Limits on the Power of Garbling Techniques
for Public-Key Encryption

Sanjam Garg[*]   Mohammad Hajiabadi[†]
Mohammad Mahmoody[‡]   Ameer Mohammed[§]

Abstract

Understanding whether public-key encryption can be based on one-way functions is a fundamental open problem in cryptography. The seminal work of Impagliazzo and Rudich [STOC'89] shows that black-box constructions of public-key encryption from one-way functions are impossible. However, this impossibility result leaves open the possibility of using non-black-box techniques for achieving this goal.

One of the most powerful classes of non-black-box techniques, which can be based on one-way functions alone, is Yao's garbled circuit technique [FOCS'86]. As for the non-black-box power of this technique, the recent work of Döttling and Garg [CRYPTO'17] shows that the use of garbling allows us to circumvent known black-box barriers in the context of identity-based encryption.

We prove that garbling of circuits that have one-way function (or even random oracle) gates in them are insufficient for obtaining public-key encryption. Additionally, we show that this model also captures (non-interactive) zero-knowledge proofs for relations with one-way function gates. This indicates that currently known one-way function based non-black-box techniques are perhaps insufficient for realizing public-key encryption.

Contents

1 Introduction

Public-key encryption (PKE) [DH76, RSA78] is a fundamental primitive in cryptography and understanding what assumptions are sufficient for realizing it is a foundational goal. Decades of research have provided us with numerous constructions of PKE from a variety of assumptions; see a recent survey by Barak [Bar17]. However, all known constructions of PKE require computational assumptions that rely on rich structure and are stronger than what is necessary and sufficient for *private*-key cryptography, namely the mere existence of one-way functions (OWF). The seminal work of Impagliazzo and Rudich [IR89] provides some evidence that this gap between the assumption complexity of private-key and public-key encryption may be inherent. In particular, the work of [IR89] shows that there is no *black-box* construction of PKE from OWFs.[1]

When studying the *impossibility* of basing PKE on OWFs, focusing on a class of constructions (e.g., black-box constructions as in [IR89]) is indeed necessary. The reason is that to rule out "OWFs implying PKE" in a *logical* sense, we have to first prove the existence of OWFs unconditionally (thus, proving $\mathbf{P} \neq \mathbf{NP}$) and then rule out the existence of PKE altogether (thus breaking all assumptions under which PKE exists). That is why this line of separation results focuses on ruling out the possibility of using certain *techniques* or generic *proof methods* (here black-box proofs/techniques) as possible natural paths from OWFs to PKE.

Garbled circuits. Over the past few decades, garbling techniques [Yao86, LP09, BHR12, App17] (or randomized encodings [IK00] more generally) have been extensively used to build many cryptographic schemes. Roughly speaking, in a circuit garbling mechanism, a PPT encoder $\mathsf{Garb}(\mathsf{C})$ takes a circuit C as input, and outputs a *garbled circuit* $\widetilde{\mathsf{C}}$ and a set of *input labels* $\{\mathsf{label}_{i,b}\}_{i\in[m],b\in\{0,1\}}$ where m is the number of input wires of C. Using another algorithm $\mathsf{Eval}(\cdot)$, one can use the garbled circuit $\widetilde{\mathsf{C}}$ and input labels $\{\mathsf{label}_{i,x_i}\}_{i\in[m]}$ for an input $x = (x_1, ..., x_m)$, to compute $\mathsf{C}(x)$ without learning any other information. Note that if the original circuit C needs to run a cryptographic primitive f internally (e.g., a circuit C for a pseudorandom generator built from a OWF f), this use of garbling leads to a *non-black-box* construction. This is because the algorithm Garb needs to work with an actual circuit description of C, whose circuit description is in turn obtained by the circuit description of f, hence making non-black-box use of f.

Garbling, as a primitive, may itself be realized using one-way functions [Yao86, LP09]. This puts forward the intriguing possibility of basing PKE solely on OWFs by making *black-box* use of garbling mechanisms over circuits that can run the one-way function. As stated above, such constructions will make non-black-box use of the underlying OWF (caused by garbling circuits that run the OWF internally) and hence the impossibility result of Impagliazzo and Rudich [IR89] has no bearing on such potential constructions. In fact, such non-black-box garbling techniques, combined with the Computational Diffie-Hellman assumption, have recently been used by Döttling and Garg [DG17] to circumvent black-box impossibility results [BPR+08, PRV12] in the context of identity-based encryption (IBE). Thus, it is natural to ask:

Can non-black-box garbling techniques be used to realize PKE from OWFs?

Our model. We study the above question in the model of Brakerski, Katz, Segev and Yerukhimovich [BKSY11] (see also follow up works [AS15, AS16, BDV17]) which gives a general way of

---

[1]A (fully) black-box construction is one that treats the OWF as an oracle, and the security proof uses the OWF and the adversary both as oracles; see the surveys of [RTV04, BBF13] for formal definitions.

capturing non-black-box techniques via circuits with cryptographic gates (e.g., OWF gates). More formally, we will model the above-stated garbling-based non-black-box use of one-way functions as black-box use of garbling mechanisms that can take as input circuits C with one-way function (or even random oracle) gates planted in them. Such constructions are indeed *non-black-box* according to the taxonomy of [RTV04] if viewed as standalone constructions solely based on the OWF itself. We stress that the allowed access to the garbling mechanism itself is black-box; the non-black-box feature arises from the fact that circuits with OWF gates may now be garbled.

A more sophisticated scenario is when the circuits being garbled have *garbling* gates, in addition to OWF gates, planted in them. We do not, however, consider such a recursive scenario and we leave it to future work. It is crucial to note that, to the best of our knowledge, all known constructions that make use of garbling schemes together with one-way functions (e.g., [Bea96, LO13, GLOS15]) fall into our model, and thus, understanding the limitations of such techniques towards obtaining PKE is impactful.

1.1 Our Result

In this work, we show that black-box use of garbling mechanisms that allow circuits with OWF gates to be garbled is not sufficient for constructing PKE. More precisely, we prove the following.

Theorem 1.1 (Main result – informally stated). *There exists no black-box construction of public-key encryption schemes from any one-way function (or even a random oracle) f together with a garbling mechanism that can garble oracle-aided circuits with f-gates embedded in them.*

Comparison with prior work: Impossibility from weaker garbling. The work of Asharov and Segev [AS15] showed that secret-key functional encryption with one-way function gates cannot be used (as a black-box) to obtain public-key encryption (or even key agreement). This result implies that for the special case of such weaker garbling schemes, called *non-decomposable* garbling, where the entire input is considered as a single unit (rather that as bit-by-bit input labels) is insufficient for realizing PKE.

On the other hand, throughout this work we use garbling to refer to a notion that supports bit-by-bit input labels, a notion that Bellare, Hoang and Rogaway [BHR12] refer to as *projective garbling* (a.k.a. decomposable garbling). Under projective garbling, for a circuit C of input size m, one generates two garbled label $\{\mathsf{label}_{i,b}\}_{i\in[m],b\in\{0,1\}}$ for the ith input wire of the circuit. An important property enabled by this bit-by-bit garbling is the decomposability property: one can pick a garbled label for each input wire to form a garbled input for a long string. In contrast, under non-decomposable garbling, for each input X to the circuit, one independently generates a corresponding garbled input $\widetilde{X}$. As a result, different strings have independent garbled inputs.

We note that projective garbling is crucial for many applications of garbling. For example, even the most basic application of garbling in two party secure computation based on oblivious transfer uses the projective property. We refer the reader to [BHR12, Figure 3] for a detailed list of applications that require projective garbling. As a recent example, we note that the IBE construction of Döttling and Garg [DG17] (that circumvents a black-box impossibility result of Papakonstantinou et al. [PRV12] using garbling) uses projective garbling crucially. Specifically, in [DG17] the encryptor provides a sequence of garbled circuits with no knowledge of what input each of those garbled circuits are later evaluated on by the decryptor. This *input-obliviousness* property is enabled by the encryptor sending all the bit-by-bit garbled labels in some encrypted

form to the decryptor. Later, the decryptor can open exactly one garbled label for each input wire, hence obtaining a garbled input for the whole string. This input-obliviousness technique cannot be enabled using non-decomposable garbling. This is because a whole garbled input cannot be formed there by putting together smaller pieces. As a result, the party who generates a garbled circuit must be aware of the input on which this garbled circuit is to be evaluated on, in order for him to be able to provide the corresponding garbled input.

1.2 Extensions

Extension to key agreements. Our proof extends to rule out any black-box construction of *constant-round* key-agreement protocols from OWFs and garbling schemes for oracle-aided circuits. However, the proof of the separation for key-agreement beyond the case of two message protocols (which are equivalent to PKE) becomes much more involved. Therefore, for clarity of the presentation, and because the most interesting special case of constant-round key-agreement protocols happens to be PKE itself, in this presentation, we focus on the case of separation for PKE. See Section A for more details.

Resolving an open question of [BKSY11]. The work of [BKSY11] proved non-black-box limitations for one-way functions when used as part of zero knowledge (ZK) proofs for relations with one-way function gates. They showed that key-agreement protocols with *perfect* completeness cannot be realized in a 'black-box' way from oracles that provide a one-way function f together with ZK proofs of satisfiability for f-aided circuits. They left ruling out the possibility of protocols with imperfect (e.g., negligible) completeness as an open problem, as their techniques indeed crucially relied on the perfect completeness assumption. We demonstrate the power of our new techniques in this work by resolving the open problem of [BKSY11] along the way, for the case of PKE schemes (or even constant-round key-agreement schemes). In particular, in Section B, we observe that the oracles we use for proving our separations for the case of garbling, indeed imply the existence of NIZK proofs for satisfiability of circuits with OWF-gates. The extension of the result of [BKSY11] explained above then follows from the above observation.

1.3 Related Work and Future Directions

There are quite a few results that prove limitations for a *broad* class of non-black-box techniques [Pas11, PTV11, GW11], so long as the security reduction is black-box. In other words, these results are proved against basing certain primitives on any *falsifiable* assumption. However, when it comes to the case of non-black-box constructions of PKE from OWFs, no such general separations are known (and proving such results might in fact be impossible).

As described earlier, the works of [BKSY11, AS15] proved limitations of certain non-black-box constructions of PKE from OWFs. This is indeed the direction pursued in this work. The work of Dachman-Soled [Dac16] takes yet another path, showing that certain non-black-box uses of one-way functions *in the security proof* are incapable of obtaining PKE from OWFs.

We note that we only consider a setting in which circuits with random oracle gates are garbled. We do not allow garbling of circuits which themselves include garbling gates. Such techniques are captured by the so called monolithic model of Garg, Mahmoody, and Mohammed [GMM17a, GMM17b]. We leave open the problem of ruling out such constructions.

Finally, as noted above, the extension of our results to the key-agreement setting, discussed in Section A only cover the *constant-round* case. The reason is that, during the proof of our main result, we modify the protocol iteratively, once for each round, which increases the parameters of the protocol by a polynomial factor each time. We leave the extension to general polynomial-round protocols as an interesting future direction.

Organization. In Section 2 we give an overview of our approach and techniques. In Section 3 we give some definitions and basic lemmas. In Section 4 we will prove our main impossibility result. In Sections A and B we will describe extensions of our results.

2 Technical Overview

For brevity, we refer to the primitive of a one-way function f and garbling circuits with f gates as GC-OWF. As usual in black-box separation results, we will prove our main theorem by providing an oracle O relative to which secure GC-OWF exists, but secure PKE does not.

2.1 Big Picture: Reducing the Problem to the Result of [IR89]

At a very high level, our approach is to reduce our problem to the result of [IR89]. Namely, we aim to show that one can always modify the PKE construction that is based on the GC-OWF oracle O into a new one that is almost as secure, but which no longer uses the garbling part of the oracle O. In other words, we modify the construction so that it becomes a construction from an OWF oracle alone. Our main result, then, follows from the impossibility result of [IR89] which rules out the possibility of getting PKE from one-way (or even random) functions. We call this process 'compiling out the garbling part' from the PKE construction.

As a technical remark, our transformation does not result in a normal black-box construction of PKE from OWFs, but rather results in an *inefficient* one which nonetheless makes a polynomial number of queries to the OWF oracle. The key point is that the *proof* of the work of Impagliazzo and Rudich [IR89] allows us to break any such (even inefficient, but still polynomial-query) constructions of PKE in the random oracle model using a polynomial number of queries during the attack. Our actual result follows by combining our compilation result with the result of [IR89], to get a polynomial query attack against the security of the original PKE. This will be sufficient for a black-box separation.

At a high level, our approach also bears similarities to recent impossibility results for indistinguishability obfuscation [GMM17a] as we also compile out the more powerful (and structured) parts of the oracle, ending up with a scheme that uses a much simpler oracle, for which an impossibility is known. However, there is a subtle distinction here. Unlike the results of [CKP15, MMN+16b, MMN16a, GMM17a], when we compile out the garbling-related queries from the PKE construction, we end up with an inefficient scheme that potentially runs in exponential time but nevertheless makes a polynomial number of queries. However, as mentioned above, this is fine for deriving our separation, because we can still rely on the fact that the result of [IR89] does something stronger and handles inefficient constructions as well.

2.2 Our Separating Idealized GC-OWF Oracle

In this subsection, we will first describe our oracle O that gives an intuitive way of obtaining GC-OWFs. The natural first version of this oracle is too strong as it also implies virtual black-box (VBB) obfuscation. We will then add a careful weakening subroutine to this oracle O to prevent it from implying obfuscation. In the next subsection we describe the ideas behind how to compile out the garbling-related subroutines of O from the PKE construction, while keeping the PKE construction "secure".

Our 1st oracle for GC-OWF. Our first version of the separating oracle $O = (f, L^f)$ will consist of a random oracle f (giving the OWF part) as well a garbling part $L^f = (\mathsf{gc}, \mathsf{eval}^f)$ with two subroutines. The encoding/garbling subroutine $\mathsf{gc}(s, \mathsf{C})$ is simply a random (injective) function that takes a seed s and a circuit C and maps them into a garbled circuit $\widetilde{\mathsf{C}}$ as well as labels $\{\mathsf{label}_{i,b}\}_{i \in [n], b \in \{0,1\}}$ for the input wires of C where n is the number of input wires in C.[2] The evaluation subroutine eval^f takes as input a garbled circuit $\widetilde{\mathsf{C}}$ as well as a vector of input labels $\widetilde{X} = (\widetilde{x}_1 \cdots \widetilde{x}_n)$ and *only if* they were all encoded using the *same* seed s, eval^f returns the right output $\mathsf{C}^f(x_1, \ldots, x_n)$. Note that we include f in the representation of eval^f but not in that of gc; the reason is gc is simply a random oracle (independent of f), while eval^f needs to call f in order to compute $\mathsf{C}^f(x_1, \ldots, x_n)$.

Adding the weakening subroutine rev. It is easy to see that this first version of the oracle O as described above can realize a secure GC-OWF, but it can do much more than that! In fact this oracle implies even VBB obfuscation of circuits with f gates (which in turn *does* imply PKE [SW14]). We will, therefore, weaken the power of the oracle O later on by adding an extra subroutine to it (which we will call rev), which roughly speaking allows an attacker to break the garbling scheme if she has access to two labels for the same wire. We will describe this subroutine after it becomes clear how it will be useful for our main goal of compiling out the garbling aspect of O. Also note that, since we are defining our oracle *after* the (supposed) construction of PKE from GC-OWF is fixed, without loss of generality, the PKE construction from GC-OWF does *not* call the extra subroutine rev. This separation technique was also used before in the work of [GMR01] and is reminiscent of the "two-oracle" approach that was first formalized in [HR04].

2.3 Compiling Out the Garbling Power of O from the Construction

Suppose $\mathcal{E}^O = (G^O, E^O, D^O)$ is a fully black-box construction of PKE using the oracle O described above. Our goal here is to 'reduce' our problem (of breaking $\mathcal{E}^O$ using a polynomial number of queries) to the result of [IR89] by compiling out the 'garbling power' of the oracle O from the scheme $\mathcal{E}^O$. But what subroutines do we have to compile out from $\mathcal{E}$? As it turns out, we do not have to eliminate both gc and eval^f subroutines; removing only eval^f queries will suffice.

Compiling out eval^f queries from PKE constructions $\mathcal{E}^O$. If we make sure that (the modified but "equally"-secure version of) $\mathcal{E}$ does not make any calls to the eval^f subroutine of the oracle O, it would be sufficient for our purposes, because the oracle $O' = (f, \mathsf{gc})$ is just a random oracle,

---

[2]In the main body, we will use two separate subroutines gc, gi for encoding circuits vs input labels, but for brevity here we combine them into one subroutine.

and by the result of [IR89] we know that such oracle is not enough for getting PKE in a black-box way. Therefore, in what follows, our goal reduces to (solely) removing the eval^f queries from PKE constructions $\mathcal{E}^O$ in a way that we can argue the new construction is 'as secure' as the original one.

In order to make our proof more modular, we compile out eval^f queries from the different components (i.e., key-generation G, encryption E, and decryption D) of the construction $\mathcal{E}^O = (G^O, E^O, D^O)$ one at a time. First, we may easily see that G^O does not need to call eval^f at all. This is because G^O is the first algorithm to get executed in the system, and so G^O knows all the generated garbled circuits/labels. Therefore, G^O can, instead of calling eval^f queries, simply run $\mathsf{C}^f(X)$ on its own by further calls to f.[3] Now, we proceed to compile out eval^f queries from the remaining two subroutines E and D in two steps. In each step, we assume that we are starting off with a construction that has no eval^f queries in some of its subroutines, and then we modify the construction to remove eval^f queries from the next subroutine.

- **Step 1:** Starting with $\mathcal{E} = (G^{f,\mathsf{gc}}, E^O, D^O)$, we will compile $\mathcal{E}$ into a new scheme $\dot{\mathcal{E}} = (\dot{G}^{f,\mathsf{gc}}, \dot{E}^{f,\mathsf{gc}}, D^O)$, removing eval^f queries asked by E^O. We have to make sure $\dot{\mathcal{E}}$ is 'almost as secure' as the original scheme $\mathcal{E}$. This step is detailed in Section 4.1.

- **Step 2:** Given $\mathcal{E} = (G^{f,\mathsf{gc}}, E^{f,\mathsf{gc}}, D^O)$, we compile $\mathcal{E}$ into a new $\ddot{\mathcal{E}} = (\ddot{G}^{f,\mathsf{gc}}, \ddot{E}^{f,\mathsf{gc}}, \ddot{D}^{f,\mathsf{gc}})$, removing eval^f queries asked by D^O. Again, we have to make sure $\ddot{\mathcal{E}}$ is 'almost as secure' as the original one. This step is detailed in Section 4.2.

Once we accomplish both of the steps above, we will combine them into a single compiler that removes eval^f queries from $\mathcal{E}^O$, obtaining another PKE construction that is secure in the random oracle model (which we already know is impossible by the result of [IR89]).

Overview of Step 1. Let us start by looking at eval queries of the encryption algorithm $E^O(pk, b)$. Since the subroutine gc of oracle O is just a random mapping, for any eval query on inputs $(\widetilde{\mathsf{C}}, \widetilde{X})$, denoted $qu = ((\widetilde{\mathsf{C}}, \widetilde{X}) \underset{\mathsf{eval}}{\longrightarrow} ?)$, made by E^O and whose answer is not trivially $\bot$, we must have either of the following cases. Either (a) $\widetilde{\mathsf{C}}$ was produced as a result of a gc query during the execution of E itself or (b) $\widetilde{\mathsf{C}}$ was produced during the execution of G which has led to the generation of the public key pk. If case (a) holds, then E does not need to make that particular eval query at all. If case (b) holds, then in order to allow $\dot{E}^{f,\mathsf{gc}}$ to simulate E^O without calling eval^f, the algorithm $\dot{E}^{f,\mathsf{gc}}$ will resort to some 'hint list' H attached to pk by $\dot{G}^{f,\mathsf{gc}}$. That is, a compiled public key $\dot{pk}$ produced by $\dot{G}^{f,\mathsf{gc}}$ will now contain the original pk as well as a hint list H. Below, we further explain how the hint H is formed.

How $\dot{G}^{f,\mathsf{gc}}$ forms the hint list H. A naive idea is to let H contain all the query/response pairs made by $G^{f,\mathsf{gc}}$ to generate pk. This method hurts security. A better idea is to provide in H answers to individual eval queries like $\mathsf{eval}(\widetilde{\mathsf{C}}, \widetilde{X})$ that are *likely* to be asked by $E^O(pk, b)$, and where $\widetilde{\mathsf{C}}$ was generated by $G^{f,\mathsf{gc}}$. That is, $\dot{G}^{f,\mathsf{gc}}$ would run $E^O(pk, b)$ many times and would let H contain all encountered eval queries as well as their answers. Note that $\dot{G}^{f,\mathsf{gc}}$ could simulate almost perfectly a random execution of $E^O(pk, b)$ without calling eval since $\dot{G}$ knows the randomness seeds of all the garbled circuits so far. However, this approach also fails! To see the difficulty, recall that a whole garbled input $\widetilde{X} = (\widetilde{x}_1, \ldots, \widetilde{x}_m)$ is made up of a sequence of garbled labels $\widetilde{x}_i$, one for each

---

[3]More formally, because of the huge output space of gc, calling eval^f on a garbled circuit $\widetilde{\mathsf{C}}$ that is produced on the fly is bound to be responded with $\bot$ with overwhelming probability as $\widetilde{\mathsf{C}}$ will not be an encoding of any circuit.

input wire. Now imagine that for a garbled circuit $\widetilde{\mathsf{C}}$ that was generated by $\dot{G}$, any new execution of $E^O(pk, b)$ calls the oracle eval on $\widetilde{\mathsf{C}}$ and on a new garbled input $\widetilde{X}$ that is formed by picking each of the garbled labels uniformly at random from the set of two lables for that corresponding input wire. If E behaves this way, then no matter how many (polynomial) times we sample from $E^O(pk, b)$, we cannot hope to predict the garbled-input part of the eval query of the next execution of $E^O(pk, b)$. We refer to this as the *garbled-input unperdictability* problem, which stems from the decomposability nature of our garbling oracle. This is what makes our results different from those of [AS15], which dealt with non-decomposable garbling, for which such a complication is absent.

In short, we could only hope to predict the garbled circuit part of an eval query of $E^O(pk, b)$, and not necessarily the garbled-input part. To fix this garbled-input unpredictability problem, $\dot{G}^{f,\mathsf{gc}}$ will do the following trick: while sampling many executions of $E^O(pk, b)$, if $\dot{G}^{f,\mathsf{gc}}$ comes across two different eval queries $\mathsf{eval}(\widetilde{\mathsf{C}}, \widetilde{X_1}), \mathsf{eval}(\widetilde{\mathsf{C}}, \widetilde{X_2})$ that are both answered with a value that is not $\bot$ (i.e., both are valid garbled circuits and inputs), then $\dot{G}^{f,\mathsf{gc}}$ releases the corresponding seed s and the *plain* circuit C of $\widetilde{\mathsf{C}}$. That is, if $\mathsf{gc}(s, \mathsf{C}) = (\widetilde{\mathsf{C}}, \cdots)$, then $\widetilde{G}^{f,\mathsf{gc}}$ puts the tuple $(s, \mathsf{C}, \widetilde{\mathsf{C}})$ into the hint list H. If, however, during these sampling, $\widetilde{\mathsf{C}}$ is evaluated upon *at most* one matching $\widetilde{X}$, then $\dot{G}^{f,\mathsf{gc}}$ simply provides the answer to the query $\mathsf{eval}(\widetilde{\mathsf{C}}, \widetilde{X})$ in H.

Looking ahead, the algorithm $\dot{E}^{f,\mathsf{gc}}((pk, \mathsf{H}), b)$, when facing an eval query $qu = \mathsf{eval}(\widetilde{\mathsf{C}}, \widetilde{X})$, will check whether qu is already answered in the list H, or whether the corresponding seed s and plain-circuit C of $\widetilde{\mathsf{C}}$ could be retrieved from H. If so, $\dot{E}^{f,\mathsf{gc}}$ will reply to qu accordingly; otherwise, it will reply to qu with $\bot$.

Using the weakening subroutine rev to reduce the security of $\dot{\mathcal{E}}$ to $\mathcal{E}$. Note that $\dot{G}^{f,\mathsf{gc}}$ does not query any oracle subroutines beyond f and gc in order to form the hint list H attached to pk. This is because $\dot{G}^{f,\mathsf{gc}}$ has all the (otherwise-hidden) query-answer pairs used to produce pk, and thus for any encountered valid garbled circuit $\widetilde{\mathsf{C}}$ during those sampled executions of E, $\dot{G}^{f,\mathsf{gc}}$ already knows the corresponding seed s and plain circuit C. Now we are left to show that this additional information H attached to pk does not degrade the security of the compiled scheme significantly. To this end, we will use the new weakening oracle intended to capture the natural use of garbling: the security of a garbled circuit $\widetilde{\mathsf{C}}$ is guaranteed to hold so long as $\widetilde{\mathsf{C}}$ is evaluated only on one garbled input. Capturing this, our new oracle rev takes as input a garbled circuit $\widetilde{\mathsf{C}}$ and two garbled inputs $\widetilde{X_1}$ and $\widetilde{X_2}$, and if all of $\widetilde{\mathsf{C}}$, $\widetilde{X_1}$ and $\widetilde{X_2}$ are encoded using the same seed s, then rev simply outputs (s, C), where $\mathsf{gc}(s, \mathsf{C}) = \widetilde{\mathsf{C}}$. For security, we will show that any adversary against the semantic security of $\dot{\mathcal{E}}$ may be used in a black-box way, along with oracle access to $(f, \mathsf{gc}, \mathsf{eval}, \mathsf{rev})$, to mount an attack against the original scheme $\mathcal{E}$. This shows that the leakage caused by revealing H was also attainable in the original scheme (in which all parties including the attacker do have access to eval) if, in addition, access to the oracle rev — which reflects the intuitive way in which garbled circuits are supposed to be used — was also granted to the adversary. In our security proof we will crucially make use of the rev subroutine in order to construct tuples $(s, \mathsf{C}, \widetilde{\mathsf{C}})$ to store in the simulated hint list whenever $\widetilde{\mathsf{C}}$ should be evaluated on two different inputs. Tuples of the form $(\widetilde{\mathsf{C}}, \widetilde{X}, y)$ can in turn be simulated using oracle access to eval.

Overview of Step 2. The main idea is similar to Step (1): $\ddot{E}^{f,\mathsf{gc}}(pk, b)$ would first run $E^{f,\mathsf{gc}}(pk, b)$ to get the ciphertext c and then appropriately attach a hint H to c. The idea is that H should allow the eval-free algorithm $\ddot{D}^{f,\mathsf{gc}}((sk, \mathsf{H}), c)$ to simulate $D^{f,\mathsf{gc},\mathsf{eval}}(sk, c)$ well enough. Again, since we cannot simply copy the entire private view of $\ddot{E}^{f,\mathsf{gc}}(pk, b)$ into H (as that cannot be simulated by the

security reduction, and therefore would hurt security) we should instead ensure that w.h.p. all eval queries during the execution of $D^O(sk, c)$, whose garbled circuits were generated by $\ddot{E}^{f,\text{gc}}(pk, b)$, can be answered using H. Let us call these eval queries $\ddot{E}$-*tied* queries. Unfortunately, when implementing this idea, we run into the following problem: $\ddot{E}^{f,\text{gc}}(pk, b)$ cannot simply run $D^O(sk, c)$ to get a sense of eval queries because sk is private. This challenge was absent in Step (1).

In order to resolve this new challenge, the algorithm $\ddot{E}^{f,\text{gc}}(pk, b)$ needs to do some more offline work in order to get an idea of $\ddot{E}$-tied eval queries that come up during $D^O(sk, c)$. The main idea is that although the true secret key sk is unknown to $\ddot{E}^{f,\text{gc}}(pk, b)$, in the eyes of $\ddot{E}^{f,\text{gc}}(pk, b)$, the value of sk is equally likely to be any sk' that agrees with the entire view of $\ddot{E}^{f,\text{gc}}(pk, b)$. Put differently, the probability that an $\ddot{E}$-tied garbled circuit comes up during $D^O(sk, c)$ is close to the probability that it comes up during the execution of $D^{O'}(sk', c)$, where O' is an offline oracle that agrees with all the private information of $\ddot{E}$, and also relative to which (pk, sk') is valid public-key/secret-key. As a result, such a fake sk' that is consistent with the view of $\ddot{E}^{f,\text{gc}}(pk, b)$ will be used to learn the answers of the evaluation queries asked by $D^{O'}(sk', c)$[4].

This is the main idea behind the second-step compilation. The full proof needs to take care of many subtle challenges, which we will defer to the main body.

Putting things together. Taken together, Steps (1) and (2) in conjunction with the result of Imagliazzo and Rudich [IR89] imply the following.

Lemma 2.1 (Informal). *The (claimed) semantic security of any candidate PKE construction $\mathcal{E}^{f,\text{gc},\text{eval}}$ can be broken by a poly-query adversary $\mathcal{A}^{f,\text{gc},\text{eval},\text{rev}}$.*

Moreover, we can show that the oracle rev does not break the one-wayness/garbling aspects of $(f, \text{gc}, \text{eval})$.

Lemma 2.2 (Informal). *The function f is one way against all poly-query adversaries with oracle access $(f, \text{gc}, \text{eval}, \text{rev})$. Moreover, there exists a garbling scheme $L^{f,\text{gc},\text{eval}}$ for garbling circuits with f gates that remains secure against all poly-query adversaries $\mathcal{B}^{f,\text{gc},\text{eval},\text{rev}}$.*

Now Lemmas 2.1 and 2.2 imply our main theorem, Theorem 1.1.

3 Preliminaries

We use κ for the security parameter. By PPT we mean a probabilistic polynomial time algorithm. By an *oracle* PPT/algorithm we mean a PPT that may make oracle calls. For any oracle algorithm A that has access to some oracle O, we denote a query qu asked by A to a subroutine T of O as $(qu \underset{T}{\to} ?)$. If the returned answer is β, then we denote the resulting query-answer pair as $(qu \underset{T}{\to} \beta)$. For a set S of query/answer pairs, we will use intuitive notation such as $(* \underset{T}{\to} \beta) \in$ S to mean that there exists a query qu such that $(qu \underset{T}{\to} \beta) \in$ S. We use $||$ to concatenate strings and we use "," for attaching strings in a way they could be retrieved. Namely, one can uniquely identify x and y from (x, y). For example $(00||11) = (0011)$, but $(0, 011) \neq (001, 1)$. For any given string x, we denote x_i to be the i'th string of x. For (family of) random variables $\{X_\kappa, Y_\kappa\}_\kappa$, by $X \overset{c}{\approx} Y$ we denote that they are computationally indistinguishable; namely, for any poly(κ)-time adversary A there is

---

[4]Note that the process of discovering such an sk' is what makes $\dot{E}$ an inefficient algorithm.

a negligible function $\mathrm{negl}(\kappa)$ such that $|\Pr[A(X_\kappa) = 1] - \Pr[A(Y_\kappa) = 1]| \leq \mathrm{negl}(\kappa)$. When writing the probabilities, by putting an algorithm A in the subscript of the probability (e.g., $\Pr_A[\cdot]$) we emphasize that the probability is over A's randomness. For any given probability distribution $\mathbf{D}$, we denote $x \leftarrow \mathbf{D}$ as sampling from this distribution and obtaining a sample x from the support of $\mathbf{D}$. We may also use $x \in \mathbf{D}$ to mean that x is in the support of $\mathbf{D}$. For any two random variables X, Y, we denote $\Delta(X, Y)$ to be the statistical distance between the two random variables. Throughout the paper, whenever we write $f_1(\kappa) \leq f_2(\kappa)$ we mean that this inequality holds asymptotically; i.e., there exists κ_0 such that for all $\kappa \geq \kappa_0$, $f_1(\kappa) \leq f_2(\kappa)$.

3.1 Some Useful Lemmas

The following lemma shows that hitting the image of a sparse injective random function without having called the function on the corresponding preimage happens with negligible probability.

Lemma 3.1 (Hitting the image of random injective function)**.** *Let A be an arbitrary polynomial-query algorithm with access to an oracle $O : \{0,1\}^\kappa \to \{0,1\}^{2\kappa}$ chosen uniformly at random from the set of all injective functions from $\{0,1\}^\kappa$ to $\{0,1\}^{2\kappa}$. We have*

$$\Pr[y \leftarrow A^O(1^\kappa) \mid \text{ for some } x \colon y = O(x) \wedge (* \xrightarrow[O]{} y) \notin \mathsf{Q_A}] \leq 2^{-\kappa/2},$$

where the probability is taken over the random choice of O as well as A's random coins, and where $\mathsf{Q_A}$ is the set of all A's query-answer pairs.

We will also use the following standard information theoretic lemma frequently in the paper.

Lemma 3.2. *Let $X_1, \ldots, X_{t+1}$ be independent, Bernoulli random variables, where $Pr[X_i = 1] = p$, for all $i \leq t+1$. Then*

$$\Pr[X_1 = 0 \wedge \cdots \wedge X_t = 0 \wedge X_{t+1} = 1] \leq \frac{1}{t}.$$

3.2 Public-Key Encryption

Definition 3.3. A *public-key single-bit encryption* (PKE) scheme is a triple of PPT algorithms (G, E, D), defined as follows:

- $G(1^\kappa)$ takes as input a security parameter and outputs a pair (pk, sk) of public/secret keys.

- $E(pk, b)$ takes as input a public key pk and a bit $b \in \{0,1\}$ and outputs a ciphertext c.

- $D(sk, c)$ takes as input a secret key sk and a ciphertext c and deterministically outputs a value $b' \in \{0,1\} \cup \{\bot\}$.

We define the following notions.

- **Correctness.** For a $\delta = \delta(\kappa)$ we say that (G, E, D) is $(\frac{1}{2} + \delta)$-correct if for all $b \in \{0,1\}$:

$$\Pr[D(sk, E(pk, b)) = b] \geq \frac{1}{2} + \delta,$$

where the probability is taken over the choice of $(pk, sk) \leftarrow G(1^\kappa)$ and the randomness of E.

- **Semantic security.** For $\gamma = \gamma(\kappa)$ we say that an adversary $\mathcal{A}$ γ-breaks (G, E, D) if

$$\Pr[A(1^\kappa, pk, c) = b] \geq \gamma,$$

where $(pk, sk) \leftarrow G(1^\kappa)$, $b \stackrel{\$}{\leftarrow} \{0, 1\}$ and $c \leftarrow E(pk, b)$. We say that $\mathcal{A}$ breaks the semantic security of (G, E, D) if $\mathcal{A}$ $(\frac{1}{2} + p)$-breaks the semantic security of the scheme for some polynomial p. We say that (G, E, D) is *semantically secure* if there does not exist a PPT adversary $\mathcal{A}$ that breaks the semantic security of (G, E, D).

3.3 Garbled Circuits

In this section we define the notion of garbling for oracle-aided circuits. First, we start by defining oracle-aided circuits.

Oracle aided circuits. A binary-output oracle-aided circuit C is a circuit with Boolean gates as well as oracle gates, and where the output of the circuit is a single bit. The input size, $\mathsf{inpsize}(\mathsf{C})$, is the number of input wires. The circuit size, denoted $|\mathsf{C}|$, denotes the number of gates and input wires of the circuit. For a fixed function f we write C^f to denote the circuit C when the underlying oracle is fixed to f.

Definition 3.4 (Garbling schemes for oracle-aided circuits). Fix a function f. A circuit garbling scheme for oracle-aided circuits relative to f (or with f gates) is a triple of algorithms $(\mathsf{Garb}, \mathsf{Eval}, \mathsf{Sim})$ defined as follows:

- $\mathsf{Garb}\,(1^\kappa, \mathsf{C})$: takes as input a security parameter κ, an oracle-aided circuit C and outputs a garbled circuit $\widetilde{\mathsf{C}}$ with a set of labels $\{\mathsf{label}_{i,b}\}_{i \in [m], b \in \{0,1\}}$, where $m = \mathsf{inpsize}(\mathsf{C})$.

- $\mathsf{Eval}^f\left(\widetilde{\mathsf{C}}, \{\mathsf{label}_{i,b_i}\}_{i \in [m]}\right)$: takes as input a garbled circuit $\widetilde{\mathsf{C}}$ and a garbled input as a sequence of input labels $\{\mathsf{label}_{i,b_i}\}_{i \in [m]}$, and outputs $y \in \{0,1\}^* \cup \{\bot\}$.

We define the following notions.

- **Correctness.** For any oracle-aided circuit C and input $x \in \{0,1\}^m$, where $m = \mathsf{inpsize}(\mathsf{C})$:

$$\Pr\left[\mathsf{C}^f(x) = \mathsf{Eval}^f\left(\widetilde{\mathsf{C}}, \{\mathsf{label}_{i,x_i}\}_{i \in [m]}\right)\right] = 1$$

where the probability is taken over $\mathsf{Garb}\,(1^\kappa, \mathsf{C}) \mapsto (\widetilde{\mathsf{C}}, \{\mathsf{label}_{i,b}\}_{i \in [m], b \in \{0,1\}})$.

- **Security.** For any polynomial $m = m(\kappa)$, any poly-size oracle circuit C with input size m, and any input $x \in \{0,1\}^m$:

$$\left(\widetilde{\mathsf{C}}, \{\mathsf{label}_{i,x_i}\}_{i \in [m]}\right) \stackrel{c}{\approx} \mathsf{Sim}\left(1^{|\mathsf{C}|}, m, \mathsf{C}^f(x)\right)$$

where $(\widetilde{\mathsf{C}}, \{\mathsf{label}_{i,b}\}_{i \in [m], b \in \{0,1\}}) \leftarrow \mathsf{Garb}\,(1^\kappa, \mathsf{C})$.

Remark 3.5. As mentioned before, we will consider projective garbling in our work (as defined above and in [BHR12]) where we are allowed to produce bit-by-bit garblings of inputs as opposed to only being allowed to garble the entire input, as was considered in the model of [AS15].

3.4 Black-box Constructions

Now, we recall the standard notion of black-box constructions [IR89, RTV04, BBF13]. We do so in the context of building PKE from one-way functions and garbling.

Definition 3.6 (Black-box constructions of PKE from GC-OWF)**.** A fully black-box construction of a PKE scheme from a one-way function and a garbling scheme for circuits with one-way function gates (shortly, from GC-OWF) consists of a triple of PPT oracle algorithms (G, E, D) and a PPT oracle security-reduction $S = (S_1, S_2)$ such that for any function f and any correct garbling scheme $L = (\mathsf{Garb}, \mathsf{Eval}, \mathsf{Sim})$ relative to f, both the following hold:

- **Correctness:** $\mathcal{E}^{f,L} = (G^{f,L}, E^{f,L}, D^{f,L})$ is a $(1 - \frac{1}{2^\kappa})$-correct PKE scheme. (See the remark after this definition.)

- **Security:** For any adversary any A that breaks the semantic security of the PKE scheme $\mathcal{E}^{f,L}$, either

  - $S_1^{f,L,A}$ breaks the one-wayness of f; or
  - $S_2^{f,L,A}$ breaks the security of the scheme $L = (\mathsf{Garb}, \mathsf{Eval}, \mathsf{Sim})$ relative to f. That is, for some oracle-aided circuit C and input x, $S_2^{f,L,A}$ can distinguish between the tuple $\left(\widetilde{\mathsf{C}}, \{\mathsf{label}_{i,x_i}\}_{i\in[m]}\right)$ and $\mathsf{Sim}\left(1^{|\mathsf{C}|}, 1^{|x|}, \mathsf{C}^f(x)\right)$, where $m = \mathsf{inpsize}(\mathsf{C})$ and we are given that $(\widetilde{\mathsf{C}}, \{\mathsf{label}_{i,b}\}_{i\in[m],b\in\{0,1\}}) \leftarrow \mathsf{Garb}\left(1^\kappa, \mathsf{C}\right)$.

Remark about the correctness condition in Definition 3.6. In Definition 3.6, for correctness we require that the constructed PKE be $(1 - \frac{1}{2^\kappa})$ correct. This is without loss of generality since one may easily boost correctness using standard techniques; i.e., let the new public key be a tuple of public keys under the original scheme. Encrypt a given plaintext bit under each individual public key. For decryption, we decrypt all the ciphertexts and go with the majority bit. The semantic security of this expanded scheme reduces to that of the base scheme using a hybrid argument, which is a fully-black-box reduction.

Calling the base primitives on the same security parameter. For simplicity of exposition, for any given black-box construction $\mathcal{E}^{f,L}$ we assume that $\mathcal{E}^{f,L}$ on the security parameter 1^κ always calls f and L on the same security parameter 1^κ. There are standard techniques for doing away with this restriction, but those extensions will only complicate the proofs further. Looking ahead, when we define our oracles $(O', \mathsf{rev})_{\kappa,n}$ in Definition 4.3, which are parameterized over a security parameter κ and a circuit size $n = n(\kappa)$, the above restriction means that $\mathcal{E}^{O'}$ on the security parameter 1^κ always calls O' on parameters such as (κ, n_1), (κ, n_2), etc. That is, the value of κ will be the same across all queries, but each query may use a different value for n.

4 Separating Public-Key Encryption from OWF-based Garbling

In this section, we state our main impossibility result and describe at a high-level the steps that we will take in order to prove our main theorem.

Theorem 4.1 (Main theorem). *There exists no fully black-box construction of a public-key encryption scheme from GC-OWFs; namely garbling schemes that garble circuits with one-way function gates in them (see Definition 3.6).*

Our theorem above follows from the following lemma.

Lemma 4.2. *There exists an oracle $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}, \mathsf{rev})$ for which the following holds (in what follows, let $O' = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval})$):*

1. *f is one-way relative to (O', rev). That is, f is one-way against all polynomial query (and even sub-exponential query) adversaries $\mathcal{A}^{O', \mathsf{rev}}$.*

2. *There exists a PPT GC-OWF construction $(\mathsf{Garb}^{O'}, \mathsf{Eval}^{O'}, \mathsf{Sim}^{O'})$ for f-aided circuits that is secure against any poly-query adversary $\mathcal{A}^{O', \mathsf{rev}}$.*

3. *For any PKE construction $\mathcal{E}^{O'}$ with access to the oracle O', there exists an attacker $\mathcal{A}^{O', \mathsf{rev}}$ that breaks the semantic security of $\mathcal{E}^{O'}$ using a polynomial number of queries.*

Note that Lemma 4.2 immediately implies Theorem 4.1.

Roadmap: Proof of Lemma 4.2. As common in black-box impossibility results, we will show the existence of the oracles required by Lemma 4.2 by proving results with respect to oracles chosen randomly according to a distribution. We will describe our oracle distribution below and will then outline the main steps we will take in order to prove Lemma 4.2.

Definition 4.3 (The ideal model/oracle). Let $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}, \mathsf{rev})_{\kappa, n}$ be an ensemble of oracles parameterized by (κ, n), where κ denotes the security parameter and n denotes the size of a circuit which we want to garble. We describe the distribution **O** from which these oracles are sampled for fixed values of (κ, n).

- $f : \{0, 1\}^\kappa \to \{0, 1\}^\kappa$: a uniformly chosen random function.

- $\mathsf{gc}(s, F) : \{0, 1\}^\kappa \times \{0, 1\}^n \to \{0, 1\}^{2(\kappa + n)}$: an injective random function that, given a key $s \in \{0, 1\}^\kappa$ and a single-bit-output oracle-aided circuit F, outputs an encoding $\widetilde{F}$.

- $\mathsf{gi}(s, i, x_i) : \{0, 1\}^\kappa \times \{0, 1\}^{\log n} \times \{0, 1\} \to \{0, 1\}^{2(\kappa + \log n)}$: an injective random function that, given a key $s \in \{0, 1\}^\kappa$, an index $i \in \{0, 1\}^{\log n}$, an input-wire bit value $x_i \in \{0, 1\}$, outputs an encoding $\widetilde{x}_i$. As notation, for any $X = (x_1, ..., x_n)$, we denote $\mathsf{gi}(s, X) := (\mathsf{gi}(s, i, x_i))_{i \in [n]} = \widetilde{X}$.

- $\mathsf{eval}(\widetilde{F}, \widetilde{X})$: given as input $\widetilde{F}$ and $\widetilde{X} = (\widetilde{x}_1, ..., \widetilde{x}_m)$, if there is a string $s \in \{0, 1\}^\kappa$ and circuit F such that $\mathsf{gc}(s, F) = \widetilde{F}$, that $m = \mathsf{inpsize}(F)$ and that for every $i \in [m]$ there exists $x_i \in \{0, 1\}$ such that $\mathsf{gi}(s, i, x_i) = \widetilde{x}_i$, then it outputs $F^f(x_1 || \cdots || x_m)$. Otherwise, it outputs $\bot$.

- $\mathsf{rev}(\widetilde{F}, \widetilde{X}, \widetilde{X}')$: if there exists $s \in \{0, 1\}^\kappa$ and circuit F such that $\mathsf{gc}(s, F) = \widetilde{F}$ and that there exists $X, X' \in \{0, 1\}^{\mathsf{inpsize}(F)}$ such that $X \neq X'$, $\mathsf{gi}(s, X) = \widetilde{X}$ and $\mathsf{gi}(s, X') = \widetilde{X}'$, then it outputs (s, F). Otherwise, it outputs $\bot$.

Remark 4.4. The size of a garbled circuit outputted by the gc oracle is roughly twice the size of the corresponding input circuit. Current garbled circuits constructions are not capable of achieving such a short expansion factor. We are able to do this as we model the garbling mechansim as a

14

totally random function. Nonetheless, working with such a short size expansion is without loss of generality, because a general black-box PKE construction out of GC-OWF should work with respect to any oracle that implements the GC-OWF securely. We should also mention that all our results hold (without having to make any changes) if the output of gc is bigger than the one specified in Definition 4.3.

First, we show that a random oracle $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}, \mathsf{rev})$ chosen according to the distribution **O** allows us to implement an ideal version of garbling for circuits with f gates. This is not surprising as O is indeed an idealized form of implementing this primitive.

Lemma 4.5 (Secure OWF and garbling exists relative to O)**.** *Let $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}, \mathsf{rev})$ be as in Definition 4.3 and let $O' = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval})$. Then, with probability (measure) one over the choice of O, the function f is one-way relative to O — i.e., f is one-way against any PPT oracle adversary with access to the oracle O. Moreover, there exists a PPT GC-OWF construction $(\mathsf{Garb}^{O'}, \mathsf{Eval}^{O'})$ for f-aided circuits which is secure relative to O with probability one over the choice of $O \leftarrow \mathbf{O}$.*

Proof. The fact that f is one-way relative to O with probability one over the choice of O is now standard (see [IR89]). Given any oracle $O = (O', \mathsf{eval})$, we now show how to construct a PPT garbling scheme $L^{O'} = (\mathsf{Garb}^{O'}, \mathsf{Eval}^{O'}, \mathsf{Sim}^{O'})$ for f-aided circuits. The algorithm $\mathsf{Garb}^{O'}$ on input $(1^\kappa, C)$ samples $s \leftarrow \{0,1\}^\kappa$, sets $m = \mathsf{inpsize}(C)$ and outputs the garbled circuit $\widetilde{C} = \mathsf{gc}(s, C)$ as well as a sequence of garbled inputs $(\widetilde{x}_{1,0}, \widetilde{x}_{1,1}, \ldots, \widetilde{x}_{m,0}, \widetilde{x}_{m,1})$, where for $i \in m$ and $b \in \{0,1\}$ we have $\widetilde{x}_{i,b} = \mathsf{gi}(s, i, b)$.

The algorithm $\mathsf{Eval}^{O'}(\widetilde{C}, \widetilde{x}_1 || \cdots || \widetilde{x}_m)$ simply outputs $\mathsf{eval}(\widetilde{C}, \widetilde{x}_1 || \cdots || \widetilde{x}_m)$. Correctness holds by definition of the oracle.

For security, we will define $\mathsf{Sim}^{O'}$ as follows: on input $(1^\kappa, n, m, y \in \{0,1\})$, where n denotes the size of the circuit, m denotes the number of input wires and y denotes the output value, we set C_0 to be a canonical circuit of size n and with m input wires that always outputs y. Sample $s \leftarrow \{0,1\}^\kappa$ and let $\widetilde{C} = \mathsf{gc}(s, C)$ and $\widetilde{X} = \mathsf{gi}(s, 0^m)$. Output $(\widetilde{C}, \widetilde{X})$. Simulation security follows from the random nature of the oracles. That is, for any polynomial-query distinguisher $\mathcal{A}^{O', \mathsf{rev}}$, for any n, m and any circuit C of size n and of input size m and any input $X \in \{0,1\}^m$, we have

$$\left| \Pr[\mathcal{A}^{O', \mathsf{rev}}(\widetilde{C}, \widetilde{X}) = 1] - \Pr[\mathcal{A}^{O', \mathsf{rev}}(\widetilde{C'}, \widetilde{X'}) = 1] \right| = \mathsf{negl}(\kappa), \tag{1}$$

where $s \leftarrow \{0,1\}^\kappa$, $\widetilde{C} = \mathsf{gc}(s, C)$, $\widetilde{X} = \mathsf{gi}(s, X)$ and $(\widetilde{C'}, \widetilde{X'}) \leftarrow \mathsf{Sim}^{O'}(1^\kappa, n, m, C^f(X))$. We omit the details of the proof of Equation 1 as it is a simple information theoretic argument. $\square$

We are left with proving Part 3 of Lemma 4.2. Proving this part is the main technical contribution of our paper, and is done via an oracle *reducibility* technique. In order to state this reducibility statement formally, we first need to define the notions of correctness and attack advantage in the *ideal model*.

Definition 4.6 (Correctness in the ideal model)**.** For a polynomial $p = p(\kappa)$ we say that a single-bit PKE scheme $\mathcal{E}^O = (G^O, E^O, D^O)$ is $\frac{1}{2} + \frac{1}{p}$ correct in the ideal model if for both $b \in \{0,1\}$:

$$\Pr[D^O(sk, c) = b] \geq \frac{1}{2} + \frac{1}{p}, \tag{2}$$

where the probability is taken over $O \leftarrow \mathbf{O}$, $(pk, sk) \leftarrow G^O(1^\kappa)$ and $c \leftarrow E^O(pk, b)$.

Definition 4.7 (Ideal model attack advantage)**.** We say that an adversary $\mathcal{A}$ breaks the semantic security of a single-bit PKE (G^O, E^O, D^O) in the ideal model with probability γ (or with advantage γ) if $\Pr[A(pk, c) = b] \geq \gamma$, where the probability is taken over $O \leftarrow \mathbf{O}$, $(pk, sk) \leftarrow G^O(1^\kappa)$, $b \leftarrow \{0, 1\}$, $c \leftarrow E^O(pk, b)$ and over $\mathcal{A}$'s random coins.

We are now ready to describe our oracle reducibility lemma.

Lemma 4.8 (Reducibility to the random oracle model)**.** *Let $\mathcal{E}$ be a given PKE construction possibly making use of all the oracles $O' = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval})$. There exists a compilation procedure and a polynomial-query security-reduction Red such that the compilation transforms $\mathcal{E}^{O'}$ into a new polynomial-query PKE construction $\ddot{\mathcal{E}}^{f, \mathsf{gc}, \mathsf{gi}}$, where $\ddot{\mathcal{E}}$ makes no eval queries and for which both the following hold:*

- ***Correctness**: If $\mathcal{E}^{O'}$ is $(1 - \frac{1}{2^\kappa})$ correct in the ideal model, the compiled scheme $\ddot{\mathcal{E}}^{f, \mathsf{gc}, \mathsf{gi}}$ has at least $(1 - \frac{1}{\kappa^7})$ correctness in the ideal model.*

- ***Security reduction**. For any constant c the following holds: if there exists an adversary $\mathcal{A}$ that breaks the semantic security of $\ddot{\mathcal{E}}^{f, \mathsf{gc}, \mathsf{gi}}$ in the ideal model with probability η, the algorithm $\mathrm{Red}^{O', \mathsf{rev}, \mathcal{A}}$ breaks the semantic security of $\mathcal{E}^{O'}$ in the ideal model with probability $\eta - \frac{1}{\kappa^c}$.*

Let us first show how to use lemmas 4.5 and 4.8 to establish Lemma 4.2.

Completing proof of Lemma 4.2 and Theorem 4.1. Let $\mathcal{E} = (G, E, D)$ be a candidate PKE construction. We will show that with probability one over the choice of $(O', \mathsf{rev}) \leftarrow \mathbf{O}$, the PKE construction $\mathcal{E}^{O'}$ can be broken by a polynomial number of queries to (O', rev). Let us first show how to use this claim to complete the proof of Theorem 4.1, and we will then prove this claim. By Lemma 4.5, we know that with probability one over the choice of O we have (a) f is one-way relative to (O', rev) and (b) $(\mathsf{Garb}^{O'}, \mathsf{Eval}^{O'}, \mathsf{Sim}^{O'})$ is a secure GC-OWF construction for f-aided circuits against all polynomial-query adversaries with access to the oracles (O', rev). Thus, the foregoing claim coupled with Lemma 4.5 implies Lemma 4.2. In what follows we prove the foregoing claim.

By Definition 3.6 we know that $\mathcal{E}^{O'}$ has $(1 - \frac{1}{2^\kappa})$-correctness in the ideal model. Thus, by Lemma 4.8 there exists a compiled scheme $\ddot{\mathcal{E}}^{f, \mathsf{gc}, \mathsf{gi}}$ that has at least $(1 - \frac{1}{\kappa^7})$-correctness in the ideal model. Note that the oracles f, gc and gi are nothing but three independent random oracles. By the results of [IR89, BMG09] there exists a polynomial query adversary $\mathcal{A}^{f, \mathsf{gc}, \mathsf{gi}}$ which breaks the semantic security of $\ddot{\mathcal{E}}^{f, \mathsf{gc}, \mathsf{gi}}$ in the ideal model with probability $(1 - \frac{1}{\kappa^6})$.[5] (See Definition 4.7 for the notion of "break in the ideal model.") Invoking Lemma 4.8 again and choosing the constant c appropriately, we will obtain a polynomial query adversary $\mathcal{B}^{O', \mathsf{rev}}$ which breaks the semantic security of $\mathcal{E}^{O'}$ in the ideal model with probability $(1 - \frac{1}{\kappa^5})$. That is,

$$\Pr_{O=(O', \mathsf{rev}), pk, b, c} \left[\mathcal{B}^{O', \mathsf{rev}}(pk, c) = b \right] \geq 1 - \frac{1}{\kappa^5}, \tag{3}$$

where $(pk, sk) \leftarrow G^{O'}(1^\kappa)$, $b \leftarrow \{0, 1\}$ and $c \leftarrow E^{O'}(pk, b)$.

Using a simple averaging argument we have

---

[5]The results of [IR89, BMG09] show how to break the semantic security of any key exchange (and hence PKE) construction in the random oracle model with a probability that is at most $\frac{1}{\kappa^{c'}}$ less than the correctness probability, for any arbitrary constant $c' > 0$.

$$\Pr_{O=(O',\mathsf{rev})}\left[\Pr_{pk,b,c}\left[\mathcal{B}^{O',\mathsf{rev}}(pk,c)=b\right]\geq 1-\frac{1}{\kappa^3}\right]\geq 1-\frac{1}{\kappa^2}. \qquad (4)$$

Equation 4 implies that for at most $\frac{1}{\kappa^2}$ fraction of the oracles $O=(O',\mathsf{rev})$, the adversary $\mathcal{B}^{O',\mathsf{rev}}(pk,c)$, on security parameter κ, recovers b with probability less than $1-\frac{1}{\kappa^3}$. Since $\sum_{i=1}^{\infty}\frac{1}{i^2}$ converges, by the Borel-Cantelli Lemma we have that for a measure-one fraction of oracles $O=(O',\mathsf{rev})\leftarrow\mathbf{O}$, the adversary $\mathcal{B}^{O',\mathsf{rev}}$ breaks the semantic security of $\mathcal{E}^{O'}$. The proof is now complete. □

Roadmap for the proof of Lemma 4.8. Finally, all that remains is proving Lemma 4.8 which shows that we can compile out eval queries from any PKE scheme without significantly hurting correctness or security. In the remainder of this paper, we show that such a compilation procedure exists. We obtain the compiled eval-free scheme $(\ddot{G}^{f,\mathsf{gc},\mathsf{gi}},\ddot{E}^{f,\mathsf{gc},\mathsf{gi}},\ddot{D}^{f,\mathsf{gc},\mathsf{gi}})$ in two steps. First, in Section 4.1, we show how to compile out eval queries from $E^{O'}$ only. In particular, we will prove the following lemma.

Lemma 4.9 (Compiling out eval from E)**.** *Let δ be an arbitrary polynomial and parse $O=(O',\mathsf{rev})$. There exists a compilation procedure that achieves the following for any constant c. Given any $(\frac{1}{2}+\delta)$-ideally-correct PKE scheme $\mathcal{E}=(G^{O'},E^{O'},D^{O'})$, the compiled PKE scheme $\dot{\mathcal{E}}=(\dot{G}^{f,\mathsf{gc},\mathsf{gi}},\dot{E}^{f,\mathsf{gc},\mathsf{gi}},D^{O'})$ is $(\frac{1}{2}+\delta-\frac{1}{\kappa^c})$-ideally-correct. Moreover, there exists a polynomial-query algorithm* SecRed *that satisfies the following: for any adversary $\mathcal{A}$ that breaks the semantic security of $\dot{\mathcal{E}}$ in the ideal model with advantage η, the adversary* SecRed$^{\mathcal{A},O}$ *breaks the semantic security of $\mathcal{E}$ in the ideal model with advantage at least $\eta-\frac{1}{\kappa^c}$.*

Then, in Section 4.2 we show how to compile out eval from $D^{O'}$, assuming neither of the algorithms G and E call eval. That is, we prove the following lemma.

Lemma 4.10 (Compiling out eval from D)**.** *Let δ be an arbitrary polynomial. There exists a compilation procedure that achieves the following for any constant c. Given any $(\frac{1}{2}+\delta)$-ideally-correct PKE scheme $\mathcal{E}=(G^{f,\mathsf{gc},\mathsf{gi}},E^{f,\mathsf{gc},\mathsf{gi}},D^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}})$, the compiled PKE scheme $\ddot{\mathcal{E}}=(\ddot{G}^{f,\mathsf{gc},\mathsf{gi}},\ddot{E}^{f,\mathsf{gc},\mathsf{gi}},\ddot{D}^{f,\mathsf{gc},\mathsf{gi}})$ is $(\frac{1}{2}+\delta-\frac{1}{\kappa^c})$-ideally-correct. Moreover, there exists a polynomial-query algorithm* SecRed *that satisfies the following: for any adversary $\mathcal{A}$ that breaks the semantic security of $\ddot{\mathcal{E}}$ in the ideal model with advantage η, the adversary* SecRed$^{\mathcal{A},O}$ *breaks the semantic security of $\mathcal{E}$ in the ideal model with advantage at least $\eta-\frac{1}{\kappa^c}$.*

Th proof of Lemma 4.8 immediately follows from Lemmas 4.9 and 4.10.

4.1 Removing Garbling Evaluation Queries from Encryption

In this section, we will prove Lemma 4.9. Namely, we will show how to compile the PKE scheme $\mathcal{E}=(G^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}},E^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}},D^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}})$ into a new PKE scheme $\dot{\mathcal{E}}=(\dot{G}^{f,\mathsf{gc},\mathsf{gi}},\dot{E}^{f,\mathsf{gc},\mathsf{gi}},D^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}})$ with correctness and security comparable to the original scheme $\mathcal{E}$ but where $\dot{E}$ would not ask any eval queries. First, we may assume without loss of generality that G does not make queries to eval — it can predict the answer itself. Thus, we will focus on removing eval queries from E assuming that G does not make any eval queries.

4.1.1 Definitions

Before we describe the compilation process, we will first present some definitions that will be used throughout this section.

Definition 4.11 (Valid outputs)**.** For any oracle $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}, \mathsf{rev})$, we say that $\widetilde{F}$ is a valid garbled circuit with respect to O if there exists (s, F) such that $\mathsf{gc}(s, F) = \widetilde{F}$. Similarly, we say that $\widetilde{X}$ is a valid garbled input with respect to O if there exists (s, X) such that $\mathsf{gi}(s, X) = \widetilde{X}$.

We also define the notion of normal form with respect to oracle-aided algorithms. At a high-level, a normal form algorithm avoids asking any redundant queries if it already knows the answer to such queries.

Definition 4.12 (Normal form)**.** Let A be an oracle algorithm that accepts as input a query-answer set $\mathsf{Q_S}$ and let $\mathsf{Q_A}$ be the query-answer pairs that A has asked so far. We say that A is in *normal-form* if it satisfies the following conditions:

1. A never asks duplicate queries.

2. Before it asks an $((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?)$ query qu, A first checks if there exists a query-answer pair $((s, F) \xrightarrow[\mathsf{gc}]{} \widetilde{F})$ in $\mathsf{Q_A} \cup \mathsf{Q_S}$. If that is the case then it would not issue qu to the oracle but would instead run $F^f(X)$ on its own where X can be obtained bit-by-bit by searching $\mathsf{gi}(s, i, x_i)$ for every index position $i \in n$ and every bit $x_i \in \{0, 1\}$.

Recall that our goal is to remove eval queries from E to obtain an eval-free algorithm $\dot{E}$. To make this transformation possible, the new algorithm $\dot{E}$ needs some help from its associated key generation algorithm $\dot{G}$ so as to make up for its lack of access to eval. This help is sent to $\dot{E}$ as part of a hint list H, attached to the public key, by the key generation algorithm $\dot{G}$. The following definition describes how $\dot{G}$ forms the hint list H based on its inside information Aux and based on some information Q that G has collected about random executions of E.

Definition 4.13 (Constructing helper tuples)**.** We define a function `ConstHelp` that takes as input a query-answer set Q along with some query-answer set Aux and outputs a set H as follows:

- If there exists $((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} y \neq \bot) \in \mathsf{Q}$ such that for no $\widetilde{X}' \neq \widetilde{X}$ do we have $((\widetilde{F}, \widetilde{X}') \xrightarrow[\mathsf{eval}]{} y' \neq \bot) \in \mathsf{Q}$, then add $((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} y)$ to H.

- If for two distinct $\widetilde{X}_1$ and $\widetilde{X}_2$ we have $((\widetilde{F}, \widetilde{X}_1) \xrightarrow[\mathsf{eval}]{} y_1 \neq \bot) \in \mathsf{Q}$ and $((\widetilde{F}, \widetilde{X}_2) \xrightarrow[\mathsf{eval}]{} y_2 \neq \bot) \in \mathsf{Q}$, then if for some (s, F) we have $((s, F) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{Aux}$, add $((s, F) \xrightarrow[\mathsf{gc}]{} \widetilde{F})$ to H.

Having a hint list H, we give the following definition that describes the idea of how the receiving algorithm $\dot{E}$ may use it to avoid making eval queries. In the following definition one may think of Q as a hint list.

Definition 4.14 (Emulating eval queries)**.** For any oracle $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}, \mathsf{rev})$, we define the function `HandleEval`$^{f, \mathsf{gi}}$ to be a subroutine that takes as input a set Q of query-answer pairs to O and a query qu of the form $((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?)$ then performs the following steps to answer qu:

- If there exists a tuple $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} y)$ in Q, then output y.

- If there exists $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{Q}$, then find X such that $\mathsf{gi}(s, X) = \widetilde{X}$ and output $y = F^f(X)$.

- If neither of the above cases happen, then return $\perp$ as the answer to qu.

We will also define the notion of a mixed oracle that uses O on non-eval queries but uses $\texttt{HandleEval}$ to answer eval queries without resorting to O. This oracle is constructed and used in the newly compiled algorithms when we want to avoid asking eval queries to O.

Definition 4.15 (Mixed oracle). For an oracle $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval})$ and a set of query-answer pairs S, we denote $O[S]$ to be an *Eval-mixed* oracle that answers all f, gc, and gi queries by forwarding them to the real oracle O, but for any eval query qu it will emulate the answer by calling and returning $y = \texttt{HandleEval}^{f,\mathsf{gi}}(S, qu)$.

4.1.2 Compilation Procedure

Let $\mathcal{E} = (G^{f,\mathsf{gc},\mathsf{gi}}, E^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}}, D^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}})$ be the give construction for which we want to remove eval queries from E. Without loss of generality, we assume that all the algorithms of $\mathcal{E}$ are in normal form (see Definition 4.12).

For simplicity, we keep O as a superscript to all the algorithms of $\mathcal{E}$, but it should be understood that the actual oracle access is of the form above.

We need the following definition as we will need to choose parameters in the compilation construction based on the query complexity of the given construction.

Definition 4.16 (Parameter $q = q(\kappa)$: size-upperbound). Throughout this section, fix $q = q(\kappa)$ to be an arbitrary polynomial that satisfies the following.

1. $q \geq \kappa$;

2. q is greater than the total number of queries that each of the algorithms (G^O, E^O, D^O) make on inputs corresponding to the security parameter 1^κ and on $O \leftarrow \mathbf{O}$; and

3. q is greater than the size of any query made by any of (G^O, E^O, D^O) on inputs corresponding to the security parameter 1^κ and on $O \leftarrow \mathbf{O}$.

Construction 4.17 (The compiled scheme $\dot{\mathcal{E}}$.). The compiled scheme $(\dot{G}, \dot{E}, D)$ is parameterized over a function $t = t(\kappa)$, which we will instantiate later.

- $\dot{G}(1^\kappa)$: Perform the following steps:

  1. Run $(pk, sk) \leftarrow G^O(1^\kappa)$. Add all query-answer pairs generated in this step to OrigG.
  2. **Generating helper set H for $\dot{E}$:** Set $\mathsf{LocalE} = \varnothing$.
     (a) For $i = [1, t]$, do the following: Run $E^{O[\mathsf{OrigG}]}(pk, 0)$ and $E^{O[\mathsf{OrigG}]}(pk, 1)$ and keep adding all the resulting query-answer pairs to LocalE.
     (b) Set $\mathsf{H} := \texttt{ConstHelp}(\mathsf{LocalE}, \mathsf{OrigG} \cup \mathsf{LocalE})$.
  3. Output $\dot{pk} = (pk, \mathsf{H})$ and $\dot{sk} = sk$.

- $\dot{E}(\dot{pk}, b)$: Parse $\dot{pk} = (pk, \mathsf{H})$. Run $\dot{c} \leftarrow E^{O[\mathsf{H}]}(pk, b)$ and add all the query-response pairs to OrigE. Return $\dot{c}$.

19

Remark about $\dot{E}$. We note that, by the definition of $O[H]$, all the `eval` queries of $E^{O[H]}(pk, b)$ will be emulated using H. Thus, $\dot{E}$ will not issue any `eval` queries.

Query complexity of $\dot{\mathcal{E}}$. It is immediate to see that the query complexity of each of the compiled algorithms is polynomial in q and t, where q the query complexity of (G, E, D).

Lemma 4.18. *Let q be the size-upperbound of (G, E, D) as given in Definition 4.16. The query complexity of $\dot{\mathcal{E}} = (\dot{G}, \dot{E}, D)$ is at most $q + (2q^2)t \leq 3tq^2$.*

4.1.3 Correctness and Security

In this section we give the correctness and security statements regarding the compiled scheme $\dot{\mathcal{E}} = (\dot{G}, \dot{E}, D)$ and prove them. By doing so, we complete the proof of Lemma 4.9.

Lemma 4.19 (Correctness of $\dot{\mathcal{E}}$). *Suppose the original scheme (G, E, D) is $(\frac{1}{2} + \delta)$ correct in the ideal model. The compiled scheme $(\dot{G}^O, \dot{E}^O, D^O)$ has at least $(\frac{1}{2} + \delta - \frac{2q}{t} - \mathrm{negl}(\kappa))$ correctness in the ideal model, where t is the number of iterations performed in $\dot{G}$.*

In particular, for any constant $c > 0$ by taking $t = q^{c+2}$, the compiled scheme $(\dot{G}^O, \dot{E}^O, D^O)$ has at least $(\frac{1}{2} + \delta - \frac{1}{\kappa^c})$ correctness.

Lemma 4.20 (Security of $\dot{\mathcal{E}}$). *There exists a polynomial-query algorithm `SecRed` that satisfies the following. For any adversary $\mathcal{A}$ that breaks the semantic security of $(\dot{G}^O, \dot{E}^O, \dot{D}^O)$ in the ideal model with probability at least γ, the algorithm `SecRed`$^{\mathcal{A},O}$ breaks the semantic security of (G^O, E^O, D^O) with probability at least $\gamma - \frac{1}{2^{\kappa/4}} - \frac{1}{\kappa^c}$ for any constant $c > 0$.*

We prove Lemma 4.19 in Section 4.1.4 and Lemma 4.20 in Section 4.1.5.

4.1.4 Proof of Correctness for $\dot{\mathcal{E}}$

In this section, we will prove Lemma 4.19, which states that $\dot{\mathcal{E}} = (\dot{G}, \dot{E}, D)$ is still a correct PKE after having removed the `eval` queries from E.

Parsing $\dot{pk} = (pk, H)$, recall that $\dot{E}^O(\dot{pk}, b)$ simply runs $E^{O[H]}(pk, b)$. With this in mind, to prove Lemma 4.19, we give the following lemma, which shows that the probability that the outputs of $E^{O[H]}(pk, b)$ and $E^O(pk, b)$ are different is small.

Lemma 4.21. *For $b \in \{0, 1\}$ we have*

$$\Pr_{O,r,pk,H}[E^O(pk, b; r) \neq E^{O[H]}(pk, b; r)] \leq \frac{2q}{t} + \frac{1}{2^{\kappa/3}}$$

where $O \leftarrow \mathbf{O}$, $((pk, H), sk) \leftarrow \dot{G}^O(1^\kappa)$ and $r \leftarrow \{0, 1\}^*$.

We first show how to derive Lemma 4.19 from Lemma 4.21.

Proof of Lemma 4.19. Parse $\dot{pk} = (pk, H)$. All the probabilities below are taken over the random choices of $\dot{pk}$, O and r. We have

$$
\begin{aligned}
\Pr[D^O(sk, \dot{E}^O(\dot{pk}, b; r)) \neq b] &= \Pr[D^O(sk, E^{O[H]}(pk, b; r)) \neq b] \\
&\leq \Pr[D^O(sk, E^O(pk, b; r)) \neq b] + 2q/t + \mathrm{negl}(\kappa) \\
&\leq \frac{1}{2} - \delta + 2q/t + \mathrm{negl}(\kappa)
\end{aligned}
$$

where the first inequality follows from Lemma 4.21. $\qquad\square$

We now focus on proving Lemma 4.21. Fix $b \in \{0, 1\}$. For compactness, we define the following experiment that outputs some random variables that will be later used to define some events.

Experiment $\mathbf{Expr}(1^\kappa)$ for fixed $b \in \{0, 1\}$: Output $\mathrm{Vars} = (pk, \mathsf{OrigG}, \mathsf{LocalE}, \mathsf{H}, r)$, where $pk, \mathsf{OrigG}, \mathsf{LocalE}$ and H are sampled as in $\dot{G}(1^\kappa)$ and $r \leftarrow \{0, 1\}^*$ is the randomness to $\dot{E}(pk, b)$.

We define the following bad events. Note that all these bad events as well as those that appear later are defined based on the output of Vars, and so we make this dependence implicit henceforth.

- Bad_1: The event that $E^O(pk, b; r)$ makes a query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?)$, where $((*, *) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \notin \mathsf{OrigG}$ and $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$.

- Bad_2: The event that the execution of $E^O(pk, b; r)$ queries $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?)$ for which we have $((*, *) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG}$, $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$ and $O[\mathsf{H}](qu) = \texttt{HandleEval}(\mathsf{H}, qu) = \bot$.

Roadmap for the proof of Lemma 4.21. The proof of Lemma 4.21 now follows from the following lemmas.

Lemma 4.22. $\Pr_{O, \mathrm{Vars}}[E^O(pk, b; r) \neq E^{O[\mathsf{H}]}(pk, b; r)] \leq \Pr[\mathsf{Bad}_1 \vee \mathsf{Bad}_2]$ *where* $O \leftarrow \mathbf{O}$ *and* $\mathrm{Vars} = (pk, \mathsf{OrigG}, \mathsf{LocalE}, \mathsf{H}, r) \leftarrow \mathbf{Expr}(1^\kappa)$.

Lemma 4.23. $\Pr_{O, \mathrm{Vars}}[\mathsf{Bad}_1] \leq \frac{1}{2^{\kappa/3}}$ *where* $O \leftarrow \mathbf{O}$ *and* $\mathrm{Vars} \leftarrow \mathbf{Expr}(1^\kappa)$.

Lemma 4.24. $\Pr_{O, \mathrm{Vars}}[\mathsf{Bad}_2 \wedge \overline{\mathsf{Bad}_1}] \leq \frac{2q}{t}$ *where* $O \leftarrow \mathbf{O}$ *and* $\mathrm{Vars} \leftarrow \mathbf{Expr}(1^\kappa)$

The proof of Lemma 4.21 follows immediately from Lemmas 4.22, 4.23, and 4.24. We now prove all these lemmas below.

Proof of Lemma 4.22. Let Bad be the event $E^O(pk, b; r) \neq E^{O[\mathsf{H}]}(pk, b; r)$. We show that whenever Bad holds, then either Bad_1 happens or Bad_2 happens, hence proving the lemma. Notice that the only difference between the executions of $E^O(pk, b; r)$ and $E^{O[\mathsf{H}]}(pk, b; r)$ is how eval queries are handled. Specifically, in $E^{O[\mathsf{H}]}(pk, b; r)$, the eval queries are simulated with respect to the set H whereas in $E^O(pk, b; r)$, the real oracle O is used to reply to these queries. All of f, gc, and gi queries will be handled identically in both experiments by forwarding them to O. Thus, we only need to consider what happens in either execution when a new query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?)$ is asked.

Suppose Bad holds and let $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?)$ be the first eval query that will be answered differently between the two executions. That is, qu will be replied to with $\bot$ under $O[\mathsf{H}]$, but receives an answer $y \neq \bot$ from the real oracle O. We will now show that either Bad_1 or Bad_2 must hold. Consider two cases:

1. $((*, *) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \notin \mathsf{OrigG}$: In this case, the fact that $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$ implies that Bad_1 holds.

2. $((*, *) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG}$: In this case the facts that $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$, that qu is a query during the execution of $E^O(pk, b; r)$, and that $O[H](qu) = \bot$ imply that Bad_2 holds.

$\square$

Proof of Lemma 4.23. The proof of this lemma follows by a simple reduction to Lemma 3.1. Letting $\alpha = \Pr[\mathsf{Bad}_1]$, we will show how to build an adversary $\mathcal{A}^{f,\mathsf{gc},\mathsf{gi}}(1^\kappa)$ in the sense of Lemma 3.1 that will win with probability $\alpha \cdot \frac{1}{\mathrm{poly}(\kappa)}$.

Let i be the index of the first query qu during the execution of $E^O(pk, b; r)$ for which the event Bad_1 holds. Note that up to the query index i, the executions of $E^O(pk, b; r)$ and $E^{O[\mathsf{OrigG}]}(pk, b; r)$ are identical. With this in mind, we build the adversary $\mathcal{A}^{f,\mathsf{gc},\mathsf{gi}}(1^\kappa)$ as follows.

The adversary $\mathcal{A}^{f,\mathsf{gc},\mathsf{gi}}(1^\kappa)$ samples $(pk, sk) \leftarrow G^{f,\mathsf{gc},\mathsf{gi}}(1^\kappa)$, forming the set of query/response pairs OrigG. Then $\mathcal{A}^{f,\mathsf{gc},\mathsf{gi}}$ guesses $i \leftarrow [q]$ and runs $E^{O[\mathsf{OrigG}]}(pk, b; r)$ for a random r. Notice that $\mathcal{A}$ makes no queries to eval whatsoever, as it handles eval queries using OrigG. If the ith query of this execution is $((\widetilde{F}, *) \xrightarrow{\mathsf{eval}} ?)$ for some $\widetilde{F}$, then $\mathcal{A}^{f,\mathsf{gc},\mathsf{gi}}$ returns $\widetilde{F}$; otherwise, $\mathcal{A}$ returns $\bot$.

$\mathcal{A}^{f,\mathsf{gc},\mathsf{gi}}(1^\kappa)$ wins with probability at least $\alpha \cdot \frac{1}{q}$. On the other hand, by Lemma 3.1 we know $\mathcal{A}$'s success probability is at most $\frac{1}{2^{\kappa/2}}$. Thus, we have $\alpha \leq \frac{1}{2^{\kappa/3}}$, and the proof is complete. $\qquad\square$

Proof of Lemma 4.24. We claim that whenever the event $\mathsf{Bad}_2 \wedge \overline{\mathsf{Bad}_1}$ holds then the event Miss, defined below, must necessarily hold. Miss is the event that during the execution of $E^{O[\mathsf{OrigG}]}(pk, b; r)$ there is a query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow{\mathsf{eval}} ?)$, such that

1. $((*, *) \xrightarrow{\mathsf{gc}} \widetilde{F}) \in \mathsf{OrigG}$;

2. $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$;

3. $((*, *) \xrightarrow{\mathsf{gc}} \widetilde{F}) \notin \mathsf{H}$ and $((\widetilde{F}, \widetilde{X}) \xrightarrow{\mathsf{eval}} *) \notin \mathsf{H}$.

The reason for the above claim is that if $\mathsf{Bad}_2 \wedge \overline{\mathsf{Bad}_1}$ holds, then $\overline{\mathsf{Bad}_1}$ must necessarily hold, and thus the two executions $E^O(pk, b; r)$ and $E^{O[\mathsf{OrigG}]}(pk, b; r)$ are identical. The rest follows by the definition of the event Bad_2. We will prove

$$\Pr[\mathsf{Miss}] \leq \frac{2q}{t}, \tag{5}$$

which will yield the proof of this lemma. Thus, in the sequel we focus on proving Equation 5.

We break the event Miss into smaller events. For that, we need some notation. Let $i \in [n]$, $d \in \{0, 1\}$, F be circuit with input size n and let $\widetilde{F} = \mathsf{gc}(s, F)$, for some s. We say that a garbled input $\widetilde{X} = (\widetilde{x}_1, \ldots, \widetilde{x}_n)$ is an (i, d)-*match* for $\widetilde{F}$ if $\widetilde{X}$ is a valid garbled input of $\widetilde{F}$ and the ith garbled bit of $\widetilde{X}$ corresponds to the bit d. Formally,

- for all $j \in [n]$ and $j \neq i$: $\widetilde{x}_j = \mathsf{gi}(s, j, 0)$ or $\widetilde{x}_j = \mathsf{gi}(s, j, 1)$;

- $\widetilde{x}_i = \mathsf{gi}(s, i, d)$.

We say that a set of query/response pairs U *contains an* (i, d)-*match for* $\widetilde{F}$ if there exists $((\widetilde{F}, \widetilde{X}) \xrightarrow{\mathsf{eval}} *) \in \mathsf{U}$ such that $\widetilde{X}$ is an (i, d)-match for $\widetilde{F}$.

We also give the following notation. Recalling the way in which LocalE is constructed in $\dot{G}$ through t iterations, for $i \in [t]$ let LocalE_i be the set formed after the i-th iteration. Also, let OrigE^* be the set of all query/response pairs during the execution of $E^{O[\mathsf{OrigG}]}(pk, b; r)$.

We now define a series of events, $\mathsf{Miss}_{i,d}$, for $i \in [q]$ and $d \in \{0, 1\}$, and will show that if Miss holds then for some i and d the event $\mathsf{Miss}_{i,d}$ must hold.

Event $\mathsf{Miss}_{i,d}$ is the event that for some $\widetilde{F}$ that $((*, *) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG}$, both the following hold:

1. OrigE^* contains an (i, d)-match for $\widetilde{F}$;

2. none of the sets $\mathsf{LocalE}_1, \cdots, \mathsf{LocalE}_t$ do contain an (i, d)-match for $\widetilde{F}$.

We claim that if Miss holds then $\mathsf{Miss}_{i,d}$ must hold, for some $i \in [q]$ and $d \in \{0, 1\}$. Suppose the event Miss holds for the query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?)$ (see above for the definition of Miss). We consider all possible cases:

- For no (i, d) does the set $\mathsf{LocalE} = \mathsf{LocalE}_1 \cup \cdots \cup \mathsf{LocalE}_t$ contain an (i, d)-match for $\widetilde{F}$. In this case, since the set OrigE^* contains $((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} *)$, there exists an (i, d)-match for $\widetilde{F}$ for all $i \in [q]$ and so $\mathsf{Miss}_{i,d}$ holds for some d and all $i \in [q]$.

- There is one and only one garbled input $\widetilde{X_1}$ which is valid for $\widetilde{F}$ and for which we have $((\widetilde{F}, \widetilde{X_1}) \xrightarrow[\mathsf{eval}]{} ?) \in \mathsf{LocalE}$. In this case, we must have $\widetilde{X_1} \neq \widetilde{X}$, because otherwise we would have $((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} *) \in \mathsf{H}$, a contradiction to the fact that Miss holds. Thus, for some (i, d) both the following must hold: (A) $\widetilde{X}$ is an (i, d)-match for $\widetilde{F}$ and (B) $\widetilde{X_1}$ is not an (i, d)-match for $\widetilde{F}$. Thus, for some i and d, the event $\mathsf{Miss}_{i,d}$ must hold.

- There are at least two different garbled inputs $\widetilde{X_1}$ and $\widetilde{X_2}$ which both are valid for $\widetilde{F}$ and which $((\widetilde{F}, \widetilde{X_1}) \xrightarrow[\mathsf{eval}]{} ?) \in \mathsf{LocalE}$ and $((\widetilde{F}, \widetilde{X_2}) \xrightarrow[\mathsf{eval}]{} ?) \in \mathsf{LocalE}$: This case cannot happen because otherwise we would have $(*, * \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{H}$, a contradiction to the fact that Miss holds.

Having proved $\Pr[\mathsf{Miss}] \leq \sum_{i,d} \Pr[\mathsf{Miss}_{i,d}]$, we bound the probability of each individual $\mathsf{Miss}_{i,d}$. To bound the probability of the event $\mathsf{Miss}_{i,d}$, note that since all of $\mathsf{LocalE}_1, \ldots, \mathsf{LocalE}_t$ and OrigE^* are obtained via independent and identical processes, by Lemma 3.1 we have

$$\Pr[\mathsf{Miss}_{i,d}] \leq \frac{1}{t}.$$

Using a union bound, $\Pr[\mathsf{Miss}] \leq \frac{2q}{t}$, and Equation 5 is now proved. This completes the proof. $\qquad \square$

4.1.5 Proof of Security for $\dot{\mathcal{E}}$

Proof of Lemma 4.20. To define the reduction algorithm `SecRed` we need to introduce the following procedure, overloading the definition of `ConstHelp` (Definition 4.13). In Definition 4.13 the procedure `ConstHelp` was given as input an auxiliary information set Aux which helps the procedure in finding answers to the eval queries provided in the given set Q. In the definition below, however, there is no auxiliary information set, but the procedure could use the oracle rev.

Definition 4.25. Procedure `ConstHelp`:

- **Input:** A set of query/answer pairs Q.

- **Oracle:** $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}, \mathsf{rev})$.

- **Output:** A "hint" set H formed as follows:

  - If there exists $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} y \neq \bot) \in \mathsf{Q}$ such that for no $\widetilde{X}' \neq \widetilde{X}$ do we have $((\widetilde{F}, \widetilde{X}') \xrightarrow[\text{eval}]{} y' \neq \bot) \in \mathsf{Q}$, then add $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} y)$ to H.

  - If for two distinct $\widetilde{X}_1$ and $\widetilde{X}_2$ we have $((\widetilde{F}, \widetilde{X}_1) \xrightarrow[\text{eval}]{} y_1 \neq \bot) \in \mathsf{Q}$ and $((\widetilde{F}, \widetilde{X}_2) \xrightarrow[\text{eval}]{} y_2 \neq \bot) \in \mathsf{Q}$, then add $(((s, F)) \xrightarrow[\text{gc}]{} \widetilde{F})$ to H, where $(s, F) = \mathsf{rev}(\widetilde{F}, \widetilde{X}_1, \widetilde{X}_2)$.

We will now describe the attack oracle-aided algorithm $\mathtt{SecRed}$ against the semantic security of (G^O, E^O, D^O). The input to $\mathtt{SecRed}$ is pair of challenge (pk, c) sampled under $\mathcal{E}^O$. Moreover, $\mathtt{SecRed}$ has oracle access to O as well as an adversary against $\dot{\mathcal{E}}^O$.

Description of $\mathtt{SecRed}^{\mathcal{A}, O}(pk, c)$:

1. Initialize $\mathsf{LocalE}^* = \varnothing$. For $i = [1, t]$, do the following: Run $E^O(pk, 0)$ and $E^O(pk, 1)$ and add all the resulting query-answer pairs to LocalE^*.

2. Set $\mathsf{H}^* \leftarrow \mathtt{ConstHelp}^O(\mathsf{LocalE}^*)$.

3. Return $b' \leftarrow \mathcal{A}(pk, \mathsf{H}^*, c)$.

We will now show that the following holds for both $b = 0$ and $b = 1$: The distribution $\mathrm{Dist1} = (pk, \mathsf{H}^*, c)$ is statistically close to $\mathrm{Dist2} = (\dot{pk}, \dot{c})$, where $(pk, sk) \leftarrow G^O(1^\kappa)$, $c \leftarrow E^O(pk, b)$, and $(\dot{pk}, *) \leftarrow \dot{G}^O(1^\kappa)$ and $\dot{c} \leftarrow \dot{E}^O(\dot{pk}, b)$. Also, H^* is sampled as in the execution of the security reduction $\mathtt{SecRed}^{\mathcal{A}, O}(pk, c)$. Let all the variables that appear below be sampled as in the above. First, it is easy to show that

$$\Delta((pk, \mathsf{H}^*), \dot{pk}) \leq \mathrm{poly}(\kappa) \times \frac{1}{2^{\kappa/2}} \leq \frac{1}{2^{\kappa/3}}.$$

Moreover, by Lemma 4.22 we have

$$\Delta(c, \dot{c}) \leq \frac{2q}{t} + \frac{1}{2^{\kappa/3}}. \tag{6}$$

Thus, $\mathtt{SecRed}^{\mathcal{A}, O}(pk, c)$ breaks the semantic security of (G^O, E^O, D^O) with probability at least $\gamma - \frac{2q}{t} - \frac{1}{2^{\kappa/4}}$.

$\square$

4.2 Removing Garbling Evaluation Queries from Decryption

In this section, we will prove Lemma 4.10. Namely, we will show the existence of a compilation procedure that compiles a PKE scheme $\mathcal{E} = (G^{f, \mathsf{gc}, \mathsf{gi}}, E^{f, \mathsf{gc}, \mathsf{gi}}, D^{f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}})$ into a new PKE scheme $\ddot{\mathcal{E}} = (\ddot{G}^{f, \mathsf{gc}, \mathsf{gi}}, \ddot{E}^{f, \mathsf{gc}, \mathsf{gi}}, \ddot{D}^{f, \mathsf{gc}, \mathsf{gi}})$ with correctness and security comparable to the original scheme $\mathcal{E}$, but where $\ddot{D}$ will not ask any eval queries.

Again, for simplicity we use the following convention where we keep the entire oracle O as a superscript to all the algorithms (G^O, E^O, D^O) as well as $(\ddot{G}^O, \ddot{E}^O, \ddot{D}^O)$ with the understanding that the actual oracle access is of the form given above. We also make the following assumption without loss of generality.

Assumption 4.26. We assume that all the algorithms (G, E, D) are in normal form (Definition 4.12). Also, we assume w.l.o.g. that the secret key outputted by G contains all the query-response pairs made by G.

4.2.1 Definitions

Here, we give some notation and definitions used in the compilation procedure.

Definition 4.27 (Query set). For an oracle-aided algorithm A^O we write $\text{Query}(A^O(x; r))$ when referring to the set of all queries asked during the execution of A^O on the input x and randomness r. We write $\text{Query}(A^O(x))$ to indicate the random variable formed by returning $\text{Query}(A^O(x; r))$ for $r \leftarrow \{0, 1\}^*$.

Definition 4.28 (Valid partial oracles). We say that a partial oracle O_1 is valid if for some $O_2 \in \text{Supp}(\mathbf{O})$: $O_1 \subseteq O_2$.

Definition 4.29 (Oracle consistency/sampling notation). We say a partial oracle O_1 is consistent with a set of query/response pairs S if $O_1 \cup \mathsf{S}$ is valid.

For a partial oracle O_1 and randomness r we say that (O_1, r) agrees with a public key pk if (1) $G^{O_1}(r) = (pk, *)$ and (2) all the queries in $\text{Query}(G^{O_1}(r))$ are defined in O_1. We say that (O_1, r) minimally agrees with pk if (1) (O_1, r) agrees with pk and (2) O_1 is defined only on the queries that occur during the execution and nothing more: namely, $O_1(qu)$ is defined iff $qu \in \text{Query}(G^{O_1}(r))$.

We let $\text{Partial}(pk, \mathsf{S})$ denote the set of all (O_1, r) where (1) (O_1, r) minimally agrees with pk and (2) O_1 agrees with S. We sometimes abuse notation and write $(O_1, sk) \leftarrow \text{Partial}(pk, \mathsf{S})$ to mean the following sampling: $(O_1, r) \leftarrow \text{Partial}(pk, \mathsf{S})$ and $(pk, sk) = G^{O_1}(r)$.

Definition 4.30 (Composed oracle). Given a partial oracle O_p and full oracle O (an oracle that is defined on all points in its domain) we define $O_p \lozenge O$ to be the *composed oracle* that uses O_p to reply if the corresponding query is defined there, and uses O otherwise. Note that $O_p \lozenge O$ is not necessarily in $\text{Supp}(O)$.

4.2.2 Compilation Procedure

Construction 4.31. The scheme $\ddot{\mathcal{E}} = (\ddot{G}, \ddot{E}, \ddot{D})$ is parameterized over two functions $\varepsilon = \varepsilon(\kappa)$ and $t = t(\kappa)$, which we will instantiate later.

- $\ddot{G}(1^\kappa)$: Do the following steps:

  1. Set $\mathsf{OrigG} = \varnothing$. Run $(pk, sk) \leftarrow G^O(1^\kappa)$, and add all query-answer pairs that are encountered during this execution to OrigG.

  2. Set $\mathsf{LearnG} = \varnothing$. While there exists a query $qu \notin \mathsf{LearnG}$ such that

  $$\Pr_{O' \leftarrow \mathbf{O}}[qu \in \text{Query}(G^{O'}(1^\kappa)) \mid pk, \mathsf{LearnG}] \geq \varepsilon,$$

  then choose the lexicographically first such qu and add $(qu \xrightarrow[T]{} O'(qu))$ to LearnG. Note that $T \in \{f, \mathsf{gc}, \mathsf{gi}\}$.

  3. Output $\ddot{pk} = (pk, \mathsf{LearnG})$ and $\ddot{sk} = sk$. (By Assumption 4.26, $\ddot{sk}$ contains OrigG.)

25

- $\ddot{E}(\ddot{pk}, b)$: Given $\ddot{pk} = (pk, \mathsf{LearnG})$ and $b \in \{0, 1\}$ do the following:

  1. Set $\mathsf{OrigE} = \varnothing$. Run $c \leftarrow E^O(pk, b)$ and add all the query-answer pairs that are observed during this execution to OrigE.

  2. **Generating helper set H for $\ddot{D}$:** Sample $t' \leftarrow [1, t]$. Set $\mathsf{S} = \mathsf{OrigE} \cup \mathsf{LearnG}$. For $i \in [1, t']$, do the following:

     (a) **Offline phase:** Sample $(\widehat{O}_i, \widehat{sk}_i) \leftarrow \mathrm{Partial}(pk, S)$.

     (b) **Semi-online phase:** Execute $D^{\widehat{O}_i \lozenge O[\mathsf{S}]}(\widehat{sk}_i, c)$ and add all query/response pairs made to the oracle O to the set S. Let $\widehat{\mathsf{OrigD}}_i$ be the set of all query-answer pairs made by this execution.

     After all iterations, set $\mathsf{H} := \mathtt{ConstHelp}(\widehat{\mathsf{OrigD}}, \mathsf{S})$, where $\widehat{\mathsf{OrigD}} = \widehat{\mathsf{OrigD}}_1 \cup \cdots \cup \widehat{\mathsf{OrigD}}_{t'}$.

  3. Output $\ddot{c} = (c, \mathsf{H})$.

- $\ddot{D}(\ddot{pk}, \ddot{sk}, \ddot{c})$: Given $\ddot{pk} = (pk, \mathsf{LearnG})$, $\ddot{sk}$, and $\ddot{c} = (c, \mathsf{H})$, Output $\widetilde{b} \leftarrow D^{O[\mathsf{H} \cup \mathsf{LearnG}]}(\ddot{sk}, c)$.

Query complexity of $\ddot{\mathcal{E}}$. The following lemma follows from the description of the compilation procedure of Construction 4.31.

Lemma 4.32. *Let q be as in Definition 4.16. Assuming $\varepsilon = \frac{1}{\mathrm{poly}(\kappa)}$ and $t = \mathrm{poly}(\kappa)$, all the algorithms of $\ddot{\mathcal{E}}^O$ make $q^{O(1)}$ queries. Concretely, the algorithm $\ddot{E}$ makes at most $\nu := 4tq^2$ queries.*

We note that by taking $\varepsilon = \frac{1}{\mathrm{poly}(\kappa)}$ the learning process of $\ddot{G}$ (i.e., for sampling LearnG) could be done by making a polynomial number of queries [BMG07].

4.2.3 Correctness and Security

In this section we give the correctness and security statements regarding the compiled scheme $\ddot{\mathcal{E}} = (\ddot{G}, \ddot{E}, \ddot{D})$ and prove them. By doing so, we complete the proof of Lemma 4.10. The following lemma gives the correctness bound for the compiled scheme. We prove the lemma in Section 4.2.4.

Lemma 4.33 (Correctness of $\ddot{\mathcal{E}}$)**.** *Suppose the original PKE scheme (G, E, D) is $(\frac{1}{2} + \delta)$-correct in the ideal model. The compiled PKE scheme $(\ddot{G}, \ddot{E}, \ddot{D})$ has at least $(\frac{1}{2} + \delta - \eta)$ correctness in the ideal model, where*

$$\eta = \frac{1}{2^{\kappa/5}} + \frac{2q}{t} + 3\varepsilon\nu t.$$

That is,

$$\Pr[\ddot{D}^O(sk, \ddot{c}) \neq b] \leq \tfrac{1}{2} - \delta + \eta \tag{7}$$

where the probability is taken over $O \leftarrow \mathbf{O}$, $(\ddot{pk}, sk) \leftarrow \ddot{G}^O(1^\kappa)$, $b \leftarrow \{0, 1\}$ and $\ddot{c} = (c, \mathsf{H}) \leftarrow \ddot{E}^O(\ddot{pk}, b)$. Here t and ε are the underlying parameters of the compilation procedure, and ν is defined in Lemma 4.32. In particular, for any constant $c > 0$ by taking $t = 2q^{c+2}$ and $\varepsilon = \frac{1}{q^{3c+8}}$, the compiled scheme $(\ddot{G}, \ddot{E}, \ddot{D})$ has at least $(\frac{1}{2} + \delta - 1/\kappa^c)$ correctness in the ideal model.

The following lemma gives the security loss for the compiled scheme $\ddot{\mathcal{E}} = (\ddot{G}, \ddot{E}, \ddot{D})$. We give the proof in Section 4.2.5.

Lemma 4.34 (Security of $\ddot{\mathcal{E}}$). *Let p be an arbitrary polynomial which satisfies*

$$8tq^2\varepsilon + \tfrac{1}{2^{\kappa/2-1}} \le \tfrac{1}{p}.$$

There exists a polynomial-query algorithm SecRed *that satisfies the following. For any adversary $\mathcal{A}$ that breaks the semantic security of $(\ddot{G}^O, \ddot{E}^O, \ddot{D}^O)$ in the ideal model O with probability at least γ, the adversary* SecRed$^{\mathcal{A},O}$ *breaks the semantic security of (G^O, E^O, D^O) with probability at least $\gamma - \beta$ where*

$$\beta = t \cdot \left(\tfrac{1}{p-1} + 4tq^2\varepsilon + \tfrac{1}{q^{2c+4}} + \tfrac{1}{2^{\kappa/2-1}} \right).$$

In particular, for any constant c by taking $t = 2q^{c+2}$ and $\varepsilon = \tfrac{1}{q^{3c+8}}$ we will have the following: For any adversary $\mathcal{A}$ breaking the semantic security of $(\ddot{G}^O, \ddot{E}^O, \ddot{D}^O)$ in the ideal model with probability at least γ, the (polynomial-query) adversary SecRed$^{\mathcal{A},O}$ *breaks the semantic security of (G^O, E^O, D^O) with probability at least $\gamma - \tfrac{22}{\kappa^{2+c}}$.*

4.2.4 Proving Lemma 4.33: Correctness for $\ddot{\mathcal{E}}$

In this section, we will prove Lemma 4.33, which states that $\ddot{\mathcal{E}} = (\ddot{G}, \ddot{E}, \ddot{D})$ is still a correct PKE after having removed the eval queries from D.

Proof roadmap. Consider the variables sampled in Lemma 4.33. Parse $\ddot{pk} = (pk, \mathsf{LearnG})$. Notice that $\ddot{D}^O(sk, \ddot{c})$ simply runs $D^{O[\mathsf{H}\cup\mathsf{LearnG}]}(sk, c)$.[6] Thus, to prove Lemma 4.33 we prove the following lemma, which shows that the probability that $D^{O[\mathsf{H}\cup\mathsf{LearnG}]}(sk, c) \ne D^O(sk, c)$ is small.

Lemma 4.35. *For any $b \in \{0, 1\}$*

$$\Pr[D^{O[\mathsf{H}\cup\mathsf{LearnG}]}(sk, c) \ne D^O(sk, c)] \le \frac{1}{2^{\kappa/5}} + \frac{2q}{t} + 3\varepsilon\nu t, \tag{8}$$

where $O \leftarrow \mathbf{O}$, $((pk, \mathsf{LearnG}), sk) \leftarrow \ddot{G}^O(1^\kappa)$ and $(c, \mathsf{H}) \leftarrow \ddot{E}^O(\ddot{pk}, b)$.

We first show how to derive Lemma 4.33 from Lemma 4.35.

Proof of Lemma 4.33. Let all the variables below be sampled as in Lemma 4.33. We have

$$\Pr[\ddot{D}^O(sk, \ddot{c}) \ne b] = \Pr[D^{O[\mathsf{H}\cup\mathsf{LearnG}]}(sk, c) \ne b]$$
$$\le \Pr[D^O(sk, c) \ne b] + \eta$$
$$\le \frac{1}{2} - \delta + \eta,$$

where the first inequality follows from Lemma 4.35 and the second inequality follows from the correctness bound assumed for the original scheme (G, E, D). Notice that for the second inequality we made use of the fact that (pk, sk, c) is identically distributed to a random triple (pk', sk', c') produced under the original scheme: That is, $(pk', sk') \leftarrow G^O(1^\kappa)$ and $c' \leftarrow E^O(pk', b)$. $\qquad\square$

To prove Lemma 4.35 we define the following two undesirable events that are defined based on the variables sampled in Lemma 4.35. We will then show that the probability of the event in Lemma 4.35 is at most the probability of the union of the following two events.

---

[6]Note that a secret key under the compiled scheme has the same format as the one under the original scheme.

Definition 4.36 (Event Miss)**.** Let OrigD be the set of query/response pairs for the execution of $D^O(sk, c)$. The event Miss holds if there is a query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} ?) \in \mathsf{OrigD}$ such that:

1. $\mathsf{eval}(\widetilde{F}, \widetilde{X}) = y \neq \bot$, and

2. $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} *) \notin \mathsf{H}$ and $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \notin \mathsf{H}$, and

3. $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{OrigE}$ for some s and F.

Definition 4.37 (Event Surprise)**.** There exists a query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} ?) \in \mathsf{OrigD}$ such that:

- $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \notin \mathsf{OrigG} \cup \mathsf{LearnG} \cup \mathsf{OrigE}$

The proof of Lemma 4.35 now follows from the following three lemmas.

Lemma 4.38. *For any $b \in \{0, 1\}$*

$$\Pr[D^{O[\mathsf{H} \cup \mathsf{LearnG}]}(sk, c) \neq D^O(sk, c)] \leq \Pr[\mathsf{Miss} \vee \mathsf{Surprise}],$$

where $O \leftarrow \mathbf{O}$, $((pk, \mathsf{LearnG}), sk) \leftarrow \ddot{G}^O(1^\kappa)$ and $(c, \mathsf{H}) \leftarrow \ddot{E}^O(\ddot{pk}, b)$.

Lemma 4.39. $\Pr[\mathsf{Surprise}] \leq \frac{1}{2^{\kappa/2}}$*, where $O \leftarrow \mathbf{O}$, $((pk, \mathsf{LearnG}), sk) \leftarrow \ddot{G}^O(1^\kappa)$ and $(c, \mathsf{H}) \leftarrow \ddot{E}^O(\ddot{pk}, b)$.*

Lemma 4.40. $\Pr[\mathsf{Miss}] \leq \frac{1}{2^{\kappa/4}} + \frac{2q}{t} + 3\varepsilon\nu t$*, where $O \leftarrow \mathbf{O}$, $((pk, \mathsf{LearnG}), sk) \leftarrow \ddot{G}^O(1^\kappa)$ and $(c, \mathsf{H}) \leftarrow \ddot{E}^O(\ddot{pk}, b)$.*

Proof of Lemma 4.35. The proof follows directly from Lemmas 4.38, 4.39 and 4.40. □

We first prove Lemma 4.38.

Proof of Lemma 4.38. Let Bad be the event that $D^{O[\mathsf{H} \cup \mathsf{LearnG}]}(sk, c) \neq D^O(sk, c)$. We show if $\overline{\mathsf{Miss}}$ and $\overline{\mathsf{Surprise}}$ hold, then $\overline{\mathsf{Bad}}$ must necessarily hold, hence proving the lemma. Suppose $\overline{\mathsf{Miss}} \wedge \overline{\mathsf{Surprise}}$ hold and to the contrary that Bad holds and consider the first query qu on which the two executions $D^{O[\mathsf{H} \cup \mathsf{LearnG}]}(sk, c)$ and $D^O(sk, c)$ are different. Since all non-eval queries are handled the same in both executions, we must have $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} ?)$ for some $\widetilde{F}$ and $\widetilde{X}$. Now since $\overline{\mathsf{Surprise}}$ holds we have $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{OrigG} \cup \mathsf{LearnG} \cup \mathsf{OrigE}$. On the other hand, by Assumption 4.26 and Definition 4.12 we cannot have $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{OrigG}$ because the secret key sk contains OrigG and thus by the normal-form restriction the query qu will not be issued to the oracle in both executions. Also, we cannot have $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{LearnG}$ because in that case we will have $O[\mathsf{H} \cup \mathsf{LearnG}](qu) = O(qu)$. Thus, $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{OrigE}$. Now a simple inspection shows that Miss holds, a contradiction. □

Proof of Lemma 4.39. As in the proof of Lemma 4.23 we can easily show that whenever the event Surprise happens, we can win in the sense of Lemma 3.1. The bound now follows from Lemma 3.1. We omit the details. □

Proving Lemma 4.40. We now focus on proving Lemma 4.40. We first start by giving a general lemma that we will use during the proof.

Lemma 4.41 (Hitting unlearned queries). *Let $\ddot{pk} = (pk, \mathsf{LearnG}) \leftarrow \ddot{G}(1^\kappa)$ and let OrigG be the query-answer pairs asked by G during $\ddot{G}$. Let $A^{f,\mathsf{gc},\mathsf{gi}}$ be an arbitrary randomized q-query algorithm that takes as input $(1^\kappa, \ddot{pk})$ and define $\mathsf{Q_A}$ to be the query-answer pairs asked by A. Then:*

$$\Pr_{O,A}[\exists qu \in \mathsf{Q_A} \cap \mathsf{OrigG} \setminus \mathsf{LearnG}] \leq q\varepsilon$$

Proof. Consider an alternative experiment where we replace A with another algorithm A' that, instead of issuing its queries to the oracle, would instead simulate the answers to these queries consistently with LearnG. That is, if the answer of a query is included in LearnG then it will use this answer, otherwise it will sample a new uniformly random answer (consistently with LearnG) and store this answer in $\mathsf{Q_{A'}}$. Note that the execution of $A^{f,\mathsf{gc},\mathsf{gi}}(1^\kappa, \ddot{pk})$ and $A'(1^\kappa, \ddot{pk})$ would proceed statistically the same up until the point when a query qu is asked such that $qu \in \mathsf{OrigG} \setminus \mathsf{LearnG}$. In that case, A would answer consistently with OrigG, whilst A' would use a uniformly random answer. Since these two executions proceed the same up until this bad event, the probability of this event happening is the same in both experiments, so we will bound this event in the A' experiment.

Since the query-answer pairs of A' depend only on pk and LearnG (and hence $\mathsf{Q_{A'}}$ is independent of OrigG given LearnG and pk), we can in fact sample LearnG first conditioned on pk followed by sampling $\mathsf{Q_{A'}}$ conditioned on LearnG and pk. Finally, we will sample OrigG conditioned on LearnG and pk. Now we compute the following probability that any query qu in $\mathsf{Q_{A'}}$ is sampled in $\mathsf{OrigG} \setminus \mathsf{LearnG}$ as follows:

$$
\begin{aligned}
\Pr_{O,A'}[\exists\, qu \in \mathsf{Q_{A'}} \cap (\mathsf{OrigG} \setminus \mathsf{LearnG})] &\leq \sum_{qu \in \mathsf{Q_{A'}}} \Pr[qu \in \mathsf{Q_{A'}} \wedge qu \in (\mathsf{OrigG} \setminus \mathsf{LearnG})] \\
&\leq \sum_{qu \in \mathsf{Q_{A'}}} \Pr[qu \in \mathsf{OrigG} \setminus \mathsf{LearnG}] \\
&\leq \sum_{qu \in \mathsf{Q_{A'}}} \Pr[qu \in \mathrm{Query}(G^O) \mid pk, \mathsf{LearnG}] \\
&\leq |\mathsf{Q_{A'}}| \times \varepsilon \leq q\varepsilon
\end{aligned}
$$

$\square$

We give Definitions 4.42 and 4.43 below to describe more compactly the random variables sampled during the compilation. We will then give some notation and terminology. We will then give a few lemmas from which Lemma 4.40 is easily derived. We will then prove those lemmas.

Definition 4.42. Procedure $\mathrm{DistGen}(pk, c, \mathsf{S})$

- **Offline phase:** Sample $(\widehat{\mathsf{OrigG}}, \widehat{sk}) \leftarrow \mathrm{Partial}(pk, \mathsf{S})$

- **Semi-online phase:** Execute $D^{\widehat{\mathsf{OrigG}} \Diamond O[\mathsf{S}]}(\widehat{sk}, c)$ and update S by adding all query/response pairs made to the oracle O. Let $\widehat{\mathsf{OrigD}}$ be the set of all query/response pairs during this decryption execution.

- **Return** $(\mathsf{S}, \widehat{\mathsf{OrigG}}, \widehat{sk}, \widehat{\mathsf{OrigD}})$

Definition 4.43. Procedure $\text{DistGen}_1(1^\kappa)$:

- **Output:** a tuple

$$\mathbf{Dist} := (\mathsf{OrigG}, sk, pk, \mathsf{LearnG}, \mathsf{OrigE}, \mathsf{OrigD}),$$

  as well as a sequence of tuples of the form

$$\mathbf{Dist}_i := (\widehat{\mathsf{OrigG}}_i, \widehat{sk}_i, pk, \mathsf{LearnG}, \mathsf{OrigE}, \widehat{\mathsf{OrigD}}_i).$$

- **Operations**

  1. Sample $(pk, sk) \leftarrow G^O(1^\kappa)$ and let OrigG be the set of all query/response pairs.
  2. Sample LearnG as in $\ddot{G}$.
  3. Sample $c \leftarrow E^O(pk, b)$ and let OrigE be the set of all query/response pairs
  4. Let OrigD be the set of all query/response pairs during the execution of $D^O(sk, c)$.
  5. Let $\mathsf{S}_1 = \mathsf{LearnG} \cup \mathsf{OrigE}$. For $i \geq 1$ sample

$$(\mathsf{S}_{i+1}, \widehat{\mathsf{OrigG}}_i, \widehat{sk}_i, \widehat{\mathsf{OrigD}}_i) \leftarrow \text{DistGen}(pk, c, \mathsf{S}_i).$$

Invalid samples. We say that a value

$$(\widehat{\mathsf{OrigG}}_i, \widehat{sk}_i, pk, \mathsf{LearnG}, \mathsf{OrigE}, \widehat{\mathsf{OrigD}}_i)$$

of the random variable $\mathbf{Dist}_i$ is *invalid* if

$$(\mathsf{LearnG} \cup \widehat{\mathsf{OrigG}}_i \cup \mathsf{OrigE} \cup \widehat{\mathsf{OrigD}}_i)$$

is not a valid partial oracle (Definition 4.28). We let $\mathsf{invalid}_i$ be the event that $\mathbf{Dist}_i$ is invalid. We also define

$$\mathsf{invalid} = \mathsf{invalid}_1 \vee \cdots \vee \mathsf{invalid}_{t'} \tag{9}$$

Let $\widehat{\mathbf{Dist}}_{i,\mathrm{good}}$ to be the induced distribution on $\widehat{\mathbf{Dist}}_i$ conditioned on $\overline{\mathsf{invalid}_i}$.

Lemma 4.44. *Let t' be as in the compilation procedure. The distributions $\mathbf{Dist}$ and $\widehat{\mathbf{Dist}}_{t'+1,\mathrm{good}}$ are identically distributed.*

Proof. The proof can be easily done by inspection. $\qquad\square$

The proof of Lemma 4.40 now follows from the following two lemmas.

Lemma 4.45. *We have $\Pr[\mathsf{Miss} \mid \overline{\mathsf{invalid}}] \leq \frac{2q}{t}$, where the probability is taken over the variables sampled by $\text{DistGen}_1(1^\kappa)$ (see Definition 4.43).*

Lemma 4.46. *We have $\Pr[\mathsf{invalid}] \leq \frac{1}{2^{\kappa/4}} + 3\varepsilon\nu t$, where the probability is taken over the variables sampled by $\text{DistGen}_1(1^\kappa)$.*

Proof of Lemma 4.40. The proof follows immediately from Lemmas 4.45 and Lemma 4.46.

$\qquad\square$

Proving Lemmas 4.45 and Lemma 4.46.

Proof of Lemma 4.45. The proof is done similarly to the proof of Lemma 4.24. Relying on the notation (i, d)-*match* defined in the proof of Lemma 4.24, as in that proof we break the event Miss into smaller events $\mathsf{Miss}_{h,d}$, for $h \in [q]$ and $d \in \{0, 1\}$, defined as follows:

Event $\mathsf{Miss}_{h,d}$. The event that for some $\widetilde{F}$ that $((*, *) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigE}$, both the following hold:

1. OrigD contains an (h, d)-match for $\widetilde{F}$;

2. none of the sets $\widehat{\mathsf{OrigD}}_1, \ldots, \widehat{\mathsf{OrigD}}_{t'}$ do contain an (h, d)-match for $\widetilde{F}$.

As in the proof of Lemma 4.24, we can show that if Miss holds then $\mathsf{Miss}_{h,d}$ must hold, for some $h \in [q]$ and $d \in \{0, 1\}$. In the sequel we show

$$\Pr[\mathsf{Miss}_{h,d} \mid \overline{\mathsf{invalid}}] \leq \frac{1}{t},$$

which will imply the desired bound for the value of $\Pr[\mathsf{Miss} \mid \overline{\mathsf{invalid}}]$.

Fix $h \in [q]$ and $d \in \{0, 1\}$ for which we want to bound the event $\mathsf{Miss}_{h,d}$. Some notation first. For a sequence of tuples $(\mathrm{Dist}_1, \mathrm{Dist}_2, \ldots,)$ sampled from $(\mathbf{Dist}_1, \mathbf{Dist}_2, \ldots,)$ given in Definition 4.43, we define a random variable $\mathbf{First}_{h,d}$ which takes a value in $\{1, \ldots, t+1\} \cup \{\bot\}$ as follows:

- $\mathbf{First}_{h,d} = i$ for $i \in [1, t+1]$ if

  1. $\widehat{\mathsf{OrigD}}_i$ contains an (h, d)-match for $\widetilde{F}$;
  2. none of the sets $\widehat{\mathsf{OrigD}}_1, \ldots, \widehat{\mathsf{OrigD}}_{i-1}$ do contain an (h, d)-match for $\widetilde{F}$.

- $\mathbf{First}_{h,d} = \bot$ if none of the sets $\widehat{\mathsf{OrigD}}_1, \ldots, \widehat{\mathsf{OrigD}}_{t+1}$ do contain an (h, d)-match for $\widetilde{F}$.

Now applying Lemma 4.44 we have that, for any (h, d):

$$\Pr[\mathsf{Miss}_{h,d} \mid \overline{\mathsf{invalid}}] = \Pr[\mathbf{First}_{h,d} = t' + 1]$$

where the random variable $\mathbf{First}_{h,d}$ is defined over the sequence $(\mathrm{Dist}_1, \mathrm{Dist}_2, \ldots, \mathrm{Dist}_{t+1})$, where the i'th tuple is sampled as $\mathrm{Dist}_i \leftarrow \widehat{\mathbf{Dist}}_{i,\mathsf{good}}$.

The random variable $\mathbf{First}_{h,d}$ takes a value in $[1, t+1] \cup \{\bot\}$ according to some arbitrary distribution. However, since t' is chosen uniformly at random from $[1, t]$ we have:

$$\Pr[\mathbf{First}_{h,d} = t' + 1] \leq \frac{1}{t}$$

$\square$

Proof of Lemma 4.46. We show that for all $i \in [t]$,

$$\Pr[\mathsf{invalid}_i] \leq \frac{1}{2^{\kappa/3}} + 3\varepsilon\nu. \tag{10}$$

The proof of the lemma then follows by the union bound. Fix i. We want to bound the probability that the tuple Tup_i sampled as

$$\mathsf{Tup_i} = (\mathsf{LearnG} \cup \widehat{\mathsf{OrigG}}_i \cup \mathsf{OrigE} \cup \widehat{\mathsf{OrigD}}_i) \leftarrow \mathbf{Dist}_i$$

is invalid. Note that by design the set

$$\mathsf{W} = \mathsf{LearnG} \cup \widehat{\mathsf{OrigG}}_i \cup \mathsf{OrigE}$$

makes up a valid partial oracle. This is because $\widehat{\mathsf{OrigG}}_i$ is chosen in a manner consistent with $\mathsf{LearnG} \cup \mathsf{OrigE}$. We first claim that if Tup_i is invalid then $\widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i$ must be invalid. To prove this, suppose $\widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i$ is valid and suppose to the contrary that Tup_i is invalid. Noting that W is valid and also, by assumption, that $\widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i$ is valid, we consider all possible cases for the inconsistency of Tup_i:

1. For some $T \in \{f, \mathsf{gc}, \mathsf{gi}\}$ we have $(qu \xrightarrow[T]{} ans_1) \in \widehat{\mathsf{OrigD}}_i$ and $(qu \xrightarrow[T]{} ans_2) \in \mathsf{LearnG} \cup \mathsf{OrigE}$ and $ans_1 \neq ans_2$: Since $(qu \xrightarrow[T]{} ans_1) \in \widehat{\mathsf{OrigD}}_i$ this means that either $(qu \xrightarrow[T]{} ans_1) \in \widehat{\mathsf{OrigG}}_i$ or $O(qu) = ans_1$. The first case contradicts the fact that $\widehat{\mathsf{OrigG}}_i$ agrees with $\mathsf{LearnG} \cup \mathsf{OrigE}$ and the second case contradicts the fact that $\mathsf{LearnG} \cup \mathsf{OrigE}$ agree with O.

2. For some $T \in \{f, \mathsf{gc}, \mathsf{gi}\}$ we have $(qu_1 \xrightarrow[T]{} ans) \in \widehat{\mathsf{OrigD}}_i$ and $(qu_2 \xrightarrow[T]{} ans) \in \mathsf{LearnG} \cup \mathsf{OrigE}$ and $qu_1 \neq qu_2$: Like above, it can be proved this case cannot happen either.

3. For some and $\widetilde{F}$ and $\widetilde{X}$, $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} \bot) \in \widehat{\mathsf{OrigD}}_i$ and qu is inconsistent with $\mathsf{LearnG} \cup \mathsf{OrigE}$: In order to have inconsistency, for some s we must at least have $((s, F) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{LearnG} \cup \mathsf{OrigE}$. Now by design in that case for any S we will have $O[\mathsf{LearnG} \cup \mathsf{OrigE} \cup \mathsf{S}](qu) = O(qu)$ and thus we will also have $O(qu) = \bot$. Thus, there cannot be any inconsistencies.

Thus, to prove Equation 10 it suffices to show

$$\Pr[\widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i \text{ is invalid}] \leq \frac{1}{2^{\kappa/3}} + 3\varepsilon\nu. \tag{11}$$

We now focus on proving Equation 11. As the first observation, we note that we cannot have any input inconsistencies in $\widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i$: namely, a query qu for which we have $(qu \xrightarrow[T]{} an_1) \in \widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i$ and $(qu \xrightarrow[T]{} an_2) \in \widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i$ for $an_1 \neq an_2$. The reason that this cannot happen is because of the normal-form restriction and the definition of composed oracles (Definition 4.30). Thus, it suffices to bound the probability of output collisions in $\widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i$. That is, we consider the inconsistencies between $\widehat{\mathsf{OrigG}}_i$ and $\widehat{\mathsf{OrigD}}_i$ that could happen as a result of having two different queries that have the same answer. We consider all possible case that an output collision may occur.

1. Col_1: There exists $((s_1, F_1) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \widehat{\mathsf{OrigG}}_i$ such that $((s, F) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \widehat{\mathsf{OrigD}}_i \setminus \widehat{\mathsf{OrigG}}_i$ and that $(s, F) \neq (s_1, F_1)$.

2. Col_2: There exists $((s_1, x_1) \underset{\mathsf{gi}}{\to} \widetilde{x}) \in \widehat{\mathsf{OrigG}}_i$ such that $((s, x) \underset{\mathsf{gi}}{\to} \widetilde{x}) \in \widehat{\mathsf{OrigD}}_i \setminus \widehat{\mathsf{OrigG}}_i$ and that $(s, x) \neq (s_1, x_1)$.

3. Col_3: There exists $((\widetilde{F}, \widetilde{X}) \underset{\mathsf{eval}}{\longrightarrow} \bot) \in \widehat{\mathsf{OrigD}}_i$ such that $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$ and $((s, F) \underset{\mathsf{gc}}{\to} \widetilde{F}) \in \widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i$ for some s and F.

We mention that the event Col_3 may not necessarily cause an inconsistency by itself, but if some query $qu = ((\widetilde{F}, \widetilde{X}) \underset{\mathsf{eval}}{\longrightarrow} \bot) \in \widehat{\mathsf{OrigD}}_i$ does cause an inconsistency in $\widehat{\mathsf{OrigG}}_i \cup \widehat{\mathsf{OrigD}}_i$, then the conditions in Case 3 must hold.

We will now show that each of the events Col_1, Col_2 and Col_3 may occur with probability at most $\frac{1}{2^{\kappa/2}} + \varepsilon\nu$, proving Equation 11.

Bounding the event Col_1: We show

$$\Pr[\text{Col}_1] \leq \tfrac{1}{2^{\kappa/2}} + \nu\varepsilon. \tag{12}$$

To prove this, first note that Col_1 implies that $\widetilde{F}$ is a valid garbled circuit, and that we indeed have $\mathsf{gc}(s, F) = \widetilde{F}$. This is because the real oracle O agrees with $\widehat{\mathsf{OrigD}}_i \setminus \widehat{\mathsf{OrigG}}_i$. We claim

$$\Pr[((s, F) \underset{\mathsf{gc}}{\to} \widetilde{F}) \in \mathsf{OrigG}] \geq 1 - \tfrac{1}{2^{\kappa/2}}. \tag{13}$$

We first show how to derive Equation 12 from Equation 13. Notice that $((s, F) \underset{\mathsf{gc}}{\to} \widetilde{F}) \notin \mathsf{LearnG}$. This is because $\widehat{\mathsf{OrigG}}_i$ agrees with LearnG and so we could not have had $((s_1, F_1) \underset{\mathsf{gc}}{\to} \widetilde{F}) \in \widehat{\mathsf{OrigG}}_i$. Thus, from Equation 13 we have $\Pr[((s, F) \underset{\mathsf{gc}}{\to} \widetilde{F}) \in \mathsf{OrigG} \setminus \mathsf{LearnG}] \geq 1 - \tfrac{1}{2^{\kappa/2}}$. Now recall that, by the definition of the event Col_1 we have $((s, F) \underset{\mathsf{gc}}{\to} \widetilde{F}) \in \widehat{\mathsf{OrigD}}_i$. Now applying Equation 13 along with Lemma 4.41 and Lemma 4.32, Equation 12 is established.

We now prove Equation 13. Let Surp' be the event that $((s, F) \underset{\mathsf{gc}}{\to} \widetilde{F}) \notin \mathsf{OrigG}$. We will show $\Pr[\mathsf{Surp}'] \leq \tfrac{1}{2^{\kappa/2}}$ by reducing it to Lemma 3.1. First, note that $\widetilde{F}$ is valid garbed circuit under gc, and that it is constructible using $\widehat{\mathsf{OrigG}}_i$ and that

$$(\widehat{\mathsf{OrigG}}_i, \widehat{sk}_i) \leftarrow (\mathbf{O}, \mathbf{sk} \mid pk, \mathsf{LearnG} \cup \mathsf{S}_{i-1}).$$

for some set S_i which completely agrees with O and so we cannot have the query/response $((*, *) \underset{\mathsf{gc}}{\to} \widetilde{F}) \in \mathsf{S}_i$. Thus, the occurrence of the event Surp' implies we could forge a valid $\widetilde{F}$ "on the fly" and so by Lemma 3.1 we have $\Pr[\mathsf{Surp}'] \leq \tfrac{1}{2^{\kappa/2}}$.

Bounding the event Col_2. As in Col_1, we can show that $\Pr[\text{Col}_2] \leq \tfrac{1}{2^{\kappa/2}} + \varepsilon\nu$.

Bounding the event Col_3 . Suppose Col_3 holds: Namely, there exists $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} \bot) \in \widehat{\text{OrigD}}_i$ such that $\text{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$ and $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \widehat{\text{OrigG}}_i \cup \widehat{\text{OrigD}}_i$ for some s and F. First, we argue that we cannot have $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \widehat{\text{OrigG}}_i$. This is because, otherwise, the query/response $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F})$ will be included in $\widehat{sk}_i$, and so by Assumption 4.12, during $D^{\widehat{O}_i \lozenge \overline{O}[S]}(\widehat{sk}_i, c)$ the query $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} ?)$ will never be asked.

Now since $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \notin \widehat{\text{OrigG}}_i$, we have $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \widehat{\text{OrigD}}_i$ and since $\widehat{\text{OrigD}}_i \neq \widehat{\text{OrigG}}_i$ agrees with O on all non-eval queries we have $\text{gc}(s, F) = \widetilde{F}$. Thus, $\widetilde{F}$ is a valid garbled circuit.

Thus, it suffices to bound the probability of the following event that we call Col': there exists $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} \bot) \in \widehat{\text{OrigD}}_i$ such that $\widetilde{F}$ is valid and $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \widehat{\text{OrigD}}_i$ for some s and F. Notice that by the normal-form restriction, the event Col' happens only when the query/response $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} *) \in \widehat{\text{OrigD}}_i$ is added to $\widehat{\text{OrigD}}_i$ *before* the query/response $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F})$ is created. Now exactly as in Col_1 we can show that the probability that Col' happens is at most $\frac{1}{2^{\kappa/2}} + \nu\varepsilon$.

The proof of Lemma 4.46 is now complete.

$\square$

4.2.5 Proving Lemma 4.34: Security for $\ddot{\mathcal{E}}$

In this section, we will prove Lemma 4.34, which states that the security of the compiled scheme $\ddot{\mathcal{E}} = (\ddot{G}, \ddot{E}, \ddot{D})$ can be reduced to the security of $\mathcal{E}$.

First, note that for a random public key $\ddot{pk} = (pk, \text{LearnG})$ produced as $((pk, \text{LearnG}), *) \leftarrow \ddot{G}(1^\kappa)$, the variable LearnG can be perfectly sampled solely based on pk and that pk is identically distributed to a random public key produced under the original $G(1^\kappa)$. With this in mind, given an attacker $\mathcal{A}$ against $\mathcal{E}$, we give a poly-query attacker $\text{SecRed}^{\mathcal{A}, O}$ that, given (pk, LearnG, c), where (pk, c) are produced under the original scheme (G^O, E^O, D^O) will sample H^* then run $\mathcal{A}(pk, \text{LearnG}, c, \mathsf{H}^*)$. We will then show that for both $b \in \{0, 1\}$ the following holds: the joint distribution of $(pk, \text{LearnG}, c, \mathsf{H})$ is close to $(pk, \text{LearnG}, c, \mathsf{H}^*)$, where $\ddot{pk} = (pk, \text{LearnG}) \leftarrow \dot{G}(1^\kappa)$, $(c, \mathsf{H}) \leftarrow \ddot{E}^O(\ddot{pk}, b)$ and the hint list H^* is generated during the security reduction $\text{SecRed}^{\mathcal{A}, O}(pk, \text{LearnG}, c)$.

Challenges and overview of techniques. Let us first start with a high level description of the main challenges and how we overcome them. Consider a random public key $\ddot{pk} = (pk, \text{LearnG}) \leftarrow \ddot{G}(1^\kappa)$ and a random ciphertext $(c, \mathsf{H}) \leftarrow \ddot{E}^O(\ddot{pk}, b)$ for some unknown plaintext bit b. Notice that c is identically distributed to a random ciphertext outputted by $E^O(pk, b)$. Thus, the main challenge is to sample a simulated H^* which is close enough to the real H, and to do this solely based on (pk, LearnG, c), and in particular without knowing OrigE, the set of query/response pairs used to sample $c \leftarrow E^O(pk, b)$.

Recall the random variables $(\widehat{O}_i, \widehat{sk}_i)$ sampled during the real execution of $\ddot{E}^O(\ddot{pk}, b)$. In order to be able to sample H^* we should be able to sample (O_i^*, sk_i^*) that is close to $(\widehat{O}_i, \widehat{sk}_i)$. The main challenge in doing so is that $(\widehat{O}_i, \widehat{sk}_i)$ is conditioned on $\text{LearnG} \cup \text{OrigE}$, and OrigE is not available to $\text{SecRed}^O(pk, \text{LearnG}, c)$. We will show that if we sample (O_i^*, sk_i^*) based on (pk, LearnG), then (O_i^*, sk_i^*) and $(\widehat{O}_i, \widehat{sk}_i)$ will be reasonably close enough, assuming that LearnG has captured all ε heavy queries for small enough values of ε.

Now assuming that (O_i^*, sk_i^*) is close to $(\widehat{O}_i, \widehat{sk}_i)$, letting $S = \mathsf{LearnG} \cup \mathsf{OrigE}$, the algorithm $\mathsf{SecRed}^{\mathcal{A},O}(pk, \mathsf{LearnG}, c)$ needs to be able to run the execution of $D^{O_i^* \lozenge O[S]}(sk_i^*, c)$ in order to be able to simulate that of $D^{\widehat{O}_i \lozenge O[S]}(\widehat{sk}_i, c)$. For simplicity of exposition, in the following assume that (O_i^*, sk_i^*) and $(\widehat{O}_i, \widehat{sk}_i)$ are identically distributed, and so we show how $\mathsf{SecRed}^{\mathcal{A},O}$ simulates the execution of $D^{\widehat{O}_i \lozenge O[S]}(\widehat{sk}_i, c)$, assuming that it knows $(\widehat{sk}_i, \widehat{O}_i)$. The main challenge in this simulation is that $\mathsf{SecRed}^{\mathcal{A},O}(pk, \mathsf{LearnG}, c)$ does not know what S is. A simple way around for $\mathsf{SecRed}^{\mathcal{A},O}$ is to use its full access to the oracle eval to make up for its lack of access to S. That is, $\mathsf{SecRed}^{\mathcal{A},O}$ would perform $D^{\widehat{O}_i \lozenge O}(\widehat{sk}_i, c)$. But this idea runs into the following problem: The oracle $\widehat{O}_i \lozenge O[S]$ may reply to an encountered query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow{\mathsf{eval}} ?)$ with $\bot$, but we may have that $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$, in which case we will have $\widehat{O}_i \lozenge O[S](qu) = \bot$ but $\widehat{O}_i \lozenge O(qu) \neq \bot$. To overcome this problem, it seems that the algorithm $\mathsf{SecRed}^{\mathcal{A},O}$, during the execution of $D^{O_i^* \lozenge S}(sk_i^*, c)$ should only resort to its eval oracle for a query $\mathsf{eval}(\widetilde{F}, \widetilde{X})$ if indeed $O[S](qu) \neq \bot$. But checking for this seems difficult, mainly because $\mathsf{SecRed}^{\mathcal{A},O}(pk, \mathsf{LearnG}, c)$ does not know what S is. Let us call this event *bad*. We resolve this bad event as follows. Notice that whenever *bad* happens for a query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow{\mathsf{eval}} ?)$, then w.h.p. $\widetilde{F}$ must have been generated in OrigG as a result of a gc query. Moreover, if this $\widetilde{F}$ appears during the decryption execution of $D^{\widehat{O}_i \lozenge O}(\widehat{sk}_i, c)$, it should appear during many other random executions of the same kind. Thus, we perform the following experiment to collect all such $\widetilde{F}$: Sample OrigE' by running $c' \leftarrow E^O(pk, b')$ for both $b' \in \{0, 1\}$; let $S' = \mathsf{LearnG} \cup \mathsf{OrigE}'$; and sample $(\widehat{O}, \widehat{sk})$ consistent with $(pk, \mathsf{LearnG} \cup \mathsf{OrigE}')$ and run the execution of $D^{\widehat{O} \lozenge O[S']}(sk', c')$ and record in a set HidG all $\widetilde{F}$ such that we see a query $((\widetilde{F}, \widetilde{X}) \xrightarrow{\mathsf{eval}} ?)$ for which we have $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$ but $\widehat{O} \lozenge O[S'](\widetilde{F}, \widetilde{X}) = \bot$.

Then, in the simulated execution of $D^{\widehat{O}_i \lozenge O}(\widehat{sk}_i, c)$ whenever $\mathsf{SecRed}^{\mathcal{A},O}$ sees a query $((\widetilde{F}, \widetilde{x}) \xrightarrow{\mathsf{eval}} ?)$, it will forward this query to O only if $\widetilde{F} \notin \mathsf{HidG}$. If $\widetilde{F} \in \mathsf{HidG}$ then $\mathsf{SecRed}^{\mathcal{A},O}$ replies to that query with $\bot$.

Procedure $\mathsf{SecRed}^{\mathcal{A},O}(pk, \mathsf{LearnG}, c)$ with output bit b':

1. **Finding hidden garbled circuits in OrigG:** Let $\mathsf{HidG} := \emptyset$. Do the following q^τ times.

   (a) For both $b \in \{0, 1\}$ do the following:
      i. Sample $c' \leftarrow E^O(pk, b)$ and let OrigE' be the set of all query/response pairs during this execution.
      ii. Let $\widehat{\mathsf{OrigD}'_1}, \ldots, \widehat{\mathsf{OrigD}'_t}$ be sampled as in $\mathrm{DistGen}_1(pk, \mathsf{LearnG}, \mathsf{OrigE}')$ and let
      $$\widehat{\mathsf{OrigD}} = \widehat{\mathsf{OrigD}'_1} \cup \cdots \cup \widehat{\mathsf{OrigD}'_t}.$$
      iii. If there is $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow{\mathsf{eval}} \bot) \in \widehat{\mathsf{OrigD}}$ but $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$, then add $\widetilde{F}$ to HidG.

2. **Sampling H^*:**

   (a) Sample $t' \leftarrow [1, t]$. For $i \in [t']$:
      i. Sample $(O_i^*, sk_i^*) \leftarrow (\mathbf{O}, \mathbf{sk} \mid pk \cup \mathsf{LearnG})$.

ii. Execute $D^{O_{lim}}(sk_i^*, c)$, where $O_{lim}(qu)$ replies exactly as $O_i^* \lozenge O(qu)$, except for the following case in which O_{lim} replies with $\bot$:

- $qu = ((\widetilde{F}, \widetilde{X}) \underset{\text{eval}}{\longrightarrow} ?)$ and $\widetilde{F} \in \mathsf{HidG}$

When the execution is over,

- Let $\mathsf{EvalD}_i^* := \emptyset$. For any eval query $qu = ((\widetilde{F}, \widetilde{X}) \underset{\text{eval}}{\longrightarrow} ?)$ during the execution for which $\widetilde{F} \notin \mathsf{HidG}$ and $\mathsf{eval}(\widetilde{F}, \widetilde{X}) = y \neq \bot$, add $((\widetilde{F}, \widetilde{X}) \underset{\text{eval}}{\longrightarrow} y)$ to EvalD_i^*

(b) Set $\mathsf{H}^* \leftarrow \mathtt{ConstHelp}^O(\mathsf{EvalD}_1^* \cup \cdots \cup \mathsf{EvalD}_{t'}^*)$, where the procedure $\mathtt{ConstHelp}^O$ is defined in Definition 4.25.

3. Output $b' \leftarrow \mathcal{A}(pk, \mathsf{LearnG}, c, \mathsf{H}^*)$

To describe our statement about the closeness of H and H^*, we define the following experiment outputting random variables over which we later state some events and probability statements.

Experiment $\mathbf{Expr}_1(1^\kappa)$ for bit $b \in \{0, 1\}$:   Output

$$\mathrm{Vars} = (\mathsf{OrigG}, \mathsf{OrigE}, c, (\mathsf{S}_i, \widehat{O}_i, \widehat{sk}_i)_{i \in [t]}, \mathsf{H}, (O_i^*, sk_i^*)_{i \in [t]}, \mathsf{HidG}, \mathsf{H}^*),$$

where

1. $((pk, \mathsf{LearnG}), sk) \leftarrow G^O(1^\kappa)$.

2. $(c, \mathsf{H}) \leftarrow \ddot{E}((pk, \mathsf{LearnG}), b)$ with OrigE and $(\widehat{O}_i, \widehat{sk}_i)$ being the random variables sampled inside this execution, and S_i is the value of S in the ith iteration based on which $(\widehat{O}_i, \widehat{sk}_i)$ is sampled — i.e., $(\widehat{O}_i, \widehat{sk}_i) \leftarrow \mathrm{Partial}(pk, \mathsf{S}_i)$.

3. H^*, HidG and (O_i^*, sk_i^*) are sampled as in the execution of $\mathtt{SecRed}^{\mathcal{A}, O}(pk, \mathsf{LearnG}, c)$.

We will prove the following lemma which shows that our proposed reduction $\mathtt{SecRed}$ does a good job of simulating H.

Lemma 4.47. *Let p be an arbitrary polynomial which satisfies*

$$8tq^2\varepsilon + \frac{1}{2^{\kappa/2-1}} \leq \frac{1}{p}.$$

We have

$$\Delta(\mathsf{H}, \mathsf{H}^*) \leq t\left(\frac{1}{p-1} + 4tq^2\varepsilon + \frac{1}{q^{\tau-1}} + \frac{1}{2^{\kappa/2-1}}\right),$$

where H and H^* are sampled as in $\mathbf{Expr}_1(1^\kappa)$.

The proof of Lemma 4.47 follows from Lemmas 4.48 and 4.49, given below.

Lemma 4.48. *Let p be a polynomial as in Lemma 4.47. For any $i \in [t]$ we have*

$$\Delta((O_i^*, sk_i^*, \mathsf{LearnG}, \mathsf{OrigE}, c), (\widehat{O}_i, \widehat{sk}_i, \mathsf{LearnG}, \mathsf{OrigE}, c)) \leq \frac{1}{p-1},$$

where all the variables are sampled as in $\mathbf{Expr}_1(1^\kappa)$.

Lemma 4.49. *Fix $i \in [t]$. Let O'_{lim} be defined as follows: O'_{lim} replies to a query qu exactly as $\widehat{O}_i \Diamond O$ except for the following case in which O'_{lim} replies with $\bot$: $qu = ((\widetilde{F}, *) \xrightarrow{\text{eval}} ?)$ and $\widetilde{F} \in \mathsf{HidG}$.*

Let Evnt: be the event that $\mathrm{Trace}(D^{O'_{lim}}(\widehat{sk}_i, c)) = \mathrm{Trace}(D^{\widehat{O}_i \Diamond O[\mathsf{OrigE} \cup \mathsf{LearnG} \cup \mathsf{S}_i]}(\widehat{sk}_i, c))$. Then

$$\Pr[\overline{\mathsf{Evnt}}] \leq 4tq^2\varepsilon + \frac{1}{q^{\tau-1}} + \frac{1}{2^{\kappa/2-1}},$$

where all the variables inside Evnt are sampled as in $\mathbf{Expr}_1(1^\kappa)$.

We will first how to prove of Lemma 4.47 using Lemmas 4.48 and 4.49. We will then prove Lemma 4.48 and Lemma 4.49.

Proof of Lemma 4.47. Fix $i \in [t]$. Let O_{lim} be defined as in the execution of the security reduction $\mathsf{SecRed}^{\mathcal{A},O}(pk, \mathsf{LearnG}, c)$: Namely, O_{lim} replies to a query qu exactly as $O_i^* \Diamond O$ except for the following case in which O_{lim} replies with $\bot$: $qu = ((\widetilde{F}, *) \xrightarrow{\text{eval}} ?)$ and $\widetilde{F} \in \mathsf{HidG}$.

Let $\mathsf{T}_1 = \mathrm{Trace}(D^{O_{lim}}(sk_i^*, c))$ and $\mathsf{T}_2 = \mathrm{Trace}(D^{\widehat{O}_i \Diamond O[\mathsf{LearnG} \cup \mathsf{OrigE} \cup \mathsf{S}_i]}(\widehat{sk}_i, c))$, where Trace denotes the set of query/response pairs made during a given execution. We show

$$\Delta(\mathsf{T}_1, \mathsf{T}_2) \leq \frac{1}{p-1} + 4tq^2\varepsilon + \frac{1}{q^{\tau-1}} + \frac{1}{2^{\kappa/2-1}}. \tag{14}$$

Using a union bound we will then obtain the desired upper bound for $\Delta(\mathsf{H}, \mathsf{H}^*)$.

We show how to obtain Equation 14 based on Lemmas 4.48 and 4.49. By Lemma 4.48 we have

$$\Delta((O_i^*, sk_i^*, \mathsf{LearnG}, \mathsf{OrigE}, c), (\widehat{O}_i, \widehat{sk}_i, \mathsf{LearnG}, \mathsf{OrigE}, c)) \leq \frac{1}{p-1}. \tag{15}$$

Note that since HidG is sampled independently of all the variables in Equation 15 we have

$$\Delta((O_i^*, sk_i^*, \mathsf{OrigE}, c, \mathsf{HidG}), (\widehat{O}_i, \widehat{sk}_i, \mathsf{OrigE}, c, \mathsf{HidG})) \leq \frac{1}{p-1}. \tag{16}$$

Define the oracle O'_{lim} as in Lemma 4.49: O'_{lim} replies to a query qu exactly as $\widehat{O}_i \Diamond O$ except for the following case in which O'_{lim} replies with $\bot$: $qu = ((\widetilde{F}, *) \xrightarrow{\text{eval}} ?)$ and $\widetilde{F} \in \mathsf{HidG}$. Define $\mathsf{T}'_1 = \mathrm{Trace}(D^{O'_{lim}}(sk_i^*, c))$

From Equation 16 we have

$$\Delta(\mathsf{T}_1, \mathsf{T}'_1) \leq \frac{1}{p-1}. \tag{17}$$

By Lemma 4.49 we have

$$\Delta(\mathsf{T}'_1, \mathsf{T}_2) \leq 4tq^2\varepsilon + \frac{1}{q^{\tau-1}} + \frac{1}{2^{\kappa/2-1}}. \tag{18}$$

Combining Equations 17 and 18 implies the bound in Equation 14, as desired. $\qquad\square$

Roadmap for the proof of Lemma 4.48. We now focus on proving Lemma 4.48. Fix an arbitrary iteration $i \in [t]$ for which we want to show $\mathbf{Dist}_1 = (O_i^*, sk_i^*, \mathsf{LearnG}, \mathsf{OrigE}, c)$ and $\mathbf{Dist}_2 = (\widehat{O}_i, \widehat{sk}_i, \mathsf{LearnG}, \mathsf{OrigE}, c)$ are statistically close. Recall that

$$(O_i^*, sk_i^*) \leftarrow \mathbf{D}_1 := (\mathbf{O}, \mathbf{sk} \mid pk \cup \mathsf{LearnG}),$$

and

$$(\widehat{O}_i, \widehat{sk}_i) \leftarrow \mathbf{D}_2 := (\mathbf{O}, \mathbf{sk} \mid pk \cup \mathsf{LearnG} \cup \mathsf{OrigE} \cup \mathsf{S}_i).$$

Define

$$\mathsf{Diff} := (\mathsf{OrigE} \cup \mathsf{S}_i \setminus \mathsf{LearnG}).$$

Comparing the ways in which (O_i^*, sk_i^*) and $(\widehat{O}_i, \widehat{sk}_i)$ are sampled, Diff is the only extra set on which $\widehat{O}_i$ is conditioned on. With this in mind, to prove that $\mathbf{Dist}_1$ and $\mathbf{Dist}_2$ are close, we first define a notion of *disjointedness* between two sets of query/response pairs. We will then show that if (a) O_i^* and Diff are disjoint and (b) $\widehat{O}_i$ and Diff are disjoint, then the induced distributions $\mathbf{Dist}_1$ and $\mathbf{Dist}_2$ are identical. We will then bound the probability that (a) or (b) does not hold.

The notion of disjointness is described below.

Definition 4.50. Let $\mathsf{S}_1, \mathsf{S}_2$ be partial oracles (i.e., two sets of query/response pairs). We say

- S_1 *hits* S_2 if there is some query qu on which both $\mathsf{S}_1(qu)$ and $\mathsf{S}_2(qu)$ are defined. We stress that even if the outputs are the same (i.e, $\mathsf{S}_1(qu) = \mathsf{S}_2(qu)$) we still call this event a hit.

- S_1 and S_2 are *disjoint* if S_1 and S_2 do not hit and that $\mathsf{S}_1 \cup \mathsf{S}_2$ is a valid oracle.

Define the following two disjoint events:

- $\mathsf{Disjoint}_1 := O_i^*$ is disjoint of Diff.

- $\mathsf{Disjoint}_2 := \widehat{O}_i$ is disjoint of Diff.

We show the following lemma.

Lemma 4.51. *Let p be as in Lemma 4.47. That is, $8tq^2\varepsilon + \frac{1}{2^{\kappa/2-1}} \leq \frac{1}{p}$. Then:*

1. $((O_i^*, sk_i^*, \mathsf{LearnG}, \mathsf{OrigE}, c) \mid \mathsf{Disjoint}_1) \equiv ((\widehat{O}_i, \widehat{sk}_i, \mathsf{LearnG}, \mathsf{OrigE}, c) \mid \mathsf{Disjoint}_2),$

2. $\Pr[\mathsf{Disjoint}_1] \geq 1 - \frac{1}{p},$

3. $\Pr[\mathsf{Disjoint}_2] \geq 1 - \frac{1}{p-1},$

where all the variables are sampled as in $\mathbf{Expr}_1(1^\kappa)$.

Proof of Lemma 4.48. The proof follows in a straightforward way from Lemma 4.51. □

We now focus on proving Lemma 4.51.

Proof of Lemma 4.51, Part 1. the proof of this part follows from the definition of disjointness. □

Roadmap for the proof of Lemma 4.51, Part 2. We first describe three events $\mathsf{Bad}_1, \mathsf{Bad}_2$ and Bad_3 in Definition 4.52. Then the proof of this part of the lemma will be divided into two lemmas. First, in Lemma 4.53 we will show that $\Pr[\mathsf{Disjoint}_1] \geq \Pr[\overline{\mathsf{Bad}_1 \vee \mathsf{Bad}_2 \vee \mathsf{Bad}_3}]$. Then, in Lemma 4.54 we will show that $\Pr[\mathsf{Bad}_1 \vee \mathsf{Bad}_2 \vee \mathsf{Bad}_3] \leq 8tq^2\varepsilon + \frac{1}{2^{\kappa/2-1}}$. These two will complete the proof of Part 2 of Lemma 4.51.

Definition 4.52. We define the following bad events. Informally, Bad_1 is the event that the output of some gc or gi query defined by O_i^* is at the range of the actual oracles gc and gi but that output does not appear in OrigG. Also, Bad_2 is the event that some non-heavy query in OrigG also does appear in Diff. Finally, Bad_3 is defined as in Bad_2 by replacing OrigG with O_i^*.

- Bad_1: the event that at least one of the following occurs:

  1. There exists $(* \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in O_i^*$ such that for some (s, F): $\mathsf{gc}(s, F) = \widetilde{F}$ and that $((s, F) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \notin (\mathsf{OrigG} \cup \mathsf{LearnG})$.

  2. There exists $(* \underset{\mathsf{gi}}{\longrightarrow} \widetilde{x}) \in O_i^*$ such that that for some (s, i, b): $\mathsf{gi}(s, i, b) = \widetilde{x}$ and that $((s, i, b) \underset{\mathsf{gi}}{\longrightarrow} \widetilde{x}) \notin (\mathsf{OrigG} \cup \mathsf{LearnG})$.

- Bad_2: the event that there exists a query that appears in both OrigG and Diff. That is, for some $T \in \{f, \mathsf{gc}, \mathsf{gi}\}$ there exists a query qu such that $(qu \underset{T}{\longrightarrow} *) \in \mathsf{OrigG}$ and $(qu \underset{T}{\longrightarrow} *) \in \mathsf{Diff}$.

- Bad_3: the event that there exists a query that appears both in O_i^* and Diff. That is, for some $T \in \{f, \mathsf{gc}, \mathsf{gi}\}$ there exists a query qu such that $(qu \underset{T}{\longrightarrow} *) \in O_i^*$ and $(qu \underset{T}{\longrightarrow} *) \in \mathsf{Diff}$.

Lemma 4.53. *We have*
$$\Pr[\mathsf{Disjoint}_1] \geq \Pr[\overline{\mathsf{Bad}_1 \vee \mathsf{Bad}_2 \vee \mathsf{Bad}_3}],$$
where all the variables are sampled as in $\mathbf{Expr}_1(1^\kappa)$.

Proof. We show whenever $\overline{\mathsf{Bad}_1 \vee \mathsf{Bad}_2 \vee \mathsf{Bad}_3}$ then the event $\mathsf{Disjoint}_1$ must necessarily hold. Suppose toward a contradiction that $\overline{\mathsf{Disjoint}_1}$ holds, or in other words O_i^* and Diff are not disjoint. Recalling the notion of disjointedness from Definition 4.50 we consider all possible cases, deriving a contradiction for each case:

1. O_i^* hits Diff or in other words there is a query that appears in both O_i^* and Diff: This contradicts the fact that $\overline{\mathsf{Bad}_3}$ holds.

2. O_i^* does not hit Diff but $O_i^* \cup \mathsf{Diff}$ is not a valid partial oracle: if this happens then at least one of the following must happen:

   (a) for some $qu_1 \neq qu_2$ we have $(qu_1 \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in O_i^*$ and $(qu_2 \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \mathsf{Diff}$: in this case we claim $(qu_2 \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \mathsf{OrigG}$, which contradicts the fact that $\overline{\mathsf{Bad}_2}$ holds. To prove the claim suppose to the contrary that $(qu_2 \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \notin \mathsf{OrigG}$. Firs, note that since $(qu_2 \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \mathsf{Diff}$ we have $\mathsf{gc}(qu_2) = \widetilde{F}$. Also, note that we must have $(qu_2 \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \notin \mathsf{LearnG}$ because we know O_i^* agrees with LearnG and so we could not have had $(qu_1 \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in O_i^*$

for $qu_1 \neq qu_2$. Now we reach a contradiction to the fact that $\overline{\mathsf{Bad}_1}$ holds, because we now have $\widetilde{F}$ is a valid garbled circuit, $(qu_1 \xrightarrow{\text{gc}} \widetilde{F}) \in O_i^*$ and $(* \xrightarrow{\text{gc}} \widetilde{F}) \notin \mathsf{OrigG} \cup \mathsf{LearnG}$.

(b) for some $qu_1 \neq qu_2$ we have $(qu_1 \xrightarrow{\text{gc}} \widetilde{x}) \in O_i^*$ and $(qu_2 \xrightarrow{\text{gc}} \widetilde{x}) \in \mathsf{Diff}$: in exactly the same way as in the previous case we can show that this case is also impossible.

$\square$

Lemma 4.54. *We have* $\Pr[\mathsf{Bad}_1 \vee \mathsf{Bad}_2 \vee \mathsf{Bad}_3] \leq 8tq^2\varepsilon + \frac{1}{2^{\kappa/2-1}}$, *where all the variables are sampled as in* $\mathbf{Expr}_1(1^\kappa)$.

Proof. We can bound the probability of any of these events using ideas discussed extensively before, and so we outline the main ideas. First, note that whenever the event Bad_1 holds, we have hit the image of a sparse random function, which in our case is either gc or gi. This is because O_i^* is sampled in offline mode while conditioning solely on $\mathsf{OrigG} \cup \mathsf{LearnG}$. Thus, by Lemma 3.1 we have $\Pr[\mathsf{Bad}_1] \leq 2 \times \frac{1}{2^{2^\kappa}}$.

Recall that $\mathsf{Diff} = \mathsf{OrigG} \setminus \mathsf{LearnG}$. The event Bad_2 is the event that there is a query $qu \in \mathsf{OrigG}$ which is at most ε heavy, but that it appears during a process which is independent of that used to produce OrigG and which consists of at most $4tq^2$ queries. (See Lemma 4.32.) Thus, we have $\Pr[\mathsf{Bad}_2] \leq 4tq^2\varepsilon$. Finally, the event Bad_3 can be bounded in exactly the same way. $\square$

Proof of Lemma 4.51, Part 2. the proof of this part follows from Lemmas 4.53 and 4.54. $\square$

Proof of Lemma 4.51, Part 3. Recall the distributions $\mathbf{D}_1$ and $\mathbf{D}_2$:

$$(O_i^*, r_i^*) \leftarrow \mathbf{D}_1 := (\mathbf{O}, \mathbf{sk} \mid pk \cup \mathsf{LearnG}).$$

$$(\widehat{O}_i, \widehat{r}_i) \leftarrow \mathbf{D}_2 := (\mathbf{O}, \mathbf{sk} \mid pk \cup \mathsf{LearnG} \cup \mathsf{OrigE} \cup \mathsf{S}_i).$$

Recall that $\mathsf{Diff} = (\mathsf{OrigE} \cup \mathsf{S}_i) \setminus \mathsf{LearnG}$.

In what follows we write $(O', r') \in \mathbf{D}_1$ to mean that (O', r') is in the support set of $\mathbf{D}_1$. Consider the following partitioning of the worlds:

- S_1: the set of all $(O', r') \in \mathbf{D}_2$ such that $G^{O'}(r')$ does not hit Diff.

- S_2: the set of all $(O', r') \in \mathbf{D}_2$ such that $G^{O'}(r')$ hits Diff.

- S_3: the set of worlds generated by $\mathbf{D}_1$ which disagree with Diff. Namely, for all $(O', r') \in \mathbf{S}_3$ we have $O' \cup \mathsf{Diff}$ is an invalid partial oracle.

Note that $\mathsf{S}_1 \cup \mathsf{S}_2 \cup \mathsf{S}_3$ covers all the worlds in the support set of $\mathbf{D}_1$. Also, $\mathsf{S}_1 \cup \mathsf{S}_2$ covers all the worlds in the support set of $\mathbf{D}_2$. By Part 2 of Lemma 4.51 we have

$$\frac{|\mathsf{S}_2| + |\mathsf{S}_3|}{|\mathsf{S}_1| + |\mathsf{S}_2| + |\mathsf{S}_3|} \leq \frac{1}{p}.$$

We want to bound $|\mathsf{S}_2|/(|\mathsf{S}_1| + |\mathsf{S}_2|)$. We have

$$\frac{|\mathsf{S}_2|}{|\mathsf{S}_1| + |\mathsf{S}_2|} \leq \frac{|\mathsf{S}_2| + |\mathsf{S}_3|}{|\mathsf{S}_1|} \leq \frac{\frac{|\mathsf{S}_2| + |\mathsf{S}_3|}{|\mathsf{S}_1| + |\mathsf{S}_2| + |\mathsf{S}_3|}}{\frac{|\mathsf{S}_1|}{|\mathsf{S}_1| + |\mathsf{S}_2| + |\mathsf{S}_3|}} \leq \frac{\frac{1}{p}}{\frac{p-1}{p}} = \frac{1}{p-1}.$$

$\square$

We now focus on proving Lemma 4.49.

Proof of Lemma 4.49. Let $\widehat{\mathsf{OrigD}}_i$ be the set of all queries issued during the execution of the procedure $D^{\widehat{O}_i \lozenge O[\mathsf{OrigE} \cup \mathsf{LearnG} \cup \mathsf{S}_i]}(\widehat{sk}_i, c)$. Define the following events:

- Miss: the event that there exists $\widetilde{F}$ such that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG}$, that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \notin \mathsf{LearnG} \cup \mathsf{OrigE} \cup \mathsf{S}_i$, that $\widetilde{F} \notin \mathsf{HidG}$ and that for some $\widetilde{X}$: $((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?) \in \widehat{\mathsf{OrigD}}_i$ and $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$.

- Bad_4: the event that there exists $\widetilde{F}$ such that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \notin \mathsf{OrigG} \cup \mathsf{LearnG} \cup \mathsf{OrigE} \cup \mathsf{S}_i$, that $((\widetilde{F}, *) \xrightarrow[\mathsf{eval}]{} ?) \in \widehat{\mathsf{OrigD}}_i$ for some s and F: $\mathsf{gc}(s, F) = \widetilde{F}$.

- Bad_5: the event that there is $\widetilde{F}$ such that $\widetilde{F} \in \mathsf{HidG}$ and $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \notin \mathsf{OrigG}$.

- Bad_6: the event that there exists a query qu such that $(qu \xrightarrow[\mathsf{gc}]{} *) \in \mathsf{OrigG} \cap \mathsf{Diff}$, where recall that $\mathsf{Diff} = (\mathsf{OrigE} \cup \mathsf{S}_i) \setminus \mathsf{LearnG}$.

We show that if $\overline{\mathsf{Evnt}}$ holds, then $\mathsf{Miss} \vee \mathsf{Bad}_4 \vee \mathsf{Bad}_5 \vee \mathsf{Bad}_6$ holds, or equivalently, if $\overline{\mathsf{Miss}} \wedge \overline{\mathsf{Bad}_4} \wedge \overline{\mathsf{Bad}_5} \wedge \overline{\mathsf{Bad}_6}$ holds then Evnt holds. We will then show that the probability that $\mathsf{Miss} \vee \mathsf{Bad}_4 \vee \mathsf{Bad}_5 \vee \mathsf{Bad}_6$ holds is at most $q\varepsilon + \frac{1}{q^{15}} + 2 \times \frac{1}{2^{\kappa/2}}$.

Suppose $\overline{\mathsf{Miss}} \wedge \overline{\mathsf{Bad}_4} \wedge \overline{\mathsf{Bad}_5} \wedge \overline{\mathsf{Bad}_6}$ holds and suppose to the contrary that $\overline{\mathsf{Evnt}}$ holds. Let $\mathsf{S} := \mathsf{LearnG} \cup \mathsf{OrigE} \cup \mathsf{S}_i$. Let $\mathsf{T}_1 = \mathrm{Trace}(D^{O_{\mathsf{ext}}}(\widehat{sk}_i, c))$ and let $\mathsf{T}_2 = \mathrm{Trace}(D^{\widehat{O}_i \lozenge O[\mathsf{S}]}(\widehat{sk}_i, c))$. Consider the first point in which T_1 and T_2 become different: this point must be an eval query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?)$, where $\mathsf{gc}(s, F) = \widetilde{F}$ for some (s, F) and for which one of the following holds: We show that each case leads to a contradiction.

1. $O_{\mathsf{ext}}(qu) = \bot$ and $\widehat{O}_i \lozenge O[\mathsf{S}](qu) \neq \bot$: The fact that $O_{\mathsf{ext}}(qu) = \bot$ implies $\widetilde{F} \in \mathsf{HidG}$. Since $\overline{\mathsf{Bad}_5}$ holds we have $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG}$. Also, by the way in which HidG is designed, we have $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \notin \mathsf{LearnG}$. Now the fact that $\widehat{O}_i \lozenge O[\mathsf{S}](qu) \neq \bot$ and that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \notin \mathsf{LearnG}$ imply that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{S} \setminus \mathsf{LearnG} = \mathsf{Diff}$. Now the facts that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{Diff}$ and that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG}$ imply that Bad_6 holds, which is a contradiction.

2. $O_{\mathsf{ext}}(qu) \neq \bot$ and $\widehat{O}_i \lozenge O[\mathsf{S}](qu) = \bot$: Since $\widehat{O}_i \lozenge O[\mathsf{S}](qu) = \bot$ and $\mathsf{eval}(qu) \neq \bot$ we have $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \notin \mathsf{S}$. Also, since $\overline{\mathsf{Bad}_4}$ holds, we have $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG} \cup \mathsf{S}$, and thus by the preceding condition we have $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG}$. Moreover, $O_{\mathsf{ext}}(qu) \neq \bot$ implies $\widetilde{F} \notin \mathsf{HidG}$. Now the facts that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG} \setminus \mathsf{S}$, that $\widetilde{F} \notin \mathsf{HidG}$, that $((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?) \in \widehat{\mathsf{OrigD}}_i$ and that $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$ imply Miss holds, which is a contradiction.

We now show how to bound each of the events above. We can easily show that each of the events Bad_4 and Bad_5 happens with probability at most $\frac{1}{2^{\kappa/2}}$ because the occurrence of each of them implies that we have obliviously hit a valid point in the image of gc. Also, the event Bad_6 is

41

a special case of the event Bad_2 (Definition 4.52), which we have shown in Lemma 4.54 that this event occurs with probability at most $4tq^2\varepsilon$.

We now bound the probability of the event Miss. To this end, we define a Bernoulli random variable for any $\widetilde{F}$ such that $((* \xrightarrow[\mathsf{gc}]{} \widetilde{F})) \in \mathsf{OrigG}$ and then show that the event Miss is the event that for some $\widetilde{F}$ where $((* \xrightarrow[\mathsf{gc}]{} \widetilde{F})) \in \mathsf{OrigG}$: the first q^τ trials of that Bernoulli variable fails, but the last one succeeds. We obtain the desired bound of $q \times \frac{1}{q^\tau}$ from Lemma 3.1 and a union bound.

Fix $\widetilde{F}$ for which $((* \xrightarrow[\mathsf{gc}]{} \widetilde{F})) \in \mathsf{OrigG}$. Define the following Bernoulli trial.

Bernoulli Trial:

1. For both $b \in \{0,1\}$ do the following:

   (a) Sample $c' \leftarrow E^O(pk, b)$ and let OrigE' be the set of all query/response pairs during this execution.

   (b) Let $\widehat{\mathsf{OrigD}}_i$ be sampled as in $\mathrm{DistGen}_1(pk, \mathsf{LearnG}, \mathsf{OrigE}')$.

   - We say the Bernoulli trial succeeds if there is a query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} \perp) \in \widehat{\mathsf{OrigD}}_i$ but $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \perp$. Otherwise, we say the Bernoulli trial fails.

Now suppose the event Miss holds. This means that there exists $\widetilde{F}$ such that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigG}$, that $\widetilde{F} \notin \mathsf{HidG}$, that $(* \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \notin \mathsf{LearnG} \cup \mathsf{OrigE} \cup \mathsf{S}_i$, and that for some $\widetilde{X}$ the query $((\widetilde{F}, \widetilde{X}) \xrightarrow[\mathsf{eval}]{} ?)$ appears during the execution of $D^{\widehat{O}_i \Diamond O[\mathsf{LearnG} \cup \mathsf{OrigE} \cup \mathsf{S}_i]}(\widehat{sk}_i, c)$ and that $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \perp$. This means that this particular Bernoulli trial (based on OrigE, $\widehat{O}_i$ and $\widehat{sk}_i$) succeeds. Moreover, $\widetilde{F} \notin \mathsf{HidG}$ means that all the previous q^τ Bernoulli trials performed to form the random variable HidG failed. The proof is now complete. $\qquad\square$

References

[App17]   Benny Applebaum. Garbled circuits as randomized encodings of functions: a primer. In *Tutorials on the Foundations of Cryptography*, pages 1–44. Springer, 2017. 3

[AS15]   Gilad Asharov and Gil Segev. Limits on the power of indistinguishability obfuscation and functional encryption. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 191–209. IEEE, 2015. 3, 4, 5, 9, 12

[AS16]   Gilad Asharov and Gil Segev. On constructing one-way permutations from indistinguishability obfuscation. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A: 13th Theory of Cryptography Conference, Part II*, volume 9563 of *Lecture Notes in Computer Science*, pages 512–541, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany. 3

[Bar17]   Boaz Barak. The complexity of public-key cryptography. In *Tutorials on the Foundations of Cryptography*, pages 45–77. Springer, 2017. 3

[BBF13]   Paul Baecher, Christina Brzuska, and Marc Fischlin. Notions of black-box reductions, revisited. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 296–315. Springer, 2013. 3, 13

[BDV17]   Nir Bitansky, Akshay Degwekar, and Vinod Vaikuntanathan. Structure vs. hardness through the obfuscation lens. In *Annual International Cryptology Conference*, pages 696–723. Springer, 2017. 3

[Bea96]   Donald Beaver. Correlated pseudorandomness and the complexity of private computations. In *28th Annual ACM Symposium on Theory of Computing*, pages 479–488, Philadephia, PA, USA, May 22–24, 1996. ACM Press. 4

[BHR12]   Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 12: 19th Conference on Computer and Communications Security*, pages 784–796, Raleigh, NC, USA, October 16–18, 2012. ACM Press. 3, 4, 12

[BKSY11]   Zvika Brakerski, Jonathan Katz, Gil Segev, and Arkady Yerukhimovich. Limits on the power of zero-knowledge proofs in cryptographic constructions. In *Theory of Cryptography Conference*, pages 559–578. Springer, 2011. 3, 5, 55, 56

[BMG07]   Boaz Barak and Mohammad Mahmoody-Ghidary. Lower bounds on signatures from symmetric primitives. In *48th Annual Symposium on Foundations of Computer Science*, pages 680–688, Providence, RI, USA, October 20–23, 2007. IEEE Computer Society Press. 26

[BMG09]   Boaz Barak and Mohammad Mahmoody-Ghidary. Merkle puzzles are optimal - an $O(n^2)$-query attack on any key exchange from a random oracle. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 374–390, Santa Barbara, CA, USA, August 16–20, 2009. Springer, Heidelberg, Germany. 16

[BPR+08]   Dan Boneh, Periklis A. Papakonstantinou, Charles Rackoff, Yevgeniy Vahlis, and Brent Waters. On the impossibility of basing identity based encryption on trapdoor permutations. In *49th Annual Symposium on Foundations of Computer Science*, pages 283–292, Philadelphia, PA, USA, October 25–28, 2008. IEEE Computer Society Press. 3

[CKP15]   Ran Canetti, Yael Tauman Kalai, and Omer Paneth. On obfuscation with random oracles. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015: 12th Theory of Cryptography Conference, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 456–467, Warsaw, Poland, March 23–25, 2015. Springer, Heidelberg, Germany. 6

[Dac16]   Dana Dachman-Soled. Towards non-black-box separations of public key encryption and one way function. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B: 14th Theory of Cryptography Conference, Part II*, volume 9986 of *Lecture Notes in Computer Science*, pages 169–191, Beijing, China, October 31 – November 3, 2016. Springer, Heidelberg, Germany. 5

[DG17]   Nico Döttling and Sanjam Garg. Identity-based encryption from the Diffie-Hellman assumption. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 537–569, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany. 3, 4

[DH76]   Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976. 3

[FLS90]   Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st Annual Symposium on Foundations of Computer Science*, pages 308–317, St. Louis, Missouri, October 22–24, 1990. IEEE Computer Society Press. 55

[GLOS15]   Sanjam Garg, Steve Lu, Rafail Ostrovsky, and Alessandra Scafuro. Garbled RAM from one-way functions. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th Annual ACM Symposium on Theory of Computing*, pages 449–458, Portland, OR, USA, June 14–17, 2015. ACM Press. 4

[GMM17a]   Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammed. Lower bounds on obfuscation from all-or-nothing encryption primitives. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 661–695, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany. 5, 6

[GMM17b]   Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammed. When does functional encryption imply obfuscation? In *TCC 2017: 15th Theory of Cryptography Conference, Part I*, Lecture Notes in Computer Science, pages 82–115. Springer, Heidelberg, Germany, March 2017. 5

[GMR01]   Yael Gertner, Tal Malkin, and Omer Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *42nd Annual Symposium on Foundations of Computer Science*, pages 126–135, Las Vegas, NV, USA, October 14–17, 2001. IEEE Computer Society Press. 7

[GW11]   Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd Annual ACM Symposium on Theory of Computing*, pages 99–108, San Jose, CA, USA, June 6–8, 2011. ACM Press. 5

[HR04]   Chun-Yuan Hsiao and Leonid Reyzin. Finding collisions on a public road, or do secure hash functions need secret coins? In Matthew Franklin, editor, *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 92–105, Santa Barbara, CA, USA, August 15–19, 2004. Springer, Heidelberg, Germany. 7

[IK00]   Yuval Ishai and Eyal Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *41st Annual Symposium*

on Foundations of Computer Science, pages 294–304, Redondo Beach, CA, USA, November 12–14, 2000. IEEE Computer Society Press. 3

[IR89]   Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st Annual ACM Symposium on Theory of Computing*, pages 44–61, Seattle, WA, USA, May 15–17, 1989. ACM Press. 2, 3, 6, 7, 8, 10, 13, 15, 16

[LO13]   Steve Lu and Rafail Ostrovsky. How to garble RAM programs. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 719–734, Athens, Greece, May 26–30, 2013. Springer, Heidelberg, Germany. 4

[LP09]   Yehuda Lindell and Benny Pinkas. A proof of security of Yao's protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, April 2009. 3

[MMN16a] Mohammad Mahmoody, Ameer Mohammed, and Soheil Nematihaji. On the impossibility of virtual black-box obfuscation in idealized models. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A: 13th Theory of Cryptography Conference, Part I*, volume 9562 of *Lecture Notes in Computer Science*, pages 18–48, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany. 6

[MMN⁺16b] Mohammad Mahmoody, Ameer Mohammed, Soheil Nematihaji, Rafael Pass, and Abhi Shelat. Lower bounds on assumptions behind indistinguishability obfuscation. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A: 13th Theory of Cryptography Conference, Part I*, volume 9562 of *Lecture Notes in Computer Science*, pages 49–66, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany. 6

[Pas11]   Rafael Pass. Limits of provable security from standard assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd Annual ACM Symposium on Theory of Computing*, pages 109–118, San Jose, CA, USA, June 6–8, 2011. ACM Press. 5

[PRV12]   Periklis A. Papakonstantinou, Charles W. Rackoff, and Yevgeniy Vahlis. How powerful are the DDH hard groups? Cryptology ePrint Archive, Report 2012/653, 2012. http://eprint.iacr.org/2012/653. 3, 4

[PTV11]   Rafael Pass, Wei-Lung Dustin Tseng, and Muthuramakrishnan Venkitasubramaniam. Towards non-black-box lower bounds in cryptography. In Yuval Ishai, editor, *TCC 2011: 8th Theory of Cryptography Conference*, volume 6597 of *Lecture Notes in Computer Science*, pages 579–596, Providence, RI, USA, March 28–30, 2011. Springer, Heidelberg, Germany. 5

[RSA78]   Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signature and public-key cryptosystems. *Communications of the Association for Computing Machinery*, 21(2):120–126, 1978. 3

[RTV04]   Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In Moni Naor, editor, *TCC 2004: 1st Theory of Cryptography Conference*, volume 2951 of *Lecture Notes in Computer Science*, pages 1–20, Cambridge, MA, USA, February 19–21, 2004. Springer, Heidelberg, Germany. 3, 4, 13

[SW14]   Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 475–484, New York, NY, USA, May 31 – June 3, 2014. ACM Press. 7

[Yao86]   Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th Annual Symposium on Foundations of Computer Science*, pages 162–167, Toronto, Ontario, Canada, October 27–29, 1986. IEEE Computer Society Press. 3

A   Extension to Constant-Round Key Agreement

Our main result showed that PKE, which is equivalent to 2-round key agreement, cannot be obtained from one-way functions and garbling. In this section, we describe how our main result can in fact extend to m-round key agreement protocols for any constant m. Unfortunately, the proof becomes more involved when handling protocols with more than 2 messages (which are equivalent to PKE schemes), and that is why in the main body of the paper we describe the full proof for the case of PKE schemes. Here we sketch the arguments needed for the extension to the constant-round KE protocols.

The reason that our techniques, at least in their current form, do not extend beyond the constant-round (in the context of KE) is that, every time that we compile the protocol into a new one (with almost the same security) and get rid of the garbling queries for a round, it blows up the parameters (e.g., the communication complexity) of the protocol by a polynomial factor. Going beyon this limitation is an interesting future direction.

We start by first introducing some notation and definitions related to key agreement protocols.

A.1   Notation and Definitions

Interactive protocols.   For any pair of interactive oracle-aided algorithms (A, B) and oracle O, we use $\langle A^O(1^\kappa), B^O(1^\kappa) \rangle$ to denote a random execution of an interactive protocol between an *initiator* A and a *successor* B, both with access to oracle O and both with no initial inputs. We denote by AMsg_i the i'th message of A and by BMsg_i the i'th message of B. We use $\mathsf{Trans}\langle A^O(1^\kappa), B^O(1^\kappa) \rangle$ to refer to the *transcript* of the protocol, which is $(\mathsf{AMsg}_1, \mathsf{BMsg}_1, \dots,)$. We may sometimes refer to A as Alice and to B as Bob.

Secret state of algorithms.   For an interactive protocol $\langle A, B \rangle$ between $A := (A_1, \dots, A_m)$ and $B := (B_1, \dots, B_m)$ we assume the following: the input to A_i, for any $i \in [m]$, is $(\mathsf{StA}, \mathsf{Tr})$, where Tr is the existing transcript at the time and StA is passed on privately to A_i from A_{i-1}, consisting of all the randomness values as well as all the query/response pairs made by $(A_1, \dots, A_{i-1})$. We use StA_{qu} to refer to all query/response pairs contained in StA and use StA_{rand} to refer to the randomness part of StA. We assume a similar convention for B.

Notation.   Let $\langle A, B \rangle$ be an m-round protocol and $m \geq 2i$. We write

$$(\mathsf{Tr}, \mathsf{state} = (\mathsf{StA}, \mathsf{StB})) \leftarrow \langle A[1, \cdots, i], B[1, \cdots, i] \rangle$$

to indicate that after running i executions of each of the algorithms (i.e., $2i$ rounds in total) the resulting transcript is Tr and the current secret states of A and B are, respectively, StA and StB.

Transcript. For any given transcript $\mathsf{Tr} = (\mathsf{AMsg}_1, \mathsf{BMsg}_1, \ldots, \mathsf{AMsg}_m, \mathsf{BMsg}_m)$, we let $\mathsf{Tr}_A :=$ $(\mathsf{AMsg}_1, \cdots, \mathsf{AMsg}_m)$ be A's portion of the transcript, and define Tr_B, B's portion, similarly.

Half-fixed executions. Let $\mathsf{AliceM} := (\mathsf{AMsg}_1, \ldots, \mathsf{AMsg}_m)$ be a given sequence of Alice's messages during the protocol. We use $\langle \mathsf{AliceM}, B^O(1^\kappa) \rangle$ to represent a random execution of the protocol in which Alice's ith message, for $i \in [m]$, is fixed to AMsg_i. We use $\langle \mathsf{AliceM}, B^O(r_1; \ldots; r_m) \rangle$ to refer to the above execution, but in which the randomness of B's i'th algorithm is fixed to r_i. We let $\mathrm{Query}(\langle \mathsf{AliceM}, B^O(1^\kappa) \rangle)$ denote the set of all Bob's query/answer pairs during the execution of $\langle \mathsf{AliceM}, B^O(1^\kappa) \rangle$.

Protocol output. At the end of an interactive protocol $\langle A, B \rangle$, each A and B output a value. We use $(\mathsf{out}_A, \mathsf{out}_B) \leftarrow \mathsf{out}(\langle A^O(1^\kappa), B^O(1^\kappa) \rangle)$ to indicate that out_A and out_B are, respectively, A's and B's output values.

Definition A.1 (Key Agreement Protocol). For any security parameter $\kappa \in \mathbb{N}$ and any $m \in \mathbb{N}$, an $2m$-round *key agreement* protocol consists of two PPT interactive algorithms (A, B) such that $\langle A(r_a), B(r_b) \rangle(1^\kappa) = (k_a, k_b)$ for (private) randomness r_a, r_b. The following completeness and security conditions must be satisfied:

Correctness. For any function $\delta(.)$, we say that a key agreement protocol is $(1 - \delta)$-correct if the following holds for any security parameter κ:

$$\Pr[k_a = k_b] \geq 1 - \delta(\kappa)$$

where the probability is over the randomness (A, B). We say that the protocol is *correct* if it is $(1 - \delta)$-correct where $\delta = 1/p(\kappa)$ for some polynomial $p(.)$.

Security. For any function $\gamma(.)$, we say that a key agreement protocol is γ-secure if, for all PPT adversaries E and any security parameter κ the following holds:

$$\Pr[E(1^\kappa, \mathsf{Tr}) = k] \leq \gamma(\kappa)$$

where $k \leftarrow \langle A(r_a), B(r_b) \rangle(1^\kappa)$, $\mathsf{Tr} \leftarrow \mathsf{Trans}\langle A(r_a), B(r_b) \rangle$ and the probability is over the randomness of E and (A, B). We say that the protocol is *secure* if there exists a negligible function $\mathrm{negl}(.)$ such that it is γ-secure for some $\gamma(\kappa) = \mathrm{negl}(\kappa)$.

Definition A.2 (Alice-consistency and Bob-consistency). Let

$$\mathsf{Tr} := (\mathsf{AMsg}_1, \mathsf{BMsg}_1, \ldots, \mathsf{AMsg}_{i-1}, \mathsf{BMsg}_{i-1})$$

be a partial transcript.

Let (O', r_B) consist of a partial oracle O' and a sequence of Bob's randomness values $r_B :=$ $(r_{B,1}, \ldots, r_{B,i-1})$. We define the following notions.

- **Bob-consistency**: We say that (O', r_B) Bob-agrees with Tr (or is Bob-consistent with Tr) if

$$\mathsf{Trans}(\langle \mathsf{Tr}_A, B^{O'}(r_{B,1}; \ldots; r_{B,i-1}) \rangle) = \mathsf{Tr}.$$

- **Bob-minimal-consistency**: We say that (O', r_B) minimally Bob-agrees with Tr if (a) (O', r_B) Bob-agrees with Tr and (b) O' is defined only on the queries that Bob makes during the execution of $\langle \mathsf{Tr}_A, B^{O'}(r_{B,1}; \ldots; r_{B,i-1}) \rangle$, and nothing more.

- **Generalized consistency**: Letting S be a set of query/answer pairs, we say (O', r_B) is Bob-consistent with $(\mathsf{Tr}, \mathsf{S})$, denoted $(O', r_B) \overset{\text{Bob}}{\vdash} \mathsf{Tr}$, if (a) O' agrees with S and (b) (O', r_B) is Bob-consistent with Tr.

Similarly, we may define the notion of Alice-consistency by $\overset{\text{Alice}}{\vdash}$.

In our compilation procedure we need to have a procedure that allows us to return a consistent sample of the next hidden state of Bob based on a given transcript and a set of query/response pairs. Recall that the secret state of each Bob and Alice at any point consists of all their previous randomness values as well as all their previous query/answer pairs. We have the following definition for our sampling purposes.

Definition A.3 (Next-Bob sampler, next-Alice sampler)**.** Let S be a set of query/answer pairs and

$$\mathsf{Tr} := (\mathsf{AMsg}_1, \mathsf{BMsg}_1, \ldots, \mathsf{AMsg}_{i-1}, \mathsf{BMsg}_{i-1}),$$

be a partial transcript. The sampler $\mathsf{NextBSamp}(\mathsf{Tr}, \mathsf{S})$ samples uniformly at random a pair $(O', r_B := (r_{B,1}, \ldots, r_{B,i-1}))$ that minimally Bob-agrees with $(\mathsf{Tr}, \mathsf{S})$ and outputs $(O', \widehat{\mathsf{StB}} := (O', r_B))$. We may similarly define the procedure $\mathsf{NextASamp}$.

A.2   Compilation procedure

Given a key agreement protocol (A, B) in the $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval})$ where the first ν steps do not make calls to eval but the $(\nu + 1)$ step might potentially do, we will show how to compile out eval from that step. We assume WLOG that the ν'th step is Alice's turn to send a message.

Parameters.   The compilation is parameterized over two polynomials $t = t(\kappa)$ and $\varepsilon = \varepsilon(\kappa)$, which we will instantiate later.

Compilation

- $\dot{A}_\nu(\mathsf{StA} := (\mathsf{StA}_{qu}, \mathsf{StA}_{rand}), \mathsf{Tr})$: Do the following steps:

  1. **Learning previous heavy queries of** Bob**:** Set $\mathsf{LearnB} := \emptyset$. While there exists a query $qu \notin \mathsf{LearnB}$ such that

     $$\Pr_{\substack{O' \leftarrow \mathbf{O} \\ r_B \leftarrow \{0,1\}^*}} \left[qu \in \mathsf{Query}(\langle \mathsf{Tr}_A, B^{O'}(r_B) \rangle) \mid (O', r_B) \overset{\text{Bob}}{\vdash} (\mathsf{Tr}, \mathsf{LearnB}) \right] \geq \varepsilon$$

     choose the lexicographically first such qu and add $(qu \underset{T}{\rightarrow} O(qu))$ to LearnB, where T is the type of the query.
  2. **Running original** Alice A_ν**:** Let $\mathsf{OrigA} = \mathsf{StA}_{qu}$. Run $\mathsf{AMsg}_\nu \leftarrow A_\nu^O(\mathsf{StA}, \mathsf{Tr})$ and add all the query-answer pairs to OrigA.

3. **Generating Helper Set H for $\dot{B}_{\nu+1}$:** Sample $t' \leftarrow [1, t]$. Set $\mathsf{S} = \mathsf{OrigA} \cup \mathsf{LearnB}$. For $i \in [1, t']$, do the following:

   (a) **Offline phase:** Sample $(\widehat{O_i}, \widehat{\mathsf{state}_i}) \leftarrow \mathsf{NextBSamp}(\mathsf{Tr}, \mathsf{S})$.

   (b) **Semi-online phase:** Execute $B_{\nu+1}^{\widehat{O_i} \lozenge O[\mathsf{S}]}(\widehat{\mathsf{state}_i}, \mathsf{Tr}||\mathsf{AMsg}_\nu)$ and add all query-answer pairs made to the oracle O to the set S. Let $\widehat{\mathsf{OrigB}_i}$ be the set of all query-answer pairs made by this execution.

   When all the iterations are over, set $\mathsf{H} \leftarrow \mathtt{ConstHelp}(\widehat{\mathsf{OrigB}}, \mathsf{S})$, where $\widehat{\mathsf{OrigB}} = \widehat{\mathsf{OrigB}_1} \cup \cdots \cup \widehat{\mathsf{OrigB}_{t'}}$.

4. Output $\dot{\mathsf{AMsg}}_\nu = (\mathsf{LearnB}, \mathsf{AMsg}_\nu, \mathsf{H})$.

- $\dot{B}_{\nu+1}(\mathsf{StB}, \mathsf{Tr}||\dot{\mathsf{AMsg}}_\nu)$: Parse

$$\dot{\mathsf{AMsg}}_\nu := (\mathsf{LearnB}, \mathsf{AMsg}_\nu, \mathsf{H}).$$

Run $B_{\nu+1}^{O[\mathsf{H} \cup \mathsf{LearnB}]}(\mathsf{StB}, \mathsf{Tr}||\mathsf{AMsg}_\nu)$ and output whatever this outputs.

A.3 Correctness and Security

In this section we give the correctness and security statements about our compiled protocol.

The proof of security follows almost exactly the same as in the PKE case. The proof of correctness require us to deal with some new cases, that we did not have before, and so we give a sketch of the proof below.

A.4 Proof of Correctness

The following lemma will help us establish correctness for the compiled protocol by relating the correctness of the compiled protocol to that of the original one.

Lemma A.4. *We have*

$$\Pr\left[B_{\nu+1}^{O[\mathsf{H} \cup \mathsf{LearnB}]}(\mathsf{StB}, \mathsf{Tr}||\mathsf{AMsg}_\nu) \neq B_{\nu+1}^O(\mathsf{StB}, \mathsf{Tr}||\mathsf{AMsg}_\nu)\right] \leq \mathrm{negl}(\kappa) \tag{19}$$

where $O \leftarrow \mathbf{O}$,
$$(\mathsf{Tr}, state = (\mathsf{StA}, \mathsf{StB})) \leftarrow \langle A[1 \cdots \nu - 1], B[1 \cdots \nu - 1] \rangle,$$
and $(\mathsf{LearnB}, \mathsf{AMsg}_\nu, \mathsf{H}) \leftarrow \dot{A}_\nu^O(\mathsf{StA}, \mathsf{Tr})$.

To prove Lemma A.4 we define the following two undesirable events that are defined based on the variables sampled in Lemma A.4. We will then show that the probability of the event in Lemma A.4 is at most the probability of the union of the following two events.

Definition A.5 (Event Miss). Let OrigB be the set of all query/response pairs during the execution of $B_{\nu+1}^O(\mathsf{StB}, \mathsf{Tr}||\mathsf{AMsg}_\nu)$.

We define the event Miss to hold if there is a query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow{\mathsf{eval}} ?) \in \mathsf{OrigB}$ such that:

1. $\mathsf{eval}(\widetilde{F}, \widetilde{X}) = y \neq \bot$, and

49

2. $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} *) \notin \mathsf{H}$ and $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \notin \mathsf{H}$, and

3. $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{OrigA}$ for some s and F.

Definition A.6 (Event Surprise)**.** There exists a query $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} ?) \in \mathsf{OrigB}$ such that:

- $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \notin \mathsf{StB} \cup \mathsf{LearnB} \cup \mathsf{OrigA}$.

The proof of Lemma A.4 now follows from the following three lemmas.

Lemma A.7. $\Pr\left[B_{\nu+1}^{O[\mathsf{H} \cup \mathsf{LearnB}]}(\mathsf{StB}, \mathsf{Tr}\|\mathsf{AMsg}_\nu) \neq B_{\nu+1}^O(\mathsf{StB}, \mathsf{Tr}\|\mathsf{AMsg}_\nu) \right] \leq \Pr[\mathsf{Miss} \vee \mathsf{Surprise}]$, *where* $O \leftarrow \mathbf{O}$,
$$(\mathsf{Tr}, state = (\mathsf{StA}, \mathsf{StB})) \leftarrow \langle A[1 \cdots \nu - 1], B[1 \cdots \nu - 1] \rangle,$$
and $(\mathsf{LearnB}, \mathsf{AMsg}_\nu, \mathsf{H}) \leftarrow \dot{A}_\nu^O(\mathsf{StA}, \mathsf{Tr})$.

Lemma A.8. $\Pr[\mathsf{Surprise}] \leq \frac{1}{2^{\kappa/2}}$, *where* $O \leftarrow \mathbf{O}$,
$$(\mathsf{Tr}, state = (\mathsf{StA}, \mathsf{StB})) \leftarrow \langle A[1 \cdots \nu - 1], B[1 \cdots \nu - 1] \rangle,$$
and $(\mathsf{LearnB}, \mathsf{AMsg}_\nu, \mathsf{H}) \leftarrow \dot{A}_\nu^O(\mathsf{StA}, \mathsf{Tr})$.

Lemma A.9. $\Pr[\mathsf{Miss}] \leq \frac{1}{2^{\kappa/4}} + \frac{2q}{t} + 3\varepsilon\nu t$, *where* $O \leftarrow \mathbf{O}$,
$$(\mathsf{Tr}, state = (\mathsf{StA}, \mathsf{StB})) \leftarrow \langle A[1 \cdots \nu - 1], B[1 \cdots \nu - 1] \rangle,$$
and $(\mathsf{LearnB}, \mathsf{AMsg}_\nu, \mathsf{H}) \leftarrow \dot{A}_\nu^O(\mathsf{StA}, \mathsf{Tr})$.

Proof of Lemma A.4. The proof follows in a straightforward from Lemmas A.7, A.8 and A.9. $\square$

We first prove Lemma A.7.

Proof of Lemma A.7. Let Bad be the event that
$$B_{\nu+1}^{O[\mathsf{H} \cup \mathsf{LearnB}]}(\mathsf{StB}, \mathsf{Tr}\|\mathsf{AMsg}_\nu) \neq B_{\nu+1}^O(\mathsf{StB}, \mathsf{Tr}\|\mathsf{AMsg}_\nu).$$

For notational convenience we let Ext_1 and Ext_2 be the whole executions traces of the following two executions:
$$\mathrm{Ext}_1 := B_{\nu+1}^{O[\mathsf{H} \cup \mathsf{LearnB}]}(\mathsf{StB}, \mathsf{Tr}\|\mathsf{AMsg}_\nu) \text{ and } \mathrm{Ext}_2 := B_{\nu+1}^O(\mathsf{StB}, \mathsf{Tr}\|\mathsf{AMsg}_\nu).$$

We show if $\overline{\mathsf{Miss}}$ and $\overline{\mathsf{Surprise}}$ hold, then $\overline{\mathsf{Bad}}$ must necessarily hold, hence proving the lemma. Suppose $\overline{\mathsf{Miss}} \wedge \overline{\mathsf{Surprise}}$ hold and to the contrary that Bad holds and consider the first query qu on which the two executions Ext_1 and Ext_2 are different. Since all non-eval queries are handled the same in both executions, we must have $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} ?)$ for some $\widetilde{F}$ and $\widetilde{X}$. Now since $\overline{\mathsf{Surprise}}$ holds we have $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{StB} \cup \mathsf{LearnB} \cup \mathsf{OrigA}$. On the other hand, by Assumption 4.26 and Definition 4.12 we cannot have $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{StB}$, because by the normal-form restriction the query qu will not be issued to the oracle in both executions. Also, we cannot have $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{LearnB}$ because in that case we will have $O[\mathsf{H} \cup \mathsf{LearnB}](qu) = O(qu)$. Thus, $(* \xrightarrow[\text{gc}]{} \widetilde{F}) \in \mathsf{OrigA}$. Now a simple inspection shows that Miss holds, a contradiction. $\square$

Proof of Lemma A.8. As in the proof of Lemma 4.23 we can easily show that whenever the event Surprise happens, we can win in the sense of Lemma 3.1. The bound now follows from Lemma 3.1. We omit the details. □

A.4.1 Proof of Lemma A.9.

We now focus on proving Lemma A.9. We first give Definitions A.10 and A.11 below to describe more compactly the random variables sampled during the compilation. We will then give some notation and terminology. We will then give a few lemmas from which Lemma A.9 is easily derived. We will then prove each of those lemmas.

Definition A.10. $\mathrm{DistGen}(\mathsf{Tr}, \mathsf{AMsg}_\nu, \mathsf{S})$

- **Offline phase:** Sample $(\widehat{O}, \widehat{\mathsf{state}}) \leftarrow \mathsf{NextBSamp}(\mathsf{Tr}, \mathsf{S})$

- **Semi-online phase:** Execute $B_{\nu+1}^{\widehat{O} \Diamond O[\mathsf{S}]}(\widehat{\mathsf{state}}, \mathsf{Tr}||\mathsf{AMsg}_\nu)$ and update S by adding all query/response pairs made to the oracle O. Let $\widehat{\mathsf{OrigB}}$ be the set of all query/response pairs during this decryption execution.

- **Return** $(\mathsf{S}, \widehat{\mathsf{state}}, \widehat{\mathsf{OrigB}})$

Definition A.11. Procedure $\mathrm{DistGen}_1(1^\kappa)$:

- **Output:** a tuple $\mathbf{Dist} := (\mathsf{StB}, \mathsf{Tr}, \mathsf{LearnB}, \mathsf{OrigA}, \mathsf{OrigB})$, as well as a sequence of tuples of the form $\mathbf{Dist_i} := (\widehat{\mathsf{StB}}_i, \mathsf{Tr}, \mathsf{LearnB}, \mathsf{OrigA}, \widehat{\mathsf{OrigB}}_i)$, which are sampled as follows.

- **Operations**

    1. Sample $(\mathsf{Tr}, state = (\mathsf{StA}, \mathsf{StB})) \leftarrow \langle A[1 \cdots \nu - 1], B[1 \cdots \nu - 1] \rangle$,
    2. Sample LearnB as in $\dot{A}_\nu(\mathsf{StA}, \mathsf{Tr})$.
    3. Sample $\mathsf{AMsg}_\nu \leftarrow A_\nu^O(\mathsf{StA}, \mathsf{Tr})$ and let OrigA contain all query/response pairs in this execution as well as all query/response pairs in StA.
    4. Let OrigB be the set of all query/response pairs during the execution of $B_{\nu+1}^O(\mathsf{StB}, \mathsf{Tr}||\mathsf{AMsg}_\nu)$.
    5. Let $\mathsf{S}_1 = \mathsf{LearnB} \cup \mathsf{OrigA}$. For $i \geq 1$ sample

    $$(\mathsf{S}_{i+1}, \widehat{\mathsf{StB}}_i, \widehat{\mathsf{OrigB}}_i) \leftarrow \mathrm{DistGen}(\mathsf{Tr}, \mathsf{AMsg}_\nu, \mathsf{S}_i).$$

Invalid samples. We say that a value $(\widehat{\mathsf{StB}}_i, \mathsf{Tr}, \mathsf{LearnB}, \mathsf{OrigA}, \widehat{\mathsf{OrigB}}_i)$ of the random variable $\mathbf{Dist_i}$ is *invalid* if $(\widehat{\mathsf{StB}}_i \cup \mathsf{LearnB} \cup \mathsf{OrigA} \cup \widehat{\mathsf{OrigB}}_i)$ is not a valid partial oracle (Definition 4.28). We let $\mathsf{invalid}_i$ be the event that $\mathbf{Dist_i}$ is invalid. We also define

$$\mathsf{invalid} = \mathsf{invalid}_1 \vee \cdots \vee \mathsf{invalid}_{t'} \tag{20}$$

Let $\widehat{\mathbf{Dist}}_{i,good}$ to be the induced distribution on $\widehat{\mathbf{Dist}}_i$ conditioned on $\overline{\mathsf{invalid}_i}$.

Lemma A.12. *Let t' be as in the compilation procedure. The two distributions $\mathbf{Dist}$ and $\widehat{\mathbf{Dist}}_{t'+1,good}$ are identically distributed.*

51

Proof. The proof can be easily done by inspection. $\qquad\square$

The proof of Lemma A.9 now follows from the following two lemmas.

Lemma A.13. *We have* $\Pr[\mathsf{Miss} \mid \overline{\mathsf{invalid}}] \leq \frac{2q}{t}$, *where the probability is taken over the variables sampled by* $\mathrm{DistGen}_1(1^\kappa)$ *(Definition A.11).*

Lemma A.14. *We have* $\Pr[\mathsf{invalid}] \leq \frac{1}{2^{\kappa/4}} + 3\varepsilon\nu t$, *where the probability is taken over the variables sampled by* $\mathrm{DistGen}_1(1^\kappa)$.

Proof of Lemma A.9. The proof follows immediately from Lemmas A.13 and Lemma A.14. $\qquad\square$

Proof of Lemma A.13. The proof is done similarly to the proof of Lemma 4.24. Relying on the notation (i, d)-*match* defined in the proof of Lemma 4.24, as in that proof we break the event Miss into smaller events $\mathsf{Miss}_{h,d}$, for $h \in [q]$ and $d \in \{0, 1\}$, defined as follows:

Event $\mathsf{Miss}_{h,d}$. The event that for some $\widetilde{F}$ that $((*, *) \xrightarrow[\mathsf{gc}]{} \widetilde{F}) \in \mathsf{OrigA}$, both the following hold:

1. OrigB contains an (h, d)-match for $\widetilde{F}$;

2. none of the sets $\widehat{\mathsf{OrigB}}_1, \ldots, \widehat{\mathsf{OrigB}}_{t'}$ do contain an (h, d)-match for $\widetilde{F}$.

As in the proof of Lemma 4.24, we can show that if Miss holds then $\mathsf{Miss}_{h,d}$ must hold, for some $h \in [q]$ and $d \in \{0, 1\}$. In the sequel we show

$$\Pr[\mathsf{Miss}_{h,d} \mid \overline{\mathsf{invalid}}] \leq \frac{1}{t},$$

which will imply the desired bound for the value of $\Pr[\mathsf{Miss} \mid \overline{\mathsf{invalid}}]$.

Fix $h \in [q]$ and $d \in \{0, 1\}$ for which we want to bound the event $\mathsf{Miss}_{h,d}$. Some notation first. For a sequence of tuples $(\mathrm{Dist}_1, \mathrm{Dist}_2, \ldots,)$ sampled from $(\mathbf{Dist}_1, \mathbf{Dist}_2, \ldots,)$ given in Definition A.11, we define a random variable $\mathbf{First}_{h,d}$ which takes a value in $\{1, \ldots, t+1\} \cup \{\bot\}$ as follows:

- $\mathbf{First}_{h,d} = i$ for $i \in [1, t+1]$ if

   1. $\widehat{\mathsf{OrigB}}_i$ contains an (h, d)-match for $\widetilde{F}$; and
   2. none of the sets $\widehat{\mathsf{OrigB}}_1, \ldots, \widehat{\mathsf{OrigB}}_{i-1}$ do contain an (h, d)-match for $\widetilde{F}$.

- $\mathbf{First}_{h,d} = \bot$ if none of the sets $\widehat{\mathsf{OrigB}}_1, \ldots, \widehat{\mathsf{OrigB}}_{t+1}$ do contain an (h, d)-match for $\widetilde{F}$.

Now applying Lemma A.12 we have that, for any (h, d):

$$\Pr[\mathsf{Miss}_{h,d} \mid \overline{\mathsf{invalid}}] = \Pr[\mathbf{First}_{h,d} = t' + 1]$$

where the random variable $\mathbf{First}_{h,d}$ is defined over the sequence of sampled tuples $(\mathrm{Dist}_1, \mathrm{Dist}_2, \ldots, \mathrm{Dist}_{t+1})$, where the ith tuple is sampled as $\mathrm{Dist}_i \leftarrow \widehat{\mathbf{Dist}}_{i,good}$.

The random variable $\mathbf{First}_{h,d}$ takes a value in $[1, t+1] \cup \{\bot\}$ according to some arbitrary distribution. However, since t' is chosen uniformly at random from $[1, t]$ we have:

$$\Pr[\mathbf{First}_{h,d} = t' + 1] \leq \frac{1}{t}.$$

$\qquad\square$

Proof of Lemma A.14. We show that for all $i \in [t]$,

$$\Pr[\mathsf{invalid}_i] \leq \frac{1}{2^{\kappa/3}} + 3\varepsilon\nu. \tag{21}$$

The proof of the lemma then follows by the union bound. Fix i.

We want to bound the probability that the tuple $\mathsf{Tup_i}$ sampled as

$$\mathsf{Tup_i} = (\mathsf{LearnB} \cup \widehat{\mathsf{StB}}_i \cup \mathsf{OrigA} \cup \widehat{\mathsf{OrigB}}_i) \leftarrow \mathbf{Dist}_i$$

is invalid. Note that by design the set

$$\mathsf{W} = \mathsf{LearnB} \cup \widehat{\mathsf{StB}}_i \cup \mathsf{OrigA}$$

makes up a valid partial oracle. This is because $\widehat{\mathsf{StB}}_i$ is sampled in a manner consistent with $\mathsf{LearnB} \cup \mathsf{OrigA}$.

We first claim that if Tup_i is invalid then $\widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i$ must be invalid. To prove this, suppose $\widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i$ is valid and suppose to the contrary that Tup_i is invalid. Noting that W is valid and also by assumption that $\widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i$ is valid, we consider all cases for the inconsistency of Tup_i:

1. For some $T \in \{f, \mathsf{gc}, \mathsf{gi}\}$ we have $(qu \xrightarrow{T} ans_1) \in \widehat{\mathsf{OrigB}}_i$ and $(qu \xrightarrow{T} ans_2) \in \mathsf{LearnB} \cup \mathsf{OrigA}$ and $ans_1 \neq ans_2$: Since $(qu \xrightarrow{T} ans_1) \in \widehat{\mathsf{OrigB}}_i$ this means that either $(qu \xrightarrow{T} ans_1) \in \widehat{\mathsf{StB}}_i$ or $O(qu) = ans_1$. The first case contradicts the fact that $\widehat{\mathsf{StB}}_i$ agrees with $\mathsf{LearnB} \cup \mathsf{OrigA}$ and the second case contradicts the fact that $\mathsf{LearnB} \cup \mathsf{OrigA}$ agree with O.

2. For some $T \in \{f, \mathsf{gc}, \mathsf{gi}\}$ we have $(qu_1 \xrightarrow{T} ans) \in \widehat{\mathsf{OrigB}}_i$ and $(qu_2 \xrightarrow{T} ans) \in \mathsf{LearnB} \cup \mathsf{OrigA}$ and $qu_1 \neq qu_2$: Like above, it can be proved this case cannot happen either.

3. For some and $\widetilde{F}$ and $\widetilde{X}$, $qu = ((\widetilde{F}, \widetilde{X}) \xrightarrow{\mathsf{eval}} \bot) \in \widehat{\mathsf{OrigB}}_i$ and qu is inconsistent with $\mathsf{LearnB} \cup \mathsf{OrigA}$: In order to have inconsistency, for some s we must at least have $((s, F) \xrightarrow{\mathsf{gc}} \widetilde{F}) \in \mathsf{LearnB} \cup \mathsf{OrigA}$. Now by design in that case for any S we will have $O[\mathsf{LearnB} \cup \mathsf{OrigA} \cup \mathsf{S}](qu) = O(qu)$ and thus we will also have $O(qu) = \bot$. Thus, there cannot be any inconsistencies.

Thus, to prove Equation 21 it suffices to show

$$\Pr[\widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i \text{ is invalid}] \leq \frac{1}{2^{\kappa/3}} + 3\varepsilon\nu. \tag{22}$$

We now focus on proving Equation 22. As the first observation, we note that we cannot have any input inconsistencies in $\widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i$: namely, a query qu for which we have $(qu \xrightarrow{T} ans_1) \in \widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i$ and $(qu \xrightarrow{T} ans_2) \in \widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i$ for $ans_1 \neq ans_2$. The reason that this cannot happen is because of the normal-form restriction and the definition of composed oracles (Definition 4.30).

Thus, it suffices to bound the probability of output collisions in $\widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i$. That is, we consider the inconsistencies between $\widehat{\mathsf{StB}}_i$ and $\widehat{\mathsf{OrigB}}_i$ that could happen as a result of having two different queries that have the same answer. We consider all possible cases that an output collision may occur.

1. Col_1: There exists $((s_1, F_1) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \widehat{\mathsf{StB}}_i$ such that $((s, F) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \widehat{\mathsf{OrigB}}_i \setminus \widehat{\mathsf{StB}}_i$ and that $(s, F) \neq (s_1, F_1)$.

2. Col_2: There exists $((s_1, x_1) \underset{\mathsf{gi}}{\longrightarrow} \widetilde{x}) \in \widehat{\mathsf{StB}}_i$ such that $((s, x) \underset{\mathsf{gi}}{\longrightarrow} \widetilde{x}) \in \widehat{\mathsf{OrigB}}_i \setminus \widehat{\mathsf{StB}}_i$ and that $(s, x) \neq (s_1, x_1)$.

3. Col_3: There exists $((\widetilde{F}, \widetilde{X}) \underset{\mathsf{eval}}{\longrightarrow} \bot) \in \widehat{\mathsf{OrigB}}_i$ such that $\mathsf{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$ and $((s, F) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i$ for some s and F.

We mention that the event Col_3 may not necessarily cause an inconsistency by itself, but if some query $qu = ((\widetilde{F}, \widetilde{X}) \underset{\mathsf{eval}}{\longrightarrow} \bot) \in \widehat{\mathsf{OrigB}}_i$ does cause an inconsistency in $\widehat{\mathsf{StB}}_i \cup \widehat{\mathsf{OrigB}}_i$, then the conditions in Case 3 must hold.

We will now show that each of the events Col_1, Col_2 and Col_3 may occur with probability at most $\frac{1}{2^{\kappa/2}} + \varepsilon\nu$, proving Equation 22.

Bounding the event Col_1:   We show

$$\Pr[\mathsf{Col}_1] \leq \tfrac{1}{2^{\kappa/2}} + \nu\varepsilon. \tag{23}$$

To prove this, first note that Col_1 implies that $\widetilde{F}$ is a valid garbled circuit, and that we indeed have $\mathsf{gc}(s, F) = \widetilde{F}$. This is because the real oracle O agrees with $\widehat{\mathsf{OrigB}}_i \setminus \widehat{\mathsf{StB}}_i$. We claim

$$\Pr[((s, F) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \mathsf{StB}] \geq 1 - \tfrac{1}{2^{\kappa/2}}. \tag{24}$$

We first show how to derive Equation 23 from Equation 24. Notice that $((s, F) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \notin \mathsf{LearnB}$. This is because $\widehat{\mathsf{StB}}_i$ agrees with LearnB and so we could not have had $((s_1, F_1) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \widehat{\mathsf{StB}}_i$. Thus, from Equation 24 we have $\Pr[((s, F) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \mathsf{StB} \setminus \mathsf{LearnB}] \geq 1 - \tfrac{1}{2^{\kappa/2}}$. Now recall that, by the definition of the event Col_1 we have $((s, F) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \widehat{\mathsf{OrigB}}_i$. Now applying Equation 24 along with Lemma 4.41 and Lemma 4.32, Equation 23 is established.

We now prove Equation 24. Let Surp' be the event that $((s, F) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \notin \mathsf{StB}$. We will show $\Pr[\mathsf{Surp}'] \leq \tfrac{1}{2^{\kappa/2}}$ by reducing it to Lemma 3.1. First, note that $\widetilde{F}$ is valid garbled circuit under gc, and that it is constructible using $\widehat{\mathsf{StB}}_i$ and that

$$\widehat{\mathsf{StB}}_i \leftarrow \mathsf{NextBSamp}(\mathsf{Tr}, \mathsf{LearnB} \cup \mathsf{S}_{i-1}).$$

for some set S_i which completely agrees with O and so we cannot have the query/response $((*, *) \underset{\mathsf{gc}}{\longrightarrow} \widetilde{F}) \in \mathsf{S}_i$. Thus, the occurrence of the event Surp' implies we could forge a valid $\widetilde{F}$ "on the fly" and so by Lemma 3.1 we have $\Pr[\mathsf{Surp}'] \leq \tfrac{1}{2^{\kappa/2}}$.

Bounding the event Col_2.   Exactly as in Col_1, we can show that $\Pr[\mathsf{Col}_2] \leq \tfrac{1}{2^{\kappa/2}} + \varepsilon\nu$.

Bounding the event Col_3. Suppose Col_3 holds: Namely, there exists $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} \bot) \in \widehat{\text{OrigB}}_i$ such that $\text{eval}(\widetilde{F}, \widetilde{X}) \neq \bot$ and $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \widehat{\text{StB}}_i \cup \widehat{\text{OrigB}}_i$ for some s and F. First, we argue that we cannot have $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \widehat{\text{StB}}_i$. This is because, otherwise, the query/response $((s, F) \xrightarrow[gc]{} \widetilde{F})$ will be included in $\widehat{\text{StB}}_i$, and so by Assumption 4.12, during $B^{\widehat{O_i} \Diamond \overline{O}[S]}_{\nu+1}(\widehat{\text{StB}}_i, \text{Tr}||\text{AMsg}_\nu)$ the query $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} ?)$ will never be asked.

Now since $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \notin \widehat{\text{StB}}_i$, we have $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \widehat{\text{OrigB}}_i$ and since $\widehat{\text{OrigB}}_i \neq \widehat{\text{StB}}_i$ agrees with O on all non-eval queries we have $\text{gc}(s, F) = \widetilde{F}$. Thus, $\widetilde{F}$ is a valid garbled circuit.

Thus, it suffices to bound the probability of the following event that we call Col': there exists $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} \bot) \in \widehat{\text{OrigB}}_i$ such that $\widetilde{F}$ is valid and $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F}) \in \widehat{\text{OrigB}}_i$ for some s and F. Notice that by the normal-form restriction, the event Col' happens only when the query/response $((\widetilde{F}, \widetilde{X}) \xrightarrow[\text{eval}]{} *) \in \widehat{\text{OrigB}}_i$ is added to $\widehat{\text{OrigB}}_i$ *before* the query/response $((s, F) \xrightarrow[\text{gc}]{} \widetilde{F})$ is created. Now exactly as in Col_1 we can show that the probability that Col' happens is at most $\frac{1}{2^{\kappa/2}} + \nu\varepsilon$.

The proof of Lemma A.14 is now complete.

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $\square$

B NIWI from Ideal Garbling

Brakerski et al. [BKSY11] show that a certain class of non-black-box techniques, when applied to one-way functions, is not sufficient to yield PKE. More specifically, they show that there exists no fully black-box constructions of perfectly-complete PKE (and even perfectly-complete constant-round key exchange protocols) from OWFs and non-interactive witness indistinguishability (NIWI) protocols for relations with OWF gates.[7] We will show that the class of garbling-based non-black-box techniques that we rule out (i.e., GC-OWF, Definition 3.6) does in fact already capture NIWI based techniques.

To proceed formally, we will first review the primitive of NIWI-OWF, which, as shown by Brakerski et al. [BKSY11], does not imply perfectly-complete PKE in a black-box way. Afterward, we will show that NIWI-OWF can be instantiated in a black-box way using our GC-OWF oracle (Definition 4.3).

By doing this we will extend the result of [BKSY11] by ruling out all fully-black-box constructions of PKE (and constant-round key exchange protocols) — which may not necessarily be perfectly complete — from NIWI-OWF.

We first start with some definitions to formalize the notion of PKE constructions from NIWI and one-way functions.

Definition B.1 (Oracle circuit satisfiability). Fix an oracle f. The language Sat^f contains oracle-aided binary-output circuits C^f, where $C^f \in \text{Sat}^f$ iff there exists an input x such that $C^f(x) = 1$. We call x a witness for C^f.

Definition B.2 (NIWI for Sat^f). Fix an oracle f. We say that a non-interactive protocol $\langle \text{NIP}, \text{NIV} \rangle$ is complete, sound and witness indistinguishable for Sat^f if all the following hold:

---

[7]Whatever we say about NIWI also holds about NIZK, because one may use standard techniques [FLS90] to go from NIWI to NIZK (in the CRS model).

- **Completeness**: for any circuit $C^f \in \mathsf{Sat}^f$ and witness x for C^f we have

$$\Pr[\mathsf{NIV}(1^\kappa, C^f, \pi) = 1] = 1,$$

where $\pi \leftarrow \mathsf{NIP}(1^\kappa, C^f, x)$.

- **Soundness**: for any circuit $C^f \notin \mathsf{Sat}^f$ and any (alleged) proof π^*, we have $\Pr[\mathsf{NIV}(1^\kappa, C^f, \pi^*) = 1] = \mathrm{negl}(\kappa)$.

- **Witness Indistinguishability**: For any $C^f \in \mathsf{Sat}^f$ and any two witnesses x_1 and x_2 for C^f, the two distributions $(x_1, x_2, \mathsf{NIP}(C^f, x_1))$ and $(x_1, x_2, \mathsf{NIP}(C^f, x_1))$ are indistinguishable.

We now formally define a black-box construction of PKE from NIWI-OWF.

Definition B.3 (PKE from NIWI-OWF)**.** A fully black-box construction of PKE from one-way functions and NIWI protocols for circuits with one-way function gates (shortly, from NIWI-OWF) consists of two oracle-aided PPT algorithms $(\mathrm{PKE}, \mathrm{Red} := (\mathrm{Red}_1, \mathrm{Red}_2))$ for which the following holds: for any function f and any sound-and-complete non-interactive protocol $\mathsf{NIWI} = (\mathsf{NIP}, \mathsf{NIV})$ for Sat^f, we have

- $\mathrm{PKE}^{f, \mathsf{NIWI}}$ is a correct implementation of PKE;

- For any adversary $\mathcal{A}$ that breaks the semantic security of $\mathrm{PKE}^{f, \mathsf{NIWI}}$, either $\mathrm{Red}_1^{f, \mathsf{NIWI}, \mathcal{A}}$ breaks the one-wayness of f or $\mathrm{Red}_2^{f, \mathsf{NIWI}, \mathcal{A}}$ breaks the witness indistinguishability of NIWI for Sat^f.

We are now ready to state our extension of the result of Brakerski et al [BKSY11].

Theorem B.4 (Impossibility of PKE from NIWI+OWF)**.** *There exists no fully black-box construction of PKE from OWFs and NIWI protocols for circuits with one-way functions gates.*

Brakerski et al. [BKSY11] proved the above theorem for the case of perfectly complete PKE. Our extension relaxes this restriction to also rule out PKE possibly with decryption error.

Roadmap. We will show the existence of NIWI-OWF with respect to measure one of our oracles $O = (f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}, \mathsf{rev})$. To this end, in Section B.1 we first show how to prove this for the case of interactive WI protocols (which we simply refer to as WI protocols). Then, in Section B.2 we will show how to make the protocol of Section B.1 non-interactive using our oracles.

B.1 WI-OWF Using Garbling

As mentioned earlier we will first show in this section how to build WI protocols for Sat^f. This construction does not require the random nature of our oracles, and can in fact be realized using any garbling scheme with natural specific properties, as we discuss below.

WI-OWF. Fix f. We give a three round WI protocol for Sat^f, built generically using oracle access to $(f, \mathcal{GS})$, where $\mathcal{GS} = (\mathsf{Garb}, \mathsf{Eval})$ is a garbling scheme relative to f gates (Definition 3.4) satisfying a specific property. We describe this property below.

Recall that correctness of a garbling scheme implies that the garbling of a non-satisfiable circuit can never be evaluated to 1 if the corresponding garbled labels were generated honestly. The

following definition extends this by demanding that even maliciously generated labels (i.e., those that are not outputted by the garbling algorithm) cannot make the garbled circuit evaluate to one. Moreover, we require that maliciously generated garbled circuits can never be evaluated to anything but $\bot$.

Definition B.5 (Non-ambiguous garbling schemes). Fix a function f and let $\mathcal{GS} = (\mathsf{Garb}, \mathsf{Eval})$ be a garbling scheme relative to f. We say that $\mathcal{GS} = (\mathsf{Garb}, \mathsf{Eval})$ is non-ambiguous if both the following properties hold.

1. For any $\widetilde{\mathsf{C}}$ for which there does not exist a seed $s \in \{0,1\}^*$ and circuit C satisfying $\mathsf{Garb}(s, \mathsf{C}) = (\widetilde{\mathsf{C}}, *)$, we have that $\mathsf{Eval}^f(\widetilde{\mathsf{C}}, \widetilde{X}) = \bot$, for any (arbitrary generated) $\widetilde{X}$.

2. For any oracle-aided circuit $\mathsf{C}^f \notin \mathsf{Sat}$, any randomness $s \in \{0,1\}^*$, assuming $\mathsf{Garb}(s, \mathsf{C}) = (\widetilde{\mathsf{C}}, *)$, for no (arbitrary generated) $\widetilde{X}$ do we have $\mathsf{Eval}^f(\widetilde{\mathsf{C}}, \widetilde{X}) = 1$.

Construction B.6 (WI protocol $\langle \mathsf{P}, \mathsf{V} \rangle$ for Sat^f using non-ambiguous garbling). We will now describe our construction of a WI protocol for Sat^f using blackbox access to a non-ambiguous garbling scheme $\mathcal{GS} := (\mathsf{Garb}, \mathsf{Eval}, \mathsf{Sim})$ for f-gate circuits.

- **Common inputs of P and V**: $(1^\kappa, r, \mathsf{C})$: a security parameter 1^κ, a polynomial $r = r(\kappa)$ and a circuit C. The value of r will determine the soundness error. Also, let $n := |\mathsf{C}|$ and let m denote the input size of C.

- **Private input of P**: $x \in \{0,1\}^m$, where $\mathsf{C}^f(x) = 1$.

- **Interactive Phase**:

    - $\mathsf{P}^{f,\mathcal{GS}:=(\mathsf{Garb},\mathsf{Eval})}(\mathsf{C}^f, x)$: for all $i \in [r]$, choose randomness $\mathsf{seed}_i \leftarrow \{0,1\}^*$ and make the query $((\mathsf{seed}_i, \mathsf{C}^f) \xrightarrow[\mathsf{Garb}]{} ?)$ to obtain

    $$(\widetilde{\mathsf{C}}_i, \mathsf{label}_{1,0}^{(i)}, \mathsf{label}_{1,1}^{(i)}, \ldots, \mathsf{label}_{m,0}^{(i)}, \mathsf{label}_{m,1}^{(i)}) := \mathsf{Garb}(\mathsf{seed}_i, \mathsf{C}).$$

    Return
    $$\mathsf{msg} := (\widetilde{\mathsf{C}}_1, \ldots, \widetilde{\mathsf{C}}_r). \tag{25}$$

    - $\mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f, \mathsf{msg})$: return a challenge message $\mathsf{ch} \leftarrow \{0,1\}^r$.
    - $\mathsf{P}(\mathsf{ch})$: send $\mathsf{msg}' := (\mathsf{msg}_1', \ldots, \mathsf{msg}_r')$ to V, where for $i \in [r]$, msg_i' is formed as:
        * if $\mathsf{ch}_i = 0$, then $\mathsf{msg}_i' := (\mathsf{label}_{1,x_1}^{(i)}, \cdots, \mathsf{label}_{m,x_m}^{(i)})$
        * if $\mathsf{ch}_i = 1$, then $\mathsf{msg}_i' := \mathsf{seed}_i$

- **Decision Phase**: $\mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f, \mathsf{msg}, \mathsf{ch}, \mathsf{msg}_1', \ldots, \mathsf{msg}_r')$: Parse $\mathsf{msg} := (\widetilde{\mathsf{C}}_1, \ldots, \widetilde{\mathsf{C}}_r)$. Return 1 if the following condition holds and return 0 otherwise: for all $i \in [r]$ we must have:

    - if $\mathsf{ch}_i := 0$, then $\mathsf{Eval}(\widetilde{\mathsf{C}}_i, \mathsf{msg}_i') = 1$;
    - if $\mathsf{ch}_i := 1$, then $\mathsf{Garb}(\mathsf{msg}_i', \mathsf{C}^f) = (\widetilde{\mathsf{C}}_i, *)$.

We will now prove all the properties required by the WI scheme.

Perfect completeness. By simple inspection, we can see that for all circuits C^f, input x where $\mathsf{C}^f(x) = 1$ we have

$$\Pr[\langle \mathsf{P}^{f,\mathcal{GS}}(\mathsf{C}^f, x), \mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f) \rangle = 1] = 1.$$

Statistical soundness. Assuming $\mathcal{GS} = (\mathsf{Garb}, \mathsf{Eval})$ is non-ambiguous, the following holds: for all unbounded provers P^* and for all $\mathsf{C}^f \notin \mathsf{Sat}$:

$$\Pr[\langle \mathsf{P}^*(\mathsf{C}^f), \mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f) \rangle = 1] \leq \frac{1}{2^r}. \tag{26}$$

Proof of statistical soundness. Fix P^* and $\mathsf{C}^f \notin \mathsf{Sat}$. We prove Equation 26. We first give some terminologies. Call a garbled circuit $\widetilde{\mathsf{C}}$ matching if there exists some randomness seed for which we have $\mathsf{Garb}(\mathsf{seed}, \mathsf{C}^f) = (\widetilde{\mathsf{C}}, *)$. Call $\widetilde{\mathsf{C}}$ non-matching if $\widetilde{\mathsf{C}}$ is not matching.

Let $(\widetilde{\mathsf{C}}_1, \ldots, \widetilde{\mathsf{C}}_n)$ be the first message that P^* sends to $\mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f)$. Let match be the set that contains all indices $i \in [r]$ where $\widetilde{\mathsf{C}}_i$ is matching. Let $\mathsf{nonmatch}$ contains all indices $i \in [r]$ where $\widetilde{\mathsf{C}}_i$ is non-matching. Note that $\mathsf{match} \cup \mathsf{nonmatch} = \{1, \ldots, r\}$.

Let $\mathsf{ch} \leftarrow \{0, 1\}^r$ be the random challenge that $\mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f)$ sends. We now claim that the probability that $\mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f)$ outputs 1 at the end of the protocol is the probability that $\mathsf{ch}_i = 1$ for all $i \in \mathsf{match}$ and that $\mathsf{ch}_j = 0$ for all $j \in \mathsf{nonmatch}$. Assuming this claim is true, our desired probability bound of $\frac{1}{2^r}$ follows. Thus in the sequel we will show why this claim holds.

To prove the above claim, we show if for some $i \in \mathsf{match}$, $\mathsf{ch}_i = 0$ then $\mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f)$ will definitely reject. Similarly, if for some $i \in \mathsf{nonmatch}$, $\mathsf{ch}_i = 1$ then $\mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f)$ will definitely reject.

- For some $i \in \mathsf{match}$ we have $\mathsf{ch}_i = 0$: In order for $\mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f)$ not to catch P^* on this index i, P^* must send some string $\widetilde{X}$ for which we have $\mathsf{Eval}(\widetilde{\mathsf{C}}, \widetilde{X}) = 1$. This is however impossible for P^* to do because $i \in \mathsf{match}$ means that $\widetilde{\mathsf{C}}$ is a valid garbled circuit for C^f, and since $\mathsf{C}^f \notin \mathsf{Sat}$, by non-ambiguity of $\mathcal{GS}$ the adversary P^* cannot find such $\widetilde{X}$.

- For some $i \in \mathsf{nonmatch}$ we have $\mathsf{ch}_i = 1$: In order for $\mathsf{V}^{f,\mathcal{GS}}(\mathsf{C}^f)$ not to catch P^* on this index i, P^* must send some randomness seed which satisfies $\mathsf{Garb}(\mathsf{seed}, C) = (\widetilde{\mathsf{C}}_i, *)$. However, this is impossible to do, because $i \in \mathsf{nonmatch}$ means that there exists no seed which satisfies $\mathsf{Garb}(\mathsf{seed}, \mathsf{C}^f) = (\widetilde{\mathsf{C}}, *)$.

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Witness Indistinguishability. For any circuit $\mathsf{C}^f \in \mathsf{Sat}$, any two witnesses x_0 and x_1 for C^f, and any PPT (but not necessarily honest) $\mathsf{V}_1^{f,\mathcal{GS}}$ we have

$$\Pr[\mathbf{Expr}(\mathsf{P}, \mathsf{V}_1, \mathsf{C}^f, x_0, x_1) = 1] \leq \frac{1}{2} + \mathsf{negl}(\kappa),$$

where the experiment $\mathbf{Expr}(\mathsf{P}, \mathsf{V}_1, \mathsf{C}^f, x_0, x_1)$ is defined as follows.

1. $b \leftarrow \{0, 1\}$;

2. $\mathsf{msg} \leftarrow \mathsf{P}^{f,\mathcal{GS}}(\mathsf{C}^f, x_b)$

3. $\mathsf{ch} \leftarrow \mathsf{V}_1^{f,\mathcal{GS}}(\mathsf{C}^f, x_0, x_1, \mathsf{msg})$

4. $\mathsf{msg}' \leftarrow \mathsf{P}^{f,\mathcal{GS}}(\mathsf{ch})$

5. $b' \leftarrow \mathsf{V}_1^{f,\mathcal{GS}}(\mathsf{msg}')$

6. output 1 if $b = b'$ and 0 otherwise.

Proof of witness indistinguishability. Let msg and msg' be as in the above experiment. First, note that msg is completely independent of the underlying challenge witness x_b. This is because: to generate msg, we sample r randomness values $\mathsf{seed}_1, \ldots, \mathsf{seed}_r$, we set for $i \in [r]$

$$(\widetilde{\mathsf{C}}_i, \mathsf{label}_{1,0}^{(i)}, \mathsf{label}_{1,1}^{(i)}, \ldots, \mathsf{label}_{m,0}^{(i)}, \mathsf{label}_{m,1}^{(i)}) := \mathsf{Garb}(\mathsf{seed}_i, \mathsf{C}),$$

and we output $\mathsf{msg} := (\widetilde{\mathsf{C}}_1, \ldots, \widetilde{\mathsf{C}}_n)$. Parse $\mathsf{msg}' = (\mathsf{msg}_1', \ldots, \mathsf{msg}_r')$. For any $i \in [r]$ where $\mathsf{ch}_i = 1$, we have $\mathsf{msg}_i' = \mathsf{seed}_i$, and thus msg_i' is again completely independent of the challenge witness x_b. Thus, we only focus on msg_i' where $\mathsf{ch}_i = 0$. Suppose for an index i we have $\mathsf{ch}_i = 0$. For rotational simplicity, set $y := x_b$. By the design of the protocol we have

$$\mathsf{msg}_i' := \mathbf{lb} := (\mathsf{label}_{1,y_1}^{(i)}, \cdots, \mathsf{label}_{m,y_m}^{(i)}).$$

We now claim that msg_i' is computationally independent of the challenge witness x_b. The reason is that by the simulation security of the garbling scheme we have $(\widetilde{\mathsf{C}}_i, \mathbf{lb}) \equiv^c \mathsf{Sim}(1^\kappa, n, m, \mathsf{C}^f(y))$, and since $\mathsf{C}^f(y) = 1$ we obtain that $\mathbf{lb}$ is computationally independent of x_b. The proof is complete. $\square$

B.2 Making the WI protocol Non-Interactive

We will now show that NIWI-OWF exists relative to measure one of our GC-OWF ideal oracles (Lemma B.8). We will then use this to derive the following impossibility: there exists no fully black-box construction of PKE from NIWI-OWF (Theorem B.4).

Construction B.7 (NIWI-OWF relative to our oracles). We now describe a non-interactive protocol $\mathsf{NIWI} := \langle \mathsf{NIP}, \mathsf{NIV} \rangle$ for the language Sat^f, which is non-interactive version of the WI protocol $\langle \mathsf{P} := (\mathsf{P}_1, \mathsf{P}_2), \mathsf{V} \rangle$ in Construction B.6.

- $\mathsf{NIP}^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}}(\mathsf{C}^f, x)$: send $(\mathsf{msg}, \mathsf{msg}')$ to NIV, where

$$\mathsf{msg} \leftarrow \mathsf{P}_1^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}}(\mathsf{C}^f, x)$$
$$\mathsf{msg}' \leftarrow \mathsf{P}_2^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}}(\mathsf{ch}), \text{ where } \mathsf{ch} = f(\mathsf{msg})$$

- $\mathsf{NIV}^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}}(\mathsf{msg}, \mathsf{msg}')$: call $(\mathsf{msg} \underset{f}{\rightarrow} ?)$ to get ch and return the decision output bit of $\mathsf{V}^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}}(\mathsf{msg}, \mathsf{ch}, \mathsf{msg}')$.

Lemma B.8 (NIWI-OWF exists relative to our oracles). *Let $\mathbf{O}$ be our garbling oracle distribution (Definition 4.3). For measure one of oracles $(f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval}, \mathsf{rev})$ outputted by $\mathbf{O}$ the following holds: the non-interactive protocol $\mathsf{NIWI}^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}}$ above is witness indistinguishable for the language Sat^f.*

Proof. We give a sketch of the proof. We need to prove completeness, soundness and witness indistinguishability for $(\mathsf{NIP}^{f,\mathsf{gc},\mathsf{gi}}, \mathsf{NIV}^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}})$.

- **Perfect completeness**: Immediate by perfect completeness of $\mathsf{WI}^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}}$.

- **Statistical soundness**: This follows by the fact that our oracle $(f, \mathsf{gc}, \mathsf{gi}, \mathsf{eval})$ provides non-ambiguous garbling. Also, recall that the statistical soundness of $\mathsf{WI}^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval}}$ depends on the condition that the verifier's challenge message ch is chosen uniformly at random and independently of msg_1. Now since under our NIWI protocol we have $\mathsf{ch} = f(\mathsf{msg}_1)$, for a random f we will have this required condition.

- **Witness indistinguishability**: For any circuit $\mathsf{C}^f \in \mathsf{Sat}$, any two witnesses x_0 and x_1 for C^f, no PPT adversary $\mathcal{A}^{f,\mathsf{gc},\mathsf{gi},\mathsf{eval},\mathsf{rev}}$ can distinguish between the outputs of $\mathsf{NIP}^{f,\mathsf{gc},\mathsf{gi}}(\mathsf{C}, x_1)$ and $\mathsf{NIP}^{f,\mathsf{gc},\mathsf{gi}}(\mathsf{C}, x_2)$. We stress that this result holds even if the distinguisher $\mathcal{A}$ has oracle access to rev.

$\square$

Proof of Theorem B.4. Combine our main impossibility (that states PKE does not exist relative to our oracles) with Lemma B.8. $\square$