

Adaptively Secure Garbled Circuits from One-Way Functions

Brett Hemenway¹, Zahra Jafargholi², Rafail Ostrovsky^{3,*},
Alessandra Scafuro^{2,4,**}, and Daniel Wichs^{2,***}

¹ University of Pennsylvania, fbrett@cis.upenn.edu

² Northeastern University, {zahra|wichs}@ccs.neu.edu

³ University of California, Los Angeles, rafail@cs.ucla.edu

⁴ Boston University, scafuro@bu.edu

Abstract. A garbling scheme is used to garble a circuit C and an input x in a way that reveals the output $C(x)$ but hides everything else. In many settings, the circuit can be garbled *off-line* without strict efficiency constraints, but the input must be garbled very efficiently *on-line*, with much lower complexity than evaluating the circuit. Yao’s garbling scheme [31] has essentially optimal on-line complexity, but only achieves *selective security*, where the adversary must choose the input x prior to seeing the garbled circuit. It has remained an open problem to achieve *adaptive security*, where the adversary can choose x after seeing the garbled circuit, while preserving on-line efficiency.

In this work, we modify Yao’s scheme in a way that allows us to prove adaptive security under one-way functions. In our main instantiation we achieve on-line complexity only proportional to the width w of the circuit. Alternatively we can also get an instantiation with on-line complexity only proportional to the depth d (and the output size) of the circuit, albeit incurring in a $2^{O(d)}$ security loss in our reduction. More broadly, we relate the on-line complexity of adaptively secure garbling schemes in our framework to a certain type of *pebble* complexity of the circuit. As our main tool, of independent interest, we develop a new notion of *somewhere equivocal* encryption, which allows us to efficiently equivocate on a small subset of the message bits.

Keywords: Adaptive Security, Garbled Circuits, Online/Offline Two-party Computation.

* Supported in part by NSF grants 09165174, 1065276, 1118126 and 1136174, US-Israel BSF grant 2008411, OKAWA Foundation Research Award, IBM Faculty Research Award, Xerox Faculty Research Award, B. John Garrick Foundation Award, Teradata Research Award, Lockheed-Martin Corporation Research Award and by DARPA Safeware program. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

** Supported by NSF grants 1012798, CNS-1414119.

*** Supported by NSF grants CNS-1347350, CNS-1314722, CNS-1413964. This work was done in part while some of the authors were visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467.

1 Introduction

Garbled Circuits. A *garbling scheme* (also referred to as a randomized encoding) can be used to garble a circuit C and an input x to derive a garbled circuit \tilde{C} and a garbled input \tilde{x} . It’s possible to evaluate \tilde{C} on \tilde{x} and get the correct output $C(x)$. However, the garbled values \tilde{C}, \tilde{x} should not reveal anything else beyond this. In particular, there is a simulator that can simulate \tilde{C}, \tilde{x} given only $C(x)$.

The notion of garbled circuits was introduced by Yao in (oral presentations of) [31,32], and can be instantiated based on one-way functions. Garbled circuits have since found countless applications in diverse areas of cryptography, most notably to secure function evaluation (SFE) starting with Yao’s work, but also in parallel cryptography [5,6], verifiable computation [16,7], software protection [20,22], functional encryption [30,21,19], key-dependent message security [9,3], obfuscation [4] and many others. These applications rely on various efficiency/functionality properties of garbled circuits and a comprehensive study of this primitive is explored in the work of Bellare, Hoang and Rogaway [12].

On-line Complexity. In many applications, the garbled circuit \tilde{C} can be computed in an *off-line* pre-processing phase before the input is known, and therefore the efficiency of this procedure may not be of paramount importance. On the other hand, once the input x becomes available in the *on-line* phase, creating the garbled input \tilde{x} should be extremely efficient. Therefore, the main efficiency measure that we consider here is the *on-line complexity*, which is the time it takes to garble an input x , and hence also a bound on the size of \tilde{x} . Ideally, the on-line complexity should only be linear in the input size $|x|$ and independent of the potentially much larger circuit size $|C|$.⁵

Yao’s Scheme. Yao’s garbling scheme already achieves essentially optimal on-line complexity, where the time to garble an input x and the size of \tilde{x} are only linear in the input size $|x|$, independent of the circuit size.⁶ However, it only realizes a weak notion of security called *selective security*, which corresponds to a setting where adversary must choose the input x before seeing the garbled circuit \tilde{C} . In particular, the adversary first chooses both C and x and then gets the garbled values \tilde{C}, \tilde{x} which are either correctly computed using the “real” garbling scheme or “simulated” using only $C(x)$. The adversary should not be able to distinguish between the real world and the simulated world.

Selective vs. Adaptive Security. Selective security is often unsatisfactory in precisely the scenarios envisioned in the off-line/on-line setting, where the garbled circuit \tilde{C} is given out first and the garbled input \tilde{x} is only given out later. In such settings, the adversary may be able to (partially) influence the choice of the

⁵ Note that, without any other restrictions on the structure of the garbling scheme, there is a trivial scheme where \tilde{C} is empty and $\tilde{x} = C(x)$, whose on-line complexity is proportional to $|C|$.

⁶ More precisely, in Yao’s garbled circuits, the garbled input is of size $|x| \cdot \text{poly}(\lambda)$ where λ is the security parameter. The work of [8] shows how to reduce this to $|x| + \text{poly}(\lambda)$ assuming stronger assumptions such as DDH, RSA or LWE.

input x after seeing the garbled circuit \tilde{C} . Therefore, we need a stronger notion called *adaptive security*, defined via the following two stage game:

1. The adversary chooses a circuit C and gets the garbled circuit \tilde{C} .
2. After seeing \tilde{C} the adversary adaptively chooses an input x and gets the garbled input \tilde{x} .

In the real world \tilde{C}, \tilde{x} are computed correctly using the garbling scheme, while in the simulated world they are created by a simulator who only gets the output $C(x)$ in step (2) of the game but does not get the input x . The adversary should not be able to distinguish these two worlds.

The work of Bellare, Hoang and Rogaway [11] gave the first thorough treatment of adaptively secure garbling schemes and showed that this notion is crucial in many applications. They point out that it remains unknown whether Yao’s garbling scheme or any of its many variants can satisfy adaptive security, and the proof techniques that work in the selective security setting do not extend to the adaptive setting. They left it as the main open problem to construct adaptively secure garbling schemes where the on-line complexity is smaller than the circuit size.⁷

Finally we emphasize that the problem of achieving adaptively secure garbled circuits is different from the problem of achieving adaptively secure two-party computation (with constant rounds) using an approach based on garbled circuits. The latter means that the adversary can *corrupt the players* adaptively. It is not known whether either problem can be reduced to the other.

1.1 Prior Approaches to Adaptive Security

Lower Bound and Yao’s scheme. The work of Applebaum et al. [8] (see also [24]) gives a lower bound on the on-line complexity of circuit garbling in the adaptive setting, showing that the size of the garbled input \tilde{x} must exceed the *output size* of the circuit. This is in contrast to the selective security setting, where Yao’s garbling scheme achieves on-line complexity that depends only on the input size and not the output size. In particular, this shows that Yao’s garbling scheme cannot directly be adaptively secure.

Complexity Leveraging. It turns out that there is a simple and natural modification of Yao’s garbling scheme (i.e., by withholding the mapping of output-wire keys to output bits until the on-line phase) that would match the above lower bound and could plausibly be conjectured to provide adaptive security. In fact, one can prove that the above variant of Yao’s scheme is secure in the adaptive

⁷ The adaptive security notion we described, is denoted `prv1` by [11]. They also consider a stronger variant called `prv2`, where the adversary adaptively chooses bits of the input x one at a time and gets the corresponding bits of the garbled input \tilde{x} . They show that there is an efficiency preserving transformation from `prv1` to `prv2` following the ideas from [20]. Therefore, in this work we can focus solely on achieving `prv1`.

setting using *complexity leveraging*, but only at a 2^n security loss in the reduction, where n is the input size. There is no known proof of security that avoids this loss.⁸

One-Time Pad and Random-Oracles. An alternate approach, suggested by [11], is to use one-time pad encryption to encrypt a Yao garbled circuit in the off-line phase and then provide the decryption key with the garbled input in the on-line phase. Intuitively, since a one-time pad encryption is “non-committing” and the ciphertext can be *equivocated* to any possible message by providing a corresponding key, the adversary does not gain any advantage in seeing such a ciphertext in the off-line phase. Unfortunately, this solution blows up the on-line complexity to be at least as large as the circuit size.

The work of [11] also noted that one can replace the one-time pad encryption in the above solution with a random-oracle based encryption scheme, which can be equivocated by programming random oracle outputs. This gives an adaptively secure garbled circuit construction with optimal parameters in the random oracle model. In fact, this approach can even be used to prove security in parameter regimes that beat the lower bound of [8], and therefore we should be suspicious about its implications in the standard model, when the random oracle is replaced by a hash function. In particular, the construction is using the random oracle for equivocation in ways that we know to be uninstantiable in the standard model [29].

UCE-Security. Bellare et al. [10] show that a variant of Yao garbled circuits (which does not violate the lower bound of [8]) can be proven secure when instantiated with a hash function that satisfies a security notion called *Universal Computational Extractor (UCE)* security. However, UCE is a strong, non-standard and non-falsifiable assumption.

Heavy Hammers. Lastly, we mention two approaches that get adaptively secure garbled circuits with good on-line complexity under significantly stronger assumptions than one-way functions. The work of Boneh et al. [13] implicitly provides such schemes where the on-line complexity is proportional to the input/output size and the depth d of the circuit, under the *learning with errors* assumption with a modulus-to-noise ratio of $2^{\text{poly}(d)}$. This translates to assuming the hardness of lattice problems with $2^{\text{poly}(d)}$ approximation factors. The work of Ananth and Sahai [2] shows how to get an essentially optimal scheme, where the on-line complexity is only proportional to the input/output size of the circuit, assuming *indistinguishability obfuscation*. In terms of both assumptions and practical efficiency, these schemes are a far cry from Yao’s original scheme.

1.2 Our results

In this work, we construct the first adaptively secure garbling scheme whose on-line complexity is smaller than the circuit size and which only relies on the

⁸ Even if we’re willing to assume exponentially secure primitives, the use of complexity leveraging blows up parameter sizes so that the garbled input must be of size at least $n^2 \cdot \text{poly}(\lambda)$ where λ is the security parameter to get any meaningful security.

existence of one-way functions. Our construction is an adaptation of Yao’s scheme that maintains essentially all of its desirable properties, such as having highly parallelizable circuit garbling and projective/decomposable input garbling.⁹ In particular, our construction simply encrypts a Yao garbled circuit with a *some-where equivocal* symmetric-key encryption scheme, which is a new primitive that we define and construct from one-way functions. The encrypted Yao garbled circuit is sent in the off-line phase, and the Yao garbled input along with the decryption key is sent in the on-line phase. We get various provably secure instantiations of the above approach depending on how we set the parameters of the encryption scheme.

As our main instantiation, we get a garbling scheme whose on-line complexity is $w \cdot \text{poly}(\lambda)$ where w is the *width* of the circuit and λ is the security parameter, but is otherwise independent of the depth d of the circuit.¹⁰ Note that, if we think of the circuit as representing a Turing Machine or RAM computation, then the width w of the circuit corresponds to the maximum of the input size n , output size m , and space complexity s of the computation, meaning that our on-line complexity is $(n + m + s) \cdot \text{poly}(\lambda)$, but otherwise independent of the run-time of the computation.

Alternately, we also get a different instantiation where the on-line complexity is only $(n + m + d) \cdot \text{poly}(\lambda)$, where n is the input size, m is the output size, and d is the depth of the circuit, but is otherwise independent of the circuit’s width w . In this case, we also incur a $2^{O(d)}$ security loss in our reduction, but this can be a significant improvement over the naive complexity-leveraging approach which incurs a 2^n security loss, where n is the input size. In particular, in the case of NC^1 circuits where $d = O(\log n)$, we get a polynomial reduction and achieve optimal on-line complexity of $(n + m) \cdot \text{poly}(\lambda)$.¹¹

More broadly, we develop a connection between constructing adaptively secure schemes in our framework and a certain type of *pebble complexity* of the given circuit. The size of the garbled input is proportional to the maximal number of pebbles and the number of hybrids in our reduction is proportional to the number of moves needed to pebble the circuit.

1.3 Applications of Our Results

We briefly mention how our results can be used to get concrete improvements in several applications of garbled circuits in prior works.

⁹ Each bit of the garbled input only depends on one bit of the original input.
¹⁰ We consider circuits made up of fan-in 2 gates with arbitrary fan-out. The circuit is composed of levels and wires can only connect gates in level i with those at the next level $i + 1$. The width of the circuit is the maximal number of gates in any level and the depth is the number of levels.
¹¹ For NC^1 circuits, there are perfectly (information theoretically) secure variants of Yao [25,26] which also achieve adaptive security. However, the on-line complexity in these schemes grows *exponentially* in the circuit depth d whereas ours is only linear in d . For example, for a boolean NC^1 circuit with depth $d = 100 \log n$, the on-line complexity of those schemes is $O(n^{100})$ whereas ours would be $O(n)$.

On-line/Off-line Two-Party Computation. One of the main uses of garbled circuits is in two-party secure computation protocols. In this setting, Alice holds an input x_A , Bob holds an input x_B and they wish to compute $f(x_A, x_B)$. To do so, Alice creates a garbled circuit \tilde{C}_f for the function f and sends \tilde{C}_f along with her portion of the garbled input \tilde{x}_A to Bob. Bob runs an oblivious transfer (OT) protocol to get the garbled version of his input \tilde{x}_B without revealing x_B to Alice. This can be done if the garbling scheme is projective/decomposable (see footnote 9) so that each bit of the garbled input only depends on one bit of the original input. Security against fully malicious parties can be obtained via zero-knowledge proofs or cut-and-choose techniques. It is possible to instantiate the above construction with selectively secure garbled circuits, by having Bob commit to x_B before he gets the garbled circuit \tilde{C}_f . This ensures that the choice of the input cannot depend on the garbled circuit.

However, in many cases, creating the garbled circuit \tilde{C}_f for the function f is expensive and we would like to do this off-line before the inputs x_A, x_B are known to Alice and Bob. Once the inputs become known, the on-line phase should be extremely efficient, and ideally much smaller than the size of the circuit of f . This setting was recently explored in the work of Lindell and Riva [28] who showed how to solve this problem very efficiently using cut-and-choose techniques, given an adaptively secure garbling scheme with low on-line complexity. To instantiate the latter primitive, they relied on the random oracle model. Using our construction of adaptively secure garbled circuit, we can instantiate the scheme of [28] in the standard model, where the on-line complexity of the two-party computation protocol would match that of our garbling schemes.

One-Time Programs and Verifiable Computation. As noted by [11], two prior works from the literature on one-time programs [20] and verifiable computation [16] implicitly require adaptively secure garbling.¹² In both cases, we can plug in our construction of adaptively secure garbling to these constructions.

In the case of one-time programs, the on-line complexity of the garbling scheme translates to the number of hardware tokens needed to create the one-time program. In the case of verifiable computation, the on-line complexity of the garbling scheme translates to the complexity of the verification protocol – it is essential that this is smaller than the circuit size to make the verification protocol non-trivial.

Compact Functional Encryption. The recent work of [1] shows how to convert any selectively secure functional encryption (FE) scheme into an adaptively secure FE. However, their transformation is not compact and the ciphertext size is as large as the maximum circuit size of the allowed functions. This is true even if the selectively secure FE that they start with is compact. Implicitly, the main bottleneck in the transformation is having adaptively secure garbled circuits with low on-line complexity. The work of [2] gives an alternate and modular transformation from a selectively secure compact FE to an adaptively secure

¹² The work of [20] requires an even stronger notion of adaptivity called `prv2` but this can be generically achieved given an adaptively secure scheme in our sense. See footnote 7.

one using adaptively secure garbled circuits (actually, their main construction is for Turing Machines and relies on garbling TMs which require heavier machinery – however, it can be scaled down to work for circuits to get the above result). This transformation applies to both bounded-collusion schemes and unbounded-collusion schemes. By plugging in our construction of adaptively secure garbled circuits into the above result we get a transformation from compact selectively secure FE to adaptive FE where the ciphertext size is only proportional to the on-line complexity of our garbling scheme.

1.4 Our Techniques

In order to explain our techniques, we must first explain the difficulties in proving the adaptive security of Yao’s garbling schemes. Since these difficulties are subtle, we begin with a description of the scheme and the proof of selective security, following Lindell and Pinkas [27]. This allows us to fix a precise notation and terminology which will be needed to also explain our new construction and proof. We expect that the reader is already familiar with the basics of Yao circuits and refer to [27] for further details.

Yao’s Scheme and The Challenge of Adaptive Security *Yao’s Scheme.*

For each wire w in the circuit, we pick two keys k_w^0, k_w^1 for a symmetric-key encryption scheme. For each gate in the circuit computing a function $g : \{0, 1\}^2 \rightarrow \{0, 1\}$ and having input wires a, b and output wire c we create a *garbled gate* consisting of 4 randomly ordered ciphertexts created as:

$$\begin{aligned} c_{0,0} &= \text{Enc}_{k_a^0}(\text{Enc}_{k_b^0}(k_c^{g(0,0)})) & c_{1,0} &= \text{Enc}_{k_a^1}(\text{Enc}_{k_b^0}(k_c^{g(1,0)})), \\ c_{0,1} &= \text{Enc}_{k_a^0}(\text{Enc}_{k_b^1}(k_c^{g(0,1)})) & c_{1,1} &= \text{Enc}_{k_a^1}(\text{Enc}_{k_b^1}(k_c^{g(1,1)})) \end{aligned} \quad (1)$$

where (Enc, Dec) is a CPA-secure encryption scheme. The garbled circuit \tilde{C} consists of all of the garbled gates, along with an *output mapping* $\{k_w^0 \rightarrow 0, k_w^1 \rightarrow 1\}$ which gives the correspondence between the keys and the bits they represent for each output wire w . To garble an n -bit value $x = x_1x_2 \cdots x_n$, the garbled input \tilde{x} consists of the keys $k_{w_i}^{x_i}$ for the n input wires w_i .

To evaluate the garbled circuit on the garbled input, it’s possible to decrypt (exactly) one ciphertext in each garbled gate and get the key $k_w^{v(w)}$ corresponding to the bit $v(w)$ going over the wire w during the computation $C(x)$. Once the keys for the output wires are computed, it’s possible to recover the actual output bits by looking them up in the output mapping.

Selective Security Simulator. To prove the selective security of Yao’s scheme, we need to define a simulator that gets the output $y = y_1y_2 \cdots y_m = C(x)$ and must produce \tilde{C}, \tilde{x} . The simulator picks random keys k_1^0, k_w^1 for each wire w just like the real scheme, but it creates the garbled gates as follows:

$$\begin{aligned} c_{0,0} &= \text{Enc}_{k_a^0}(\text{Enc}_{k_b^0}(k_c^0)) & c_{1,0} &= \text{Enc}_{k_a^1}(\text{Enc}_{k_b^0}(k_c^0)), \\ c_{0,1} &= \text{Enc}_{k_a^0}(\text{Enc}_{k_b^1}(k_c^0)) & c_{1,1} &= \text{Enc}_{k_a^1}(\text{Enc}_{k_b^1}(k_c^0)) \end{aligned} \quad (2)$$

where all four ciphertexts encrypt the same key k_c^0 . It creates the output mapping $\{k_w^0 \rightarrow y_w, k_w^1 \rightarrow 1 - y_w\}$ by “programming it” so that the key k_w^0 corresponds to the correct output bit y_w for each output wire w . This defines the simulated garbled circuit \tilde{C} . To create the simulated garbled input \tilde{x} the simulator simply gives out the keys k_w^0 for each input wire w . Note that, when evaluating the simulated garbled circuit on the simulated garbled input, the adversary only sees the keys k_w^0 for every wire w .

Selective Security Hybrids. To prove indistinguishability between the real world and the simulation, there is a series of carefully defined hybrid games that switch the distribution of one garbled gate at a time, starting with the input level and proceeding up the circuit level by level. In each step we switch the distribution of the ciphertexts in the targeted gate to:

$$\begin{aligned} c_{0,0} &= \text{Enc}_{k_a^0}(\text{Enc}_{k_b^0}(k_c^{v(c)})) & c_{1,0} &= \text{Enc}_{k_a^1}(\text{Enc}_{k_b^0}(k_c^{v(c)})), \\ c_{0,1} &= \text{Enc}_{k_a^0}(\text{Enc}_{k_b^1}(k_c^{v(c)})) & c_{1,1} &= \text{Enc}_{k_a^1}(\text{Enc}_{k_b^1}(k_c^{v(c)})) \end{aligned} \quad (3)$$

where $v(c)$ is the correct *value* of the bit going over the wire c during the computation of $C(x)$.

Let us give names to the three modes for creating garbled gates that we defined above: (1) is called **RealGate** mode, (2) is called **SimGate** mode, and (3) is called **InputDepSimGate** mode, since the way that it is defined depends adaptively on the choice of the input x .

We can switch a gate from **RealGate** to **InputDepSimGate** mode if the gates in the previous level are in **InputDepSimGate** mode (or we are in the input level) by CPA security of encryption. In particular, we are *not* changing the value contained in ciphertext $c_{v(a),v(b)}$ encrypted under the keys $k_a^{v(a)}, k_b^{v(b)}$ that the adversary obtains during evaluation, but we *can* change the values contained in all of the other ciphertexts since the keys $k_a^{1-v(a)}, k_b^{1-v(b)}$ do not appear anywhere inside the garbled gates in the previous level.

At the end of the above sequence of hybrid games, all gates are switched from **RealGate** to **InputDepSimGate** mode and the output mapping is computed as in the real world. The resulting distribution is *statistically identical* to the simulation where all the gates are in **SimGate** mode and the output mapping is programmed. This is because, at any level that’s not the output, the keys k_c^0, k_c^1 are used completely identically in the subsequent level so there is no difference between always encrypting $k_c^{v(c)}$ (**InputDepSimGate**) and k_c^0 (**SimGate**). At the output level there is no difference between encrypting $k_c^{v(c)}$ and giving the real mapping $k_c^{v(c)} \rightarrow y_c$ or encrypting k_c^0 and giving the programmed mapping $k_c^0 \rightarrow y_c$ where y_c is the output bit on wire c .

Challenges in Achieving adaptive security. There are two issues in using the above strategy in the adaptive setting: an immediate but easy to fix problem and a more subtle but difficult to overcome problem.

The first immediate issue is that the selective simulator needs to know the output $y = C(x)$ to create the garbled circuit \tilde{C} and in particular to program the output mapping $\{k_w^0 \rightarrow y_w, k_w^1 \rightarrow 1 - y_w\}$ for the output wires w . However, the

adaptive simulator does not get the output y until *after* it creates the garbled circuit \tilde{C} . Therefore, we cannot (even syntactically) use the selective security simulator in the adaptive setting. This issue turns out to be easy to fix by modifying the construction to send the output-mapping as part of the garbled input \tilde{x} in the on-line phase, rather than as part of the garbled circuit \tilde{C} in the off-line phase. This modification raises on-line complexity to also being linear in the output size of the circuit, which we know to be necessary by the lower bound of [8]. With this modification, the adaptive simulator can program the output mapping after it learns the output $y = C(x)$ in the on-line phase and therefore we get a syntactically meaningful simulation strategy in the adaptive setting.

The second problem is where the true difficulty lies. Although we have a syntactically meaningful simulation strategy, the previous proof of indistinguishability of the real world and the simulation completely breaks down in the adaptive setting. Recall that the proof consisted of a sequence of hybrids where we changed one garbled gate at a time (starting from the input level) from `RealGate` mode to the `InputDepSimGate` mode. In the latter mode, the gate is created in a way that depends on the input x , but in the adaptive setting the input x is chosen adaptively after the garbled circuit is created, leading to a circularity. In other words, the distribution of `InputDepSimGate` as specified in equation (3) doesn't even syntactically make sense in the adaptive setting. Therefore, *although we have a syntactically meaningful simulation strategy for the adaptive setting, we do not have any syntactically meaningful sequence of intermediate hybrids to prove indistinguishability between the real world and the simulated world.*

(One could hope to bypass `InputDepSimGate` mode altogether and define the hybrids by changing a gate directly from `RealGate` mode to `SimGate` mode. Unfortunately, this change is easily distinguishable already for the very first gate we would hope to change at the input level – the output value on the gate would no longer be $v(w)$ but 0 which may result in an overall incorrect output since we have not programmed the output map yet. On the other hand, we cannot immediately jump to a hybrid where we program the output map since all of the keys and their semantics are contained under encryption in prior levels of the circuit and we haven't argued about the security of the ciphertexts in these levels yet.)

Our Solution *Outer Encryption Layer.* Our construction starts with the approach of [11] which is to encrypt the entire garbled circuit with an additional outer encryption layer in the off-line phase (this is unrelated to the encryption used to construct the garbled gates). Then, in the on-line phase, we give out the secret key for this outer encryption scheme. The approach of [11] required a symmetric-key, one-time encryption scheme which is *equivocal*, meaning that the ciphertext doesn't determine the message and it is possible to come up with a secret key that can open the ciphertext to any possible message. Unfortunately, any fully equivocal encryption scheme where a ciphertext can be opened to any message (e.g., the one-time pad) must necessarily have a secret key size which is as large as the message size. In our case, this is the entire garbled circuit and

therefore this ruins the on-line efficiency of the scheme. Our main idea is to use a new type of *partially* equivocal encryption scheme, we call *somewhere equivocal*. *Somewhere Equivocal Encryption*. Intuitively, a somewhere equivocal encryption scheme allows us to create a simulated ciphertext which contains “holes” in some small subset of the message bit positions I chosen by the simulator, but all other message bits are fixed. The simulator can later equivocate this ciphertext and “plug the holes” with any bits it wants by deriving a corresponding secret key. An adversary cannot distinguish between seeing a real encryption of some message $m = m_1 m_2 \cdots m_n$ and the real secret key, from seeing a simulated encryption created using only $(m_i)_{i \notin I}$ with “holes” in positions I and an equivocated secret key that later plugs the holes to the correct bits $(m_i)_{i \in I}$. We show how to construct somewhere equivocal encryption using one-way functions. The size of the secret key is only proportional to the maximum number of holes $t = |I|$ that we allow, which we call the “equivocation parameter”, but can be much smaller than the message size.¹³

Our proof of security departs significantly from that of [11]. In particular, our simulator does *not* take advantage of the equivocation property of the encryption scheme at all, and in fact, our simulation strategy is identical to the adaptive simulator we outlined above for the variant of Yao’s garbling where the output map is sent in the on-line phase. However, we crucially rely on the equivocation property to carefully define a meaningful sequence of hybrids that allows us to prove the indistinguishability of the real and simulated worlds.

Hybrids for adaptive security. We define hybrid distributions where various garbled gates will be created in one of three modes discussed above: **RealGate**, **SimGate** and **InputDepSimGate**. However, to make the last option meaningful (even syntactically) in the adaptive setting, we rely on the somewhere equivocal encryption scheme. For these hybrids, when we create the encrypted garbled circuit in the off-line phase, we will simulate the outer encryption layer with a ciphertext that contains “holes” in place of all gates that are in **InputDepSimGate** mode. Only when we open the outer encryption in the on-line phase after the input x is chosen, we will “plug the holes” by sampling these gates correctly in **InputDepSimGate** mode in a way that depends on the input x . Our equivocation parameter t for the somewhere equivocal encryption scheme therefore needs to be large enough to support the maximum number of gates in **InputDepSimGate** mode that we will have in any hybrid.

Sequence of hybrids. For our main result, we use the following sequence of hybrids to prove indistinguishability of real and simulated worlds. We start by switching the first two levels of gates (starting with the input level) to **InputDepSimGate** mode. We then switch the first level of gates to **SimGate** mode and switch the third level **InputDepSimGate** mode. We continue this process, where in each step i we maintain level i in **InputDepSimGate** mode but switch the previous level

¹³ A different notion of partially equivocal encryption, called *somewhat non-committing* encryption, was introduced in [15]. The latter notion allows a ciphertext to be opened to some small, polynomial size, set of messages which can be chosen arbitrarily by the simulator at encryption time. The two notions are incomparable.

$i - 1$ from `InputDepSimGate` to `SimGate` and then switch the next level $i + 1$ from `RealGate` to `InputDepSimGate`. Eventually all gates will be in `SimGate` mode as we wanted. We can switch a level $i - 1$ from `InputDepSimGate` to `SimGate` mode when the subsequent level i is in `InputDepSimGate` mode since the keys k_c^0, k_c^1 for wires c crossing from level $i - 1$ to i are used identically in level i and therefore there is statistically no difference between encrypting the key $k_c^{v(c)}$ (`InputDepSimGate`) and k_c^0 (`SimGate`). We can also switch a level $i + 1$ from `RealGate` to `InputDepSimGate` when the previous level i is `InputDepSimGate` (or $i + 1$ is the input level) by CPA security following the same argument as in the selective setting. With this strategy, at any point in time we have at most two levels in `InputDepSimGate` mode and therefore our equivocation parameter only needs to be proportional to the circuit width w .

Connection to pebbling. We can generalize the above idea and get other meaningful sequences of hybrids with different parameters and implications. We can think of the process of switching between `RealGate`, `SimGate` and `InputDepSimGate` modes as a new kind of *graph pebbling game*, where pebbles can be placed on the graph representing the circuit according to certain rules. Initially, all gates are in `RealGate` mode, which we associate with *not having any pebble* on them. We associate `InputDepSimGate` mode with having a *black pebble* and `SimGate` mode with having a *gray pebble*. The rules of the game go as follows:

- We can place or remove a black pebble on a gate as long as both predecessors of that gate have black pebbles on them (or the gate is an input gate).
- We can replace a black pebble with a gray pebble on a gate as long as all successors of that gate have black or gray pebbles on them (or the gate is an output gate).

The goal of the game is to end up with a gray pebble on every gate. Any such pebbling strategy leads to a sequence of hybrids that shows the indistinguishability between the real world and the simulation. The number of moves needed to complete the pebbling corresponds to the number of hybrids in our proof, and therefore the security loss of our reduction. The maximum number of black pebbles that are in play at any given time corresponds to the equivocation parameter needed for our somewhere equivocal encryption scheme.

For example, the sequence of hybrids discussed above corresponds to a pebbling strategy where the number of black pebbles used is linear in the circuit width w (but independent of the depth) and the number of moves is linear in the circuit size. We give an alternate recursive pebbling strategy where the number of black pebbles used is linear in the circuit depth d (but independent of the width) and the number of moves is $2^{O(d)}$ times the circuit size.

Constructing somewhere equivocal encryption. Lastly, we discuss our construction of somewhere equivocal encryption from one-way functions, which may be of independent interest. Recall that a somewhere equivocal encryption provides a method for equivocating some small number (t out of n) of bits of the message.

Our construction is based on the techniques developed in recent constructions of *distributed point functions* [17,14]. These techniques give us a way to

construct a pseudorandom function (PRF) family f_k with the following equivocation property: for any input x , we can create two PRF keys k_0, k_1 that each individually look uniformly random but such that $f_{k_0}(x') = f_{k_1}(x')$ for all $x' \neq x$ and $f_{k_0}(x) \neq f_{k_1}(x)$. The construction is based on a clever adaptation of the Goldreich-Goldwasser-Micali (GGM) PRF [18].

Using distributed point functions, we can immediately create a somewhere equivocal encryption with equivocation parameter $t = 1$. We rely on a PRF family f_k with the above equivocation property and with one-bit output. To encrypt a message $m = m_1 m_2 \cdots m_n \in \{0, 1\}^n$ we create a ciphertext $c = f_k(1) \oplus m_1 || f_k(2) \oplus m_2 || \cdots || f_k(n) \oplus m_n$ using the PRF outputs as a one-time pad. To create a simulated encryption with a hole in position i , the simulator samples two PRF keys k_0, k_1 that only differ on input $x = i$. The simulator encrypts the n -bit message by setting the unknown value in position i to $m_i := 0$ and using k_0 . If it later wants to open this value to 0, it sets the decryption key to k_0 else k_1 .

We can extend the above approach to an arbitrarily large equivocation parameter t , by using the XOR of t independently chosen PRFs with the above equivocation property. The key size will be $t \cdot \text{poly}(\lambda)$.

2 Preliminaries

General Notation. For a positive integer n , we define the set $[n] := \{1, \dots, n\}$. We use the notation $x \leftarrow X$ for the process of sampling a value x according to the distribution X . For a vector $\bar{m} = (m_1, m_2, \dots, m_n)$, and a subset $P \subset [n]$, we use $(m_i)_{i \in P}$ to denote a vector containing only the values m_i in positions $i \in P$ and \perp symbols in all other positions. We use $(m_i)_{i \notin P}$ as shorthand for $(m_i)_{i \in [n] \setminus P}$.

Circuit Notation. A boolean circuit C consists of gates $\text{gate}_1, \dots, \text{gate}_q$ and wires w_1, w_2, \dots, w_p . A gate is defined by the tuple $\text{gate}_i = (g, w_a, w_b, w_c)$ where $g : \{0, 1\}^2 \rightarrow \{0, 1\}$ is the function computed by the gate, w_a, w_b are the incoming wires, and w_c is the outgoing wire. Although each gate has a unique outgoing wire w_c , this wire can be used as an incoming wire to several different gates and therefore this models a circuit with fan-in 2 and unbounded fan-out. We let q denote the number of gates in the circuit, n denotes the number of input wires and m denote the number of output wires. The total number of wires is $p = n + q$ (since each wire can either be input wire or an outgoing wire of some gate). For convenience, we denote the n input wires by $\text{in}_1, \dots, \text{in}_n$ and the m output wires by $\text{out}_1, \dots, \text{out}_m$. For $x \in \{0, 1\}^n$ we write $C(x)$ to denote the output of evaluating the circuit C on input x .

We say C is leveled, if each gate has an associated level and any gate at level l has incoming wires only from gates at level $l - 1$ and outgoing wires only to gates at level $l + 1$. We let the *depth* d denote the number of levels and the *width* w denote the maximum number of gates in any level.

A circuit C is fully specified by a list of gate tuples $\text{gate}_i = (g, w_a, w_b, w_c)$. We use $\Phi_{\text{topo}}(C)$ to refer to the topology of a circuit— which indicates how gates

are connected, without specifying the function implemented by each gate. In other words, $\Phi_{\text{topo}}(C)$ is the list of *sanitized gate tuples* $\widehat{\text{gate}}_i = (\perp, w_a, w_b, w_c)$ where the function g that the gate implements is removed from the tuple.

3 Garbling Scheme

We now give a formal definition of a garbling scheme. There are many variants of such definitions in the literature, and we refer the reader to [12] for a comprehensive treatment.

Definition 1. *A Garbling Scheme is a tuple of PPT algorithms $\text{GC} = (\text{GCircuit}, \text{GInput}, \text{Eval})$ such that:*

- $(\tilde{C}, k) \stackrel{\$}{\leftarrow} \text{GCircuit}(1^\lambda, C)$: takes as input a security parameter λ , a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$, and outputs the garbled circuit \tilde{C} , and key k .
- $\tilde{x} \leftarrow \text{GInput}(k, x)$: takes as input $x \in \{0, 1\}^n$, and key k and outputs \tilde{x} .
- $y = \text{Eval}(\tilde{C}, \tilde{x})$: given a garbled circuit \tilde{C} and a garbled input \tilde{x} output $y \in \{0, 1\}^m$.

Correctness *There is a negligible function ν such that for any $\lambda \in \mathbb{N}$, any circuit C and input x it holds that $\Pr[C(x) = \text{Eval}(\tilde{C}, \tilde{x})] = 1 - \nu(\lambda)$, where $(\tilde{C}, k) \leftarrow \text{GCircuit}(1^\lambda, C)$, $\tilde{x} \leftarrow \text{GInput}(k, x)$.*

Adaptive Security. *There exists a PPT simulator $\text{Sim} = (\text{SimC}, \text{SimIn})$ such that, for any PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr[\text{Exp}_{\mathcal{A}, \text{GC}, \text{Sim}}^{\text{adaptive}}(1^\lambda, 0) = 1] - \Pr[\text{Exp}_{\mathcal{A}, \text{GC}, \text{Sim}}^{\text{adaptive}}(1^\lambda, 1) = 1] \leq \nu(\lambda)$$

where the experiment $\text{Exp}_{\mathcal{A}, \text{GC}, \text{Sim}}^{\text{adaptive}}(1^\lambda, b)$ is defined as follows:

1. The adversary \mathcal{A} specifies C and gets \tilde{C} where \tilde{C} is created as follows:
 - if $b = 0$: $(\tilde{C}, k) \leftarrow \text{GCircuit}(1^\lambda, C)$,
 - if $b = 1$: $(\tilde{C}, \text{state}) \leftarrow \text{SimC}(1^\lambda, \Phi_{\text{topo}}(C))$, where $\Phi_{\text{topo}}(C)$ reveals the topology of C .
2. The adversary \mathcal{A} specifies x and gets \tilde{x} created as follows:
 - if $b = 0$, $\tilde{x} \leftarrow \text{GInput}(k, x)$,
 - if $b = 1$, $\tilde{x} \leftarrow \text{SimIn}(C(x), \text{state})$.
3. Finally, the adversary outputs a bit b' , which is the output of the experiment.

On-line Complexity. The time it takes to garble an input x , (i.e., time complexity of $\text{GInput}(\cdot, \cdot)$) is the *on-line complexity* of the scheme. Clearly the on-line complexity of the scheme gives a bound on the size of the garbled input \tilde{x} . Ideally, the on-line complexity should be much smaller than the circuit size $|C|$.

Projective Scheme. A garbling scheme is *projective* if each bit of the garbled input \tilde{x} only depends on one bit of the actual input x . In other words, each bit of the input, is garbled independently of other bits of the input. Projective schemes are essential for two-party computation where the garbled input is transmitted using an oblivious transfer (OT) protocol. Our constructions will be projective.

Hiding Topology. A garbling scheme that satisfies the above security definition may reveal the topology of the circuit C . However, there is a way to transform any such garbling scheme into one that hides everything, including the topology of the circuit, without a significant asymptotic efficiency loss. More precisely, we rely on the fact that there is a function $\text{HideTopo}(\cdot)$ that takes a circuit C as input and outputs a functionally equivalent circuit C' , such that for any two circuits C_1, C_2 of equal size, if $C'_1 = \text{HideTopo}(C_1)$ and $C'_2 = \text{HideTopo}(C_2)$, then $\Phi_{\text{topo}}(C'_1) = \Phi_{\text{topo}}(C'_2)$. An easy way to construct such function HideTopo is by setting C' to be a universal circuit, with a hard-coded description of the actual circuit C . Therefore, to get a topology-hiding garbling scheme, we can simply use a topology-revealing scheme but instead of garbling the circuit C directly, we garble the circuit $\text{HideTopo}(C)$.

4 Somewhere Equivocal Symmetric-Key Encryption

We introduce a new cryptographic primitive called *somewhere* equivocal encryption scheme. Intuitively, a somewhere equivocal encryption scheme allows one to create a simulated ciphertext which contain “holes” in some small subset of the messages in positions I chosen by the simulator, but all other messages are fixed. The simulator can later equivocate this ciphertext and “plug the holes” with any message it wants by deriving a corresponding secret key.

In more detail, encryptions can be computed in two modes: real mode and simulated mode. In the real mode, a key $\text{key} \leftarrow \text{KeyGen}(1^\lambda)$ is generated using the honest key generation procedure and a vector of n messages $\bar{m} = m_1, \dots, m_n$ is encrypted using the honest encryption procedure $\bar{c} \leftarrow \text{Enc}(\text{key}, \bar{m})$.

In the simulated mode, there is an encryption procedure SimEnc that given a set I (set of holes) and only a subset of messages $(m_i)_{i \notin I}$, outputs simulated ciphertext \bar{c} that is equivocal in positions I . In a later stage, upon learning the remaining messages $(m_i)_{i \in I}$, there exists a procedure SimKey that plugs the holes by generating a key key' that will decrypt \bar{c} correctly according to \bar{m} .

The security property that we require is that the distributions of $\{\bar{c}, \text{key}\}$ generated in the two modes are indistinguishable. To capture this property, one could envision a non-adaptive security game where an adversary \mathcal{A} first selects the full vector \bar{m} and the set I , then it receives the tuple (\bar{c}, key) and needs to distinguish which distribution it belongs to. However, such security definition is not sufficient for our indistinguishability proof where instead we need an adversary to decide on the missing messages *after* she receives the ciphertext \bar{c} . Therefore, we consider an adaptive security definition where the security game is defined in two stages: in the first stage, the adversary chooses I , an *incomplete*

vector of messages $(m_i)_{i \notin I}$, and a *challenge* index $j \notin I$ and receives the ciphertext \bar{c} . In the second stage, the adversary sends the remaining messages $(m_i)_{i \in I}$ and gets key. The adversary knows that all positions in I are equivocal and are plugged to the values $(m_i)_{i \in I}$ chosen in the second stage. The challenge is to distinguish whether the position j is also equivocal or not. Note that this two-stage (adaptive) security definition is stronger than the non-adaptive security definition sketched above. For completeness, we give the simpler non-adaptive definition and prove the above implication in the full version [23].

Definition 2. A somewhere equivocal encryption scheme with block-length s , message-length n (in blocks), and equivocation-parameter t (all polynomials in the security parameter) is a tuple of probabilistic polynomial algorithms $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{SimEnc}, \text{SimKey})$ such that:

- The key generation algorithm KeyGen takes as input the security parameter 1^λ and outputs a key: $\text{key} \leftarrow \text{KeyGen}(1^\lambda)$.
- The encryption algorithm Enc takes as input a vector of n messages $\bar{m} = m_1, \dots, m_n$, with $m_i \in \{0, 1\}^s$, and a key key , and outputs ciphertext $\bar{c} \leftarrow \text{Enc}(\text{key}, \bar{m})$.
- The decryption algorithm Dec takes as input ciphertext \bar{c} and a key key and outputs a vector of messages $\bar{m} = m_1, \dots, m_n$. Namely, $\bar{m} \leftarrow \text{Dec}(\text{key}, \bar{c})$.
- The simulated encryption algorithm SimEnc takes as input a set of indexes $I \subset [n]$, such that $|I| \leq t$, and a vector of $n - |I|$ messages $(m_i)_{i \notin I}$ and outputs ciphertext \bar{c} , and a state state . Namely, $(\text{state}, \bar{c}) \leftarrow \text{SimEnc}((m_i)_{i \notin I}, I)$.
- The simulated key algorithm SimKey , takes in the variable state and messages $(m_i)_{i \in I}$ and outputs a key key' . Namely, $\text{key}' \leftarrow \text{SimKey}(\text{state}, (m_i)_{i \in I})$.

and satisfies the following properties:

Correctness. For every $\text{key} \leftarrow \text{KeyGen}(1^\lambda)$, $\bar{m} \in \{0, 1\}^{s \times n}$ it holds that:

$$\text{Dec}(\text{key}, (\text{Enc}(\text{key}, \bar{m}))) = \bar{m}$$

Simulation with No Holes. We require that the distribution of (\bar{c}, key) computed via $(\bar{c}, \text{state}) \leftarrow \text{SimEnc}(\bar{m}, \emptyset)$ and $\text{key} \leftarrow \text{SimKey}(\text{state}, \emptyset)$ to be identical to $\text{key} \leftarrow \text{KeyGen}(1^\lambda)$ and $\bar{c} \leftarrow \text{Enc}(\text{key}, \bar{m})$. In other words, simulation when there are no holes (i.e., $I = \emptyset$) is identical to honest key generation and encryption.

Security. For any PPT adversary \mathcal{A} , there is a negligible function $\nu(\lambda)$ s.t.:

$$\Pr[\text{Exp}_{\mathcal{A}, \Pi}^{\text{simenc}}(1^\lambda, 0) = 1] - \Pr[\text{Exp}_{\mathcal{A}, \Pi}^{\text{simenc}}(1^\lambda, 1) = 1] \leq \nu(\lambda)$$

where the experiment $\text{Exp}_{\mathcal{A}, \Pi}^{\text{simenc}}$ is defined as follows:

Experiment $\text{Exp}_{\mathcal{A}, \Pi}^{\text{simenc}}(1^\lambda, b)$

1. The adversary \mathcal{A} on input 1^λ outputs a set $I \subseteq [n]$ s.t. $|I| < t$, vector $(m_i)_{i \notin I}$, and a challenge index $j \in [n] \setminus I$. Let $I' = I \cup j$.
2. – If $b = 0$, compute \bar{c} as follows: $(\text{state}, \bar{c}) \leftarrow \text{SimEnc}((m_i)_{i \notin I}, I)$.

- If $b = 1$, compute \bar{c} as follows: $(\text{state}, \bar{c}) \leftarrow \text{SimEnc}((m_i)_{i \notin I'}, I')$.
- 3. Send \bar{c} to the adversary \mathcal{A} .
- 4. The adversary \mathcal{A} outputs the set of remaining messages $(m_i)_{i \in I}$.
 - If $b = 0$, compute key as follows: $\text{key} \leftarrow \text{SimKey}(\text{state}, (m_i)_{i \in I})$.
 - If $b = 1$, compute key as follows: $\text{key} \leftarrow \text{SimKey}(\text{state}, (m_i)_{i \in I'})$.
- 5. Send key to the adversary \mathcal{A} .
- 6. \mathcal{A} outputs b' which is the output of the experiment.

In the full version of this paper, [23], we construct somewhere equivocal encryption from one-way functions, proving the following theorem.

Theorem 1. *Assuming the existence of one-way functions, there exists a somewhere equivocal encryption scheme for any polynomial message-length n , block-length s , and equivocation parameter t , having key size $t \cdot s \cdot \text{poly}(\lambda)$ and ciphertext of size $n \cdot s$ bits.*

5 Adaptively Secure Garbling Scheme and Simulator

In this section we describe our garbling scheme and the simulation strategy.

5.1 Construction

Our adaptively-secure garbling scheme consists in two simple steps: (1) garble the circuit using Yao’s garbling scheme; (2) hide the garbled circuit (without the output tables) under an **outer** layer of encryption instantiated with a *somewhere-equivocal* encryption scheme. In the on-line phase, the garbled input consists of Yao’s garbled input plus the output tables. Next we provide the formal description of our scheme that contains the details of Yao’s garbling scheme.

Let C be a leveled boolean circuit with fan-in 2 and unbounded fan-out, with inputs size n , output size m , depth d and width w . Let q denote the number of gates in C . Recall that wires are uniquely identified with labels w_1, w_2, \dots, w_p , and a circuit C is specified by a list of gate tuples $\text{gate} = (g, w_a, w_b, w_c)$. To simplify the description of our construction, we first describe the procedure for garbling a single gate, that we denote by GarbleGate . Let $\Gamma = (G, E, D)$ be a CPA-secure symmetric-key encryption scheme satisfying the special correctness property, that is, the decryption procedure will abort if an incorrect key is used. $\text{GarbleGate}(g, \{k_a^\sigma, k_b^\sigma, k_c^\sigma\}_{\sigma \in \{0,1\}})$ computes 4 ciphertexts $c_{\sigma_0, \sigma_1} : \sigma_0, \sigma_1 \in \{0, 1\}$ as defined below and outputs them in a random order as $\tilde{g} = [c_1, c_2, c_3, c_4]$.

$$\begin{aligned} c_{0,0} &\leftarrow E_{k_a^0}(E_{k_b^0}(k_c^{g(0,0)})) & c_{0,1} &\leftarrow E_{k_a^0}(E_{k_b^1}(k_c^{g(0,1)})) \\ c_{1,0} &\leftarrow E_{k_a^1}(E_{k_b^0}(k_c^{g(1,0)})) & c_{1,1} &\leftarrow E_{k_a^1}(E_{k_b^1}(k_c^{g(1,1)})) \end{aligned}$$

Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{SimEnc}, \text{SimKey})$ be a somewhere-equivocal symmetric-encryption scheme as defined in Section 4. Recall that in this primitive the plaintext is a vector of n blocks, each of which has s bits. In our construction we use the following parameters: the vector size $n = q$ is the number of gates and the block size $s = |\tilde{g}|$ is the size of a single garbled gate. The equivocation parameter t is defined by the strategy used in the security proof and will be specified later. The garbling scheme is formally described in Figure 1.

| | |
|--|--|
| <p><u>GCircuit</u>($1^\lambda, C$)</p> <ol style="list-style-type: none"> 1. Garble Circuit (Yao's scheme) <ul style="list-style-type: none"> – (Wires) $k_{w_i}^\sigma \leftarrow G(1^\lambda)$, $i \in [p]$, $\sigma \in \{0, 1\}$. (Input wires) $K = (k_{in_i}^0, k_{in_i}^1)_{i \in [n]}$. – (Gates) For $\text{gate}_i = (g, w_a, w_b, w_c)$ in C: $\tilde{g}_i \leftarrow \text{GarbleGate}(g, \{k_{w_a}^\sigma, k_{w_b}^\sigma, k_{w_c}^\sigma\}_{\sigma \in \{0,1\}})$ – (Output tables) For each output $j \in [m]$: $\tilde{d}_j := [(k_{out_j}^0 \rightarrow 0), (k_{out_j}^1 \rightarrow 1)]$. 2. Outer Encryption <ul style="list-style-type: none"> – $\text{key} \xleftarrow{\\$} \text{KeyGen}(1^\lambda)$. – $\tilde{C} \leftarrow \text{Enc}(\text{key}, (\tilde{g}_1, \dots, \tilde{g}_q))$. <p>Output \tilde{C}, $k = (K, \text{key}, (\tilde{d}_j)_{j \in [m]})$.</p> <p><u>Glinput</u>($x, k$)</p> <ul style="list-style-type: none"> – (Select input keys) $K^x = (k_{in_1}^{x_1}, \dots, k_{in_n}^{x_n})$. – Output $\tilde{x} = (K^x, \text{key}, (\tilde{d}_j)_{j \in [m]})$. | <p><u>Eval</u>(\tilde{C}, \tilde{x})</p> <ol style="list-style-type: none"> 1. Parse $\tilde{x} = (K, \text{key}, (\tilde{d}_j)_{j \in [m]})$. 2. Decrypt Outer Encryption $(\tilde{g}_i)_{i \in q} \leftarrow \text{Dec}(\text{key}, \tilde{C})$. 3. Evaluate Circuit. Parse $K = (k_{in_1}, \dots, k_{in_n})$. For each level $j = 1, \dots, d$, For $\widehat{\text{gate}}_i = (\perp, w_a, w_b, w_c)$ at level j: <ul style="list-style-type: none"> – Let $\tilde{g}_i = [c_1, c_2, c_3, c_4]$; – For $\delta \in [4]$ let $k'_{w_c} \leftarrow D_{k_{w_a}}(D_{k_{w_b}}(c_\delta))$ If $k'_{w_c} \neq \perp$ then set $k_{w_c} := k'_{w_c}$. 4. Decrypt output. For $j \in [m]$: <ul style="list-style-type: none"> – Parse $\tilde{d}_j = [(k_{out_j}^0 \rightarrow 0), (k_{out_j}^1 \rightarrow 1)]$. – Set $y_j = b$ iff $k_{out_j} = k_{out_j}^b$. <p>Output y_1, \dots, y_m.</p> |
|--|--|

Fig. 1: Adaptively-secure Garbling Scheme.

5.2 Adaptive Simulator

The adaptive security simulator for our garbling scheme is essentially the same as the static security simulator for Yao's scheme (as in [27]), with the only difference that the output table is sent in the on-line phase, and is computed adaptively to map to the correct output. Note that the garbled circuit simulator does not rely on the simulation properties of the somewhere equivocal encryption scheme - these are only used in the proof of indistinguishability.

More specifically, the adaptive simulator ($\text{SimC}, \text{SimIn}$) works as follows. In the off-line phase, SimC computes the garbled gates using procedure GarbleSimGate , that generates 4 ciphertexts that encrypt the same output key. More precisely, $\text{GarbleSimGate}(\{k_{w_a}^\sigma, k_{w_b}^\sigma\}_{\sigma \in \{0,1\}}, k'_{w_c})$ takes both keys for input wires w_a, w_b and a single key for the output wire w_c , that we denote by k'_{w_c} . It then outputs $\tilde{g}_c = [c_1, c_2, c_3, c_4]$ where the ciphertexts, arranged in random order, are computed as follows.

$$\begin{aligned} c_{0,0} &\leftarrow E_{k_a^0}(E_{k_b^0}(k'_c)) & c_{0,1} &\leftarrow E_{k_a^0}(E_{k_b^1}(k'_c)) \\ c_{1,0} &\leftarrow E_{k_a^1}(E_{k_b^0}(k'_c)) & c_{1,1} &\leftarrow E_{k_a^1}(E_{k_b^1}(k'_c)) \end{aligned}$$

The simulator invokes GarbleSimGate on input $k'_c = k_c^0$. It then encrypts the garbled gates so obtained by using the honest procedure for the somewhere equivocal encryption.

In the on-line phase, `SimIn`, on input $y = C(x)$ adaptively computes the output tables so that the evaluator obtains the correct output. This is easily achieved by associating each bit of the output, y_j , to the only key encrypted in the output gate g_{out_j} , which is $k_{out_j}^0$. For the input keys, `SimIn` just sends keys $k_{in_i}^0$ for each $i \in [n]$. The detailed definition of `(SimC, SimIn)` is provided in Figure 2.

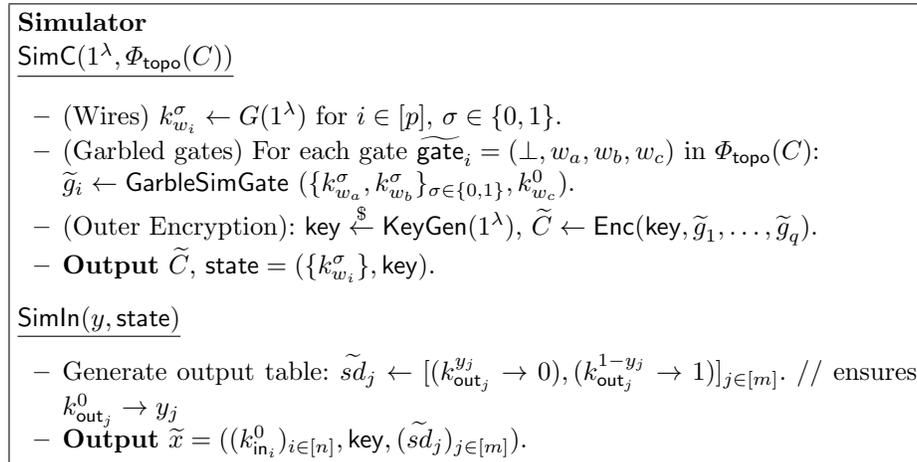


Fig. 2: Simulator for Adaptive Security.

5 Hybrid Games

We now need to prove the indistinguishability of our garbling scheme and the simulation. We devise a modular approach for proving indistinguishability using different strategies that result in different parameters. We first provide a template for defining hybrid games, where each such hybrid game is parametrized by a *circuit configuration*, that is, a vector indicating the way the gates are garbled and encrypted. Then we define the rules that allow us to indistinguishably move from one configuration to another. With this framework in place, an indistinguishability proof consists of a strategy to move from the circuit configuration of the real game to the circuit configuration of the simulated game, using the allowed rules.

5.1 Template for Defining Hybrid Games

Gate/Circuit Configuration. We start by defining a *gate configuration*. A gate configuration is a pair (**outer mode**, **garbling mode**) indicating the way a gate is computed. The outer encryption mode can be $\{\text{EquivEnc}, \text{BindEnc}\}$ depending on whether the outer encryption contains a “hole” in place of that gate

or whether it is binding on that gate. The garbling mode can be $\{\text{RealGate}, \text{SimGate}, \text{InputDepSimGate}\}$ which corresponds to the distributions outlined in Figure 3. We stress that, if the garbling mode of a gate is InputDepSimGate then we require that the outer encryption mode is EquivEnc . This means that there are 5 valid gate configurations for each gate.

| RealGate | SimGate | InputDepSimGate |
|---|--|---|
| $c_{0,0} \leftarrow E_{k_a^0}(E_{k_b^0}(k_c^{g(0,0)}))$ | $c_{0,0} \leftarrow E_{k_a^0}(E_{k_b^0}(k_c^0))$ | $c_{0,0} \leftarrow E_{k_a^0}(E_{k_b^0}(k_c^{v(c)}))$ |
| $c_{0,1} \leftarrow E_{k_a^0}(E_{k_b^1}(k_c^{g(0,1)}))$ | $c_{0,1} \leftarrow E_{k_a^0}(E_{k_b^1}(k_c^0))$ | $c_{0,1} \leftarrow E_{k_a^0}(E_{k_b^1}(k_c^{v(c)}))$ |
| $c_{1,0} \leftarrow E_{k_a^1}(E_{k_b^0}(k_c^{g(1,0)}))$ | $c_{1,0} \leftarrow E_{k_a^1}(E_{k_b^0}(k_c^0))$ | $c_{1,0} \leftarrow E_{k_a^1}(E_{k_b^0}(k_c^{v(c)}))$ |
| $c_{1,1} \leftarrow E_{k_a^1}(E_{k_b^1}(k_c^{g(1,1)}))$ | $c_{1,1} \leftarrow E_{k_a^1}(E_{k_b^1}(k_c^0))$ | $c_{1,1} \leftarrow E_{k_a^1}(E_{k_b^1}(k_c^{v(c)}))$ |

Fig. 3: Garbling Gate modes: RealGate (left), SimGate (center), InputDepSimGate (right). The value $v(c)$ depends on the input x and corresponds to the bit going over the wire c in the computation $C(x)$.

A *circuit configuration* simply consists of the gate configuration for each gate in the circuit. More specifically, we represent a circuit configuration by a tuple $(I, (\text{mode}_i)_{i \in [q]})$ where

- Set $I \subseteq [q]$ contains the indices of the gates i whose outer mode is EquivEnc .
- The value $\text{mode}_i \in \{\text{RealGate}, \text{SimGate}, \text{InputDepSimGate}\}$ describes the garbling mode of gate i .

A *valid circuit configuration* is one where all indexes i such that $\text{mode}_i = \text{InputDepSimGate}$ satisfy $i \in I$.

The *Hybrid Game* $\text{Hyb}(I, (\text{mode}_i)_{i \in [q]})$. Every valid circuit configuration $I, (\text{mode}_i)_{i \in [q]}$ defines a hybrid game $\text{Hyb}(I, (\text{mode}_i)_{i \in [q]})$ as specified formally Figure 4 and described informally below. The hybrid game consists of two procedures: $\text{GCircuit}'$ for creating the garbled circuit \tilde{C} and GInput' for creating the garbled input \tilde{x} respectively. The garbled circuit is created by picking random keys $k_{w_j}^\sigma$ for each wire w_j . For each gate i , such that $\text{mode}_i \in \{\text{RealGate}, \text{SimGate}\}$ it creates a garbled gate \tilde{g}_i using the corresponding distribution as described in Figure 3. The garbled circuit \tilde{C} is then created by simulating the outer encryption using the values \tilde{g}_i in locations $i \notin I$ and “holes” in the locations I . The garbled input is created by first sampling the garbled gates \tilde{g}_i for each i such that $\text{mode}_i = \text{InputDepSimGate}$ using the corresponding distribution in Figure 3 and using knowledge of the input x . Then the decryption key key is simulated by plugging in the holes in locations I with the correctly sampled garbled gates \tilde{g}_i . There is some subtlety about how the input labels $K[i]$ and the output label maps \tilde{d}_j are created when computing \tilde{x} :

- If all of the gates having in_i as an input wire are in SimGate mode, then $K[i] := k_{\text{in}_i}^0$ else $K[i] := k_{\text{in}_i}^{x_i}$.

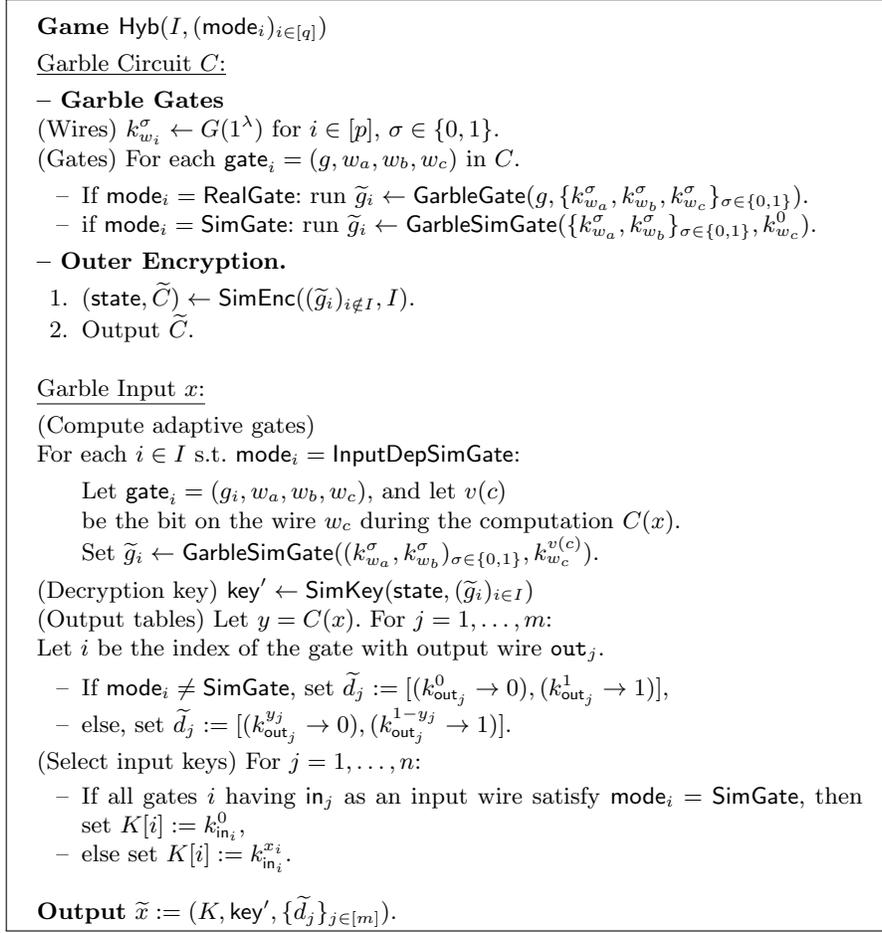


Fig. 4: The Hybrid Game.

- If the unique gate having out_j as an output wire is in SimGate mode, then we give the simulated output map $\tilde{d}_j := [(k_{\text{out}_j}^{y_j} \rightarrow 0), (k_{\text{out}_j}^{1-y_j} \rightarrow 1)]$ else the real one $\tilde{d}_j := [(k_{\text{out}_j}^0 \rightarrow 0), (k_{\text{out}_j}^1 \rightarrow 1)]$.

Real game and Simulated Game. By definition of adaptively secure garbled circuits (Definition 1), the real game $\text{Exp}_{\mathcal{A}, \text{GC}, \text{Sim}}^{\text{adaptive}}(1^\lambda, 0)$ is equivalent to $\text{Hyb}(I = \emptyset, (\text{mode}_i = \text{RealGate})_{i \in [q]})$ and the simulated game $\text{Exp}_{\mathcal{A}, \text{GC}, \text{Sim}}^{\text{adaptive}}(1^\lambda, 1)$ is equivalent to $\text{Hyb}(I = \emptyset, (\text{mode}_i = \text{SimGate})_{i \in [q]})$. Therefore, the main aim is to show that these hybrids are indistinguishable.¹⁴

¹⁴ Note that, the games $\text{Hyb}(\dots)$ use the simulated encryption and key generation procedures of the somewhere equivocal encryption, while the games $\text{Exp}_{\mathcal{A}, \text{GC}, \text{Sim}}^{\text{adaptive}}(1^\lambda, b)$

5.2 Rules for Indistinguishable Hybrids

Next, we provide rules that allow us to move from one configuration to another and prove that the corresponding hybrid games are indistinguishable. We define three rules that allow us to do this. We define $\text{mode} \stackrel{\text{def}}{=} (\text{mode}_i)_{i \in [q]}$.

Indistinguishability Rule 1: Changing the Outer Encryption Mode BindEnc \leftrightarrow EquivEnc. This rule allows to change the outer encryption of a single gate. It says that one can move from a valid circuit configuration (I, mode) to a circuit configuration (I', mode) where $I' = I \cup j$. Thus one more gate is now computed equivocally (and vice versa).

Lemma 1. *Let (I, mode) be any valid circuit configuration, let $j \in [q] \setminus I$ and let $I' = I \cup j$. Then $\text{Hyb}(I, \text{mode}) \stackrel{\text{comp}}{\approx} \text{Hyb}(I', \text{mode})$ are computationally indistinguishable as long as $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{SimEnc}, \text{SimKey})$ is a somewhere equivocal encryption scheme with equivocation parameter t such that $|I'| \leq t$.*

Proof. Towards a contradiction, assume there exists a PPT distinguisher \mathcal{A} that distinguishes the distributions $H_0 = \text{Hyb}(I, \text{mode})$ and $H_1 = \text{Hyb}(I', \text{mode})$ as defined in the Lemma.

We construct a distinguisher B for the security of somewhere equivocal encryption scheme as follows. Informally, adversary B is playing in experiment $\text{Exp}_{B, \Pi}^{\text{simenc}}(1^\lambda, b)$ and uses her oracle access to SimEnc to reproduce the distribution of H_b . B , on input I, j and $\text{mode} = \text{mode}_1, \dots, \text{mode}_q$ computes each garbled gate \tilde{g}_i on its own exactly as in H_0 / H_1 accordingly to mode_i . B computes the outer encryptions of the gates by sending the gates, along with sets I, j to $\text{Exp}^{\text{simenc}}$.

In the on-line phase, after obtaining x from \mathcal{A} , B computes the values for the missing gates $(\tilde{g}_i)_{i \in I}$ and send them to $\text{Exp}^{\text{simenc}}$, and obtain a key key' . B uses such key to compute the garbled inputs \tilde{x} .

Now, if B is playing the game $\text{Exp}_{B, \Pi}^{\text{simenc}}(1^\lambda, b)$ with a bit b , then the view generated by B is distributed identically to H_b . Thus, B distinguishes whether it is playing the game with $b = 0$ or $b = 1$ with the same probability that \mathcal{A} distinguishes H_0 from H_1 . A more detailed description of adversary B is provided in the full version [23].

Indistinguishability Rule 2. Changing the Garbling Mode RealGate \leftrightarrow InputDepSimGate This rule allows us to change the mode of a gate j from RealGate to InputDepSimGate as long as $j \in I$ and that $\text{gate}_j = (g, w_a, w_b, w_c)$ has incoming wires w_a, w_b that are either input wires or are the outgoing wires of some predecessor gates both of which are in InputDepSimGate mode.

only use the real key generation and encryption procedures. However, by definition, these are equivalent when $I = \emptyset$ (no “holes”).

Double Encryption Encryption Security. For convenience, we use the notion of double encryption security, following [27]. This notion is implied by standard CPA security but is more convenient to use in our security proof of garbled circuit security. See the full version [23] for more details.

Definition 3 (Predecessor/Successor/Sibling Gates). *Given a circuit C and a gate $j \in [q]$ of the form $\text{gate}_j = (g, w_a, w_b, w_c)$ with incoming wires w_a, w_b and outgoing wire w_c :*

- We define the predecessors of j , denoted by $\text{Pred}(j)$, to be the set of gates whose outgoing wires are either w_a or w_b . If w_a, w_b are input wires then $\text{Pred}(j) = \emptyset$, else $|\text{Pred}(j)| = 2$.
- We define the successors of j , denoted by $\text{Succ}(j)$ to be the set of gates that contain w_c as an incoming wire. If w_c is an output wire then $\text{Succ}(j) = \emptyset$.
- We define the siblings of j , denoted by $\text{Siblings}(j)$ to be the set of gates that contain either w_a or w_b as an incoming wire.

Lemma 2. *Let $(I, \text{mode} = (\text{mode}_i)_{i \in [q]})$ be a valid circuit configuration and let $j \in I$ be an index such that $\text{mode}_j = \text{RealGate}$ and for all $i \in \text{Pred}(j)$: $\text{mode}_i = \text{InputDepSimGate}$. Let $\text{mode}' = (\text{mode}'_i)_{i \in [q]}$ be defined by $\text{mode}'_i = \text{mode}_i$ for all $i \neq j$ and $\text{mode}'_j = \text{InputDepSimGate}$. Then the games $\text{Hyb}(I, \text{mode}) \stackrel{\text{comp}}{\approx} \text{Hyb}(I, \text{mode}')$ are computationally indistinguishable as long as $\Gamma = (G, E, D)$ is an encryption scheme secure under chosen double encryption.*

Proof. Let I, mode, j and mode' be as in the statement of the Lemma. Towards a contradiction, assume that there exists a PPT adversary \mathcal{A} distinguishing distributions generated in $H^0 := \text{Hyb}(I, \text{mode})$ and $H^1 := \text{Hyb}(I, \text{mode}')$.

We construct an adversary B that breaks the CPA-security of the inner encryption scheme $\Gamma = (G, E, D)$ which is used to garble gates. More specifically, we show that B wins the chosen double encryption security game which is implied by CPA security. Informally, B , on input mode, I and target gate j aims to use her CPA-oracle access in $\text{Exp}^{\text{double}}(1^\lambda, b)$ to generate a distribution H^b . Recall that the only difference between H^0 and H^1 is in the way that the garble gate \tilde{g}_j is computed. On a high level, the reduction B will compute all garbled gates \tilde{g}_i for $i \neq j$, according to experiment $\text{Hyb}(I, \text{mode})$, and will compute the garbled gate \tilde{g}_j using the ciphertexts obtained as a challenge in the experiment $\text{Exp}^{\text{double}}(1^\lambda, b)$.

In more detail, let $\text{gate}_j = (g, w_a, w_b, w_c)$ be the target gate. Recall $j \in I$ and therefore the value \tilde{g}_j is only needed in the on-line phase. If the values going over the wires w_a, w_b during the computation $C(x)$ are α, β respectively, the reduction B will know all wire keys *except* for $k_{w_a}^{1-\alpha}, k_{w_b}^{1-\beta}$. To create the garbled gate \tilde{g}_j it will create the ciphertext $c_{\alpha, \beta}$ as an encryption of $k_{w_c}^{g(\alpha, \beta)}$ on its own, but the remaining three ciphertexts $c_{\alpha', \beta'}$ will come from the experiment $\text{Exp}^{\text{double}}(1^\lambda, b)$ as either encryptions of different values $k_{w_c}^{g(\alpha', \beta')}$ (real) or of the same value $k_{w_c}^{g(\alpha, \beta)}$.

The one subtlety is that reduction needs to create encryptions under the keys $k_{w_a}^{1-\alpha}, k_{w_b}^{1-\beta}$ to create garbled gates \tilde{g}_i for gates i that are siblings of gate j . It

can do that by using the encryption oracles which are given to it as part of the experiment $\text{Exp}^{\text{double}}(1^\lambda, b)$. However, since some of the sibling gates i might be in RealGate or SimGate modes, the reduction needs to create these encryptions already in the offline phase and therefore needs to know the values of α, β in the offline phase before the input x is chosen. To deal with this, we simply have the reduction *guess* the bits α, β randomly in the offline phase. If in the online phase it finds out that the guess is incorrect it outputs a random bit and aborts, else continues. See the full version [23], for a detailed description of the reduction B .

Let Correct be the event that B guesses α and β correctly. Then

$$\begin{aligned} & |\Pr[\text{Exp}_B^{\text{double}}(1^\lambda, 0) = 1] - \Pr[\text{Exp}_B^{\text{double}}(1^\lambda, 1) = 1]| \\ &= \frac{1}{4} |\Pr[\text{Exp}_B^{\text{double}}(1^\lambda, 0) = 1 | \text{Correct}] - \Pr[\text{Exp}_B^{\text{double}}(1^\lambda, 1) = 1 | \text{Correct}]| \\ &= \frac{1}{4} |\Pr[H_{\mathcal{A}}^0(1^\lambda) = 1] - \Pr[H_{\mathcal{A}}^1(1^\lambda) = 1]| \\ &\implies |\Pr[H_{\mathcal{A}}^0(1^\lambda) = 1] - \Pr[H_{\mathcal{A}}^1(1^\lambda) = 1]| \\ &\leq 4 |\Pr[\text{Exp}_B^{\text{double}}(1^\lambda, 0) = 1] - \Pr[\text{Exp}_B^{\text{double}}(1^\lambda, 1) = 1]| \leq \text{negl}(\lambda) \end{aligned}$$

which proves the Lemma.

Indistinguishability Rule 3. Changing the Garbling Mode:

InputDepSimGate \leftrightarrow **SimGate**. This rule allows us to change the mode of a gate j from InputDepSimGate to SimGate under the condition that all successor gates $i \in \text{Succ}(j)$ satisfy that $\text{mode}_i \in \{\text{InputDepSimGate}, \text{SimGate}\}$.

Lemma 3. *Let $(I, \text{mode} = (\text{mode}_i)_{i \in [q]})$ be a valid circuit configuration and let $j \in I$ be an index such that $\text{mode}_j = \text{InputDepSimGate}$ and for all $i \in \text{Succ}(j)$ we have $\text{mode}_i \in \{\text{SimGate}, \text{InputDepSimGate}\}$. Let $\text{mode}' = (\text{mode}'_i)_{i \in [q]}$ be defined by $\text{mode}'_i = \text{mode}_i$ for all $i \neq j$ and $\text{mode}'_j = \text{SimGate}$. Then the games $\text{Hyb}(I, \text{mode}) \equiv \text{Hyb}(I, \text{mode}')$ are identically distributed.*

Proof. Define $H_0 := \text{Hyb}(I, \text{mode})$ and $H_1 := \text{Hyb}(I, \text{mode}')$. Let $\text{gate}_j = (g, w_a, w_b, w_c)$, and let $v(c)$ be the bit on the wire w_c during the computation $C(x)$, which is defined in the on-line phase.

The main difference between the hybrids is how the garbled gate \tilde{g}_j is created:

- In H_0 , we set $\tilde{g}_j \leftarrow \text{GarbleSimGate}((k_{w_a}^\sigma, k_{w_b}^\sigma)_{\sigma \in \{0,1\}}, k_{w_c}^{v(c)})$.
- In H_1 , we set $\tilde{g}_j \leftarrow \text{GarbleSimGate}((k_{w_a}^\sigma, k_{w_b}^\sigma)_{\sigma \in \{0,1\}}, k_{w_c}^0)$.

If j is not an output gate, and all successor gates $i \in \text{Succ}(j)$ are in $\{\text{SimGate}, \text{InputDepSimGate}\}$ modes then the keys $k_{w_c}^0$ and $k_{w_c}^1$ are treated symmetrically everywhere in the game other than in \tilde{g}_j . Therefore, by symmetry, there is no difference between using $k_{w_c}^0$ and $k_{w_c}^{v(c)}$ in \tilde{g}_j .

If j is an output gate then the keys $k_{w_c}^0$ and $k_{w_c}^1$ are only used in \tilde{g}_j and in the output map \tilde{d}_j . Therefore, by symmetry, there is no difference between using

$k_{w_c}^{y_j}$ in \tilde{g}_j and setting $\tilde{d}_j := [(k_{\text{out}_j}^0 \rightarrow 0), (k_{\text{out}_j}^1 \rightarrow 1)]$ (in H_0) versus using $k_{w_c}^0$ in \tilde{g}_j and setting $\tilde{d}_j := [(k_{\text{out}_j}^{y_j} \rightarrow 0), (k_{\text{out}_j}^{1-y_j} \rightarrow 1)]$ (in H_1).

One last difference between the hybrids occurs if some wire in_i becomes only connected to gates that are in `SimGate` in H_1 . In this case, when we create the garbled input \tilde{x} , then in H_0 we give $K[i] := k_{\text{in}_i}^{x_i}$ but in H_1 we give $K[i] := k_{\text{in}_i}^0$. Since the keys $k_{\text{in}_i}^0, k_{\text{in}_i}^1$ are treated symmetrically everywhere in the game (both in H_0 and H_1) other than in $K[i]$, there is no difference between setting $K[i] := k_{\text{in}_i}^0$ versus $K[i] := k_{\text{in}_i}^{x_i}$.

6 Pebbling and Sequences of Hybrid Games

In the last section we defined hybrid games parameterized by a configuration (I, mode) . We also gave 3 rules, which describe ways that allow us to indistinguishably move from one configuration to another. Now our goal is to use the given rules so as to define a *sequence of indistinguishable hybrid games* that takes us from the *real game* $\text{Hyb}(I = \emptyset, (\text{mode}_i = \text{RealGate})_{i \in [q]})$ to the simulation $\text{Hyb}(I = \emptyset, (\text{mode}_i = \text{SimGate})_{i \in [q]})$.

Pebbling Game. We show that the problem of finding such sequences of hybrid games can be captured by a certain type of *pebbling game* on the circuit C . Each gate can either have *no pebble*, a *black pebble*, or a *gray pebble* on it (this will correspond to `RealGate`, `InputDepSimGate` and `SimGate` modes respectively). Initially, the circuit starts out with no pebbles on any gate. The game consist of the following possible moves:

Rule A. We can place or remove a black pebble on a gate as long as both predecessors of that gate have black pebbles (or the gate is an input gate).

Rule B. We can replace a black pebble with a gray one, only if successors of that gate have black or gray pebbles on them (or the gate is an output gate).

A *pebbling* of a circuit C is a sequence of γ moves that follow rules A and B and that end up with a gray pebble on every gate. We say that a pebbling uses t black pebbles if this is the maximal number of black pebbles on the circuit at any point in time during the game.

From Pebbling to Sequence of Hybrids. In our next theorem we prove that any pebbling of a circuit C results in a sequence of hybrids that shows indistinguishability of the real and simulated games. The number of hybrids is proportional to the number of moves in the pebbling and the equivocation parameter is proportional to the number of black pebbles it uses.

Theorem 2. *Assume that there is a pebbling of the circuit C in γ moves. Then there is a sequence of $2 \cdot \gamma + 1$ hybrid games, starting with the real game $\text{Hyb}(I = \emptyset, (\text{mode}_i = \text{RealGate})_{i \in [q]})$ and ending with the simulated game $\text{Hyb}(I = \emptyset, (\text{mode}_i = \text{SimGate})_{i \in [q]})$ such that any two adjacent hybrid games in the sequence are indistinguishable by rules 1, 2 or 3 from the previous section. Furthermore if*

pebbling uses t^* black pebbles then every hybrid $\text{Hyb}(I, \text{mode})$ in the sequence satisfies $|I| \leq t^*$. In particular, indistinguishability holds as long as the equivocation parameter is at least t^* .

Proof. A pebble configuration specifies whether each gate contains no pebble, a black pebble, or a gray pebble. A pebbling in γ moves gives rise to a sequence of $\gamma + 1$ pebble configurations starting with no pebbles and ending with a gray pebble on each gate. Each pebble configuration follows from the preceding one by a move that satisfies pebbling rules A or B.

We let each pebble configuration define a hybrid $\text{Hyb}(I, \text{mode})$ where:

- For every gate $i \in [q]$, we set $\text{mode}_i = \text{RealGate}$ if gate i has no pebble, $\text{mode}_i = \text{InputDepSimGate}$ if gate i has a black pebble, and $\text{mode}_i = \text{SimGate}$ if gate i has a gray pebble.
- We set I to be the set of gates with black pebbles on them.

Therefore a pebbling defines a sequence of hybrids $\text{Hyb}_\alpha = \text{Hyb}(I^\alpha, \text{mode}^\alpha)$ for $\alpha = 0, \dots, \gamma$ where $\text{Hyb}_0 = \text{Hyb}(\emptyset, (\text{mode}_i^0 = \text{RealGate})_{i \in [q]})$ is the real game and $\text{Hyb}_\gamma = \text{Hyb}(\emptyset, (\text{mode}_i^\gamma = \text{SimGate})_{i \in [q]})$ is the simulated game, and each Hyb_α is induced by the pebbling configuration after α moves. We will need to add additional intermediate hybrids (which we call “half steps”) to ensure that each pair of consecutive hybrids is indistinguishable by rules 1,2 or 3. We do this as follows:

- Assume that move $\alpha + 1$ of the pebbling applies rule A to place a black pebble on gate j .
Let $\text{Hyb}_\alpha = \text{Hyb}(I^\alpha, \text{mode}^\alpha)$ and $\text{Hyb}_{\alpha+1} = \text{Hyb}(I^{\alpha+1}, \text{mode}^{\alpha+1})$. Then $I^{\alpha+1} = I^\alpha \cup \{j\}$, $\text{mode}_i^{\alpha+1} = \text{mode}_i^\alpha$ for all $i \neq j$, and $\text{mode}_j^\alpha = \text{RealGate}$, $\text{mode}_j^{\alpha+1} = \text{InputDepSimGate}$.
Define the intermediate “half-step” hybrid $\text{Hyb}_{\alpha+\frac{1}{2}} := \text{Hyb}(I^{\alpha+1}, \text{mode}^\alpha)$.
It holds that $\text{Hyb}_\alpha \stackrel{\text{comp}}{\approx} \text{Hyb}_{\alpha+\frac{1}{2}}$ by rule 1, and $\text{Hyb}_{\alpha+\frac{1}{2}} \stackrel{\text{comp}}{\approx} \text{Hyb}_{\alpha+1}$ by rule 2. The conditions needed to apply rule 2 are implied by pebbling rule A.
- Assume that move $\alpha + 1$ of the pebbling applies rule A to remove a black pebble from gate j .
Let $\text{Hyb}_\alpha = \text{Hyb}(I^\alpha, \text{mode}^\alpha)$ and $\text{Hyb}_{\alpha+1} = \text{Hyb}(I^{\alpha+1}, \text{mode}^{\alpha+1})$. Then $I^{\alpha+1} = I^\alpha \setminus \{j\}$, $\text{mode}_i^{\alpha+1} = \text{mode}_i^\alpha$ for all $i \neq j$, and $\text{mode}_j^\alpha = \text{InputDepSimGate}$, $\text{mode}_j^{\alpha+1} = \text{RealGate}$.
Define the intermediate “half-step” hybrid $\text{Hyb}_{\alpha+\frac{1}{2}} := \text{Hyb}(I^\alpha, \text{mode}^{\alpha+1})$.
It holds that $\text{Hyb}_\alpha \stackrel{\text{comp}}{\approx} \text{Hyb}_{\alpha+\frac{1}{2}}$ by rule 2, and $\text{Hyb}_{\alpha+\frac{1}{2}} \stackrel{\text{comp}}{\approx} \text{Hyb}_{\alpha+1}$ by rule 1. The conditions needed to apply rule 2 are implied by pebbling rule A.
- Assume that move $\alpha + 1$ of the pebbling applies rule B to replace a black pebble with a gray pebble on gate j .
Let $\text{Hyb}_\alpha = \text{Hyb}(I^\alpha, \text{mode}^\alpha)$ and $\text{Hyb}_{\alpha+1} = \text{Hyb}(I^{\alpha+1}, \text{mode}^{\alpha+1})$. Then $I^{\alpha+1} = I^\alpha \setminus \{j\}$, $\text{mode}_i^{\alpha+1} = \text{mode}_i^\alpha$ for all $i \neq j$, and $\text{mode}_j^\alpha = \text{InputDepSimGate}$, $\text{mode}_j^{\alpha+1} = \text{SimGate}$.

Define the intermediate “half-step” hybrid $\text{Hyb}_{\alpha+\frac{1}{2}} := \text{Hyb}(I^\alpha, \text{mode}^{\alpha+1})$.

It holds that $\text{Hyb}_\alpha \stackrel{\text{comp}}{\approx} \text{Hyb}_{\alpha+\frac{1}{2}}$ by rule 3, and $\text{Hyb}_{\alpha+\frac{1}{2}} \stackrel{\text{comp}}{\approx} \text{Hyb}_{\alpha+1}$ by rule 1. The conditions needed to apply rule 3 are implied by pebbling rule B.

Therefore the sequence $\text{Hyb}_0, \text{Hyb}_{\frac{1}{2}}, \text{Hyb}_1, \text{Hyb}_{1+\frac{1}{2}}, \text{Hyb}_2, \dots, \text{Hyb}_\gamma$ consisting of $2\gamma + 1$ hybrids satisfies the conditions of the theorem.

Combining Theorem 2 and Theorem 1 we obtain the following corollary.

Corollary 1. *There exists an adaptively secure garbling scheme such that the following holds. Assuming the existence of one-way functions, there is an instantiation of the garbling scheme that has on-line complexity $(n + m + t^*)\text{poly}(\lambda)$ for any circuit C that admits a pebbling with $\gamma = \text{poly}(\lambda)$ moves and t^* black pebbles. Furthermore, assuming the existence of sub-exponentially secure one-way functions, there is an instantiation of the garbling scheme that has on-line complexity $(n + m + t^*)\text{poly}(\lambda, \log \gamma)$ for any circuit C that admits a pebbling strategy with $\gamma = 2^{\text{poly}(\lambda)}$ moves and t^* black pebbles.*

Proof. We instantiate our construction from Section 5 with a CPA-secure “inner encryption” Γ having special correctness, and a somewhere-equivocal “outer encryption” Π from Section 4 using an equivocation parameter $t = t^*$. Both components can be instantiated from one-way functions. Assuming that $\gamma = \text{poly}(\lambda)$, Theorem 2 tells us that the resulting garbling scheme is adaptively secure as long as Γ, Π are. The on-line complexity consists of $n + m$ keys for Γ along with the key of Π for a total of $(n + m)\text{poly}(\lambda) + t^*\text{poly}(\lambda)$ as claimed.

When $\gamma = 2^{\text{poly}(\lambda)}$, then Theorem 2 tells us that the resulting garbling scheme is adaptively secure as long as the schemes Γ, Π provide a higher level of security so as to survive $2\gamma + 1$ hybrids, meaning that the distinguishing advantage for each of the schemes needs to be $2^{-(2\gamma+1)}\text{negl}(\lambda)$. This can be accomplished assuming sub-exponentially secure one-way functions by setting the security parameter of Γ, Π to some $\lambda' = \text{poly}(\lambda, \log \gamma)$ and results in on-line complexity $(n + m)\text{poly}(\lambda, \log \gamma) + t^*\text{poly}(\lambda, \log \gamma)$ as claimed.

6.1 Pebbling Strategies

In this section we give two pebbling strategies for arbitrary circuit with width w , depth d , and q gates. The first strategy uses $O(q)$ moves and $O(w)$ black pebbles. The second strategy uses $O(q2^d)$ moves and $O(d)$ black pebbles.

Strategy 1 To pebble the circuit proceed as follows:

Pebble(C):

1. Put a black pebble on each gate at the input level (level 1).
2. For $i = 1$ to $d - 1$, repeat:
 - (a) Put a black pebble on each gate at level $i + 1$.
 - (b) For each gate at level i , replace the black pebble with a gray pebble.

(c) $i \leftarrow i + 1$ 3. For each gate at level d , replace the black pebble with a gray pebble.

This strategy uses $\gamma = 2q$ moves and $t^* = 2w$ black pebbles. By instantiating Corollary 1 with this strategy, we obtain the following corollary.

Corollary 2. *Assuming the existence of one-way functions there exists an adaptively secure garbling scheme with on-line complexity $w \cdot \text{poly}(\lambda)$, where w is the width of the circuit.*

Strategy 2 This is a recursive strategy defined as follows.

- **Pebble(C):**
 - For each gate i in C starting with the gates at the top level moving to the bottom level:
 1. **RecPutBlack(C, i)**
 2. Replace the black pebble on gate i with a gray pebble.
- **RecPutBlack(C, i):** // Let **LeftPred(C, i)** and **RightPred(C, i)** are the two predecessors of gate i in C .
 1. If gate i is an input gate, put a black pebble on i and **return**.
 2. Run **RecPutBlack($C, \text{LeftPred}(C, i)$)**, **RecPutBlack($C, \text{RightPred}(C, i)$)**
 3. Put a black pebble on gate i .
 4. Run **RecRemoveBlack($C, \text{LeftPred}(C, i)$)**
and **RecRemoveBlack($C, \text{RightPred}(C, i)$)**
- **RecRemoveBlack(C, i):** This is the same as **RecPutBlack**, except that instead of putting a black pebble on gate i , in steps 1 and 3, we remove it.

To analyze the correctness of this strategy, we note the following invariants: if the circuit C is in a configuration where it does not contain any pebbles at any level below that of gate i , then (1) the procedure **RecPutBlack(C, i)** results in a configuration where a single black pebble is added to gate i , but nothing else changes, (2) the procedure **RecRemoveBlack(C, i)** results in a configuration where a single black pebble is removed from gate i , but nothing else changes. Using these two invariants the correctness of of the entire strategy follows.

To calculate the number of black pebbles used and the number of moves that the above strategy takes to pebble C , we use the following simple recursive equations. Let **#PebPut(d)** and **#PebRem(d)** be the number of black pebbles on gate i and below it used to execute **RecPutBlack** and **RecRemoveBlack** on a gate at level d , respectively. We have,

$$\begin{aligned} \#PebPut(1) &= 1, & \#PebPut(d) &\leq \max(\#PebPut(d-1), \#PebRem(d-1)) + 2 \\ \#PebRem(1) &= 1, & \#PebRem(d) &\leq \max(\#PebPut(d-1), \#PebRem(d-1)) + 2 \end{aligned}$$

Therefore the strategy requires at most $2d$ black pebbles to pebble the circuit.

To calculate the number of moves it takes run **Pebble(C)**, we use the following recursive equations. Let **#Moves(d)** be the number of moves it takes to put a black pebble on, or remove a black pebble from, a gate at level d . Then

$$\#Moves(1) = 1, \quad \#Moves(d) = 4(\#Moves(d-1)) + 1$$

Hence, each call of `RecPutBlack` takes at most 4^d moves, and the total number of moves to pebble the circuit is at most $q4^d$.

In summary, the above gives us a strategy to pebble any circuit with at most $\gamma = q4^d$ moves and $t^* = 2d$ black pebbles. By instantiating Corollary 1 with the above strategy, we obtain the following corollary.

Corollary 3. *Assuming the existence of (standard) one-way functions, there exists an adaptively secure garbling schemes that has on-line complexity $(n + m)\text{poly}(\lambda)$ for all circuits having depth $d = O(\log \lambda)$.*

Assuming the existence of sub-exponentially secure one-way functions, there exists an adaptively secure garbling scheme that has on-line complexity $(n + m)\text{poly}(\lambda, d)$, for arbitrary circuits of depth $d = \text{poly}(\lambda)$.

7 Conclusions

We have shown how to achieve adaptively secure garbling schemes under one-way functions by augmenting Yao’s construction with an additional layer of somewhere-equivocal encryption. The on-line complexity in our constructions can be significantly smaller than the circuit size. In our main instantiation, the on-line complexity only scales with the width w of the circuit, which corresponds to the space complexity of the computation.

It remains as an open problem to get the optimal on-line complexity $(n + m)\text{poly}(\lambda)$ which does not depend on the circuit depth or width. Currently, this is only known assuming the existence of indistinguishability obfuscation and therefore it remains open to achieve the above under one-way functions or even stronger assumptions such as DDH or LWE. It also remains open if Yao’s scheme (or more precisely, a variant of it where the output map is sent in the on-line phase) can already achieve adaptive security without relying on somewhere-equivocal encryption. We have no proof nor a counter-example. It would be interesting to see if there is some simple-to-state standard-model security assumption that one could make on the encryption scheme used to create the garbled gates in Yao’s construction (e.g., circular security, key-dependent message security, etc.), under which one could prove that the resulting garbling scheme is adaptively secure.

References

1. P. Ananth, Z. Brakerski, G. Segev, and V. Vaikuntanathan. From selective to adaptive security in functional encryption. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 657–677. Springer, Heidelberg, Aug. 2015.
2. P. Ananth and A. Sahai. Functional encryption for turing machines. Cryptology ePrint Archive, Report 2015/776, 2015. <http://eprint.iacr.org/>.
3. B. Applebaum. Key-dependent message security: Generic amplification and completeness. In K. G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 527–546. Springer, Heidelberg, May 2011.

4. B. Applebaum. Bootstrapping obfuscators via fast pseudorandom functions. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 162–172. Springer, Heidelberg, Dec. 2014.
5. B. Applebaum, Y. Ishai, and E. Kushilevitz. Cryptography in NC^0 . In *45th FOCS*, pages 166–175. IEEE Computer Society Press, Oct. 2004.
6. B. Applebaum, Y. Ishai, and E. Kushilevitz. Computationally private randomizing polynomials and their applications. In *20th Annual IEEE Conference on Computational Complexity (CCC 2005), 11-15 June 2005, San Jose, CA, USA*, pages 260–274. IEEE Computer Society, 2005.
7. B. Applebaum, Y. Ishai, and E. Kushilevitz. From secrecy to soundness: Efficient verification via secure computation. In S. Abramsky, C. Gavoille, C. Kirchner, F. Meyer auf der Heide, and P. G. Spirakis, editors, *ICALP 2010, Part I*, volume 6198 of *LNCS*, pages 152–163. Springer, Heidelberg, July 2010.
8. B. Applebaum, Y. Ishai, E. Kushilevitz, and B. Waters. Encoding functions with constant online rate or how to compress garbled circuits keys. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 166–184. Springer, Heidelberg, Aug. 2013.
9. B. Barak, I. Haitner, D. Hofheinz, and Y. Ishai. Bounded key-dependent message security. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 423–444. Springer, Heidelberg, May 2010.
10. M. Bellare, V. T. Hoang, and S. Keelveedhi. Instantiating random oracles via UCEs. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 398–415. Springer, Heidelberg, Aug. 2013.
11. M. Bellare, V. T. Hoang, and P. Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In X. Wang and K. Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 134–153. Springer, Heidelberg, Dec. 2012.
12. M. Bellare, V. T. Hoang, and P. Rogaway. Foundations of garbled circuits. In T. Yu, G. Danezis, and V. D. Gligor, editors, *ACM CCS 12*, pages 784–796. ACM Press, Oct. 2012.
13. D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Heidelberg, May 2014.
14. E. Boyle, N. Gilboa, and Y. Ishai. Function secret sharing. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 337–367. Springer, Heidelberg, Apr. 2015.
15. J. A. Garay, D. Wichs, and H.-S. Zhou. Somewhat non-committing encryption and efficient adaptively secure oblivious transfer. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 505–523. Springer, Heidelberg, Aug. 2009.
16. R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 465–482. Springer, Heidelberg, Aug. 2010.
17. N. Gilboa and Y. Ishai. Distributed point functions and their applications. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 640–658. Springer, Heidelberg, May 2014.
18. O. Goldreich, S. Goldwasser, and S. Micali. On the cryptographic applications of random functions. In G. R. Blakley and D. Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 276–288. Springer, Heidelberg, Aug. 1984.

19. S. Goldwasser, Y. T. Kalai, R. A. Popa, V. Vaikuntanathan, and N. Zeldovich. Reusable garbled circuits and succinct functional encryption. In D. Boneh, T. Roughgarden, and J. Feigenbaum, editors, *45th ACM STOC*, pages 555–564. ACM Press, June 2013.
20. S. Goldwasser, Y. T. Kalai, and G. N. Rothblum. One-time programs. In D. Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 39–56. Springer, Heidelberg, Aug. 2008.
21. S. Gorbunov, V. Vaikuntanathan, and H. Wee. Functional encryption with bounded collusions via multi-party computation. In R. Safavi-Naini and R. Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 162–179. Springer, Heidelberg, Aug. 2012.
22. V. Goyal, Y. Ishai, A. Sahai, R. Venkatesan, and A. Wadia. Founding cryptography on tamper-proof hardware tokens. In D. Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 308–326. Springer, Heidelberg, Feb. 2010.
23. B. Hemenway, Z. Jafarholi, R. Ostrovsky, A. Scafuro, and D. Wichs. Adaptively secure garbled circuits from one-way functions. *IACR Cryptology ePrint Archive*, 2015:1250, 2015.
24. P. Hubacek and D. Wichs. On the communication complexity of secure function evaluation with long output. In T. Roughgarden, editor, *ITCS 2015*, pages 163–172. ACM, Jan. 2015.
25. Y. Ishai and E. Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 294–304, 2000.
26. Y. Ishai and E. Kushilevitz. Perfect constant-round secure computation via perfect randomizing polynomials. In *Automata, Languages and Programming, 29th International Colloquium, ICALP 2002, Malaga, Spain, July 8-13, 2002, Proceedings*, pages 244–256, 2002.
27. Y. Lindell and B. Pinkas. A proof of security of Yao’s protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, Apr. 2009.
28. Y. Lindell and B. Riva. Cut-and-choose Yao-based secure computation in the online/offline and batch settings. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 476–494. Springer, Heidelberg, Aug. 2014.
29. J. B. Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 111–126. Springer, Heidelberg, Aug. 2002.
30. A. Sahai and H. Seyalioglu. Worry-free encryption: functional encryption with public keys. In E. Al-Shaer, A. D. Keromytis, and V. Shmatikov, editors, *ACM CCS 10*, pages 463–472. ACM Press, Oct. 2010.
31. A. C.-C. Yao. Protocols for secure computations (extended abstract). In *23rd FOCS*, pages 160–164. IEEE Computer Society Press, Nov. 1982.
32. A. C.-C. Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, Oct. 1986.