widehat{B}_{i,x_{\text{inp}_0(i)},x_{\text{inp}_1(i)}}$  matrices. Moreover, the variable from the matrix  $\widehat{B}_{1,x_{\text{inp}_0(1)},x_{\text{inp}_1(1)}}$  comes from row  $s$ , and the variable from the matrix  $\widehat{B}_{\ell,x_{\text{inp}_0(\ell)},x_{\text{inp}_1(\ell)}}$  comes from column  $t$ .

**Lemma 7.2.** *Each single-input element  $p_x$  can be efficiently decomposed into a sum*

$$p_x = \sum_{s \in [d_0], t \in [d_1]} \beta_s \gamma_t p_{x,s,t}$$

where  $p_{x,s,t}$  are single input/single output elements for  $x, s, t$

**Proof.** Write the  $\mathbf{C}$  in terms of the  $\widehat{\mathbf{B}}$  and  $\mathbf{S}, \mathbf{T}$ , where  $\mathbf{S}$  and  $\mathbf{T}$  are the diagonal matrices with  $\beta_s$  and  $\gamma_t$  on the diagonal, respectively. That is,

$$\mathbf{C}_{1,b_1,b_2} = \mathbf{S} \cdot \widehat{\mathbf{B}}_{1,b_1,b_2} \quad \mathbf{C}_{\ell,b_1,b_2} = \widehat{\mathbf{B}}_{1,b_1,b_2} \cdot \mathbf{T} \quad \mathbf{C}_{i,b_1,b_2} = \widehat{\mathbf{B}}_{i,b_1,b_2} \text{ for each } i \in [2, \ell - 1]$$

For each  $s \in [d_0], t \in [d_\ell]$ , set  $\beta_s = 1, \gamma_t = 1$  and  $\beta_{s'} = \gamma_{t'} = 0$  for all  $s' \neq s, t' \neq t$ . Let  $p_{x,s,t}$  be the polynomial remaining. Then  $p_{x,s,t}$  is exactly a single-input/single-output element for  $x, s, t$ . Moreover, after doing this for all  $s, t$ , we have that

$$p_x = \sum_{s \in [d_0], t \in [d_1]} \beta_s \gamma_t p_{x,s,t}$$

as desired.  $\square$

Next, for each  $x \in D$ , we query the function oracle to learn  $BP(x)$ , and use  $BP(x)$  to determine an input distribution on which we test the various polynomials  $p_{x,s,t}$ . Starting at this point, our simulation and analysis departs from previous works. Existing works rely on Kilian [Kil88] simulation to argue that the distribution of test inputs matches the distribution in the actual obfuscator. This allows them to determine whether  $p_x$  should evaluate to zero or not with overwhelming probability.

Unfortunately for us, Kilian simulation only applies to square invertible matrices. Therefore, we need to modify the simulation and/or analysis to handle this.

Fix  $x, s, t$ , and let  $b_c^i = x_{\text{inp}_c(i)}$ . For  $i \in [0, \ell - 1]$ , let  $\widehat{\mathbf{A}}_i$  denote  $\widehat{\mathbf{B}}_{i, b_0^i, b_1^i}$  and  $\mathbf{A}_i = \mathbf{B}_{i, b_0^i, b_1^i}$ . Let  $\widehat{\mathbf{A}}_1$  be row  $s$  of  $\widehat{\mathbf{B}}_{1, b_0^1, b_1^1}$  and  $\widehat{\mathbf{A}}_\ell$  be row  $t$  of  $\widehat{\mathbf{B}}_{\ell, b_0^\ell, b_1^\ell}$ <sup>6</sup>. Then  $p_{x,s,t}$  is an allowable polynomial in the  $\widehat{\mathbf{A}}_i$

For each polynomial  $p_{x,s,t}$ , we determine whether  $p_{x,s,t}$  evaluates to zero. We do this as follows:

- If  $BP(x)[s, t] = 1$ , we choose totally random matrices  $\widehat{\mathbf{A}}_i$ , and test if  $p_{x,s,t}$  evaluates to zero on these matrices. If the result is zero, we say  $p_{x,s,t}$  evaluates to zero, and if the result is non-zero, we say  $p_{x,s,t}$  evaluates to non-zero. There are two cases:
  - $p_{x,s,t}$  is identically zero. Then our test will give zero with probability 1, and  $p_{x,s,t}$  evaluates to zero in the actual scheme with probability 1. Therefore, we correctly determine if  $p_{x,s,t}$  evaluates to zero.
  - $p_{x,s,t}$  is not identically zero. Then, by Schwartz-Zippel, our test will, with overwhelming probability, obtain non-zero, and we will report non-zero. In the actual scheme, since  $BP(x)[s, t] = 1$ , the  $\mathbf{A}_i$  satisfy the first set of requirements of Theorem 4.2. Therefore, since  $p_{x,s,t}$  is allowable, Theorem 4.2 shows that  $p_{x,s,t}$  is also *not* identically zero as a polynomial over the randomization matrices  $\mathbf{R}_i$ . Schwartz-Zippel then shows that in the actual scheme, with overwhelming probability,  $p_{x,s,t}$  will evaluate to non-zero. Thus we correctly guess whether  $p_{x,s,t}$  evaluates to non-zero with overwhelming probability.
- If  $BP(x)[s, t] = 0$ , we choose random matrices  $\widehat{\mathbf{A}}_i$  subject to the restriction that their product is zero, and test  $p_{x,s,t}$  on these matrices. This is done as follows. Choose random  $\mathbf{A}_i$  for  $i \in [\ell - 1]$ . Let  $\mathbf{v} = (v_1, \dots, v_{d_{\ell-1}})$  be the row vector  $\prod_{i=0}^{\ell-1} \mathbf{A}_i$ . Now we sample values  $w_2, \dots, w_{d_{\ell-1}}$  at random, and let  $\widehat{\mathbf{A}}_\ell$  be the column vector

$$\widehat{\mathbf{A}}_\ell = \begin{pmatrix} -\sum_{i=2}^{d_{\ell-1}} v_i w_i \\ v_1 w_2 \\ v_1 w_3 \\ \vdots \\ v_1 w_{d_{\ell-1}} \end{pmatrix} \quad (7.1)$$

---

<sup>6</sup>The simulator does not actually compute the  $\mathbf{A}_i$ ; we are just using them for the analysis.



Then  $\mathbf{v} \cdot \widehat{\mathbf{A}}_\ell = 0$ . We now make the following claim:

**Claim 7.3.** *If a polynomial  $p$ , after making the substitution in Equation 7.1, becomes identically zero, then  $p$  was originally a multiple of the matrix product polynomial.*

**Proof.**  $p$  being identically zero in the substitution in Equation 7.1 is equivalent to  $p$  being identically zero after making the substitution

$$\widehat{A}_{\ell,1,1} \leftarrow \frac{-\sum_{i=2}^{d_{\ell-1}} v_i \widehat{A}_{\ell,i,1}}{v_1}$$

If this substitution gives a zero polynomial, it must be that

$$\widehat{A}_{\ell,1,1} + \frac{\sum_{i=2}^{d_{\ell-1}} v_i \widehat{A}_{\ell,i,1}}{v_1}$$

divides  $p$ . Since  $p$  is a polynomial, we can remove the  $v_1$  in the denominator and conclude that, in fact,

$$v_1 \widehat{A}_{\ell,1,1} + \sum_{i=2}^{d_{\ell-1}} v_i \widehat{A}_{\ell,i,1} = \sum_{i=1}^{d_{\ell-1}} v_i \widehat{A}_{\ell,i,1}$$

divides  $p$ . But the polynomial above is exactly the matrix product polynomial, as desired.  $\square$

Now we test  $p_{x,s,t}$  on the samples  $\widehat{\mathbf{A}}_i$ . If the result is zero, we say  $p_{x,s,t}$  evaluates to zero, and if the result is non-zero, we say  $p_{x,s,t}$  evaluates to non-zero. There are two cases:

- $p_{x,s,t}$  is a multiple of the matrix product polynomial. Then our test will give zero with probability 1, and  $p_{x,s,t}$  evaluates to zero in the actual scheme with probability 1. Therefore, we correctly determine if  $p_{x,s,t}$  evaluates to zero.
- $p_{x,s,t}$  is not a multiple of the matrix product polynomial. Claim 7.3 then shows  $p_{x,s,t}$  must be not identically zero after making the substitutions in Equation 7.1. Therefore, Schwartz-Zippel shows that the polynomial evaluates to non-zero with overwhelming probability. Therefore, we will say the value is non-zero. In the actual scheme, since  $BP(x)[s,t] = 0$  and  $BP$  is non-shortcutting, the  $\mathbf{A}_i$  satisfy the second set of requirements for Theorem 4.2. Since  $p_{x,s,t}$  is not a multiple of the matrix product polynomial but is allowable, Theorem 4.2 shows that the polynomial is not identically zero as a polynomial over the randomization matrices  $\mathbf{R}_i$ . Schwartz-Zippel then shows that in the actual scheme, with overwhelming probability,  $p_{x,s,t}$  will evaluate to non-zero. Thus we correctly guess whether  $p_{x,s,t}$  evaluates to non-zero with overwhelming probability.

Therefore, we will correctly determine whether  $p_{x,s,t}$  evaluates to 0 for each  $x, s, t$  with overwhelming probability. Now recall that

$$p = \sum_{x \in D, s \in [d_0], t \in [d_\ell]} \left( \prod_{i \in [\ell]} \alpha_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}} \right) \beta_s \gamma_t p_{x,s,t}$$

If any of the  $p_{x,s,t}$  evaluate to non-zero, we respond to the zero-test with non-zero. If all evaluate to zero, we respond to the zero-test with zero. Since the number of  $p_{x,s,t}$  is polynomial

(namely  $|D| \times d_0 \times d_\ell$ ), we can test each  $p_{x,s,t}$  efficiently. In the case where any of the  $p_{x,s,t}$  are non-zero, we again appeal to Schwartz-Zippel (this time on the  $\alpha$ s,  $\beta$ s, and  $\gamma$ s) to see that with overwhelming probability the polynomial  $p$  evaluates to non-zero. If all of the  $p_x$  are zero, then with probability 1  $p$  will evaluate to zero. Therefore, we correctly guess the value of  $p$  with overwhelming probability.  $\square$

## References

- [AB15] Benny Applebaum and Zvika Brakerski. Obfuscating circuits via composite-order graded encoding. In *12th Theory of Cryptography Conference (TCC)*, pages 528–556, 2015.
- [AGIS14] Prabhakaran 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/>.
- [AHKM14] Daniel Apon, Yan Huang, Jonathan Katz, and Alex J. Malozemoff. Implementing cryptographic program obfuscation. Cryptology ePrint Archive, Report 2014/779, 2014. <http://eprint.iacr.org/>.
- [App14] Benny Applebaum. Bootstrapping obfuscators via fast pseudorandom functions. In *ASIACRYPT*, 2014.
- [Bar86] David A. Mix Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC1. In *STOC*, 1986.
- [BBC<sup>+</sup>14] Boaz Barak, Nir Bitansky, Ran Canetti, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Obfuscation for evasive functions. In *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, pages 26–51, 2014.
- [BGH<sup>+</sup>15] Zvika Brakerski, Craig Gentry, Shai Halevi, Tancrede Lepoint, Amit Sahai, and Mehdi Tibouchi. Cryptanalysis of the quadratic zero-testing of GGH. Cryptology ePrint Archive, Report 2015/845, 2015. <http://eprint.iacr.org/>.
- [BGI<sup>+</sup>01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *CRYPTO*, pages 1–18, 2001.
- [BGK<sup>+</sup>14] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. In *EUROCRYPT*, 2014.
- [BLR<sup>+</sup>14] Dan Boneh, Kevin Lewi, Mariana Raykova, Amit Sahai, Mark Zhandry, and Joe Zimmerman. Semantically secure order revealing encryption: Multi-input functional encryption without obfuscation, 2014.
- [BR14] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In *TCC*, 2014.
- [BWZ14] Dan Boneh, David J. Wu, and Joe Zimmerman. Immunizing multilinear maps

- against zeroizing attacks. Cryptology ePrint Archive, Report 2014/930, 2014. <http://eprint.iacr.org/>.
- [BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In *Advances in Cryptology - CRYPTO*, pages 480–499, 2014.
- [CGH<sup>+</sup>15] Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancrede 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.
- [CHL<sup>+</sup>15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In *EUROCRYPT*, 2015.
- [Cle91] Richard Cleve. Towards optimal simulations of formulas by bounded-width programs. *computational complexity*, 1(1):91–105, 1991.
- [CLR15] Jung Hee Cheon, Changmin Lee, and Hansol Ryu. Cryptanalysis of the new clt multilinear maps. Cryptology ePrint Archive, Report 2015/934, 2015. <http://eprint.iacr.org/>.
- [CLT13] Jean-Sebastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In *CRYPTO*, pages 476–493, 2013.
- [CLT15] Jean-Sebastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. New multilinear maps over the integers. In *CRYPTO*, 2015.
- [GGG<sup>+</sup>14] Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. In *Advances in Cryptology - EUROCRYPT*, pages 578–602, 2014.
- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, pages 1–17, 2013.
- [GGH<sup>+</sup>13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, pages 40–49, 2013.
- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In *12th Theory of Cryptography Conference (TCC)*, pages 498–527, 2015.
- [GGHR14] Sanjam Garg, Craig Gentry, Shai Halevi, and Mariana Raykova. Two-round secure MPC from indistinguishability obfuscation. In *11th Theory of Cryptography Conference (TCC)*, pages 74–94, 2014.
- [GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In *STOC*, 2013.
- [Gie01] Oliver Giel. Branching program size is almost linear in formula size. *J. Computer and System Sciences*, pages 222–235, 2001.
- [GIS<sup>+</sup>10] Vipul Goyal, Yuval Ishai, Amit Sahai, Ramarathnam Venkatesan, and Akshay Wadia.

- Founding cryptography on tamper-proof hardware tokens. In *TCC*, pages 308–326, 2010.
- [GLSW14] Craig Gentry, Allison Lewko, Amit Sahai, and Brent Waters. Indistinguishability obfuscation from the multilinear subgroup elimination assumption. Cryptology ePrint Archive, Report 2014/309, 2014. <http://eprint.iacr.org/2014/309>.
- [GLW14] Craig Gentry, Allison B. Lewko, and Brent Waters. Witness encryption from instance independent assumptions. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, pages 426–443, 2014.
- [Hal15] Shai Halevi. Graded encoding, variations on a scheme. *IACR Cryptology ePrint Archive*, 2015:866, 2015.
- [HJ15] Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. *IACR Cryptology ePrint Archive*, 2015:301, 2015.
- [Kil88] Joe Kilian. Founding cryptography on oblivious transfer. In *STOC*, pages 20–31, 1988.
- [MF15] Brice Minaud and Pierre-Alain Fouque. Cryptanalysis of the new multilinear map over the integers. Cryptology ePrint Archive, Report 2015/941, 2015. <http://eprint.iacr.org/>.
- [MSW14] Eric Miles, Amit Sahai, and Mor Weiss. Protecting obfuscation against arithmetic attacks. *IACR Cryptology ePrint Archive*, 2014:878, 2014.
- [PST14a] Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In *CRYPTO*, pages 500–517, 2014.
- [PST14b] Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In JuanA. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014*, volume 8616 of *Lecture Notes in Computer Science*, pages 500–517. Springer Berlin Heidelberg, 2014.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *Symposium on Theory of Computing (STOC)*, pages 475–484, 2014.
- [SZ14] Amit Sahai and Mark Zhandry. Obfuscating low-rank matrix branching programs. *IACR Cryptology ePrint Archive*, 2014:773, 2014.
- [Zha14] Mark Zhandry. How to avoid obfuscation using witness prfs. Cryptology ePrint Archive, Report 2014/301, 2014. <http://eprint.iacr.org/>.
- [Zim14] Joe Zimmerman. How to obfuscate programs directly. Cryptology ePrint Archive, Report 2014/776, 2014. <http://eprint.iacr.org/>.

## A Building Low-Rank Branching Programs

In this section, we describe some procedures to be carried out on branching programs. We will use these procedures to show how to make any branching program non-shortcutting, and how to convert formulas into matrix branching programs.

In particular, we will describe a way to convert any formula of size  $s$  over NOT, AND, XOR gates into an exact matrix branching program of length  $s + 1$  and maximum width  $\lceil \log_2(s + 2) \rceil$ . Our result gives a qualitative improvement to a result of Cleve [Cle91], which achieves similar asymptotics, but only for balanced formula, and only for formula over NOT, AND gates.

First, we give some further definitional details on branching programs, building on Section 2.3.

### A.1 Branching Programs

We will actually define three notions of branching programs. The first, layered (graphical) branching programs, corresponds to the standard notion of branching programs found in the literature. Second, we define the notion of a matrix branching program, which can be seen as a generalization of graphical branching programs. Finally, we define a matrix branching program sampler, which is again a generalization of matrix branching programs.

**Layered Graphical Branching Programs.** Our notion of a layered graphical branching program corresponds to the traditional notion of branching programs.

**Definition A.1.** A (*graphical*) *branching program* is a finite directed acyclic graph with two special nodes, a source node and a sink node, also referred to as an “accept” node. Each non-sink node is labeled with a variable  $x_i$  and can have arbitrary out-degree. Each of the out-edges is either labeled with  $x_i = 0$  or  $x_i = 1$ . The sink node has out-degree 0. We denote a branching program by  $BP$  and denote the restriction of the branching program consistent with input  $x$  by  $BP|_x$ .  $BP$  accepts an input  $x \in \{0, 1\}^n$  if and only if there is at least one path from the source node to the accept node in  $BP|_x$ . The *length*  $\ell$  of  $BP$  is the maximum length of any such path in the graph. The *node size*  $t$  of the branching program is the total number of nodes in the graph.

A *layered (graphical) branching program* is a branching program such that nodes can be partitioned into a sequence of layers  $L_0$  through  $L_\ell$  where all the nodes in  $L_{i-1}$  have only outgoing edges to  $L_i$ , and all of the outgoing edges from  $L_{i-1}$  are labeled with the same input variable, denoted  $x_{\text{inp}(i)}$  where  $\text{inp} : [\ell] \rightarrow [n]$ . We can assume without loss of generality that  $L_0$  contains only the source node and  $L_\ell$  contains only the sink node. The *length* of a layered branching program is  $\ell$ , and the *shape*  $(d_0, \dots, d_\ell)$  counts the number of nodes in each layer:  $d_i = |L_i|$ . Finally, the *width*  $w$  is the maximum  $d_i$ . The *node size*  $t$  of  $BP$  is still the total number of nodes in the graph,  $\sum_{i=0}^{\ell} d_i$ . We also define an additional quantity, the *total size*  $u$  is the sum of the products of sizes of adjacent layers  $\prod_{i=1}^{\ell} d_{i-1}d_i$ . Consider a slight modification to  $BP$  where there is an edge from *every* node in  $L_i$  to *every* node in  $L_{i+1}$ , and the edges are labeled with either  $x_i = 0$ ,  $x_i = 1$ ,  $x_i = 0, 1$  (representing that this edge is always used), or  $x_i = \perp$  (representing that this edge is never used). Then  $u$  counts the total number of edges in  $BP$ , and therefore represents the actual size of the description of  $BP$ .

**Matrix Branching Programs.** Note that our definition of a matrix branching program will depart in several ways from the standard definitions of matrix branching programs in the literature.

At a high level, a matrix branching program consists of a sequence of pairs of matrices  $(B_{i,0}, B_{i,1})$ . To evaluate the branching program on an input  $x$ , select one matrix from each pair based on the input, and multiply all of the matrices together. The matrices are shaped so that the products are valid and the end result is a scalar; otherwise, there are no restrictions on the shapes of the matrices. The branching program evaluates to 0 if and only if the product of all the matrices is 0, and otherwise it evaluates to 1. We can also easily generalize to multi-bit outputs by having the final matrix product be an actual matrix, and test each component independently for zero.

**Definition A.2.** A *generalized matrix branching program* of length  $\ell$  and shape  $(d_0, d_1, \dots, d_\ell) \in (\mathbb{Z}^+)^{\ell+1}$  for  $n$ -bit inputs is given by a sequence

$$BP = \left( \text{inp}, (\mathbf{B}_{i,0}, \mathbf{B}_{i,1})_{i \in [\ell]} \right)$$

where  $\mathbf{B}_{i,b} \in \mathbb{Z}^{d_{i-1} \times d_i}$  are  $d_{i-1} \times d_i$  matrices, and  $\text{inp} : [\ell] \rightarrow [n]$  is the evaluation function of  $BP$ .  $BP$  defines the following three functions:

- $BP_{arith} : \{0, 1\}^n \rightarrow \mathbb{Z}^{d_0 \times d_\ell}$  computed as

$$BP_{arith}(x) = \prod_{i=1}^n \mathbf{B}_{i, \text{inp}(i)}$$

- $BP_{bool} : \{0, 1\}^n \rightarrow \{0, 1\}^{d_0 \times d_\ell}$  computed as

$$BP_{bool}(x)_{j,k} = \begin{cases} 0 & \text{if } BP_{arith}(x)_{j,k} = 0 \\ 1 & \text{if } BP_{arith}(x)_{j,k} \neq 0 \end{cases}$$

- $BP_{bool(q)} : \{0, 1\}^n \rightarrow \{0, 1\}^{d_0 \times d_\ell}$  computed as

$$BP_{bool(q)}(x)_{j,k} = \begin{cases} 0 & \text{if } BP_{arith}(x)_{j,k} = 0 \pmod{q} \\ 1 & \text{if } BP_{arith}(x)_{j,k} \neq 0 \pmod{q} \end{cases}$$

A matrix branching program is  $t$ -bounded if  $|BP_{arith}(x)_{j,k}| \leq t$  for all  $x, j, k$ . In other words,  $t$  bounds the possible output values of  $BP_{arith}$ .

We define the width  $w = \max_{i \in [0, \ell]} d_i$ , node size  $t = \sum_{i=0}^{\ell} d_i$ , and total size  $u = \sum_{i=1}^{\ell} d_{i-1} d_i$ .

**Fact A.3.** Any graphical layered branching program  $BP$  of length  $\ell$  and shape  $(d_0, \dots, d_\ell)$  can be converted into a generalized matrix branching program  $BP'$  of length  $\ell$ , shape  $(d_0, \dots, d_\ell)$ , and bound  $t = \prod_{i=1}^{\ell-1} d_i \leq w^{\ell-1}$  such that  $BP'_{bool}(x) = BP(x)$  for all  $x$ .

For our obfuscator, similar to existing works, we will need to actually consider *dual-input* generalized matrix branching programs:

**Definition A.4.** A *dual-input generalized matrix branching program* of length  $\ell$  and shape  $(d_0, d_1, \dots, d_\ell) \in (\mathbb{Z}^+)^{\ell+1}$  for  $n$ -bit inputs is given by a sequence

$$BP = \left( \text{inp}_0, \text{inp}_1, \{\mathbf{B}_{i,b_0,b_1}\}_{i \in [\ell], b_0, b_1 \in \{0,1\}} \right)$$

where  $\mathbf{B}_{i,b_0,b_1} \in \mathbb{Z}^{d_{i-1} \times d_i}$  are  $d_{i-1} \times d_i$  matrices, and  $\text{inp} : [\ell] \rightarrow [n]$  is the evaluation function of  $BP$ .  $BP$  defines the following three functions:

- $BP_{arith} : \{0, 1\}^n \rightarrow \mathbb{Z}^{d_0 \times d_\ell}$  computed as

$$BP_{arith}(x) = \prod_{i=1}^n \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}}$$

- $BP_{bool} : \{0, 1\}^n \rightarrow \{0, 1\}^{d_0 \times d_\ell}$  computed as

$$BP_{bool}(x)_{j,k} = \begin{cases} 0 & \text{if } BP_{arith}(x)_{j,k} = 0 \\ 1 & \text{if } BP_{arith}(x)_{j,k} \neq 0 \end{cases}$$

- $BP_{bool(q)} : \{0, 1\}^n \rightarrow \{0, 1\}^{d_0 \times d_\ell}$  computed as

$$BP_{bool(q)}(x)_{j,k} = \begin{cases} 0 & \text{if } BP_{arith}(x)_{j,k} = 0 \pmod{q} \\ 1 & \text{if } BP_{arith}(x)_{j,k} \neq 0 \pmod{q} \end{cases}$$

A matrix branching program is  $t$ -bounded if  $|BP_{arith}(x)_{j,k}| \leq t$  for all  $x, j, k$ .

We note that it is easy to transform any normal matrix branching program into a dual input matrix branching program of the same length, shape, and bound: set  $\text{inp}_0 = \text{inp}_1 = \text{inp}$ , and  $\mathbf{B}_{i,b,b} = \mathbf{B}_{i,b}$  (the values  $\mathbf{B}_{i,b,1-b}$  can be set arbitrarily).

Unlike previous obfuscation constructions [GGH<sup>+</sup>13b, BR14, BGK<sup>+</sup>14, PST14b, AGIS14], we allow the matrices in the branching program to be singular, and even to be rectangular. This gives us the ability to have the product matrix have any desired shape — in particular, it can be a scalar for single-bit outputs. Thus, we do not need the “bookends” used in previous works to turn the matrix product into a scalar.

We will impose one requirement on matrix branching program, called *non-shortcutting*, which will be important in the security analysis of our obfuscator. We say that a branching program *shortcuts* on an input  $x$  if there is an interval  $[j, k] \subsetneq [\ell]$  strictly smaller than  $[\ell]$  such that the sub-product

$$\prod_{i=j}^k \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}} = 0 \text{ .}$$

In other words, for the input  $x$ , it is possible to determine that  $BP(x)_{j,k} = 0$  prematurely without evaluating the entire product. We say that a branching program is *non-shortcutting* if it does not shortcut on any  $x$ :

**Definition A.5.** A dual-input generalized matrix branching program is *non-shortcutting* if, for any input  $x$ , and any  $j \in [d_0]$  and any  $k \in [d_\ell]$ , the following holds:

$$\mathbf{e}_j^T \cdot \left( \prod_{i=1}^{\ell-1} \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}} \right) \neq 0^{d_\ell-1} \quad \text{and} \quad \left( \prod_{i=2}^{\ell} \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}} \right) \cdot \mathbf{e}_k \neq 0^{d_1}$$

where  $\mathbf{e}_j$  and  $\mathbf{e}_k$  are the  $j$ th and  $k$ th standard basis vectors of the correct dimension. Equivalently, each row of the product  $\prod_{i=1}^{\ell-1} \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}}$  and each column of the product  $\prod_{i=2}^{\ell} \mathbf{B}_{i, x_{\text{inp}_0(i)}, x_{\text{inp}_1(i)}}$  has at least one non-zero entry.

A similar definition holds for regular (non-dual-input) branching programs.

We will see in the next section that it is easy to convert any matrix branching program into a non-shortcutting branching program, and only increasing its width by 2.

A final property of matrix branching programs, which we call *exactness*, says that the outputs of  $BP_{arith}$  and  $BP_{bool}$  are the same on all inputs:

**Definition A.6.** A matrix branching program  $BP$  is *exact* if, for all inputs  $x$ , it holds that

$$BP_{arith}(x) \in \{0, 1\}^{d_0 \times d_\ell}$$

In other words,

$$BP_{arith}(x) = BP_{bool}(x)$$

In this case, we simply write  $BP(x)$  to denote  $BP_{arith}(x) = BP_{bool}(x) = BP_{bool(q)}$  for all  $q \geq 2$ .

**Matrix Branching Program Samplers.** We now define a matrix branching program sampler (MBPS). Roughly, an MBPS is a procedure that takes as input a modulus  $q$ , and outputs a matrix branching program  $BP$ . However, we will be interested mainly in the function  $BP_{bool(q)}$ .

**Definition A.7.** A matrix branching program sampler (MBPS) is a possibly randomized procedure  $BP^S$  that takes as input a modulus  $q$  satisfying  $q > t$  for some bound  $t$ . It outputs a matrix branching program.

**Fact A.8.** Any matrix branching program  $BP$  with bound  $t$  can trivially be converted into a matrix branching program sampler  $BP^S$  with the same bound  $t$ , such that if  $BP' \leftarrow BP^S(q)$ , then  $BP'_{bool(q)}(x) = BP_{bool}(x)$ .

## A.2 Operations on Generalized Matrix Branching Programs

We now describe several operations on generalized matrix branching programs. We will use these operations to build our branching program for formulas.

**Transpose.** Let  $BP = \left( \text{inp}, (\mathbf{B}_{i,0}, \mathbf{B}_{i,1})_{i \in [\ell]} \right)$  be a branching program of length  $\ell$  and shape  $(d_0, d_1, \dots, d_\ell) \in (\mathbb{Z}^+)^{\ell+1}$ . The *transpose* of  $BP$ , denoted  $BP^T$ , is a branching program of length  $\ell$  and shape  $(d_\ell, \dots, d_0)$ , given by

$$BP^T = \left( \text{inp}^T, \left( \mathbf{B}_{\ell+1-i,0}^T, \mathbf{B}_{\ell+1-i,1}^T \right)_{i \in [\ell]} \right)$$

where  $\text{inp}^T(i) = \text{inp}(\ell + 1 - i)$ . Observe that  $(BP^T)_{arith/bool}(x) = (BP_{arith/bool}(x))^T$ . Note that if  $BP$  is exact, then so is  $BP^T$ .

**Augment.** Let  $BP$  be as above. The  $r$ -augmentation of  $BP$  is the branching program  $BP' = \text{Augment}(BP, r)$  of length  $\ell$  and shape  $(d_0 + r, d_1 + r, \dots, d_\ell + r)$  given by

$$BP' = \left( \text{inp}, \left( \mathbf{B}'_{i,0}, \mathbf{B}'_{i,1} \right)_{i \in [\ell]} \right) \text{ where } \mathbf{B}'_{i,b} = \begin{array}{cc} \mathbf{B}_{i,b} & \mathbf{0}^{d_{i-1} \times r} \\ \mathbf{0}^{r \times d_i} & I_r \end{array}$$

Observe that

$$BP'_{arith/bool}(x) = \begin{array}{cc} BP_{arith/bool}(x) & \mathbf{0}^{d_0 \times r} \\ \mathbf{0}^{r \times d_\ell} & I_r \end{array}$$

Moreover, if  $BP$  is exact, then so is  $BP'$ . We will define  $\text{Augment}(BP) = \text{Augment}(BP, 1)$ .



**Linear Operations.** Let  $BP$  be as above. Given a  $d'_0 \times d_0$  matrix  $\mathbf{L}$  and a  $d_\ell \times d'_\ell$  matrix  $\mathbf{R}$ , we can compute the branching program  $\mathbf{L} \cdot BP \cdot \mathbf{R}$  which has length  $\ell$ , shape  $(d'_0, d_1, \dots, d_{\ell-1}, d'_\ell)$ , and is given by

$$\mathbf{L} \cdot BP \cdot \mathbf{R} = \left( \text{inp}, \left( \mathbf{B}'_{i,0}, \mathbf{B}'_{i,1} \right)_{i \in [\ell]} \right) \text{ where } \mathbf{B}'_{i,b} = \begin{cases} \mathbf{B}_{i,b} & \text{if } i \neq 1, \ell \\ \mathbf{L} \cdot \mathbf{B}_{1,b} & \text{if } i = 1 \\ \mathbf{B}_{\ell,b} \cdot \mathbf{R} & \text{if } i = \ell \end{cases}$$

Observe that  $(\mathbf{L} \cdot BP \cdot \mathbf{R})_{arith}(x) = \mathbf{L} \cdot (BP_{arith}(x) \cdot \mathbf{R})$ , and that if  $BP$  is non-shortcutting, then  $\mathbf{L} \cdot BP \cdot \mathbf{R}$  is also non-shortcutting.

**Merge.** Let  $BP^{(0)}, BP^{(1)}$  be two branching programs of length  $\ell^{(b)}$  and shape  $(d_0^{(b)}, \dots, d_\ell^{(b)})$  for  $b \in \{0, 1\}$  with the property that  $d_\ell^{(0)} = d_1^{(1)}$ . Then we can compute the merge of  $BP^{(0)}$  and  $BP^{(1)}$  as

$$BP^{(0)} \cdot BP^{(1)} = \left( \text{inp}, \left( \mathbf{B}_{i,0}, \mathbf{B}_{i,1} \right)_{i \in [\ell']} \right)$$

where  $\ell' = \ell^{(0)} + \ell^{(1)}$ ,  $\text{inp}' : [\ell^{(0)} + \ell^{(1)}] \rightarrow [n]$  is defined as  $\text{inp}'(i) = \begin{cases} \text{inp}^{(0)}(i) & \text{if } i \leq \ell^{(0)} \\ \text{inp}^{(1)}(i - \ell^{(0)}) & \text{if } i > \ell^{(0)} \end{cases}$ , and

$$\mathbf{B}_{i,b} = \begin{cases} \mathbf{B}_{i,b}^{(0)} & \text{if } i \leq \ell^{(0)} \\ \mathbf{B}_{i-\ell^{(0)},b}^{(1)} & \text{if } i > \ell^{(0)} \end{cases}. \text{ Observe that } (BP^{(0)} \cdot BP^{(1)})_{arith}(x) = (BP_{arith}^{(0)}(x)) \cdot (BP_{arith}^{(1)}(x)).$$

### A.3 Making any branching program non-shortcutting

Let  $BP$  be an arbitrary matrix branching program of shape  $(d_0, d_1, \dots, d_\ell)$ . Define

$$BP' = \begin{pmatrix} I_{d_0} & 1^{d_0 \times 1} & 0^{d_0 \times 1} \end{pmatrix} \cdot \text{Augment}(BP, 2) \cdot \begin{pmatrix} I_{d_\ell} \\ 0^{1 \times d_\ell} \\ 1^{1 \times d_\ell} \end{pmatrix}$$

Notice that the shape of  $BP'$  is  $(d_0, d_1 + 2, \dots, d_{\ell-1} + 2, d_\ell)$ . Moreover,  $BP'_{arith}$  computes the matrix

$$\begin{aligned} BP'_{arith} &= \begin{pmatrix} I_{d_0} & 1^{d_0 \times 1} & 0^{d_0 \times 1} \end{pmatrix} \cdot \begin{pmatrix} BP_{arith} & 0^{d_0 \times 2} \\ 0^{2 \times d_\ell} & I_2 \end{pmatrix} \cdot \begin{pmatrix} I_{d_\ell} \\ 0^{1 \times d_\ell} \\ 1^{1 \times d_\ell} \end{pmatrix} \\ &= \begin{pmatrix} I_{d_0} & 1^{d_0 \times 1} & 0^{d_0 \times 1} \end{pmatrix} \cdot \begin{pmatrix} BP_{arith} \\ 0^{d_\ell \times 1} \\ 1^{d_\ell \times 1} \end{pmatrix} = BP_{arith} \end{aligned}$$

Finally, we have the following:

**Lemma A.9.**  $BP'$ , as defined above, is non-shortcutting.

**Proof.** The branching program  $BP^{(L)} = \text{Augment}(BP, 2) \cdot \begin{pmatrix} I_{d_\ell} \\ 0^{1 \times d_\ell} \\ 1^{1 \times d_\ell} \end{pmatrix}$  on input  $x$  computes the matrix

$$\begin{pmatrix} BP_{arith}(x) \\ 0^{d_\ell \times 1} \\ 1^{d_\ell \times 1} \end{pmatrix}$$

which always has all columns not identically zero. Therefore the sub-product of  $BP_{arith}^{(L)}(x)$  consisting of all matrices except the left-most matrix will always have non-zero columns. We obtain  $BP'$  from  $BP^{(L)}$  by left-multiplying the left-most matrix by a matrix. Therefore, the sub-product of  $BP'_{arith}(x)$  that drops the left-most matrix is identical to the sub-product of  $BP_{arith}^{(L)}(x)$  that drops the left-most matrix, and therefore has non-zero columns. Similarly, we can conclude that any sub-product that does not include the right-most matrix has non-zero rows. Therefore, the branching program is non-shortcutting, as desired.  $\square$

#### A.4 Arithmetic Formulas to Matrix Branching Programs

We now give our conversion of formulas to matrix branching programs. We will build a branching program for any *arithmetic* formula taking 0/1 inputs, where every gate is an arbitrary bilinear polynomial in its inputs. That is, for any such arithmetic formula  $f$ , we build a branching program  $BP$  such that  $BP_{arith}(x) = f(x)$ . We note, however, that  $BP_{bool}$  only reveals if  $f(x)$  is zero or not.

As a first step, we build a branching program  $BP$  such that  $BP_{arith}(x) = \begin{pmatrix} 1 & f(x) \end{pmatrix}$ . The final branching program  $BP'$  is given as  $BP' = BP \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix}$

For input wires  $x_i$ , the branching program is trivial:  $\text{inp}$  maps 1 to  $i$ , and  $B_{1,b} = \begin{pmatrix} 1 & b \end{pmatrix}$ .

We now build  $BP$  recursively. Suppose  $f = f_0 \text{OP} f_1$  for some bilinear  $\text{OP}$ . Write  $y_0 \text{OP} y_1 = c_0 + c_1 y_0 + c_2 y_1 + c_3 y_0 y_1$ . We observe that

$$\begin{pmatrix} 1 & y_0 \end{pmatrix} \cdot \begin{pmatrix} 0 & c_2 & 1 \\ c_1 & c_3 & 0 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ y_1 & 0 \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 0 & 1 \\ 1 & c_0 \end{pmatrix} = \begin{pmatrix} 1 & y_0 \text{OP} y_1 \end{pmatrix}$$

Therefore, we can compute the branching programs  $BP_0$  and  $BP_1$  for  $f_0$  and  $f_1$ , and let

$$BP = BP_0 \cdot \left( \begin{pmatrix} 0 & c_2 & 1 \\ c_1 & c_3 & 0 \end{pmatrix} \cdot \text{Augment}(BP_1^T) \cdot \begin{pmatrix} 0 & 1 \\ 1 & c_0 \end{pmatrix} \right)$$

It follows that  $BP_{arith}(x) = \begin{pmatrix} 1 & f(x) \end{pmatrix}$ , meaning  $BP$  computes the correct function.

Now notice that  $f_0$  and  $f_1$  are not treated symmetrically above. Indeed, the width  $w$  of the branching program  $BP$  is equal to  $\max(w_0, 1 + w_1)$ , where  $w_0$  and  $w_1$  are the widths of  $BP_0$  and  $BP_1$ , respectively. For very unbalanced formula, this could lead the total width to be linear in the formula size.

However, we can easily exchange the roles of  $f_0$  and  $f_1$ . In particular, for a gate  $\text{OP}$ , let  $\text{OP}^T$  be the operation defined as  $y_0 \text{OP}^T y_1 = y_0 \text{OP} y_1$ . Now we can write  $f = f_0 \text{OP} f_1$  or  $f = f_1 \text{OP}^T f_0$ . In the first case, the total width becomes  $w = \max(w_0, 1 + w_1)$ , and in the second case, the width

becomes  $w = \max(w_1, 1 + w_0)$ . Therefore, we can choose which recursion to perform to arrive at the minimum:

$$w = \min(\max(w_0, 1 + w_1), \max(w_1, 1 + w_0))$$

For reasons that make the proof below more straightforward, we actually choose the order based on the formula size, rather than the resulting branching program width. That is, we have  $s_0 \geq s_1$  where  $s_b$  is the size of  $f_b$ .

**Lemma A.10.** *The length  $\ell$ , width  $w$ , total nodes  $t$  and total size  $u$  of the branching program  $BP'$  satisfy*

$$\begin{aligned}\ell &= s + 1 \\ w &\leq \lceil \log_2(s + 2) \rceil \\ t &\leq \lceil (s + 1)(1 + \frac{1}{2} \log_2(s + 2)) \rceil \\ u &\leq \lceil (s + 1)(3 + \log_2(s + 2))^2 / 4 \rceil\end{aligned}$$

**Proof.** The fact that  $\ell = s + 1$  follows easily from the recurrence. We will prove that the profile of  $BP$ ,  $(d_0, \dots, d_{s+1})$ , satisfies the following:

$$\begin{aligned}d_0 &= 1 & d_{s+1} &= 2 \\ 2 &\leq d_1 \leq \lceil \log_2(s + 2) \rceil & \text{for } i &\in [s] \\ \sum_{i=0}^{s+1} d_i &\leq 1 + \lceil (s + 1)(1 + \frac{1}{2} \log_2(s + 2)) \rceil \\ \sum_{i=0}^s d_i d_{i+1} &\leq \lceil (s + 1)(3 + \log_2(s + 2))^2 / 4 \rceil\end{aligned}$$

The final branching program  $BP'$  has the same profile, except that  $d_{s+1}$  is set to 1 instead of 2. The lemma easy follows.

The base case where  $s = 0$  is trivial. For a formula  $f = f_0 \text{OP} f_1$  of size  $s$ , let the size of the subformulas  $f_0$  and  $f_1$  be  $s_0$  and  $s_1$ :  $s_0 + s_1 + 1 = s$ . Then by induction the branching programs  $BP_0$  and  $BP_1$  have profiles  $d_0^{(b)}, \dots, d_{s_b+1}^{(b)}$  for  $b = 0, 1$  where  $d_0^{(b)} = 1$ ,  $d_{s_b+1}^{(b)} = 2$ , and  $2 \leq d_i^{(b)} \leq w_b$  where  $\lceil \log_2(s_b + 2) \rceil$  for  $i \in [1, s_b]$ . Moreover,  $\sum_{i=0}^{s_b+1} d_i^{(b)} \leq t_b$  where  $t_b = 1 + \lceil (s_b + 1)(1 + \frac{1}{2} \log_2(s_b + 2)) \rceil$  and  $\sum_{i=0}^{s_b} d_i^{(b)} d_{i+1}^{(b)} \leq u_b$  where  $u_b = \lceil (s + 1)(3 + \log_2(s + 2))^2 / 4 \rceil$

For now, suppose  $s_0 \geq s_1$ . The case  $s_1 > s_0$  is handled similarly. Then  $BP$  has the profile

$$(1, d_1^{(0)}, d_2^{(0)}, \dots, d_{s_0}^{(0)}, 2, d_{s_1}^{(1)} + 1, d_{s_1-1}^{(1)} + 1, \dots, d_1^{(1)} + 1, 2)$$

If  $s_0 > s_1$ , the total width is at most  $w_0 \leq \lceil \log_2(s + 2) \rceil$ , as desired.

Now if  $s_0 = s_1$ , then the total width will be  $w_0 + 1$ . However,  $s = 2s_0 + 1$  and so

$$w_0 + 1 = 1 + \lceil \log_2(s_0 + 2) \rceil = \lceil \log_2(2s_0 + 4) \rceil = \lceil \log_2(s + 3) \rceil$$

At first, this bound looks worse than what we are trying to prove. However, we know that  $s$  is odd. Suppose  $\lceil \log_2(s + 3) \rceil > \lceil \log_2(s + 2) \rceil$ . Then it must be that  $s + 2$  is a power of 2, and  $s + 3$

is one greater. However, this is impossible since  $s$  is odd. Therefore,  $w \leq \lceil \log_2(s+2) \rceil$  in this case as well.

The linear size  $t$  satisfies  $t \leq t_0 + t_1 - 2 + s_1$ . We have that

$$t \leq \lceil (s_0 + 1)(1 + \frac{1}{2} \log_2(s_0 + 2)) \rceil + \lceil (s_1 + 1)(1 + \frac{1}{2} \log_2(s_1 + 2)) \rceil + s_1$$

We can bound this expression as

$$t \leq 1 + \lceil (s_0 + 1)(1 + \frac{1}{2} \log_2(s_0 + 2)) + (s_1 + 1)(1 + \frac{1}{2} \log_2(s_1 + 2)) + s_1 \rceil$$

Given that  $s = s_0 + s_1 + 1$  and  $s_0 \geq s_1$ , it is straightforward but tedious to bound this as

$$t \leq 1 + \lceil (s + 1)(1 + \frac{1}{2} \log_2(s + 2)) \rceil$$

Next, we observe that

$$u = \left( \sum_{i=0}^{s_0} d_i^{(0)} d_{i+1}^{(0)} \right) + 2(d_{s_1}^{(1)} + 1) + \left( \sum_{i=0}^{s_1-1} (d_i^{(1)} + 1)(d_{i+1}^{(1)} + 1) \right)$$

The term  $\left( \sum_{i=0}^{s_0} d_i^{(0)} d_{i+1}^{(0)} \right)$  is equal to  $u_0$ . The term  $\left( \sum_{i=0}^{s_1-1} (d_i^{(1)} + 1)(d_{i+1}^{(1)} + 1) \right)$  is equal to  $\left( \sum_{i=0}^{s_1-1} d_i^{(1)} d_{i+1}^{(1)} \right) + 2 \left( \sum_{i=0}^{s_1-1} d_i^{(1)} \right) + s_1 - 1 - d_0^{(1)} - d_{s_1}^{(1)}$ . Using the fact that  $d_{s_1+1}^{(1)} = 2$ ,  $d_0^{(1)} = 1$  and  $d_i^{(1)} \geq 2$  for all other  $i$ , we can therefore, we can bound

$$u \leq u_0 + u_1 + 2t_1 + s_1 - 4$$

We can then bound this as

$$u \leq \lceil (s_0 + 1)(3 + \log_2(s_0 + 2))^2 / 4 + (s_1 + 1)(3 + \log_2(s_1 + 2))^2 / 4 + 2(s_1 + 1)(1 + \frac{1}{2} \log_2(s_1 + 2)) + s_1 + 1 \rceil$$

It is straightforward but tedious to bound this as  $u \leq \lceil (s + 1)(3 + \log_2(s + 2))^2 / 4 \rceil$ , as desired.  $\square$

**Making the branching program non-shortcutting.** We can make  $BP'$  non-shortcutting by increasing the profile by 2 as in Lemma A.9. However, it turns out in this case it is sufficient to only increase the profile by 1. This is because the branching program  $BP'$  always has non-zero outputs, and we obtain  $BP'$  from  $BP$  by right-multiplying the rightmost matrix. Therefore, any sub-product of  $BP'$  that does not contain the rightmost matrix must be a non-zero row vector (so its one and only row is non-zero). Therefore,  $BP'$  is part-way to non-shortcutting already. However, there is still a possibility that sub-products that do contain the right-most matrix will be zero.

We will fix this by augmenting the branching program, and then collapsing it back to a scalar by modifying the left- and right-most matrices. We do this in such a way that leaving out the left-most matrix always gives a non-zero product. Our branching program is set to

$$BP^{final} = \begin{pmatrix} 1 & 0 \end{pmatrix} \cdot \text{Augment}(BP') \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

**Lemma A.11.**  $BP^{final}$ , as constructed above, is non-shortcutting. Moreover, it satisfies

$$\begin{aligned} w &\leq \lceil 1 + \log_2(s+2) \rceil \\ t &\leq \lceil (s+1)(4 + \log_2(s+2))/2 \rceil \\ u &\leq \lceil (s+1)(5 + \log_2(s+2))^2/4 \rceil \end{aligned}$$

**Proof.** By a similar analysis to Lemma A.9 and the discussion above about  $BP'$  being partially non-shortcutting, we have that  $BP^{final}$  is non-shortcutting.

The augment procedure only increases the profile by one in each coordinate except  $d_0$  and  $d_{s+1}$ . Therefore, the bound on  $w$  follows from the previous analysis. The linear size increases by exactly  $s-1$ , so the bound of  $t$  follows from the previous analysis. Finally, it is straightforward to see that the actual size  $u$  only increases by less than twice the old linear size, plus 1. Therefore, the bound on  $u$  also follows from the previous analysis.  $\square$

## A.5 Extensions

**Boolean Formula.** Any boolean gate can be seen as a bilinear polynomial when its inputs and outputs are treated as integers. Therefore, for any boolean formula  $f$ , we can use the conversion above to build a matrix branching program  $BP$  such that  $BP_{arith} = f(x)$ . Since  $f(x)$  is either 0 or 1, we see that  $BP_{bool}(x) = BP_{arith}(x)$ , and so the resulting matrix branching program is exact.

**Arithmetic Formula With Integer Input.** Given any arithmetic formula where the input variables are bounded in some range of size  $B$ , we can break each input variable  $x$  into  $\lceil \log_2(B+1) \rceil$  input bits  $x_i$ , which can then be assembled into  $x$  using  $\lceil \log_2(B+1) \rceil - 1$  gates. Thus we increase the size of the formula by approximately a factor of  $\log_2 B$ , but can handle large inputs. We note that while  $BP_{arith}(x) = f(x)$ ,  $BP_{bool}$  still only reveals if  $f(x)$  is 0 or not.

**Arithmetic Formula With Bounded Outputs.** One limitation of the above constructions is that for general arithmetic formula, the function  $BP_{bool}$  does not reveal  $f(x)$ , but only a single bit. When we build our obfuscators, we will see that the function  $BP_{bool}$  is what we can actually obfuscate. Therefore, we would like to build a branching program  $BP$  such that  $BP_{bool}$  reveals the entire output  $f(x)$ . Here we make partial progress towards this goal by solving the case where the output  $f(x)$  is confined to a small range. Let  $f$  be an arithmetic formula with the guarantee that  $f(x) \in [B_0, B_1]$  for some integers  $B_0, B_1$  where  $B_1 - B_0$  is relatively small. We will construct a branching program  $BP$  such that  $BP_{bool}$  (rather than  $BP_{arith}$ ) reveals the entire output of  $f$ .

First, construct the branching program  $BP$  where  $BP_{arith}(x) = \begin{pmatrix} 1 & f(x) \end{pmatrix}$  as above. Then construct the branching program

$$BP' = BP \cdot \begin{pmatrix} B_0 & B_0 + 1 & \cdots & B_1 - 1 & B_1 \\ -1 & -1 & \cdots & -1 & -1 \end{pmatrix}$$

Notice that

$$BP'_{arith}(x) = \begin{pmatrix} B_0 - f(x) & (B_0 + 1) - f(x) & \cdots & (B_1 - 1) - f(x) & B_1 - f(x) \end{pmatrix}$$

Then  $BP'_{bool}(x)$  is all 1's, except for index  $f(x) - B_0 + 2$ , which will be zero. Therefore, there is a bijective mapping between  $BP'_{bool}(x)$  and  $f(x)$ , as desired. Moreover,  $BP'_{arith}(x)$  always contains

a non-zero entry, so we can make the branching program non-shortcutting by only increasing the profile by 1 instead of 2.

While the above increases the maximum width to at least  $B_1 - B_0 + 2$ , it does not increase the total size of the branching program by much. Indeed, the above modification only increases the profile in the last position. Therefore, the vertex size increases by an additive  $B_1 - B_0$ , while the total size increases by at most an additive  $(B_1 - B_0)w$  where  $w$  was the maximum width before the conversion. If we were restricted to using square matrices, the total size would increase by approximately  $s(2(B_1 - B_0)w + (B_1 - B_0)^2)$ , which is considerably worse for large  $B_1 - B_0$  or large  $s$ .

## B Efficiency improvements

In all current constructions, applying a core obfuscation method directly to circuits requires overhead that grows exponentially with depth. This occurs for two reasons. First, and perhaps most problematically, the level of multilinearity required by known core obfuscators grows exponentially with the depth of the circuit being obfuscated, and known implementations of graded encoding schemes have complexity that grows polynomially with the level of multilinearity [GGH13a, CLT13]. Furthermore, the only known method for converting circuits to full-rank matrix branching programs requires the size of the representation to grow exponentially with the depth of the circuit [Bar86]. The recent construction of Zimmerman [Zim14], which operates directly on circuits, does not suffer from the latter blowup, but still requires a level of multilinearity that grows exponentially with circuit depth (using known implementations of graded encoding schemes).

Nevertheless, a core obfuscator can be used to obfuscate general (high-depth) circuits with a polynomial overhead via bootstrapping as mentioned above. However, attempting such bootstrapping for obfuscation [GGH<sup>+</sup>13b, GIS<sup>+</sup>10, App14] based on existing core obfuscators encounters overheads that are asymptotically polynomial but easily reach above  $2^{100}$ . Such large overheads primarily arise due to the depth of the circuit that needs to be obfuscated by the core obfuscator (even though asymptotically, this circuit has depth logarithmic in the security parameter). Indeed, similarly large overheads arise when attempting to apply the core obfuscator to other programs represented in circuit form, since few interesting and non-learnable families of circuits have depth below, say, 50.

This suggests that perhaps representing programs as circuits may not be the best approach toward efficient obfuscation. Indeed, if we can expand the classes of program representations that are amenable to direct obfuscation by our core obfuscator, then this may allow for alternative methods of bootstrapping that yield substantially better efficiency (see also [AGIS14] for a speculative discussion of one such approach). Thus, improving the capabilities of the core obfuscator is a critical goal in obfuscation research.

Our core obfuscator substantially generalizes the types of branching programs that can be directly obfuscated. In existing core obfuscators operating on branching programs [GGH<sup>+</sup>13b, BR14, BGK<sup>+</sup>14, PST14a, GLSW14, AGIS14, MSW14], the program is required to consist of square and invertible matrices. As we elaborate below, this limitation of previous work is quite restrictive. Our obfuscator does not require this condition, and instead works on *any* matrix branching program satisfying a mild natural condition called non-shortcutting<sup>7</sup>. Furthermore, our work achieves two

---

<sup>7</sup>Non-shortcutting is defined in Section 2.3. Intuitively, this condition requires that no intermediate product of

Table 1: Comparing the efficiency of obfuscation schemes for keyed formulas over different bases. We use  $\tilde{O}$  to suppress the multiplicative polynomial dependence on the security parameter and other poly-logarithmic terms and  $O_\epsilon$  to suppress multiplicative constants which depend on  $\epsilon$ . Here  $s$  is the formula size,  $\epsilon > 0$  is an arbitrarily small constant, and  $\phi$  is a constant such that for  $\kappa$ -level multilinear encodings, the size of each encoding is  $\tilde{O}(\kappa^\phi)$ . The current best known constructions have  $\phi = 2$ . Evaluation time is given in the form  $a \cdot b$ , where  $a$  denotes the number of multilinear operations (up to lower order additive terms) and  $b$  denotes the time for carrying out one multilinear operation.

Work	Levels of Multilinearity	Size of Obfuscation/ Evaluation Time
[AGIS14] (previous work) {AND, NOT}-basis	$s$	$4s^3 \cdot \tilde{O}(s^\phi)$
[AGIS14] + [Gie01] (previous work) {AND, XOR, NOT}-basis	$O(s^{1+\epsilon})$	$O(2^{(2^{1/\epsilon})} s^{(1+\epsilon)}) \cdot \tilde{O}(s^{(1+\epsilon)\phi})$
This work (direct) {AND, XOR, NOT}-basis	$s$	$\frac{1}{4}s \log_2(s)^2 \cdot \tilde{O}(s^\phi)$

other advances over previous work:

- Our construction allows for sequences of *non-square* matrices of arbitrary compatible dimensions. We show that this flexibility yields concrete efficiency gains (see Section A for details).
- Our construction allows for a single obfuscation to yield multiple bits of output contained in different entries in the output matrix  $M$ . This is in contrast to previous work that yielded only one bit of output, and required many parallel obfuscations to obtain multiple bits.

We also show how to exploit our results to yield efficiency improvements over previous obfuscations of both Boolean formulas and layered graphical branching programs. These improvements are summarized in Tables 1 and 2.

Furthermore, our analysis can also be used in settings where obfuscation is not directly used, but where low-rank matrix branching programs are considered in the context of multilinear maps. Indeed, subsequent to the initial online publication of this work, our theorems were used by [BLR<sup>+</sup>14] to yield substantial efficiency improvements in the context of multi-input functional encryption and semantically secure order-revealing encryption.

To obtain our results for Boolean formulas, in Section A, we give a simple conversion from formulas to (low-rank) matrix branching programs that achieves qualitatively better parameters than was previously known [Cle91]. This conversion may be of independent interest.

---

matrices yields the all-zero matrix.

Table 2: Comparing the efficiency of obfuscation schemes for keyed layered graphical branching programs, as defined in Section 2.3. For a layered graphical branching program,  $\ell$  denotes the length of the branching program, and  $w$  is the width. Other notation is as in Table 1.

Work	Levels of Multilinearity	Size of Obfuscation/ Evaluation Time
[AGIS14] (previous work)	$\ell$	$4\ell^3 w^2 \cdot \tilde{O}(\ell^\phi)$
This work	$\ell$	$\ell w^2 \cdot \tilde{O}(\ell^\phi)$