

Obfuscation Combiners

Marc Fischlin¹, Amir Herzberg², Hod Bin Noon², and Haya Shulman³

¹ Technische Universität Darmstadt, Germany
marc.fischlin@cryptoplexity.de

² Bar Ilan University, Israel

³ Fraunhofer SIT, Darmstadt, Germany

Abstract. Obfuscation is challenging; we currently have practical candidates with rather vague security guarantees on the one side, and theoretical constructions which have recently experienced jeopardizing attacks against the underlying cryptographic assumptions on the other side. This motivates us to study and present *robust combiners for obfuscators*, which integrate several candidate obfuscators into a single obfuscator which is secure as long as a quorum of the candidates is indeed secure.

We give several results about building obfuscation combiners, with matching upper and lower bounds for the precise quorum of secure candidates. Namely, we show that one can build 3-out-of-4 obfuscation combiners where at least three of the four combiners are secure, whereas 2-out-of-3 structural combiners (which combine the obfuscator candidates in a black-box sense) with only two secure candidates, are impossible. Our results generalize to $(2\gamma + 1)$ -out-of- $(3\gamma + 1)$ combiners for the positive result, and to 2γ -out-of- 3γ results for the negative result, for any integer γ .

To reduce overhead, we define *detecting combiners*, where the combined obfuscator may sometimes produce an error-indication instead of the desired output, indicating that some of the component obfuscators is faulty. We present a $(\gamma + 1)$ -out-of- $(2\gamma + 1)$ detecting combiner for any integer γ , bypassing the previous lower bound. We further show that γ -out-of- 2γ structural detecting combiners are again impossible.

Since our approach can be used for practical obfuscators, as well as for obfuscators proven secure (based on assumptions), we also briefly report on implementation results for some applied obfuscator programs.

1 Introduction

Software obfuscation has a long tradition in aiming at protecting against reverse engineering. For example, the first International Obfuscated C Code Contest (www.ioccc.org) has been organized in 1984 and experienced the 23rd event in this series in 2014. There are obfuscators for all popular programming languages today. For example, for Java, there are several open-source projects like ProGuard, ClassEncrypt, or JavaGuard, and an even larger number of commercial products. These approaches are usually based on heuristics and best

practices, ranging from simple renaming of function and variables names, to elaborate schemes, e.g. [19]. However, these practical obfuscators do not provide verifiable, *proven* security guarantees.

Provably secure obfuscation, in the sense that it is based on some reasonable cryptographic assumption, has long been a highly desirable yet hard-to-reach goal. Even worse, there have been devastating impossibility results for the natural notion of virtual black-box obfuscation [8] and only limited positive results for special cases like point functions [16]. A significant breakthrough came with the work by Garg et al. [28], indicating that the relaxed yet useful notion of indistinguishability obfuscation may be achievable for general circuits. This notion basically says that one cannot distinguish the obfuscated codes of two functionally equivalent circuit programs.

It is fair to say that the underlying cryptographic assumption, on which is security of the construction of Garg et al. [28] is based upon, is non-standard and not well analyzed (yet). This is also true for the alternative approach to build indistinguishability obfuscators proposed by Pass et al. [48]. This is complemented by yet other proposals of Gentry et al. [32] based on a more standard-like computational assumption about multilinear maps, and of Ananth and Jain [3] based on compact functional encryption. At the same time, recent attacks [18,23,31,21,22] on multilinear maps, albeit currently not known to break the aforementioned obfuscation candidates, testify that constructions may suddenly turn out to lack the desired security guarantees. New suggestions and attacks keep on appearing at high frequency [41,6,30,46].

The above leaves us with multiple choices of candidates for building obfuscators, both in practice as well as in theory, and it is currently difficult to determine the best choice in terms of security. For the heuristic, practical obfuscators, it may be even harder to distinguish sound constructions from weak approaches, since the design strategies may be vague. A straightforward idea to boost confidence in obfuscator candidates, both in theory as well as in practice, is to interlock multiple solutions and approaches. This idea of failure-tolerant cryptographic designs has traditionally been subsumed under the notion of robust combiners.

1.1 Robust Combiners for Obfuscation

The notion of robust combiners has been introduced by Harnik et al. [34] based on the idea of tolerant cryptographic designs by Herzberg [35,36,37]. Such combiners take several candidates for a cryptographic task and provide a secure solution if a quorum of the candidates is indeed secure. The idea has been successfully applied to several cryptographic primitives, including hash functions [14,49,25,50,47,27,45], encryption [24,34], commitments [34,35,36], and oblivious transfer [34,43,44].

A robust combiner for obfuscation would take as input a program (abstractly in form of a circuit or a Turing machine⁴) and create an obfuscated version with the help of the candidate obfuscators $\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_N$. As long as a sufficient number of candidate obfuscators is indeed secure, the combiner should also provide a secure obfuscator. In order to make formal claims about the robustness of the combiner, due to the lack of rigorous security properties for practical obfuscators, one inevitably needs to base the notion of security for the combiner on the various models in the cryptographic literature, such as virtual black-box obfuscation or indistinguishability obfuscation.⁵

What distinguishes the idea of combiners for obfuscation from the previous scenarios is that obfuscation combiners are higher-order combiners which are closely linked to the functionality of their inputs. Consider for instance the case of hash function combiners where it usually suffices that the combiner preserves the security property only, enabling solutions like the concatenation combiner $\text{Comb}^{H_1, H_2}(x) = H_1(x) || H_2(x)$ with longer output for collision resistance. Devising hash combiners with equal output size as H_1, H_2 , retaining this mild functional property, is conceivably hard [14,49,50]. An obfuscation combiner, in contrast, must provide a circuit which computes the same function as the input circuit; it cannot implement a different function with a larger output. Indeed, note that functional preservation and input hiding are conflicting requirements for obfuscation and one is easy to achieve without the other.

As a concrete example consider combiners in the context of virtual black-box obfuscation. Herzberg and Shulman [38] show that the cascading construction $\text{Comb}^{\mathcal{O}_1, \mathcal{O}_2}(\cdot) = \mathcal{O}_2(\mathcal{O}_1(\cdot))$ of two candidate obfuscators $\mathcal{O}_1, \mathcal{O}_2$ is robust for this notion *as long as functional correctness of the candidates is guaranteed*. If this is not granted and the inner obfuscator is corrupt then \mathcal{O}_1 may implement an arbitrary function, such that the combiner neither preserves functional correctness nor necessarily input hiding. The latter holds as one usually does not have any security guarantees for input circuits with diverging functionalities, even if obfuscator \mathcal{O}_2 is sound. Analogously, if the outer obfuscator is corrupt then the resulting cascade may no longer sustain functionality.

While functional correctness of an obfuscator is usually not based on unproven cryptographic assumptions, unlike the obfuscation property, there are two reasons why certifying functional correctness may still be hard. First, software implementations are error prone, and the complexity of previous theoretical proposals for obfuscation [28,48,32] seems to be inimical in this regard. Secondly, one may have little control over, or insights into, the actual obfuscation program. This is clearly true for commercial obfuscation programs; in fact, the programs of such obfuscators are often themselves obfuscated. The creation of a corrupt obfuscator, which intentionally leaks some information, is easy; to demonstrate

⁴ In our presentation of our formal results we focus on circuits instead of Turing machines, since our approach applies equally well to both settings but the state of the art of solutions is much more advanced in the circuit setting.

⁵ There are approaches to define metrics for practical obfuscators [20,5], but mainly in terms of software complexity. We discuss them in Section 7.

this, we implemented demos of different types of corrupted-obfuscators, including obfuscators which leak information even when used in cascade.

The concern about corrupt-obfuscators may also emerge in theoretical solutions. As an example for the latter, in the universal-parameter generation setting [39] a *trusted* party publishes an obfuscated program which parties can use to generate common parameters. What if we now prefer to use several potentially *untrusted* authorities and combine their obfuscated programs?

1.2 Our Results

Our goal is to provide a general combiner for obfuscation. It should satisfy the formal requirements in order to allow for sound solutions both in theory and in practice. Ideally, the combiner should tolerate a large number of corrupt obfuscators, be very efficient, and ensure various notions of obfuscation simultaneously. Note that, while virtual black-box obfuscation may be impossible in general, for some functions and attack models [15,7] the notion may still be achievable, such that our combiner should also comply with this notion.

On the positive side we present a 3-out-of-4 combiner which can tolerate a single corrupt combiner out of four candidates. It is depicted in Figure 1 and consists of two layers. In the first layer we insert the input circuit C into three combinations of three of the obfuscators each; in each combination, we output a circuit that produces the majority of the three obfuscated circuits. We only require three of the four combinations of picking three of four obfuscators. Each unit ensures that if at most one candidate is corrupt then functional correctness is still preserved. In the next layer we then run each of the first-layer majority circuits through the complementary fourth obfuscation candidate and again take the majority to ensure correctness. Obfuscation follows as either all three candidates on the first layer are sound and thus hide the input circuit, or the fourth candidate on the second layer ensures this.

Our combiner indeed works for different notions of obfuscation such as virtual black-box and grey box obfuscation,⁶ indistinguishability obfuscation, and differing-input obfuscation. In total it requires twelve calls to obfuscators and has depth 2. The latter is important as obfuscation may cause a polynomial blow-up in size. Remarkably, while most theoretical solutions currently induce a significant size expansion, with a few exceptions [13,4], obfuscators in practice only display a mild increase in code size. Note that devising combiners of depth 1 with a structure as above is impossible as the corrupt obfuscator may then leak information about the input circuit via the output.

We then show an impossibility result for 2-out-of-3 combiners. There are, of course, trivial combiners in this case, such as the combiner which simply uses the sound candidate only, and the (inefficient) combiner for indistinguishability obfuscation that evaluates the input circuit and then outputs the lexicographic smallest equivalent circuit. We thus focus on *structural* combiners that use a fixed pattern, independently of the status of the candidates, and do not semantically

⁶ With respect to dependent auxiliary inputs [33].

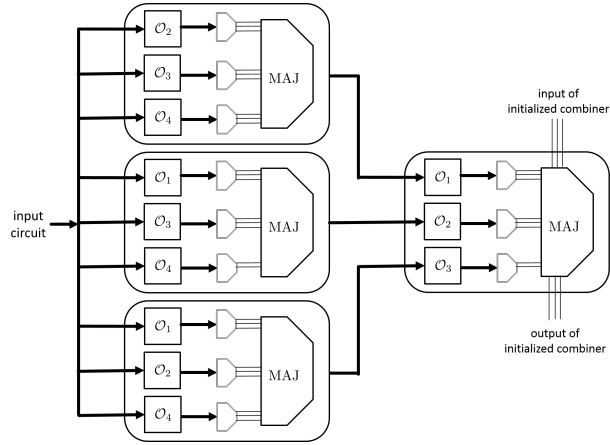


Fig. 1: Our 3-out-of-4 combiner. The MAJ circuit has three hardwired circuits C_1, C_2 and C_3 with equal input and output sizes, which also correspond to the input and output size of the MAJ circuit. For input x the MAJ circuit evaluates each of the three circuits for x and returns the bit-wise majority of the circuit’s outputs.

interpret the input circuit. Our 3-out-of-4 combiner is structural in this regard. We show that no 2-out-of-3 structural combiner may ensure both functional correctness and obfuscation. This holds for the weaker notion of indistinguishability obfuscation and therefore also for the stronger notions of black-box and grey-box obfuscation. Note that this also applies to any 1-out-of-2 combiner.⁷

We extend the positive result as well as the negative result to the case of $(2\gamma + 1)$ -out-of- $(3\gamma + 1)$ resp. 2γ -out-of- 3γ combiners. That is, we give a construction which can be seen as a less efficient generalization of our basic solution if one can corrupt at most γ out of $3\gamma + 1$ obfuscators. We then argue that one cannot have structural combiners if γ out of the 3γ obfuscators can be corrupt. For both settings we can draw on the the ideas and techniques from the basic cases.

The combiners above are correcting in the sense that they guarantee functional correctness if a quorum of input obfuscator candidates is secure. One can also envision a weaker notion of combiners, which output circuits that either compute the correct output of the input circuit, but may also output, instead, a special error indicator \perp . This error indicator should be output only when one of the component obfuscator is faulty; if all obfuscators are sound then the combiner must never output \perp . In particular, such combiners cannot output false answers. We call them *detecting combiners* in analogy to coding theory.

For detecting combiners we achieve slightly different bounds. That is, we show that one can have $(\gamma + 1)$ -out-of- $(2\gamma + 1)$ combiners for any γ . For the case $\gamma = 1$ and 2-out-of-3 combiners we can again provide an optimized version similar to our original 3-out-of-4 combiner. Concerning lower bounds, we can apply the

⁷ Every 1-out-of-2 combiner is also a 2-out-of-3 combiner if it ignores the third obfuscator.

ideas of the other combiners to show that there cannot exist structural γ -out-of- 2γ detecting combiners for any γ . The reduced overhead of detecting combiners may make them attractive option for practical implementations, where once detection ability exists, the attack-vector of providing faulty obfuscator appears unlikely.

While our main results follow the common approach in provably secure obfuscation, we stress that we view our approach to be equally well suited for practice. In Section 7 we therefore evaluate performance of our combiner when applied to practical obfuscators, and discuss the implications of our findings in this domain.

Concurrent Work. Independently of our work, Ananth et al. [2] also discuss the idea of obfuscation combiners. Their approach is fundamentally different from ours, and results in (non-structural) obfuscation combiners which are secure as long as a single candidate is secure. However, this comes at the cost of a significant overhead, and also requires additional cryptographic assumptions such as LWE or DDH, and indistinguishability obfuscation against sub-exponential adversaries.

2 Preliminaries

We exclusively treat circuits here; the approach can be transferred to the case of Turing machines straightforwardly. When speaking of circuits C from some class $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ we usually mean some arbitrary (but efficiently computable) description of the circuit. When considering specific encodings with dedicated properties, as required for our lower bounds, we usually write $\langle C \rangle$ for the encoding of the circuit under scheme $\langle \cdot \rangle$. If, on the other hand, we consider the function implemented by the circuit we usually write $C(\cdot)$ instead, and $C(x)$ for the output of circuit C on input x . When writing $C(\cdot) = C'(\cdot)$ or $C \equiv C'$ we refer to functional equality of circuits C and C' , comprising input and output length, whereas $C = C'$ or $\langle C \rangle = \langle C' \rangle$ means equal descriptions (under the encoding in question).

2.1 Obfuscators

Barak et al. [8] defined several notions of obfuscators, with *virtual black-box* (VBB) obfuscators being the strongest one. This notion says that the adversary cannot learn anything from an obfuscated circuit beyond the circuit's outputs for chosen inputs. While they also showed that this notion is in general unachievable, for specific cases such as point functions one may be able to attain this level of obfuscation. Below we mainly consider obfuscation of circuits, and we also consider the possibility that the obfuscator itself may be non-uniform and work specifically for different values of λ . The latter allows corrupt (also called malicious) obfuscators to match the algorithm class of adversaries and distinguishers. All obfuscators here, sound and corrupt ones, are nonetheless considered to be stateless.

Definition 1 (Virtual Black-Box Obfuscation). A (possibly non-uniform) PPT algorithm \mathcal{O} is a virtual black-box obfuscator for circuit class $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ if the following holds:

Functional Correctness: For any $\lambda \in \mathbb{N}$, any circuit $C \in \mathcal{C}_\lambda$, any obfuscated version $O \leftarrow \mathcal{O}(1^\lambda, C)$ we have $C \equiv O$.

VBB Obfuscation: For any (possibly non-uniform) PPT algorithm \mathcal{A} there exists a (possibly non-uniform) algorithm PPT \mathcal{S} and a negligible function $\epsilon(\lambda)$ such that for all circuits $C \in \mathcal{C}_\lambda$ we have

$$|\text{Prob}[\mathcal{A}(1^\lambda, \mathcal{O}(1^\lambda, C)) = 1] - \text{Prob}[\mathcal{S}^C(1^\lambda) = 1]| \leq \epsilon(\lambda),$$

where the probabilities are over the randomness of \mathcal{O} and \mathcal{A} resp. \mathcal{S} .

Virtual grey-box (VGB) obfuscation [10] is defined analogously, only that the simulator above is computationally unbounded but can make at most a polynomial number of queries to its oracle circuit. Clearly, VBB obfuscation implies VGB obfuscation. A stronger notion is based on the extension to (dependent) auxiliary inputs [33] where both the adversary and the simulator receive a random sample aux as additional input, where aux may depend on any circuit $C' \in \mathcal{C}_\lambda$.⁸ We will use this version for proving the security of our combiners for VBB and VGB obfuscation.

Another meaningful relaxation, implied by both notions above in the non-uniform setting, is *indistinguishability obfuscation* [8] which basically says that the obfuscations of two functional equivalent circuits are indistinguishable:

Definition 2 (Indistinguishability Obfuscator). A (possibly non-uniform) PPT algorithm $i\mathcal{O}$ is called an indistinguishability obfuscator for a circuit class $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$ if the following conditions hold:

Functional Correctness: For any $\lambda \in \mathbb{N}$, any circuit $C \in \mathcal{C}_\lambda$, any obfuscated version $O \leftarrow i\mathcal{O}(1^\lambda, C)$ we have $C \equiv O$.

Indistinguishability: For any (possibly non-uniform) PPT distinguisher \mathcal{D} , there exists a negligible function $\epsilon(\lambda)$ such that for all circuits $C_0, C_1 \in \mathcal{C}_\lambda$ with $C_0 \equiv C_1$ we have

$$\begin{aligned} & |\text{Prob}[\mathcal{D}(1^\lambda, C_0, C_1, i\mathcal{O}(1^\lambda, C_0)) = 1] \\ & - \text{Prob}[\mathcal{D}(1^\lambda, C_0, C_1, i\mathcal{O}(1^\lambda, C_1)) = 1]| \leq \epsilon(\lambda), \end{aligned}$$

where the probabilities are over the randomness of $i\mathcal{O}$ and \mathcal{D} .

There are several variations of the above definitions. For one, we can allow for a negligible error in the functional correctness (over the random choices of the obfuscator). Both our positive and our negative result are robust with respect to such a change. That is, our 3-out-of-4-combiners uses a constant number of

⁸ This slightly strengthens the original auxiliary input setting [33] where only $C' = C$ is allowed.

obfuscator calls such that the error would remain negligible; obfuscation would still hold, because the leakage due to incorrect obfuscator outputs has negligible probability. Similarly, our impossibility result about 2-out-of-3 combiners would still hold, even if the starting combiners would have perfect functional correctness, but the (fixed-size structural) combiner could have a negligible error. Alternatively, one may use the recent approach in [12] to eliminate the error first.

Finally, yet another version of obfuscation, called *differing-inputs* obfuscation [8], demands indistinguishability of two obfuscated circuits, but only if the input circuits C_0, C_1 can be sampled such that finding inputs where C_0 and C_1 differ, is infeasible. More formally, we assume that there is a PPT algorithm `Sampler` associated to the circuit family \mathcal{C} such that for any PPT algorithm \mathcal{A} there exists a negligible function $\epsilon(\lambda)$ such that the probability that $C_0(x) \neq C_1(x)$, where $(C_0, C_1, \text{aux}) \leftarrow \text{Sampler}(1^\lambda)$ and $x \leftarrow \mathcal{A}(1^\lambda, C_0, C_1, \text{aux})$, is at most $\epsilon(\lambda)$. Note that we assume that `Sampler`(1^λ) only outputs circuits $C_0, C_1 \in \mathcal{C}_\lambda$,

A differing-inputs obfuscator $\text{di}\mathcal{O}$ for \mathcal{C} and `Sampler` is now defined analogously to an indistinguishability obfuscator, only that it is infeasible to distinguish outputs $\text{di}\mathcal{O}(1^\lambda, C_0)$ and $\text{di}\mathcal{O}(1^\lambda, C_1)$ for $(C_0, C_1, \text{aux}) \leftarrow \text{Sampler}(1^\lambda)$, even if given `aux` as additional input. While the notion is also quite useful for the design of protocols [1], Garg et al. [29] argue that the notion may be hard to achieve.

2.2 Combiners for Obfuscators

Roughly, a combiner for obfuscators is a procedure which uses a set of obfuscators $\mathcal{O}_1, \mathcal{O}_2, \dots$ to turn an input circuit C into an obfuscated one, with the guarantee that if an (unspecified) quorum of the underlying obfuscators is secure, then so is the combiner. In the definition below we abstractly speak of o -obfuscators, leaving open which obfuscation category $o \in \{\text{VBB}, \text{VGB}, \text{indistinguishability}, \text{differing-inputs}\}$ we refer to.

For combiners of primitives with multiple properties, such as functional correctness and obfuscation here, there are varying levels of combiners, called weak, mild, and strong [26,27]. A strong combiner preserves security “property-wise”, i.e., for each property individually if sufficiently many candidates have this property then so does the combiner. A weak combiner only preserves all properties if there are enough candidates which are secure and thus have all properties simultaneously. The mild notion is in between where the candidates must somehow cover all properties but for each property possibly by different candidates. In [26,27] it has been discussed that strong robustness implies mild robustness which in turn implies weak robustness, and that the implications are strict in case of hash functions for some properties.

Definition 3 (Robust Combiner for o -Obfuscation). *Let `Comb` be a PPT oracle algorithm and let $\mathcal{O}_1, \dots, \mathcal{O}_N$ be o -obfuscators candidates. Then `Comb` is called a*

- *strongly robust t -out-of- N combiner* if for each of functional correctness and o -obfuscation, if at least t of the N candidates have this property, then so does the combiner $\text{Comb}^{\mathcal{O}_1, \dots, \mathcal{O}_N}$;
- *mildly robust t -out-of- N combiner* if, whenever functional correctness and o -obfuscation are each satisfied by at least t of the N candidates, then the combiner too has both properties;
- *weakly robust t -out-of- N combiner* if the combiner is a functional correct o -obfuscator if there are at least t out of N candidates which are simultaneously functionally correct and o -obfuscators.

The definition assumes that the obfuscators and combiner all work for the same class \mathcal{C} of obfuscatable circuits. This neglects an important aspect, though: If the combiner calls obfuscators recursively then the candidates need to be able to handle obfuscated circuits, too. We assume that this is indeed the case — and discuss it more explicitly for our structural combiners below— making the implicit assumption that the candidates also allow for a superclass $\mathcal{C}^{\text{Comb}}$ of circuits which is rich enough to capture intermediate circuits created by the specific combiner. Still, the task for the combiner is to obfuscate the “core” class \mathcal{C} of circuits.

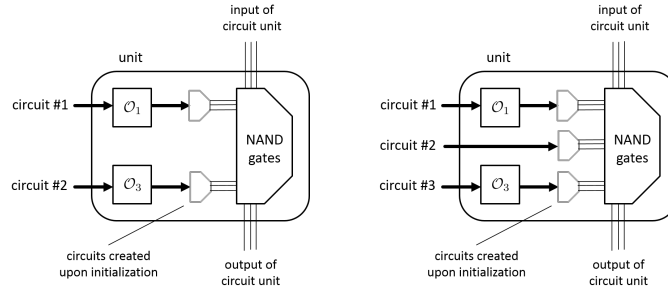


Fig. 2: Example of a unit (with pass-through version on the right-hand side).

As usual for combiners in general, there is always a secure obfuscation combiner, namely, the one which “obviously” uses the secure obfuscator \mathcal{O}_i and ignores the other ones in order to obfuscate the input circuit. However, this only provides an existential proof and says nothing about how to design an actual solution. Even worse, for indistinguishability obfuscation there is a trivial (non-efficient) combiner for obfuscators which can be described effectively [8]. The combiner takes as input (the description of) a circuit C and finds the (lexicographically) minimal circuit C_{\min} which computes the same functionality as C and outputs this circuit C_{\min} . Then any two circuits C, C' with the same functionality yield the same obfuscated circuit C_{\min} . This combiner ignores the candidate obfuscators and already constitutes an unconditionally secure obfuscator itself. It is even efficient relative to a Σ_2^P oracle. Hence, any lower bound

for combiners would need to bypass this result and therefore need to implicitly show that $\Sigma_2^P \neq P$.

One option to circumvent the first problem is to require to have an effective mean to turn attacks against the combiner into attacks for the candidate obfuscators. This option of so-called *black-box combiners* has been used for other lower bounds such as for hash function combiners [14,49,50]. Still, in our setting such black-box combiners would have to deal with the problem of the inefficient combiner.

An alternative path, which we also take here, is therefore to restrict the way how the combiner works. Whereas the above unconditional combiner approaches the circuit *semantically* by plotting its behavior, we look into what we call *structural* combiners here. Basically, these are combiners which have a prescribed structure with place-holder gates for the obfuscators, and they merely plug in the input circuit and derive the output circuit according to this fixed structure, without evaluating the circuits. It turns out that our 3-out-of-4 combiner is in fact structural.

2.3 Structural Combiners

A structural combiner for obfuscators is a circuit consisting of NAND gates and of obfuscator gates, where each one of the latter is labeled with one of the obfuscators \mathcal{O}_i . The layout is independent of the actual obfuscators and should thus work with any concrete obfuscator candidates, i.e., be black-box. The combiner is structured in so-called units. A *unit* is a sub circuit which takes as input the descriptions of circuits and itself describes a circuit. The unit first inserts the input circuits into some of the obfuscators, where we allow multiple appearances of obfuscators in a unit, and then processes the output circuits by a circuit consisting of NAND gates only. An example is given in the left part of Figure 2. If the input circuit is given to obfuscators i_1, i_2, \dots then we call this an $\{i_1, i_2, \dots\}$ -unit for the multiset $\{i_1, i_2, \dots\}$. The example in Figure 2 describes a $\{1, 3\}$ -unit. Furthermore, we can even let some input circuit be passed to the NAND-circuit completely, saying that the unit is *pass-through* in this case. Since it is irrelevant for our lower bound which circuit is passed through, we do not need to specify the identifier. The right hand side of Figure 2 shows a pass-through version of a $\{1, 3\}$ -unit.

The output of a unit can itself serve again as the input for another unit. We can therefore nest units in a tree-like structure as in Figure 3. In particular, we can analogously to the notion of depths of circuits define the depth of a unit, starting with level-1 units, as well as paths from the input circuit to the final unit. We call the path of units from level-1 units to the final unit a full path. A unit which is level-1 always receive the combiner's input circuit as inputs, but potentially also other unit circuits if it is simultaneously a higher level unit. Every unit has at least one input circuit, and a unit can of course serve as multiple inputs to other units.

To complete the description of a structural combiner we need to specify the output of our combiner for some input circuit C , once the obfuscator candidates

are determined. We call this the *initialization* of the combiner with C . Basically the output is again a circuit and it is derived by stepwise replacing the obfuscator gates in units (starting with level-1 units which receive C as input) with samples of the output of the corresponding obfuscator. Note that the structure of the combiner circuit remains, only the obfuscator gates are now filled in with concrete circuits. In case of pass-through units we additionally place the code of the unit's input circuit inside the new circuit at the corresponding position. Once a unit has been initialized we can use it as input to a higher-level unit and initialize that unit, till we have eventually initialized the final unit. Instructively, the reader may think of this as a left-to-right pass in Figure 3 to compute the final output circuit, denoted as $\text{Comb}^{\mathcal{O}_1, \mathcal{O}_2, \dots}(C)$. Note that this is a random variable, depending on the randomness of the obfuscators. A sample of this random variable can then be fed with inputs x to produce some output y .

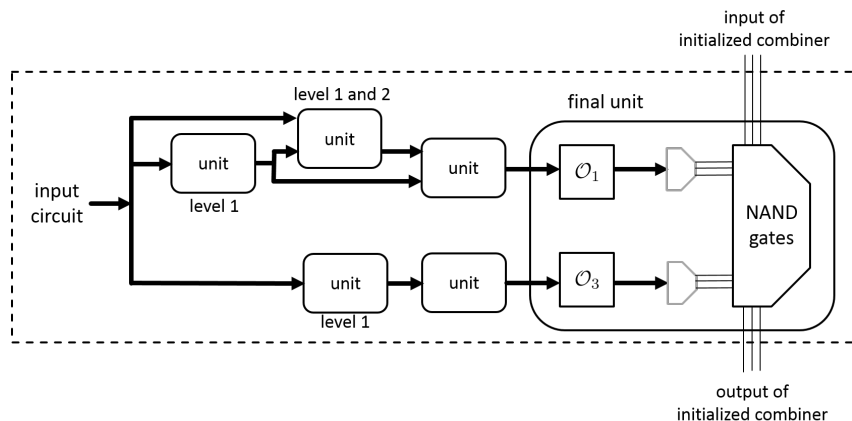


Fig. 3: Combiner circuit consisting of units.

The above assumes that the class of obfuscatable circuits for structural circuits is closed under recursive constructions of units. We note that for concrete constructions such as our 3-out-of-4 combiner in the next section it suffices that we can also obfuscate level-1 units of the original input circuits. Given a circuit class $\mathcal{C} = (\mathcal{C}_\lambda)_{\lambda \in \mathbb{N}}$, some fixed structural combiner Comb , and fixed obfuscators $\mathcal{O}_1, \mathcal{O}_2, \dots$ we denote by $\mathcal{C}^{\text{Comb}} = (\mathcal{C}_\lambda^{\text{Comb}})_{\lambda \in \mathbb{N}}$ the class of circuits which, besides all circuits $C \in \mathcal{C}_\lambda$, for any C also includes all possible initializations of all units of the combiner (except for the final unit) for the given obfuscators. It is understood that, when considering a specific combiner Comb , all candidate obfuscators $\mathcal{O}_1, \mathcal{O}_2, \dots$ must be able to handle the class $\mathcal{C}^{\text{Comb}}$, whereas the combiner only works for the “inner” class \mathcal{C} . Instructively, one may think of \mathcal{C} as the class one would like to obfuscate, although the candidate obfuscators allow for broader classes.

3 Robust 3-out-of-4 Combiner for Obfuscators

In this section we present a 3-out-of-4 (structural) combiner for obfuscation, depicted in Figure 1 on Page 5.

3.1 Construction

The idea is to first obfuscate the input circuit C by all combinations of 3 out of the 4 given obfuscators $\mathcal{O}_1, \dots, \mathcal{O}_4$ and for each combination taking the majority of the output of the three obfuscated circuits. Note that since at least 2 of the 3 obfuscators in such a combination work properly, the majority decision provides a functionally correct output. Formally, for the majority circuit MAJ combining three input circuits by evaluating each one for a given input x and taking the bit-wise majority of the outputs, we thus build the circuits

$$\mathcal{O}_{i_1, i_2, i_3} \leftarrow \text{MAJ}(\mathcal{O}_{i_1}(C), \mathcal{O}_{i_2}(C), \mathcal{O}_{i_3}(C)), \quad 1 \leq i_1 < i_2 < i_3 \leq 4$$

for all possible 4 combinations of i_1, i_2, i_3 . Since we merely need an arbitrary 3 of these 4 circuits for the next stage, we take the combinations leaving out obfuscators 1, 2 and 3 (in this order).

Of course, a corrupt obfuscator among $\mathcal{O}_{i_1}, \mathcal{O}_{i_2}, \mathcal{O}_{i_3}$ in the majority combination could still reveal information about the input circuit C . We hence add another layer where we now combine three of the majority combinations as before, by running each combination $\mathcal{O}_{i_1, i_2, i_3}$ through the complementary obfuscator \mathcal{O}_{i_4} and taking the majority of these circuits again. Put differently, we now build the circuit

$$\text{MAJ}(\mathcal{O}_1(\mathcal{O}_{2,3,4}), \mathcal{O}_2(\mathcal{O}_{1,3,4}), \mathcal{O}_3(\mathcal{O}_{1,2,4})).$$

Functional correctness of our combiner is guaranteed because each of the input circuits $\mathcal{O}_{2,3,4}, \mathcal{O}_{1,3,4}, \mathcal{O}_{1,2,4}$ computes the correct function and at least two of the level-2 obfuscators $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3$ are correct. The obfuscation property holds because if one of the level-2 obfuscators, say, \mathcal{O}_1^* , is malicious, then the level-1 obfuscators generating $\mathcal{O}_{2,3,4}$ already hide the input circuit. Furthermore, the malicious obfuscator \mathcal{O}_1^* cannot bias the functional correctness of the circuits $\mathcal{O}_{1,3,4}$ and $\mathcal{O}_{1,2,4}$ in the other branches, such that the sound second-layer obfuscators $\mathcal{O}_2, \mathcal{O}_3$ also hide $\mathcal{O}_{1,3,4}$ and $\mathcal{O}_{1,2,4}$ and thus the input circuit C , even if \mathcal{O}_1^* on the first level reveals information about C .

3.2 Security

We start by showing that the combiner is (strongly) robust for indistinguishability obfuscation. Recall that strong robustness refers to the fact that each property, functional correctness and obfuscation, is preserved individually. Note that for our combiner (and also the security proof) it suffices that the parties merely have black-box access to all obfuscators.

Theorem 1. *The combiner in Figure 1 is a strongly robust 3-out-of-4 combiner for indistinguishability obfuscation.*

Proof. Functional correctness is straightforward, given that for each unit at least two obfuscators are functionally correct and since we apply the majority of the outputs.

We next show indistinguishability. Take an arbitrary distinguisher \mathcal{D} against our combiner. We need to show that there exists a negligible function ϵ such that for an arbitrary pair $C_0, C_1 \in \mathcal{C}_\lambda$ of circuits, the distinguishing advantage of \mathcal{D} is smaller than $\epsilon(\lambda)$. The idea is to show that one can gradually replace the input circuits C_0 to the obfuscators in the combiner by circuit C_1 , taking some care with the single corrupt obfuscator.

For the gradual replacement fix the order of the nine level-1 obfuscators $\mathcal{O}_2, \mathcal{O}_3, \mathcal{O}_4, \dots, \mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_4$ according to their appearance in Figure 1 from top to down, with one exception: for a parameter $k \in \{1, 2, 3, 4\}$, a reminiscent for the index of the corrupt obfuscator \mathcal{O}_k^* , we move all occurrences of this obfuscator to the very end of the list. For instance, for $k = 2$ we would have the order $\mathcal{O}_3, \mathcal{O}_4, \mathcal{O}_1, \dots, \mathcal{O}_4, \mathcal{O}_2^*, \mathcal{O}_2^*$. Let $K = K(k) \in \{7, 8\}$ be the first index of \mathcal{O}_k^* in that list. Define now the random hybrid variables $H_i^k(C_0, C_1)$ for $i = 0, 1, \dots, 9$ as the output of our combiner if we pass circuit C_0 for the first i obfuscators (according to our order) and C_1 for the remaining $9 - i$ ones. Then, clearly $H_9^k(C_0, C_1)$ corresponds to the distribution of our combiner for input C_0 , and $H_0^k(C_0, C_1)$ to the one of our combiner for C_1 . It hence suffices to show for any i that \mathcal{D} 's probability of distinguishing adjacent H_{i-1}^k, H_i^k is negligible.

To bound the advantage of \mathcal{D} for each pair (H_{i-1}^k, H_i^k) we will wrap the algorithm into a sequence of distinguishers \mathcal{D}_i^k for $i = 1, 2, \dots, 9$. The distinguisher \mathcal{D}_i^k works in two modes, depending on the status of the i -th obfuscator in our sequence:

- If i is such that the i -th obfuscator is not corrupt, i.e., $i < K(k)$, then \mathcal{D}_i^k expects as input a pair C_0, C_1 and an obfuscated circuit O' generated by the i -th obfuscator in our order for C_b , $b \in \{0, 1\}$. Algorithm \mathcal{D}_i^k computes the output of our combiner (with the given obfuscators) but inserts C_1 as input in the first $i - 1$ level-1 obfuscators, O' as the *output* of the i -th level-1 obfuscator, and C_0 as input in the final $9 - i$ slots. It completes the output O of the combiner for these data and lets \mathcal{D} run on C_0, C_1 and O . Algorithm \mathcal{D}_i^k returns whatever \mathcal{D} outputs.
- If the i -th obfuscator is corrupt, i.e., $i \geq K(k)$, then \mathcal{D}_i^k expects as extra auxiliary input a pair C_0, C_1 and a sample O' of one of the sound obfuscator candidates. Here the obfuscator \mathcal{O}_j producing O' is determined by looking at the level-1 unit u in which the i -th (corrupt) obfuscator \mathcal{O}_k^* appears. For this unit, and its three obfuscators, there exists the fourth, complementing obfuscator \mathcal{O}_j to which the unit's output is fed to on the level-2 unit. For instance, if $k = 2$, $K = 8$, and $i = 8$, then the corresponding level-1 unit u is the top one in Figure 1, and the complementing obfuscator is \mathcal{O}_1 . The input to the complementing obfuscator \mathcal{O}_j for deriving O' is either a sample of the level-1 unit where all honest obfuscators are initialized with

C_0 and the corrupt one with C_1 , or all of them are initialized with C_0 . By assumption, both samples are in the class $\mathcal{C}_\lambda^{\text{Comb}}$ such that the sample can be passed to \mathcal{O}_j . Algorithm \mathcal{D}_i^k now evaluates our combiner, by replacing inputs to obfuscators up to index i by C_1 , for subsequent indices giving input C_0 , and replacing the output of the complementing obfuscator \mathcal{O}_j in unit u when evaluating our combiner by O' . Return \mathcal{D} 's output bit on input C_0, C_1 and the combiner's output O .

Assume i is such that the i -th obfuscator in order is still different from \mathcal{O}_k^* , i.e., $i < K$. Then if O' is the obfuscation of C_1 , then \mathcal{D}_i^k runs \mathcal{D} exactly on the distribution of the hybrid variable $H_{i-1}^k(C_0, C_1)$. In particular, for $i = 1$ algorithm \mathcal{D}_i^k runs \mathcal{D} on a sample of our combiner's output for C_1 . Analogously, for $i = 9$ and O' stemming from the complementing obfuscator for a sample of the level-1 unit with all C_0 inputs, the input to \mathcal{D} is distributed like a sample of our combiner for C_0 (and thus of $H_9^k(C_0, C_1)$).

Assume that the k -th obfuscator is indeed corrupt. For $i < K$ it follows from the indistinguishability obfuscation of the sound obfuscators that there exist negligible functions $\epsilon_i(\lambda)$ such that for any C_0, C_1 , the advantage of \mathcal{D}_i^k in distinguishing the two input cases is at most $\epsilon_i(\lambda)$. For $i \geq K$ this follows as the input circuits to the two sound obfuscators in unit u are already C_0 , such that the majority computation of the unit ensures that in both cases the unit circuit computes the function $C_0(\cdot)$. It follows that both input circuits to the complementing obfuscator \mathcal{O}_j compute the same function and we can again conclude from the security of the obfuscator that the advantage must be bounded by some function $\epsilon_i(\lambda)$. Note that here we take advantage of the fact that indistinguishability holds for all circuits and therefore in particular also for our partly combiner samples.

It therefore also holds for any i that the advantage of \mathcal{D} in distinguishing $H_{i-1}(C_0, C_1)$ and $H_i(C_0, C_1)$ for any C_0, C_1 is at most $\epsilon_i(\lambda)$, too. Hence, the overall advantage of \mathcal{D} is at most $\epsilon(\lambda) := \sum_{i=1}^9 \epsilon_i(\lambda)$ and thus negligible.⁹

The claim carries over to the case of differing-inputs obfuscation. Recall that the main difference to indistinguishability obfuscation is that, for the differing-inputs case, the circuits in question are generated by an algorithm `Sampler` such that the circuits may compute different functions, but `Sampler` ensures that finding differing inputs is infeasible. We can basically apply the same hybrid argument in this case as above. However, for the step $i \geq K$, when using the obfuscation of our level-1 unit, we need to specify sampler `Sampler'_k` with oracle access to $\mathcal{O}_1, \dots, \mathcal{O}_4$ to generate the input circuit for the complementing obfuscator. Algorithm `Sampler'_k` first runs `Sampler` to get (C_0, C_1, aux) , then generates two samples of the level-1 unit (one time using C_0 for the honest obfuscators and C_1 for \mathcal{O}_k^* , and the other time using C_0 everywhere), and finally outputs these two samples and $\text{aux}' = (C_0, C_1, \text{aux})$ as auxiliary data. Note that finding an input

⁹ Note that we do not need to know the index k of the obfuscator; unlike the construction it suffices that the proof provides an existential result.

x where the two level-1 unit samples differ is impossible, as both implement the same function.

We next show that the claim remains true with respect to virtual black-box and grey-box obfuscation. For this we assume that the adversary and the simulator receive some circuit-dependent auxiliary input \mathbf{aux} as additional input, as explained in Section 2.

Proposition 1. *The combiner in Figure 1 is a strongly robust 3-out-of-4 combiner for virtual black-box and grey-box obfuscation with respect to dependent auxiliary input.*

Proof. Functional correctness follows as in the case of indistinguishability obfuscation. We only discuss the VBB property here; the VGB property follows analogously.

Consider an adversary \mathcal{A}_0 against VBB obfuscation. This adversary receives an output sample O' of our combiner as input and some auxiliary input $\mathbf{aux}[0] = \mathbf{aux}[0](C)$. Let k be again the index of the malicious obfuscator and this time define $L = L(k) \in \{3, 5\}$ as follows. For $k = 4$ we would have the malicious obfuscator \mathcal{O}_4^* only on first-level units and we only need to look at the $L = 3$ second-level obfuscators. For $k \in \{1, 2, 3\}$, on the other hand, the malicious combiner appears in a second-level unit and we thus consider the $L = 5$ sound obfuscators, consisting of the 3 obfuscators leading to the second-level appearance of \mathcal{O}_k^* and the remaining 2 honest level-two obfuscators.

Assume now that we change the auxiliary input to include the obfuscator results of our combiner for all L sound obfuscators defined above. Denote these intermediate results, ordered according to the obfuscator application, by $O[1..L] = (O_{i_1}, O_{i_2}, O_{i_3}, \dots, O_{i_L})$, and let $O[1..i]$ denote the first i entries in $O[1..L]$. Let $\mathbf{aux}[0..i]$ denote the sample given by a sample of first i obfuscator outputs, together with the (independent) sample $\mathbf{aux}[0]$ of \mathcal{A}_0 .

Instead of considering $\mathcal{A}_0(1^\lambda, O', \mathbf{aux}[0])$ we construct an algorithm \mathcal{A}_1 which receives 1^λ and $\mathbf{aux}[0..L]$ as input, assembles a combiner output O' from $\mathbf{aux}[1..L]$ by possibly evaluating the (level-2) malicious obfuscator, and runs adversary $\mathcal{A}_0(1^\lambda, O', \mathbf{aux}[0])$. Then, clearly, the output distribution of both algorithms are identical. We can now view \mathcal{A}_1 as an algorithm which receives $\mathbf{aux}[0..L-1]$ as auxiliary input, and the obfuscated circuit $\mathbf{aux}[L]$ together with 1^λ as regular input.¹⁰ For this algorithm \mathcal{A}_1 , by assumption about the security of \mathcal{O}_{i_L} producing O_{i_L} , there exists a simulator $\mathcal{S}_1^C(1^\lambda, \mathbf{aux}[0..L-1])$ with negligibly close output distribution.

Given \mathcal{S}_1 we construct an adversary \mathcal{A}_2 which receives auxiliary input $\mathbf{aux}[0..L-2]$, and 1^λ and $\mathbf{aux}[L-1]$ as regular input. It runs $\mathcal{S}_1(1^\lambda, \mathbf{aux}[0..L-2])$ and uses $\mathbf{aux}[L-1]$ to answer oracle calls. Note that, by the functional correctness of $\mathcal{O}_{i_{L-1}}$, using $\mathbf{aux}[L-1]$ to simulate the oracle C of \mathcal{S}_1 is sound as both circuits compute the same function. We can set this argument forth to eventually obtain

¹⁰ By construction, if we shift the input of a level-2 obfuscator then this is a sample of a level-1 unit, whereas the other auxiliary inputs are based on the original and functional equivalent circuit C .

a simulator $\mathcal{S}_L^C(1^\lambda, \text{aux}[0])$, producing some output distribution which is negligibly close to the one of our initial adversary $\mathcal{A}_0(1^\lambda, O', \text{aux}[0])$. This shows VBB obfuscation.

4 Lower Bounds for Combiners

To illustrate how we use the two required security properties, function preservation and indistinguishability, against each other to derive our general result, it is useful to demonstrate our technique for some toy examples. In the examples we use an unspecified notion of indistinguishability of the obfuscators as we merely highlight the issues; the reader may think for sake of concreteness of the notion of indistinguishability obfuscation.

4.1 Simple Attempts That Fail

The first attempt to build a secure combiner consists of a single unit and is given in the left hand part of Figure 4. It uses three obfuscators $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3$ and runs the input circuit through each of them. Then it combines the three obfuscated circuits by a majority circuit. Note that this means that this part takes the input x and outputs y as the bit-wise majority of the answers. While we cannot break the functional correctness of the combiner with a single corrupt obfuscator \mathcal{O}_i^* , we can easily break the indistinguishability property. To this end we take control of obfuscator \mathcal{O}_1^* and let it simply output the input circuit in clear. Note that this means that the unit, after having been initialized, reveals the input circuit in clear as well, and this easy to distinguish.




Fig. 4: Examples of (insecure) structural combiners

In our second example we have two obfuscators $\mathcal{O}_1, \mathcal{O}_2$, let the output of them be combined arbitrarily, and then input the derived circuit into obfuscator \mathcal{O}_3 . Note that in this case it is unclear how to break the indistinguishability property by corrupting a single obfuscator only. If it is \mathcal{O}_3 then the obfuscators of the first unit already hide the input circuit; if we corrupt one of the obfuscators $\mathcal{O}_1, \mathcal{O}_2$ then the final obfuscation hides the actual circuit.

We can, nonetheless, in the second example break the functional preservation property. Namely, assume that both \mathcal{O}_1 and \mathcal{O}_2 are secure and that there are two potential input circuits D_0, D_1 computing *different* functions for the same input and output length. If we control \mathcal{O}_1 , then we let it on any input circuit rather obfuscate D_1 . Vice versa, if we control \mathcal{O}_2 then we let it always obfuscate D_0 , independently of the actual input. It follows that the initialized combiners for D_0 (with our malicious \mathcal{O}_1^* and with genuine \mathcal{O}_2) and for D_1 (with genuine \mathcal{O}_1 and our malicious \mathcal{O}_2^*) have the same distribution. For at least one of the two cases the computed function must then be incorrect, as the initialization samples for both input circuits D_0, D_1 have the same distributions in both cases.

4.2 The General Case of 2-out-of-3 Combiners

The attacks in the simple case show the path for our general impossibility result for 2-out-of-3 structural combiners. If one of the three obfuscators appears in all units on the path from some level-1 unit to the final unit, then it can pass on information about the input circuit C to the final unit. This is done by forwarding some information about the input circuit in the output of the obfuscator. This would clearly violate the indistinguishability property. Hence, on all paths there must be a unit which only uses (at most) the same two obfuscators. But then we can “confuse” the combiner as we did in the second example above. The argument, however, requires some care to deal with the fact that we have many paths. Note also that the confusion strategy fails in our 3-out-of-4 construction because the majority of the three combiners yields the correct function.

For sake of concreteness we use the notion of indistinguishability obfuscation for the obfuscators and the combiners. Recall that this means that for any functionally equivalent circuits C_0, C_1 from class \mathcal{C} the combiners initialization with these two circuits must be computationally indistinguishable. To avoid trivial cases we assume that the class \mathcal{C} contains at least two distinct but functionally equivalent circuits E_0, E_1 , and that it also contains two circuits D_0, D_1 computing *different* functions. We call such classes *non-trivial*.

Since we pass on circuits as inputs we need to fix some encoding. Let $\langle \cdot \rangle$ denote such a function mapping circuits from the class \mathcal{C} to strings. We assume that the encoding is such that given an encoding of a unit (after initialization) one can reconstruct the circuits output by the obfuscators. That is, there exists an efficient algorithm `reconstruct` such that given any initialization $V \leftarrow U(C_1, C_2, C_3, \dots)$ of a unit U , including obfuscated circuits $O_i \leftarrow \mathcal{O}_{i_j}(C_i)$ and possibly pass-through circuits $O_i = C_i$, we have `reconstruct`($\langle V \rangle$) = $(\langle O_1 \rangle, \langle O_2 \rangle, \langle O_3 \rangle, \dots)$. Furthermore, we assume that the encoding of no unit coincides with the encoding of our equivalent circuits E_0 or E_1 such that it is clear if each $\langle O_i \rangle$ is the result of a unit initialization or rather one of the circuits E_0 or E_1 . We call such encodings $\langle \cdot \rangle$ *admissible*. Note that this means that the combiner itself cannot apply any obfuscation techniques beyond the ones provided by the obfuscators placed inside the unit.

We first show that on any full path (from level-1 units to the final unit) each of the three combiners does not appear in a unit (and that a unit in which it

does not appear is not pass-through). This holds for any structural combiner, independently of the total number of obfuscators and the number of malicious ones:

Lemma 1. *Let $\text{Comb}^{\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_N}$ be a structural combiner for a non-trivial circuit class \mathcal{C} with admissible encoding $\langle \cdot \rangle$. Then for any full path of units of the combiner and for any $i \in \{1, 2, \dots, N\}$ there must be a unit which is not pass-through and which is not an $\{i, \dots\}$ -unit, or else the combiner cannot be an indistinguishable obfuscator.*

Proof. Assume that there exists a full path of units and an $i \in \{1, 2, \dots, N\}$ such that each unit on the path is an $\{i, \dots\}$ -unit or that it is pass-through (or both). Then we show how to break indistinguishability obfuscation as follows. Let E_0, E_1 be some functional equivalent circuits in the class with distinct encodings under $\langle \cdot \rangle$. We corrupt obfuscator \mathcal{O}_i and for each input circuit let it, for each call, simply output the input circuit in clear, by duplicating the input description.

Since each unit on the pass includes the i -th obfuscator (or is pass through) a distinguisher can distinguish between a combiner obfuscation of E_0 and E_1 as follows. The distinguisher receives as input the initialization of the final unit U , and runs $\text{reconstruct}(\langle U \rangle)$ to recover all (obfuscated or pass-through) input circuits O_1, O_2, \dots . Since the distinguisher knows the layout of the combiner it can recursively apply the reconstruction algorithm to outputs of the i -th combiner resp. to passed circuits; both are initialized units. Following the full path in question, the distinguisher eventually obtains either E_0 or E_1 as the input circuit, and can thus distinguish the two cases easily. \square

We next show that, given that each full path contains a unit in which, say, obfuscator \mathcal{O}_3 does not appear, we can confuse the combiner. This time, the claim only holds for 2-out-of-3 combiners:

Lemma 2. *Let $\text{Comb}^{\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3}$ be a structural combiner for a non-trivial circuit class \mathcal{C} with admissible encoding $\langle \cdot \rangle$. Then the combiner cannot be perfectly correct.*

In particular, if u denotes the number of units in the structural combiner and m the maximal number of obfuscator gates in a unit, then with probability at least $2^{-u(mu)^{-mu}}$ (over the random choices of the obfuscators) the combiner's function is different from the one of the input circuit. If u and m are constant, for instance, this means a constant error in functional preservation.

Proof. By Lemma 1 for each path from level-1 units to the final unit there exists a unit which does not contain, say, the obfuscator \mathcal{O}_3 and which is neither pass-through. Put differently, such a unit contains (at most) the obfuscators $\mathcal{O}_1, \mathcal{O}_2$, each one possibly multiple times. Let U_1, U_2, \dots be the corresponding units which we call *confusion units*. In the example in Figure 5 the confusion units on the three paths are marked by dotted lines.

We consider two cases, one time corrupting obfuscator \mathcal{O}_1 , the other time corrupting obfuscator \mathcal{O}_2 . Let us first consider the case that we corrupt obfuscator \mathcal{O}_1 . Our version \mathcal{O}_1^* of the obfuscator will internally hold, and formally




Fig. 5: Confusion units in this example are marked by dotted lines.

attributed to the non-uniformity, an initialization sample of $\text{Comb}^{\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3}(D_1)$ with the genuine obfuscators for input circuit D_1 . In particular, for each confusion unit U_i it will include the $j \leq m$ circuit codes of $\mathcal{O}_i^j[D_1]$ which the original obfuscator \mathcal{O}_1 output in unit U_i in this sample. In order to make our obfuscator state-free we will guess the right insertion positions and injected circuits. That is, for each call (about some input circuit) our malicious obfuscator \mathcal{O}_1^* tosses a coin. If it comes out as head, then the obfuscator proceeds as the genuine obfuscator would. If it is tail, then it picks one of the at most mu circuits $\mathcal{O}_i^j[D_1]$ at random, and returns this circuit. An example of a run with good guesses is given in the left part of Figure 5.

For the other case we corrupt \mathcal{O}_2 and include a sample of $\text{Comb}^{\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3}(D_0)$ of the genuine obfuscators, this time for input circuit D_0 . Analogously to the other case denote the output of \mathcal{O}_2 in the confusion unit U_i by $\mathcal{O}_i^j[D_0]$. When called, the malicious obfuscator \mathcal{O}_2^* also generates an honest answer with probability $\frac{1}{2}$, and inserts one of the pre-sampled circuits $\mathcal{O}_i^j[D_0]$, the choice made at random, in the other case.

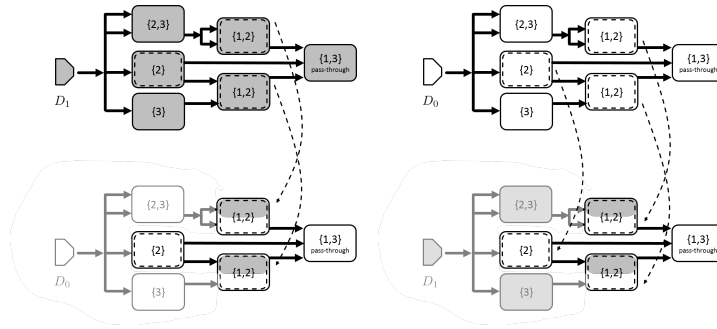


Fig. 6: Confusion strategy with malicious obfuscator \mathcal{O}_1^* injecting parts of the upper D_1 initialization sample into the lower D_0 initialization (left), and malicious obfuscator \mathcal{O}_2^* injecting parts of the upper D_0 initialization sample into the lower D_1 initialization (right).

For the analysis we start with the case of a malicious obfuscator \mathcal{O}_1^* . Note that, if we let u denote the number of units in the combiner, then with probability

2^{-u} we overwrite the obfuscator’s behavior exactly for the confusion units, since we predict the status of each unit (confusion or not) exactly with probability $\frac{1}{2}$. If so, then we also inject the hardwired circuits $\mathcal{O}_i^j[D_1]$ “correctly” in confusion unit U_i with probability at least $(mu)^{-mu}$, since we have at most u units with at most m obfuscator gates and need to guess for each unit correctly among the at most mu possibilities among all $\mathcal{O}_i^j[D_1]$ ’s. If this happens, and the combiner receives circuit D_0 as input, then in the confusion units we have consistent samples for \mathcal{O}_1 gates (if present), as if the combiner’s input had been D_1 . See the left part of Figure 6 for an example. Simultaneously, in the same unit, we have consistent samples for \mathcal{O}_2 gates (if present), as if the overall input had been circuit D_0 .

By symmetry, the same is true if we control obfuscator \mathcal{O}_2^* and the combiner’s input is D_1 . Hence, with probability at least $2^{-u}(mu)^{-mu}$ either case creates the same output distribution. If this happens, then on each path to the final unit the corresponding confusion unit produces the same distribution upon the single initialization in both cases. It follows that the combiner must implement an incorrect function in one of the cases, showing that functional preservation is not satisfied. It follows that the combiner cannot be perfectly correct.

Noting that the combiner cannot work even if 2 of the 3 obfuscators both have both properties simultaneously, the previous lemmas imply that there are not even weakly robust 2-out-of-3 combiners for indistinguishability obfuscation. It follows that there cannot exist stronger forms of structural combiners either, such as 1-out-of-2 combiners, strong combiner, or virtual grey-box combiners.

Theorem 2. *For any $o \in \{VBB, VGB, \text{indistinguishability}, \text{differing-inputs}\}$ there is no structural weakly robust 2-out-of-3 o -obfuscation combiner $\text{Comb}^{\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3}$ for non-trivial circuit classes \mathcal{C} with admissible encoding $\langle \cdot \rangle$.*

5 The General Case of $(2\gamma[+1])$ -out-of- $(3\gamma[+1])$ Combiners

In this section we present a generalization of our 3-out-of-4 combiner to the case of $(2\gamma+1)$ -out-of- $(3\gamma+1)$ combiners for any fixed integer γ . In fact, our combiner for $\gamma = 1$ in Section 3 can be seen as a special parallelized version of the general approach here. We then discuss that our lower bound for 2-out-of-3 structural combiners also carries over to the more general case of 2γ -out-of- 3γ combiners, showing that our general combiner here is optimal in this regard.

5.1 Robust $(2\gamma + 1)$ -out-of- $(3\gamma + 1)$ Combiners

Consider all sets I of subsets of $\{1, 2, \dots, 3\gamma+1\}$ of size $2\gamma+1$. For each such set I form the unit which, similar to our 3-out-of-4 case, first in parallel obfuscates the input circuit with each obfuscator \mathcal{O}_i for $i \in I$, and then compute the majority circuit over all these $2\gamma + 1$ obfuscated circuits. We write

$$O_I(\cdot) = \text{MAJ}\{\mathcal{O}_i(\cdot) \mid i \in I\}$$

for this unit. To obfuscate a circuit C compose each of these units for all the I 's sequentially, in arbitrary order. Let us denote this process by

$$\left(\prod_I O_I\right)(\cdot) = O_{I_\ell}(\cdots O_{I_3}(O_{I_2}(O_{I_1}(\cdot)))\cdots)$$

for constant $\ell = \binom{3\gamma+1}{2\gamma+1}$. Call this the *sequential-subset combiner* for γ .

Intuitively, the sequential-subset combiner guarantees robustness as there exists a subset I such that this subset only uses the $2\gamma+1$ uncorrupt obfuscators. At the same time each circuit O_I computes the correct function as the majority of the $2\gamma+1$ obfuscators faithfully computes the correct function.

Theorem 3. *For any constant γ the sequential-subset combiner is a strongly robust $(2\gamma+1)$ -out-of- $(3\gamma+1)$ combiner for indistinguishability obfuscation, for differing-inputs obfuscation, for virtual black-box obfuscation, and for grey-box obfuscation, the latter ones for dependent auxiliary inputs.*

The proof is similar to our 3-out-of-4 combiner. Functionally correctness follows from the fact that the majority computation in each O_{I_i} ensures that the at most γ corrupt obfuscators cannot bias the outcome. Obfuscation follows as before because there must exist one set I which exclusively contains non-malicious obfuscators.

5.2 Impossibility for 2γ -out-of- 3γ Combiners

In this section we discuss that our lower bound for 2-out-of-3 structural combiners carries over to the more general case of 2γ -out-of- 3γ combiners.

Theorem 4. *There is no structural weakly robust 2γ -out-of- 3γ obfuscation combiner $\text{Comb}^{\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_{3\gamma}}$ for non-trivial circuit classes \mathcal{C} with admissible encoding $\langle \cdot \rangle$.*

Proof. Recall the proof for the 2-out-of-3 case. There, in the first step we have shown that on each path from a level-1 unit to the output unit there must be a unit in which obfuscator \mathcal{O}_3 does not appear and which is not pass-through. We called these units confusion units.

The same argument now applies here as well for the γ obfuscators with indices $2\gamma+1, \dots, 3\gamma$. Else, if there was a path in which one of these obfuscators appears in each unit (or if the unit is pass-through), then we could easily corrupt these obfuscators and forward information about the input circuit through the admissible encoding $\langle \cdot \rangle$. Hence, in the case here there must be a confusion unit on each path, which only uses circuits with indices $1, 2, \dots, 2\gamma$ and which are not pass-through.

In the second step of the proof for the 2-out-of-3 case we then show that in the confusion units with obfuscators \mathcal{O}_1 and \mathcal{O}_2 we can confuse the combiner. One time we corrupt \mathcal{O}_1 and let it insert samples of circuit D_1 , and the other time we corrupt \mathcal{O}_2 and insert samples for D_0 , where D_0, D_1 compute different

functions. Then the combiner’s view when run on input D_0 in the first case, and on D_1 in the second case, has the same distribution and the combiner cannot provide functional correctness.

We apply the same argument here, one time corrupting the first γ obfuscators with indices $1, \dots, \gamma$ and inserting a sample for D_1 , and the other time corrupting obfuscators with indices between $\gamma + 1, \dots, 2\gamma$ and using a sample for D_0 . Then the combiner’s views in both cases (for input circuit D_0 in the first case, and for D_1 in the second case) are identical again such that it cannot provide a correct combiner.

As in the 2-out-of-3 case the malicious obfuscators above insert the confusion samples at random positions, such that it only achieves confusion with the same bound as in the previous case. Note also that we took advantage of the fact that corrupt combiners are coordinated centrally by the adversary.

6 Detecting Combiners

The combiners in the previous section were correcting in the sense that they guaranteed functionality correctness if a quorum of obfuscator candidates is secure. Here we consider combiners which should create circuits which either output the correct value, but may give some error output \perp . We call them detecting combiners.

For detecting combiners we require a weaker correctness property, namely that for any circuit $C \in \mathcal{C}$, for any $O \leftarrow \text{Comb}^{\mathcal{O}_1, \mathcal{O}_2, \dots}(C)$ we have that $O(x) \in \{C(x), \perp\}$ for all $x \in \{0, 1\}^*$ in the domain of C . This means that the combiner may sometimes fail to compute the correct function value but then it signals this by outputting a special symbol \perp . To prevent trivial solutions like the combiner which outputs the circuit that always returns \perp we assume that $C \equiv O$ if all obfuscators are secure. Note that our assumption about the obfuscators $\mathcal{O}_1, \mathcal{O}_2, \dots$ being able to deal with (intermediate) combiner outputs in $\mathcal{C}^{\text{Comb}}$ implies that the obfuscators may now also receive circuits which occasionally output \perp .

6.1 Robust $(\gamma + 1)$ -out-of- $(2\gamma + 1)$ Detecting Combiners

To build our $(\gamma + 1)$ -out-of- $(2\gamma + 1)$ combiner we follow the approach of our sequential-subset combiner. We can also straightforwardly give the optimized version for the case of a 1-out-of-3 combiner, akin to our 1-out-of-4 combiner, but omit this step here. To build the sequential-subset combiner consider here all sets I of subsets of $\{1, 2, \dots, 2\gamma + 1\}$ of size $\gamma + 1$. For each such set I we first obfuscate the input circuit with each obfuscator \mathcal{O}_i for $i \in I$. But now instead of completing the computation by adding a majority sub circuit, we now use the detecting version which (a) either outputs the string on which all circuits agree upon as output (even if it is \perp), or (b) returns \perp if there is no such unanimous decision. Let

$$O_I(\cdot) = \text{UNAN} \{\mathcal{O}_i(\cdot) \mid i \in I\}$$

denote this unit with the unanimity circuit at the end. For obfuscation of C now compute the sequential-subset combiner

$$\left(\prod_I O_I\right)(\cdot) = O_{I_\ell}(\cdots O_{I_3}(O_{I_2}(O_{I_1}(\cdot))) \cdots)$$

as before for constant $\ell = \binom{2\gamma+1}{\gamma+1}$.

Theorem 5. *For any constant γ the sequential-subset combiner is a strongly robust $(\gamma+1)$ -out-of- $(2\gamma+1)$ detecting combiner for indistinguishability obfuscation, for differing-inputs obfuscation, for virtual black-box obfuscation, and for grey-box obfuscation, the latter ones for dependent auxiliary inputs.*

The proof is similar to the case of correcting combiners, except that we only guarantee the weaker functional correctness. This property is given since in each unit for index set I there is at least one honest obfuscator among the $\gamma+1$ ones, the unanimity circuit either outputs the function value computed by the honest obfuscator (if all other circuits agree), which may either be the correct function value for some x or \perp , or it returns the error message \perp . It follows that the overall output of the combiner circuit can only comply with the circuit's output, or returns \perp . The obfuscation properties follow as before noting that the obfuscators are able to handle input circuits with output \perp , and that there must exist an index set I which only contains good obfuscators.

6.2 Impossibility of γ -out-of- 2γ Detecting Combiners

The idea for the lower bound for correcting combiners carries over to detecting combiners, with

Theorem 6. *There is no structural weakly robust γ -out-of- 2γ obfuscation combiner $\text{Comb}^{\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_{2\gamma}}$ for non-trivial circuit classes \mathcal{C} with admissible encoding $\langle \cdot \rangle$.*

Proof. As in the case of 2-out-of-3 combiners and 2γ -out-of- 3γ combiners, here, there must be also (non-pass-through) confusion units in each path from input units to the final unit, where none of the obfuscators with indices $\gamma+1, \dots, 2\gamma$ appears. Assume now that we corrupt the obfuscators with indices $1, \dots, \gamma$ and let these obfuscators insert intermediate samples of a circuit D_0 of the combiner's obfuscation in the confusion units, independently of the input. If the insertions happen at the right position with significant probability, then the combiner must output an obfuscated circuit *as if the combiner has been run on D_0 for honest obfuscators*. In particular, the combiner's circuit must then compute the function D_0 on every input. This holds even if the original input circuit was D_1 , computing a different function than D_0 , i.e., $D_0(z) \neq D_1(z)$ for some string z . But then the combiner's circuit produces a false output $D_0(z) \neq \perp$ for input z and cannot be detecting.

7 Implementation and Evaluation

Our formal results have been stated in terms of the common notion of circuit obfuscation. In practice, however, programs are usually considered to be better modeled for Turing machines. We stress that our results, especially for the majority-based combiner, hold for such Turing machine programs as well. Namely, our 3-out-of-4 combiner would then output the program implementing the nested majority implementations.

Concerning provably secure instantiations for Turing machine obfuscation, we note that if the running time and the input length of the Turing machine are bounded then one can in principle transform such machines into corresponding circuits, albeit at the cost of increasing the complexity significantly. A more efficient solution is to use obfuscation techniques for Turing machines directly. Given the current state of constructions this is possible if the input length can be bounded [40] and, for other constructions, if the space is also bounded beforehand [17,11].

To evaluate the suggested combiners for typical obfuscation programs in practice we implemented the PyObf python package [9] that can be used to wrap existing obfuscators and to implement new combiners. Even though the conceptual construction of a combiner is not related to the concrete implementation of the underlying obfuscators, implementations for different programming languages might differ, since some constructions introduce new run-time parts (e.g., as the MAJ circuit in our case) to the program. We chose to use JavaScript as the implementation programming language because of the relatively high number of available obfuscators.

7.1 Performance Evaluation

For performance evaluation we used Yahoo!’s YUICompressor v2.4.8¹¹ as \mathcal{O}_1 , a slightly randomized version of it as \mathcal{O}_2 , Google’s Closure Compiler v20151015¹² as \mathcal{O}_3 , and jsPacker.pl v1.00b¹³ as \mathcal{O}_4 . Note that there is no essential difference between \mathcal{O}_1 and \mathcal{O}_2 , especially in terms of obfuscation overhead, since the latter only uses different and randomized symbol selection routine. Security of combiners usually relies on somewhat independent components but since we are mainly interested in performance evaluation here we opted for using the related choice. The evaluated combiners are:

$$\begin{aligned} C_1(\cdot) &= \mathcal{O}_4(\mathcal{O}_3(\mathcal{O}_2(\mathcal{O}_1(\cdot)))) \\ C_2(\cdot) &= \mathcal{O}_2(\mathcal{O}_1(\mathcal{O}_4(\mathcal{O}_3(\cdot)))) \\ C_3(\cdot) &= \text{MAJ}(\mathcal{O}_1(\mathcal{O}_{2,3,4}), \mathcal{O}_2(\mathcal{O}_{1,3,4}), \mathcal{O}_3(\mathcal{O}_{1,2,4})) \\ C_4(\cdot) &= \text{MAJ}(\mathcal{O}_3(\mathcal{O}_{4,1,2}), \mathcal{O}_4(\mathcal{O}_{3,1,2}), \mathcal{O}_1(\mathcal{O}_{3,4,2})) \end{aligned}$$

¹¹ <https://github.com/yui/yuicompressor>

¹² <https://github.com/google/closure-compiler>

¹³ <http://dean.edwards.name/download/>

The evaluated programs (with varying input size, ranging from a few thousand bytes to a roughly million bytes) are Cookies.js v1.2.2¹⁴ (6,637 bytes), Highlight.js v9.0.0¹⁵ (22,604 bytes), jCarousel v0.3.4¹⁶ (46,007 bytes), Backbone.js v1.2.3¹⁷ (71,415 bytes), Chart.js v1.0.2¹⁸ (109,612 bytes), Epoch v0.8.4¹⁹ (115,940 bytes), Swig v1.4.2²⁰ (143,975 bytes), PhysicsJS v0.7.0²¹ (171,847 bytes), jQuery v1.6.4²² (238,166 bytes), Raphaël v2.1.4²³ (304,254 bytes), Dojo v1.10.4²⁴ (629,481 bytes), Video.js v5.4.4²⁵ (675,527 bytes) and AngularJS v1.4.5²⁶ (1,052,336 bytes).

Note that the above circuit model describes a program as a function with input and output, in contrast to the common software design of JavaScript libraries that heavily depends on the JavaScript context (e.g., the window object). But this difference is irrelevant to performance evaluation.

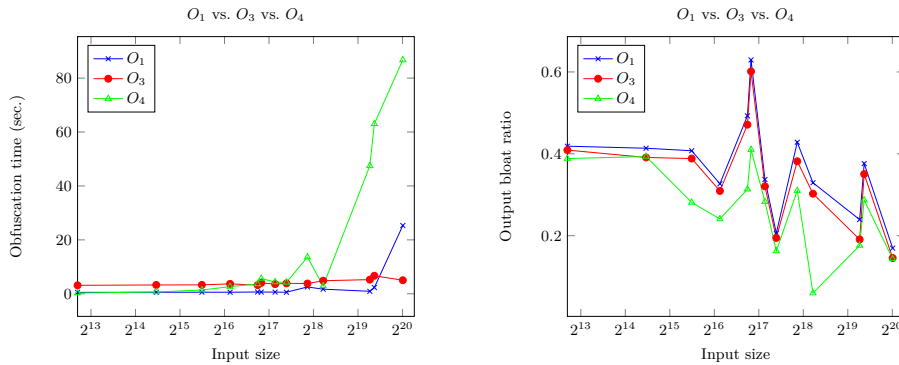


Fig. 7: Overhead of individual obfuscators (time and output bloat ratio) for the various programs (in relation to their input sizes for the obfuscator). Note that we do not display obfuscator \mathcal{O}_2 here as its performance is essentially identical to the one of \mathcal{O}_1 .

Figure 7 gives the effectiveness of the obfuscators in terms of obfuscation time and of output-bloat ratio. Here we show the figures in relation of the sizes

¹⁴ <https://github.com/ScottHamper/Cookies>
¹⁵ <https://highlightjs.org>
¹⁶ <http://sorgalla.com/jcarousel/>
¹⁷ <http://backbonejs.org/>
¹⁸ <http://www.chartjs.org/>
¹⁹ <http://epochjs.github.io/epoch/>
²⁰ <http://paularmstrong.github.io/swig/>
²¹ <http://wellcaffeinated.net/PhysicsJS/>
²² <https://jquery.com/>
²³ <https://github.com/DmitryBaranovskiy/raphael>
²⁴ <https://dojotoolkit.org/>
²⁵ <http://videojs.com/>
²⁶ <https://angularjs.org/>

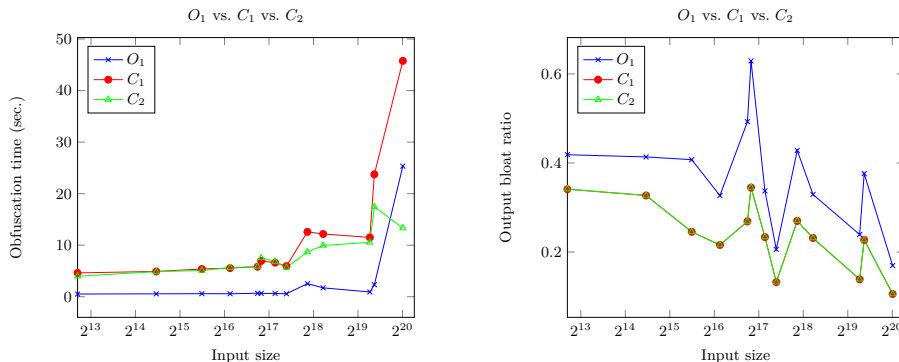


Fig. 8: Overhead of cascaded combiner (time and output bloat ration).

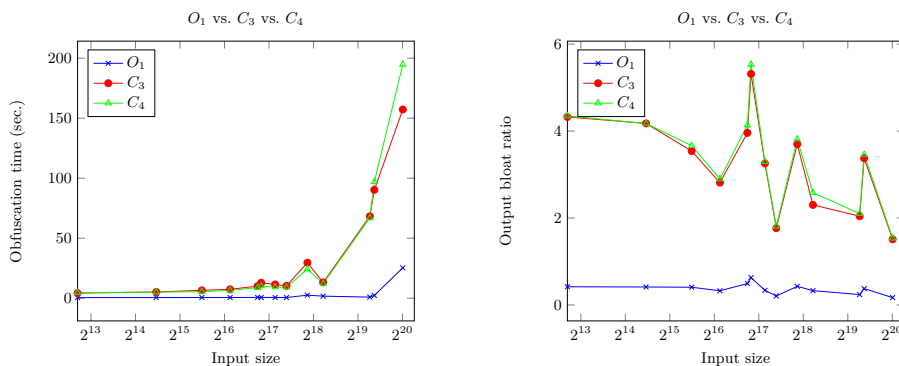


Fig. 9: Overhead of 3-out-of-4 combiner (time and output bloat ration).

of the various programs (from 6,637 bytes to 1,052,336 bytes) given as input to the obfuscators. Note that the results may depend heavily on the specific input programs such that we cannot expect perfectly monotonic behavior in the graphs. Also, as mentioned before, many practical obfuscators come with techniques for code size reduction such that the output bloat ratio can be—and often is—smaller than 1. Next, we compare these figures to the results of the suggested combiners, first to the cascade combiners in Figure 8 and then to the 3-out-of-4 combiners in Figure 9.

In summary, the proposed obfuscation combiners do not add significant run time overhead compared to a single obfuscator. The factor is roughly proportional to the number of invoked instances, with some gains presumably due to the intermediate code optimization. Due to the advanced compression techniques the code size of our cascaded combiners is in the same order as the individual obfuscators. For the 3-out-of-4 combiner we of course get an increased output size because of the tripling for each majority step, potentially also hampering some code reductions.

7.2 Security Evaluation

Due to the unclear situation about security properties of practical obfuscators we have proven robustness of our combiners with respect to the common theoretical notions of obfuscation in the literature. There are approaches to define metrics for practical obfuscators, though. A first approach is by Collberg et al. [20] who define notions for *potency* (the incomprehensibility of the transformed program for humans), *resilience* (the hardness of undoing the transformation through the joint effort of engineers and deobfuscation techniques), and *cost* (the overhead caused by the obfuscator). The measure of quality of an obfuscator is then given by a vector of these three metrics.

While the notion of cost in [20] even distinguishes the full range between exponential and constant overhead in execution resources, the metrics for potency and resilience in [20] are less rigorous. They are accompanied by suitable software complexity measures such as program length or cyclomatic complexity [42]. Anckaert et al. [5] later used software complexity measures, too, for establishing a benchmarking system for obfuscators for binary executables.

While it is beyond the scope of our work here, it may be interesting to benchmark our obfuscation combiners according to the metrics in [5]. Note that their metrics focus on resilience and somewhat neglect the overhead. Since our combiners in principle increase the software complexity at the cost of incurring additional steps, one should expect that combinations of benchmarked obfuscators yield better values in this regard.

8 Conclusion

Our positive results about combiners, and also our lower bounds, indicate how to proceed both in theory and practice. If you only have two available candidates then the best solution appears to be the sequential composition $\mathcal{O}_2(\mathcal{O}_1(\cdot))$, if one can somehow guarantee that the inner obfuscator provides functional correctness. For three candidates (out of which at least two are sound) then our 2-out-of-3 detecting combiner should be the primary choice. To ensure correct output, our 3-out-of-4 combiner provides a secure solution.

Acknowledgments

We are grateful to Christian Collberg for his feedback and encouragement. Marc Fischlin is supported by the Heisenberg grant Fi 940/3-2 and the SPP 1736 grant Fi 940/5-1 of the German Research Foundation (DFG). Amir Herzberg is supported by the Israeli Ministry of Science and Technology.

References

1. Ananth, P., Boneh, D., Garg, S., Sahai, A., Zhandry, M.: Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689 (2013), <http://eprint.iacr.org/2013/689>

2. Ananth, P., Jain, A., Naor, M., Sahai, A., Yagev, E.: Universal obfuscation and witness encryption: Boosting correctness and combining security (2016), <http://eprint.iacr.org/2016/281>
3. Ananth, P., Jain, A.: Indistinguishability obfuscation from compact functional encryption. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 308–326. Springer, Heidelberg (Aug 2015)
4. Ananth, P., Jain, A., Sahai, A.: Indistinguishability obfuscation with constant size overhead. IACR Cryptology ePrint Archive, Report 2015/1023 (2015), <http://eprint.iacr.org/2015/1023>
5. Anckaert, B., Madou, M., Sutter, B.D., Bus, B.D., Bosschere, K.D., Preneel, B.: Program obfuscation: a quantitative approach. In: Proceedings of the 3th ACM Workshop on Quality of Protection, QoP 2007, Alexandria, VA, USA, October 29, 2007. pp. 15–20. ACM (2007)
6. Badrinarayanan, S., Miles, E., Sahai, A., Zhandry, M.: Post-zeroizing obfuscation: New mathematical tools, and the case of evasive circuits. In: Advances in Cryptology - EUROCRYPT 2016. Lecture Notes in Computer Science, vol. 9666, pp. 764–791. Springer (2016)
7. Barak, B., Garg, S., Kalai, Y.T., Paneth, O., Sahai, A.: Protecting obfuscation against algebraic attacks. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 221–238. Springer, Heidelberg (May 2014)
8. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. *J. ACM* 59(2), 6 (2012)
9. Bin Noon, H.: Pyobf. <https://github.com/hodbn/pyobf> (2016)
10. Bitansky, N., Canetti, R.: On strong simulation and composable point obfuscation. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 520–537. Springer, Heidelberg (Aug 2010)
11. Bitansky, N., Garg, S., Lin, H., Pass, R., Telang, S.: Succinct randomized encodings and their applications. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC. pp. 439–448. ACM Press (Jun 2015)
12. Bitansky, N., Vaikuntanathan, V.: Indistinguishability obfuscation: from approximate to exact, <http://eprint.iacr.org/2015/704>
13. Bitansky, N., Vaikuntanathan, V.: Indistinguishability obfuscation from functional encryption. In: IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015. pp. 171–190. IEEE Computer Society (2015)
14. Boneh, D., Boyen, X.: On the impossibility of efficiently combining collision resistant hash functions. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 570–583. Springer, Heidelberg (Aug 2006)
15. Brakerski, Z., Rothblum, G.N.: Virtual black-box obfuscation for all circuits via generic graded encoding. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 1–25. Springer, Heidelberg (Feb 2014)
16. Canetti, R.: Towards realizing random oracles: Hash functions that hide all partial information. In: Kaliski Jr., B.S. (ed.) CRYPTO’97. LNCS, vol. 1294, pp. 455–469. Springer, Heidelberg (Aug 1997)
17. Canetti, R., Holmgren, J., Jain, A., Vaikuntanathan, V.: Succinct garbling and indistinguishability obfuscation for RAM programs. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC. pp. 429–437. ACM Press (Jun 2015)
18. Cheon, J.H., Han, K., Lee, C., Ryu, H., Stehlé, D.: Cryptanalysis of the multilinear map over the integers. In: EUROCRYPT 2015, Part I. LNCS, vol. 9056. Springer, Heidelberg (Apr 2015)

19. Collberg, C., Thomborson, C.: Watermarking, tamper-proofing, and obfuscation—tools for software protection. *IEEE TSE: IEEE Transactions on Software Engineering* 28 (2002)
20. Collberg, C., Thomborson, C., Low, D.: A taxonomy of obfuscating transformations. Technical Report #148, Department of Computer Science, The University of Auckland, New Zealand (1997)
21. Coron, J.S., Gentry, C., Halevi, S., Lepoint, T., Maji, H.K., Miles, E., Raykova, M., Sahai, A., Tibouchi, M.: Zeroizing without low-level zeroes: New MMAP attacks and their limitations. In: Gennaro, R., Robshaw, M.J.B. (eds.) *CRYPTO 2015, Part I. LNCS*, vol. 9215, pp. 247–266. Springer, Heidelberg (Aug 2015)
22. Coron, J.S., Lee, M., Lepoint, T., Tibouchi, M.: Cryptanalysis of GGH15 multilinear maps. *Cryptology ePrint Archive, Report 2015/1037* (2015), <http://eprint.iacr.org/2015/1037>
23. Coron, J.S., Lepoint, T., Tibouchi, M.: Cryptanalysis of two candidate fixes of multilinear maps over the integers. *Cryptology ePrint Archive, Report 2014/975* (2014), <http://eprint.iacr.org/2014/975>
24. Dodis, Y., Katz, J.: Chosen-ciphertext security of multiple encryption. In: Kilian, J. (ed.) *TCC 2005. LNCS*, vol. 3378, pp. 188–209. Springer, Heidelberg (Feb 2005)
25. Fischlin, M., Lehmann, A.: Security-amplifying combiners for collision-resistant hash functions. In: Menezes, A. (ed.) *CRYPTO 2007. LNCS*, vol. 4622, pp. 224–243. Springer, Heidelberg (Aug 2007)
26. Fischlin, M., Lehmann, A.: Multi-property preserving combiners for hash functions. In: Canetti, R. (ed.) *TCC 2008. LNCS*, vol. 4948, pp. 375–392. Springer, Heidelberg (Mar 2008)
27. Fischlin, M., Lehmann, A., Pietrzak, K.: Robust multi-property combiners for hash functions. *Journal of Cryptology* 27(3), 397–428 (Jul 2014)
28. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th *FOCS*. pp. 40–49. IEEE Computer Society Press (Oct 2013)
29. Garg, S., Gentry, C., Halevi, S., Wichs, D.: On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In: Garay, J.A., Gennaro, R. (eds.) *CRYPTO 2014, Part I. LNCS*, vol. 8616, pp. 518–535. Springer, Heidelberg (Aug 2014)
30. Garg, S., Mukherjee, P., Srinivasan, A.: Obfuscation without the vulnerabilities of multilinear maps. *Cryptology ePrint Archive, Report 2016/390* (2016), <http://eprint.iacr.org/2016/390>
31. Gentry, C., Halevi, S., Maji, H.K., Sahai, A.: Zeroizing without zeroes: Cryptanalyzing multilinear maps without encodings of zero. *Cryptology ePrint Archive, Report 2014/929* (2014), <http://eprint.iacr.org/2014/929>
32. Gentry, C., Lewko, A., Sahai, A., Waters, B.: Indistinguishability obfuscation from the multilinear subgroup elimination assumption (2015)
33. Goldwasser, S., Kalai, Y.T.: On the impossibility of obfuscation with auxiliary input. In: 46th *FOCS*. pp. 553–562. IEEE Computer Society Press (Oct 2005)
34. Harnik, D., Kilian, J., Naor, M., Reingold, O., Rosen, A.: On robust combiners for oblivious transfer and other primitives. In: Cramer, R. (ed.) *EUROCRYPT 2005. LNCS*, vol. 3494, pp. 96–113. Springer, Heidelberg (May 2005)
35. Herzberg, A.: Folklore, practice and theory of robust combiners. *Cryptology ePrint Archive, Report 2002/135* (2002), <http://eprint.iacr.org/2002/135>
36. Herzberg, A.: On tolerant cryptographic constructions. In: Menezes, A. (ed.) *CT-RSA 2005. LNCS*, vol. 3376, pp. 172–190. Springer, Heidelberg (Feb 2005)

37. Herzberg, A.: Folklore, practice and theory of robust combiners. *Journal of Computer Security* 17(2), 159–189 (2009)
38. Herzberg, A., Shulman, H.: Robust combiners for software hardening. In: *Trust and Trustworthy Computing, Third International Conference, TRUST 2010*, Berlin, Germany, June 21–23, 2010. Proceedings. *Lecture Notes in Computer Science*, vol. 6101, pp. 282–289. Springer (2010)
39. Hofheinz, D., Jager, T., Khurana, D., Sahai, A., Waters, B., Zhandry, M.: How to generate and use universal samplers. *Cryptology ePrint Archive*, Report 2014/507 (2014), <http://eprint.iacr.org/2014/507>
40. Koppula, V., Lewko, A.B., Waters, B.: Indistinguishability obfuscation for turing machines with unbounded memory. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC. pp. 419–428. ACM Press (Jun 2015)
41. Lin, H.: Indistinguishability obfuscation from constant-degree graded encoding schemes. In: *Advances in Cryptology - EUROCRYPT 2016*. *Lecture Notes in Computer Science*, vol. 9665, pp. 28–57. Springer (2016)
42. McCabe, T.J.: A complexity measure. *IEEE Trans. Software Eng.* 2(4), 308–320 (1976)
43. Meier, R., Przydatek, B.: On robust combiners for private information retrieval and other primitives. In: Dwork, C. (ed.) *CRYPTO 2006*. LNCS, vol. 4117, pp. 555–569. Springer, Heidelberg (Aug 2006)
44. Meier, R., Przydatek, B., Wullschlegel, J.: Robuster combiners for oblivious transfer. In: Vadhan, S.P. (ed.) *TCC 2007*. LNCS, vol. 4392, pp. 404–418. Springer, Heidelberg (Feb 2007)
45. Mennink, B., Preneel, B.: Breaking and fixing cryptophia’s short combiner. In: Gritzalis, D., Kiayias, A., Askoxylakis, I.G. (eds.) *CANS 14*. LNCS, vol. 8813, pp. 50–63. Springer, Heidelberg (Oct 2014)
46. Miles, E., Sahai, A., Zhandry, M.: Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. *Cryptology ePrint Archive*, Report 2016/147, <http://eprint.iacr.org/2016/147>
47. Mittelbach, A.: Cryptophia’s short combiner for collision-resistant hash functions. In: Jacobson Jr., M.J., Locasto, M.E., Mohassel, P., Safavi-Naini, R. (eds.) *ACNS 13*. LNCS, vol. 7954, pp. 136–153. Springer, Heidelberg (Jun 2013)
48. Pass, R., Seth, K., Telang, S.: Indistinguishability obfuscation from semantically-secure multilinear encodings. In: Garay, J.A., Gennaro, R. (eds.) *CRYPTO 2014, Part I*. LNCS, vol. 8616, pp. 500–517. Springer, Heidelberg (Aug 2014)
49. Pietrzak, K.: Non-trivial black-box combiners for collision-resistant hash-functions don’t exist. In: Naor, M. (ed.) *EUROCRYPT 2007*. LNCS, vol. 4515, pp. 23–33. Springer, Heidelberg (May 2007)
50. Pietrzak, K.: Compression from collisions, or why CRHF combiners have a long output. In: Wagner, D. (ed.) *CRYPTO 2008*. LNCS, vol. 5157, pp. 413–432. Springer, Heidelberg (Aug 2008)