

Adaptive Distributional Security for Garbling Schemes with $\mathcal{O}(|x|)$ Online Complexity

Estuardo Alpírez Bock¹, Chris Brzuska², Pihla Karanko², Sabine Oechsner³,
and Kirithivaasan Puniamurthy²

¹ Xiphera, Finland

² Aalto University, Finland

³ University of Edinburgh, UK

Abstract. Garbling schemes allow to garble a circuit C and an input x such that $C(x)$ can be computed while hiding both C and x . In the context of adaptive security, an adversary specifies the input to the circuit *after* seeing the garbled circuit, so that one can pre-process the garbling of C and later only garble the input x in the *online phase*. Since the online phase may be time-critical, it is an interesting question how much information needs to be transmitted in this phase and ideally, this should be close to $|x|$. Unfortunately, Applebaum, Ishai, Kushilevitz, and Waters (AIKW, *CRYPTO 2013*) show that for some circuits, specifically PRGs, achieving online complexity close to $|x|$ is impossible with simulation-based security, and Hubáček and Wichs (HW, *ITCS 2015*) show that online complexity of maliciously secure 2-party computation needs to grow with the incompressibility entropy of the function. We thus seek to understand under which circumstances optimal online complexity is feasible despite these strong lower bounds.

Our starting point is the observation that lower bounds (only) concern *cryptographic* circuits and that, when an embedded secret is not known to the adversary (distinguisher), then the lower bound techniques do not seem to apply. Our main contribution is *distributional simulation-based security* (DSIM), a framework for capturing weaker, yet meaningful simulation-based (adaptive) security which does not seem to suffer from impossibility results akin to AIKW. We show that DSIM can be used to prove security of a distributed symmetric encryption protocol built around garbling. We also establish a bootstrapping result from DSIM-security for NC^0 circuits to DSIM-security for arbitrary polynomial-size circuits while preserving their online complexity.

1 Introduction

The garbled circuits approach to secure two-party computation goes back to a seminal work of Yao [Yao82a] and allows two parties to evaluate a circuit C on their private inputs. Bellare, Hoang and Rogaway (BHR) [BHR12b] suggest to abstract out the central building block behind this approach into a *garbling*

scheme. A garbling scheme allows a party (called garbler) to *garble* a circuit C and input x into \tilde{C} and \tilde{x} such that another party (called evaluator) can derive $C(x)$ from \tilde{C} and \tilde{x} . Selective (simulation-based) security of a garbling scheme is defined as the comparison between real and simulated garbling: A garbling scheme is secure if no efficient adversary can distinguish a real garbling from the output of a simulator that is given only $C(x)$ and some leakage $\Phi(C)$ on C .

BHR further point out [BHR12a] that in some settings, a stronger adaptive security notion is needed, in which the adversary chooses the input x adaptively after seeing the garbled circuit \tilde{C} . In particular, now the simulator works in two stages and has to produce \tilde{C} first given only $\Phi(C)$, and produces \tilde{x} only after seeing $C(x)$. The second *adaptive* phase is often referred to as the *online phase*, and $|\tilde{x}|$ as the online complexity.

In the selective setting, $|\tilde{x}|$ can be linear in the size of x and independent of any other circuit dimension. For example, $|\tilde{x}| = \lambda|x|$ in the case of Yao’s garbling scheme and many of its variants, where λ denotes the security parameter. Ideally, the online complexity of *adaptively* secure garbling schemes would also match this bound. Unfortunately, Applebaum, Ishai, Kushilevitz and Waters (AIKW [AIKW13]) show that adaptively secure garbling schemes for pseudo-random generators (PRGs) require online complexity of at least the output size $|C(x)|$, and indeed, adaptively secure constructions for general circuits match this bound [BHR12a, AS16, HJO⁺16, JW16, GS18, JO20].

Circumventing the AIKW lower bound. Approaches to circumvent the AIKW lower bound include complexity leveraging over all possible values of x (with a superpolynomial-time simulator) and proving security in the programmable random oracle model [BHR12a]. Another line of work shows how to construct adaptively secure garbling schemes with low online complexity ($|\tilde{x}| = O(|x|)$) for restricted circuit classes [JSW17, KKP21, AS16] which is contained in the class of all efficiently invertible circuits and, therefore, unfortunately exclude garbling most cryptographic functions with adaptive simulation-security.

Relaxing garbling scheme security. In this work, we explore a new approach which seems suitable for *cryptographic* circuits. Concretely, we identify a relaxed (simulation-based) garbling scheme security notion that is not affected by the same lower bound and ask which circuit classes can be securely garbled with respect to it. When garbling *cryptographic* circuits, some inputs are typically generated *honestly* and *secretly*, and hence, an (efficient) adversary does not know them - an indication that the standard definition of adaptive simulation-based security might be too strong.

1.1 Summary of contributions

Generalizing adaptive garbling scheme security. Our main contribution is a security framework for garbling schemes, called *distributional simulation-based security* (DSIM) which relaxes the existing simulation-based adaptive security

notion. The idea is to model realistic restrictions on the knowledge of adversary and simulator such as partially hiding the circuit (or inputs) that is garbled from the adversary. DSIM security bypasses the AIKW lower bound and our generalization of it as they are specific to the adversary’s and simulator’s (lack of) knowledge in the adaptive simulatability game.

We then show relations between DSIM and existing adaptive security notions, in particular that adaptive simulatability implies DSIM when suitably restricting DSIM parameters. We further show how use a DSIM-secure garbling scheme to turn an authenticated encryption scheme (AE) into a two-party distributed encryption protocol. The AE scheme needs to be additionally secure under linear related-key (RK) attacks. PRFs secure under RK attacks can be achieved, e.g., based on key-homomorphic PRFs [BLMR13].

Bootstrapping distributional security. Moreover, we show a bootstrapping result for circuits which are *output indistinguishable*, i.e. for which there exists an efficient circuit sampler producing C such that $C(x)$ and $C(0^{|x|})$ are computationally indistinguishable for arbitrary adversarially chosen inputs x .

Theorem 1 (Informal). *Assume that garbling scheme G_b is DSIM-secure for garbling the class of output indistinguishable NC^0 circuits with online complexity α and assume IND-CPA secure symmetric encryption exists. Then, there exists a garbling scheme that is DSIM-secure for garbling the class of all polynomial size output indistinguishable circuits with online complexity α .*

Note that the bootstrapping result preserves the online complexity of the garbling scheme, and hence it suffices to construct a DSIM-secure garbling scheme for the class of output indistinguishable NC^0 circuits with online complexity $\mathcal{O}(|x|)$ to obtain a DSIM-secure garbling scheme for the class *all* output indistinguishable circuits with online complexity $\mathcal{O}(|x|)$. The construction of such a garbling scheme remains a very interesting open problem.

Further results about standard simulation-based security To complement our main contribution, we also include two smaller results within the context of the standard adaptive security notion. We generalize the AIKW lower bound and provide small improvements regarding the online complexity for garbling NC^1 circuits.

Bounding online complexity in terms of (pseudo-)entropy. AIKW show that garbling a pseudorandom generator (PRG) with adaptive security requires the online complexity of the garbling scheme to grow with the output size $|C(x)|$. Using a similar, more general idea, Hubáček and Wichs (HW [HW15]) show that the online complexity of maliciously secure 2-party computation is lower bounded by the Yao incompressibility entropy [Yao82b,TVZ05] of the computed function. HW also translates into a (slightly worse) analogous lower bound for the online complexity of adaptively secure garbling schemes: Assume that the

2PC protocol is realized by a garbling scheme combined with a maliciously secure two-party computation protocol to securely evaluate the input garbling (e.g., an oblivious transfer protocol in the case of a projective garbling scheme), one can deduct the communication complexity of the latter and obtain a lower bound for the online complexity of the garbling scheme.

We apply the HW approach to lower bound the online complexity of a garbling scheme by the Yao incompressibility entropy of the outputs of the garbled circuit and thereby obtain a direct proof of this result as well as slightly improved bounds compared to HW. The lower bound includes cryptographic functions such as PRGs, pseudorandom functions, and encryption schemes, all of which (when using suitable parameters) have higher incompressibility entropy than their input size. Note that the more standard notion of HILL pseudo-entropy (indistinguishability from a distribution with k bits of true entropy), if high, also implies high incompressibility entropy, but the converse is not necessarily true [HLR07,HMS23], whence HW and we state our results in terms of incompressibility entropy.

Theorem (Informal). *Let $C \leftarrow_{\S} \text{Sam}_C(1^\lambda)$ be a distribution over circuits and $x \leftarrow_{\S} \text{Sam}_x(C, 1^\lambda)$ be a distribution over inputs, depending on C . Then if Gb is a SIM_{Φ} -secure garbling scheme, its online complexity $|\tilde{x}|$ is greater or equal to the Yao incompressibility-entropy of $C(x)$ (minus a small constant) conditioned on $\Phi(C)$, where Φ is some function on the circuit.*

Here Φ is a function that depends on the flavour of SIM security we want to capture. Typically, Φ is just the topology of the circuit C , which would not contribute to the entropy, but it can also be a less trivial function.

Garbling schemes with (almost) optimal online complexity. Our final result revisits the question of constructing garbling schemes with (almost) optimal online complexity under existing security notions. To circumvent the AIKW lower bound, various works considered the weaker notion of *indistinguishability-based* adaptive security [BHR12a]. (For brevity, we will call this notion adaptive indistinguishability or adaptive IND security.) Table 1 summarizes the existing results, the circuit class they apply to, their online complexity and security assumptions. Note that any circuit with constant treewidth can be simulated in NC^1 [JS14], and that there exist NC^1 circuits that do not have low treewidth, e.g. Goldreich’s PRG [Gol11] is in $\text{NC}^0 \subseteq \text{NC}^1$ [App14] but has high expansion and hence high treewidth. In fact, low treewidth already implies invertibility [Bod88,Fre90,BK08]⁴.

Concretely, we revisit the adaptively indistinguishable garbling scheme by Jafarholi, Scafuro and Wichs (JSW) [JSW17] for garbling NC^1 circuits. The JSW construction has an online complexity that includes a small linear overhead in the *circuit depth*. We show how to modify the construction to remove this overhead.

⁴ See Appendix A for further discussion.

	Circuit class	online complexity	assumption
Ananth & Sahai [AS16]	Arbitrary	$\text{poly}(\lambda, x)$	iO + OWF
Kamath, Klein & Pietrzak [KKP21]	Constant tw	$\lambda x $	OWF
Jafargholi, Scafuro & Wichs [JSW17]	NC^1	$2\lambda x + 2\lambda^2d$	OWF
this work	NC^1	$2\lambda x $	OWF

Table 1: Overview of adaptively indistinguishable garbling schemes with online complexity linear in the input size $|x|$, by circuit class for which security is known. d denotes the circuit depth and λ the security parameter.

Theorem (Informal). *Assuming the existence of one-way functions, there exists an adaptively indistinguishable garbling scheme for NC^1 circuits with online complexity $2\lambda|x|$.*

Jafargholi, Scafuro and Wichs further observed that an adaptively indistinguishable garbling scheme (IND) is also adaptively simulatable when restricting the garbling scheme to *efficiently invertible functions*. Adaptive IND security ensures that garblings of (C_0, x_0) and (C_1, x_1) are indistinguishable, even when x_0 and x_1 are chosen adaptively, as long as $C_0(x_0) = C_1(x_1)$. Since this equality requirement is rather restrictive, it is not obvious how to use IND security for cryptographic circuits in general. However, IND secure garbling should be useful to use with all techniques that are compatible with indistinguishability obfuscation (iO) as well, because iO requires that for all x , $C_0(x) = C_1(x)$ which implies the condition $C_0(x_0) = C_1(x_1)$ for $x = x_0 = x_1$. Therefore, IND security should be useful, e.g., for puncturable PRFs. We show the following:

Corollary (Informal). *Assuming the existence of one-way functions, there exists an adaptively simulatable garbling scheme for efficiently invertible NC^1 circuits with online complexity $2\lambda|x|$.*

1.2 Outline

In Section 2, we provide a technical overview over our main results. Section 3 provides background on garbling schemes and cryptographic primitives. Section 4 introduces our notion of distributional simulation-based security (DSIM). We show how to use the notion in Section 5 on the example of distributed symmetric encryption and prove a bootstrapping results for DSIM-secure garbling schemes in Section 6. Our adaptation of the HW incompressibility approach to generalize the AIKW lower bound is presented in Section 7. Finally, we show how to garble efficiently invertible NC^1 circuits with online complexity $2\lambda|x|$ based on one-way functions in Sections 8.




Fig. 1: Simulation-based security games $\text{Sim}_{\mathcal{A}, \text{Gb}}^0$ and $\text{Sim}_{\mathcal{A}, \mathcal{S}}^1$ for garbling scheme Gb , adversary \mathcal{A} and simulator \mathcal{S} , and circuit leakage function Φ .

2 Technical overview

This section provides an overview of our core contribution, a relaxed simulation-based security notion for garbling schemes.

2.1 Defining distributional simulation-based (DSIM) security

A quick recap of the existing adaptive simulation-based security notion [BHR12a]: Figure 1 shows the real and ideal security games as interaction between adversary \mathcal{A} , the security game, and a simulator \mathcal{S} . In the ideal game $\text{Sim}_{\mathcal{A}, \mathcal{S}}^1$, \mathcal{A} chooses (and hence knows) the circuit C and input x to be garbled, while the simulator has to simulate given leakage $\Phi(C)$ and output $C(x)$ only. Intuitively, this simultaneously captures strong privacy guarantees on the circuit and input.

This security notion gives the adversary the power to choose and learn the full circuit C and the input x to be garbled. The AIKW lower bound and our generalization effectively exploit an attack vector that is specific to this security notion: Since the adversary knows both C and x , they can compute the output $C(x)$ themselves and then compare it to the result of the garbled circuit evaluation. A simulator on the other hand is asked to simulate the garbled input given only $C(x)$. As AIKW and HW show, this simulation can simply not succeed if $C(x)|\Phi(C)$ produces too much (pseudo-)entropy. However, the main purpose of using a garbling scheme is to hide at least some information about either C or x , else one could evaluate $C(x)$ in the clear. For example when garbling a cryptographic function such as an encryption function $\text{enc}(k, \cdot)$ for secret key k , the key is typically sampled as part of the outer protocol. Hence only the garbler knows k . Thus when defining and proving security against a malicious evaluator, it suffices to prove security against an adversary (i.e. the evaluator) who does not know k and thus does not know the entire circuit C . As it suffices to consider a weaker adversary, we introduce two modifications to games $\text{Sim}_{\mathcal{A}, \text{Gb}}^0$ and $\text{Sim}_{\mathcal{A}, \mathcal{S}}^1$:

Modification 1: Partially hiding the circuit from \mathcal{A} . Instead of \mathcal{A} choosing C , the game is parameterized by a sampler Sam whose description is known to the adversary, leading to a *distributional* definition style. The sampler outputs a circuit C and circuit leakage lkg_C . The latter is given to both adversary and simulator




Fig. 2: Distributional simulation-based security games $\text{DSim}_{\text{Sam}, \mathcal{A}, \text{Gb}, \mathcal{A}}^0$ and $\text{DSim}_{\text{Sam}, \mathcal{A}, \mathcal{S}, \mathcal{A}}^1$ for garbling scheme Gb , sampler Sam , adversary \mathcal{A} , simulator \mathcal{S} , and simulator leakage Λ .

and captures the partial information that \mathcal{A} receives about C . An adversary may for instance be allowed to choose and learn which *class* of functions is garbled, e.g. the particular encryption scheme, but not the random choice of key. If on the other hand Sam samples from a distribution containing a single fixed circuit, then the adversary’s knowledge is the same as in the existing security notion. We further augment the leakage so that it outputs a PPT oracle \mathcal{O} which depends on the (secret) state of Sam . Oracle \mathcal{O} models, e.g., encryption queries which depend on a secret key k that is also used in a circuit. Allowing such oracles supports composability, as we will see in Section 5, and it is similar to the interface which the ideal functionality exposes to the simulator in Canetti’s universal composability framework [Can01]. For example, in our case this allows us to prove meaningful security of garbling an \mathbb{S} -RK-AE secure encryption scheme; here the \mathbb{S} -RK-AE security allows the adversary to make multiple queries to the encryption oracle, which would be impossible to model without our “free-form” oracle \mathcal{O} that can contain the encryption oracle in this example use-case.

Modification 2: Relaxing the consistency requirement on \mathcal{S} . The relaxation to partial adversarial knowledge of C then makes it also natural to relax simulation requirements: If the partial knowledge an adversary gains about C and x does not allow to infer $C(x)$ but only some distribution of a possible output, then simulation may be with respect to this distribution instead of a particular output. We capture this observation through a new leakage function $\Lambda(C, x)$ which as before, is applied to obtain the simulator’s input in the game’s online phase. This leakage could be $C(x)$ (as in SIM), or empty (the other extreme) or anything in between, depending on what models the situation best (the less information we give to the simulator, the stronger the security guarantee by DSIM). In our \mathbb{S} -RK-AE encryption example $\Lambda(C, x) = \emptyset$, since as long as the key remains unknown to the adversary, a uniformly random output is indistinguishable from $\text{enc}(k, x)$ for any x . So, the simulator can model the output as a uniformly random string, without knowing the actual output. Interestingly, this gives the security notion a semi-adaptive flavour as the simulator does not depend on the input x at all.

DSIM security games. Figure 2 shows a slightly simplified version of our new security games as interaction between adversary \mathcal{A} , the security game, and a simulator \mathcal{S} . Note that, while consistency is relaxed in the security definition, the garbling scheme is still expected to provide the usual correctness guarantees.

Related work Our distributional security notion for garbling schemes is inspired by some existing techniques for sampling inputs in security games.

Distributional security notions. A distributional approach is the standard way to define security of deterministic public-key encryption schemes [BBO07]. In this context, an adversary that is allowed to choose the message m to be encrypted could simply re-encrypt m under the public key and compare with the challenge ciphertext, leading to a trivial attack. As discussed above, the AIKW lower bound is derived from a similar attack where the adversary can compare $C(x)$ to the output of the simulation. In both cases, the solution is to limit the adversary’s knowledge, and define and prove security only for the scenario where the adversary knows the input *distribution*.

Another context where distributional security notions have been used before is zero knowledge [Gol93,DNRS99,JKKR17,Khu21]. Distributional zero knowledge and related notions relax the zero knowledge property by choosing the statement, i.e. the verifier’s input to the computation, from an efficiently samplable distribution instead of universally quantifying over it. In the context of garbling schemes on the other hand, existing notions give the adversary the power to choose *all* inputs to the computation, and our relaxation to distributional security allows to hide some of the inputs from the adversary while maintaining a relaxed form of adversarial control over them.

Circuits samplers. Our definitional style uses a circuit sampler which is reminiscent of the treatment of Universal Computational Extractors (UCE) as abstraction of keyed hash functions by Bellare, Hoang and Keelveedhi [BHK13]. When using the DSIM definition with respect to the class of *output indistinguishable* circuits, one can view the DSIM notion as a notion of garbling schemes for *probabilistic functionalities* which is somewhat similar to the concept of probabilistic indistinguishability obfuscation (piO), a notion put forward by Canetti, Lin, Tessaro and Vaikuntanathan [CLTV15], where one should not be able to distinguish obfuscations of two computationally indistinguishable distributions \mathcal{D}_0 and \mathcal{D}_1 . Our DSIM variant is different from piO in that DSIM is simulation-based and, since we study garbling schemes, the adversary only gets to see a *single* sample of the distribution.

2.2 Application: Distributed symmetric encryption

A distributed symmetric encryption protocol (DSE) allows multiple servers to jointly perform symmetric encryption of a message, such that each server holds

only a *share* of the key. The notion was formally introduced by Agrawal, Mohassel, Mukherjee and Rindal (AMMR [AMMR18]) and a construction for an arbitrary number of parties based on distributed PRFs was proposed.

Distributed symmetric encryption for arbitrary encryption schemes. For efficiency reasons, existing constructions have focused on distributed version of special-purpose encryption schemes [Muk20,ADL⁺22]. AMMR point out that while DSE can in principle be achieved from general-purpose MPC (which would allow to distribute any existing encryption scheme), this approach would be prohibitively expensive. Nevertheless, such a construction would be preferable when the encryption scheme to be used in an application is already fixed. In this case, efficiency can be improved in other ways, e.g. through a message-independent preprocessing phase.

In Section 5, we show a construction for the two-server case based on a garbling scheme. The idea is to split encryption into two phases: First (and possibly in advance), the encryption circuit is garbled. Then in the online phase, only the message needs to be garbled. This preprocessing strategy means that the garbling scheme needs some flavour of adaptive security and we show that our new distributional simulation-based security notion suffices.

Remark. Encryption is typically length-preserving, so one could use a garbling scheme whose online complexity is proportional to its output length also. However, if the same message was, say, encrypted under t keys into t ciphertexts, then we would have a length-expanding circuit whose pseudo-entropy grows with t . For simplicity of exposition, we will focus on the single-key version.

We show security in a simplified and more restricted model in comparison to AMMR: We assume a two-server setting where only server S_2 makes encryption queries. Server S_2 acts as evaluator in the garbling scheme, and we focus on showing security against a corrupt S_2 as this is the more difficult case. Decryption is assumed to be performed by a trusted party. Finally, we restrict the adversary to multiple indirect encryption queries but only a single direct encryption query in the terminology of AMMR.

Protocol overview. For two servers S_1 and S_2 , the construction works as follows. Assume a symmetric authenticated encryption scheme se , a garbling scheme Gb with DSIM security with respect to output indistinguishable samplers such as authenticated encryption, and ideal oblivious transfer. In the basic construction, each server S_i holds an additive share k_i of the symmetric encryption key $k = k_1 \oplus k_2$. To encrypt a message m under se and k and randomness $r = r_1 \oplus r_2$, the servers act as follows: Server S_1 samples r_1 and garbles $\text{enc}(k_1 \oplus \cdot, \cdot; r_1 \oplus \cdot)$. Both servers then run an OT protocol to garble the missing inputs k_2 , m , and r_2 . We show that the AE security of the symmetric encryption scheme se still holds against a malicious server S_2 . We need to assume related-key security of the AE against linear functions since the malicious server can choose an arbitrary key k_1 to be xored on k_2 .

For the case of encrypting the message under t different keys k^1, \dots, k^t , the output of the circuit would consist of t ciphertexts. To compress the size of S_2 's input, the key shares k_2^j can be computed by a PRG as $k_2^j = \text{PRG}(s^j)$, and similarly for S_2 's randomness. Then in the online phase, we only need to garble the two PRG seeds and the message, both of which are *independent* in size of t and hence the output size. See Section 5 for details on the single-key version of the protocol.

2.3 Bootstrapping NC^0 to polynomial-size circuits

In the realm of obfuscation, *bootstrapping* is an established technique to turn positive results for a *limited* class of circuits into a positive result for arbitrary polynomial-size circuits. The bootstrapping techniques in the context of obfuscation follow the randomized encoding approach by Applebaum, Ishai, Kushilevitz [AIK04], see e.g. [GGH⁺13, Agr19] for examples. Assume we can obfuscate low-depth circuits. Instead of directly obfuscating a given high-depth circuit $C(x)$ for inputs x , a randomized encoding function is applied to the circuit. The result $\tilde{C}(x; r)$ is a new circuit that outputs a randomized encoding of $C(x)$. $\tilde{C}(x; r)$ uses additional randomness r and is chosen such that for random r , $\tilde{C}(x; r)$ does not leak more information than $y = C(x)$. If the encoding function itself has low depth, then $\tilde{C}(x; r)$ can now be obfuscated, and y can be recovered from $\tilde{C}(x; r)$ via a (high-depth, public) decoding function Dec .

Since randomized encodings and garbling schemes are different abstractions of the same underlying idea, it is natural to ask if the randomized encodings bootstrapping technique can be used to bootstrap adaptive security of garbling schemes from low-depth circuits to arbitrary polynomial-size circuits. That is, instead of directly garbling a circuit $C(x)$, we could try to (outer) garble the (inner) garbling function $\tilde{C}(x; r) := [\text{GCircuit}(C(\cdot); r), \text{GInput}(x)]$ with explicit hardcoded randomness r . Now, if the inner garbling scheme produces a low depth circuit \tilde{C} , the outer garbling scheme only needs to be secure for low-depth circuits, and, the inner garbling scheme only needs to be selectively secure, not adaptively. Now, the inner garbling scheme \tilde{C} can be simply, e.g., Yao's garbling scheme, which is useful for two reasons: (1) Yao can be proven selectively secure assuming it uses $\$$ -IND-CPA symmetric encryption [LP09, BHR12b, BO23] and (2) Yao's garbling can be implemented in constant depth (see Lemma 1 for details).

However, proving SIM security in this setting is infeasible for the following reason. When we try to reduce the SIM security of the combined (garbling of garbling) garbling scheme to the SIM security of the outer garbling scheme, the adversarial evaluator chooses the circuit to be garbled, in this case, the inner garbling function $\tilde{C}(\cdot; r)$ with the hardcoded randomness r . Adversary knowing r trivially compromises the security of the inner garbling.

Luckily, our DSIM security notion precisely allows to circumvent the randomness issue by dividing the adversary into two parts that do not share a state: first part samples the circuit from a suitable *output indistinguishable* distribution (e.g. $\$$ -IND-CPA secure encryption with a uniformly random key, see

Definition 8 for how exactly we define the natural property of output indistinguishability) and the second adversary interacts with the garbler. Hence, we can prove Theorem 1.

3 Preliminaries

All algorithms take as input the security parameter 1^λ . We write it explicitly for some algorithms, but leave it implicit for most algorithms. $a \leftarrow A(x)$ assigns the result of an execution of the deterministic algorithm A on input x to variable a . $a \leftarrow_{\$} A(x)$ denotes the execution of a *randomized* algorithm. $a||b$ denotes the concatenation of two bit strings a and b . $A_2 \circ A_1$ denotes the composition of two algorithms in the form $A_2(A_1(\cdot))$. We often write $\mathcal{A}^{\mathcal{O}_1, \mathcal{O}_2}$ for an adversary accessing oracle \mathcal{O}_1 and \mathcal{O}_2 and sometimes write $\mathcal{A}_{\mathcal{O}_2}^{\mathcal{O}_1}$ for conciseness. We sometimes use the notation $\mathcal{A}^{\mathcal{O}_1, \dots, \mathcal{O}_t} \mathbb{G}^b$ inspired from state-separating proofs [BDF⁺18] to say that the adversary \mathcal{A} is the main procedure who has access to oracles $\mathcal{O}_1, \dots, \mathcal{O}_t$ of a distinguishing game \mathbb{G}^b . $\Pr[1 = \mathcal{A}^{\mathcal{O}_1, \dots, \mathcal{O}_t} \mathbb{G}^b]$ then refers to the probability that the adversary \mathcal{A} , after interacting with the oracles $\mathcal{O}_1, \dots, \mathcal{O}_t$ of \mathbb{G}^b returns 1. We like this notation, because it makes both the name of the game and the adversary's oracle explicit, while the standard oracle subscript notation $\mathcal{A}^{\mathcal{O}_1, \dots, \mathcal{O}_t}$ only contains the oracles (but not the name of the game), and the experiment notation $\mathbb{G}_{\mathcal{A}}^b$ only contains the name of the game (but not the oracles).

3.1 Cryptographic primitives

For a symmetric encryption scheme se , $\$$ -IND-CPA-security captures that ciphertexts are indistinguishable from random strings of the same length. The related-key authenticated encryption security game $\$$ -RK-AE provides decryption queries in addition

$\underline{\$$ -IND-CPA $_{\text{se}}^b(1^\lambda)$	$\underline{\$$ -RK-AE $_{\text{se}}^b(1^\lambda)$	
ENC(m)	ENC(m, Δ)	DEC(c, Δ)
if $k = \perp$	if $k = \perp$	if $b = 1$
$k \leftarrow_{\$} \{0, 1\}^\lambda$	$k \leftarrow_{\$} \{0, 1\}^\lambda$	return \perp
$c \leftarrow_{\$} \text{enc}(k, m)$	$c \leftarrow_{\$} \text{enc}(k \oplus \Delta, m)$	if $k = \perp$
	$S \leftarrow S \cup \{c\}$	$k \leftarrow_{\$} \{0, 1\}^\lambda$
if $b = 1$	if $b = 1$	assert $c \notin S$
$c \leftarrow_{\$} \{0, 1\}^{ c }$	$c \leftarrow_{\$} \{0, 1\}^{ c }$	$m \leftarrow \text{dec}(k \oplus \Delta, c)$
return c	return c	return m

Fig. 3: Games $\$$ -IND-CPA $_{\text{se}}^b(1^\lambda)$ and $\$$ -RK-AE $_{\text{se}}^b(1^\lambda)$.

and allows adversarially chosen linear offsets on the key. Related-key security has been introduced by Bellare and Kohno [BK03].

Definition 1 ($\$$ -IND-CPA and $\$$ -RK-AE Security). *A symmetric encryption scheme $\text{se} = (\text{enc}, \text{dec})$ is indistinguishable under chosen plaintext attacks ($\$$ -IND-CPA) if for all PPT adversaries \mathcal{A} , the advantage $\text{Adv}_{\text{se}, \mathcal{A}}^{\$$ -IND-CPA}(1^\lambda) :=*

$$\left| \Pr[1 = \mathcal{A}(1^\lambda) \xrightarrow{\text{ENC}} \$\text{-IND-CPA}^0(1^\lambda)] - \Pr[1 = \mathcal{A}(1^\lambda) \xrightarrow{\text{ENC}} \$\text{-IND-CPA}^1(1^\lambda)] \right|$$

is negligible in λ . se is authenticated encryption under linear related key attacks if for all PPT adversaries \mathcal{A} , the advantage $\text{Adv}_{\text{se}, \mathcal{A}}^{\text{\$-RK-AE}}(1^\lambda) :=$

$$\left| \Pr \left[1 = \mathcal{A}(1^\lambda) \xrightarrow{\text{ENC, DEC}} \text{\$-RK-AE}^0(1^\lambda) \right] - \Pr \left[1 = \mathcal{A}(1^\lambda) \xrightarrow{\text{ENC, DEC}} \text{\$-RK-AE}^1(1^\lambda) \right] \right|$$

is negligible in λ .

3.2 Garbling schemes

Definition 2 (Garbling scheme). A Garbling scheme Gb consists of three PPT algorithms (GCircuit , GInput , GEval) with the following syntax:

- $(\tilde{C}, K, d) \leftarrow_{\text{\$}} \text{GCircuit}(C)$: the garbling algorithm takes as input a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$, and outputs a garbled circuit \tilde{C} , keys K and output decoding information d . m and n are polynomials in λ .
- $\tilde{x} \leftarrow_{\text{\$}} \text{GInput}(K, x)$: the input garbling algorithm takes an input $x \in \{0, 1\}^n$ and the keys K and outputs an encoding of the input \tilde{x} .
- $y \leftarrow_{\text{\$}} \text{GEval}(\tilde{C}, \tilde{x}, d)$: the evaluation algorithm takes as input the garbled circuit, a garbled input, the output encoding information d and returns a value $y \in \{0, 1\}^m$, where m is the output length of C .

Correctness. of the garbling scheme holds if for any λ , any circuit C and any input $x \in \{0, 1\}^n$ we have

$$\Pr \left[C(x) = \text{GEval}(\tilde{C}, \tilde{x}, d) \right] = 1 - \text{negl}(\lambda),$$

where $(\tilde{C}, K, d) \leftarrow_{\text{\$}} \text{GCircuit}(C)$, $\tilde{x} \leftarrow_{\text{\$}} \text{GInput}(K, x)$.

We first define *selective simulation security*, where the adversary chooses circuit C and input x at the same time. We define adaptive garbling scheme security in Section 4, where we compare it to the DSIM definition, which we propose.

$\text{SelSim}_{\mathcal{A}, \text{Gb}}^0(1^\lambda)$ $(C, x, \text{st}) \leftarrow_{\text{\$}} \mathcal{A}(1^\lambda)$ $(\tilde{C}, K, d) \leftarrow_{\text{\$}} \text{GCircuit}(C)$ $\tilde{x} \leftarrow_{\text{\$}} \text{GInput}(K, x)$ $b' \leftarrow_{\text{\$}} \mathcal{A}(\tilde{C}, \tilde{x}, d, \text{st})$ return b'	$\text{SelSim}_{\mathcal{A}, \mathcal{S}}^1(1^\lambda)$ $(C, x, \text{st}) \leftarrow_{\text{\$}} \mathcal{A}(1^\lambda)$ $(\tilde{C}, \tilde{x}, d) \leftarrow_{\text{\$}} \mathcal{S}(\Phi(C), C(x))$ $b' \leftarrow_{\text{\$}} \mathcal{A}(\tilde{C}, \tilde{x}, d, \text{st})$ return b'
--	--

Fig. 4: Experiments $\text{SelSim}_{\mathcal{A}, \text{Gb}}^0$ and $\text{SelSim}_{\mathcal{A}, \mathcal{S}}^1$.

Definition 3 (SelSim security). Let Φ be a leakage function. We say that Gb is *selectively simulation secure (SelSim)* if for any PPT adversary \mathcal{A} there exists a PPT simulator \mathcal{S} such that

$$\text{Adv}_{\mathcal{A}, \mathcal{S}, \text{Gb}}^{\text{SelSim}}(1^\lambda) := \left| \Pr \left[1 = \text{SelSim}_{\mathcal{A}, \mathcal{S}, \text{Gb}}^0 \right] - \Pr \left[1 = \text{SelSim}_{\mathcal{A}, \mathcal{S}}^1 \right] \right| = \text{negl}(\lambda),$$

where experiments $\text{SelSim}_{\mathcal{A}, \text{Gb}}^0$ and $\text{SelSim}_{\mathcal{A}, \mathcal{S}}^1$ are specified in Figure 4.

4 Distributional simulation-based security (DSIM)

In this section, we first define the simulation-based and indistinguishability flavours of adaptive security and then present our new definition of adaptive distributional simulation-based security (DSIM) for garbling schemes.

$\text{Sim}_{\mathcal{A},\text{Gb}}^0(1^\lambda)$	$\text{Sim}_{\mathcal{A},\mathcal{S}}^1(1^\lambda)$	$\text{Ind}_{\mathcal{A},\text{Gb}}^b(1^\lambda)$
$(C, \text{st}) \leftarrow_{\$} \mathcal{A}(1^\lambda)$	$(C, \text{st}) \leftarrow_{\$} \mathcal{A}(1^\lambda)$	$(C_0, C_1, \text{st}) \leftarrow_{\$} \mathcal{A}(1^\lambda)$
$(\tilde{C}, K, d) \leftarrow_{\$} \text{GCircuit}(C)$	$(\tilde{C}, d, \text{st}_S) \leftarrow_{\$} \mathcal{S}(\Phi(C))$	$(\tilde{C}, K, d) \leftarrow_{\$} \text{GCircuit}(C_b)$
$(x, \text{st}') \leftarrow_{\$} \mathcal{A}(\tilde{C}, d, \text{st})$	$(x, \text{st}') \leftarrow_{\$} \mathcal{A}(\tilde{C}, d, \text{st})$	$(x_0, x_1, \text{st}') \leftarrow_{\$} \mathcal{A}(\tilde{C}, d, \text{st})$
$\tilde{x} \leftarrow_{\$} \text{GInput}(K, x)$	$\tilde{x} \leftarrow_{\$} \mathcal{S}(C(x), \text{st}_S)$	$\tilde{x} \leftarrow_{\$} \text{GInput}(K, x_b)$
$b' \leftarrow_{\$} \mathcal{A}(\tilde{x}, \text{st}')$	$b' \leftarrow_{\$} \mathcal{A}(\tilde{x}, \text{st}')$	$b' \leftarrow_{\$} \mathcal{A}(\tilde{x}, \text{st}')$
return b'	return b'	return b'
$\text{DSim}_{\text{Sam},\mathcal{A},\text{Gb},\mathcal{A},\text{Filter}}^0(1^\lambda)$	$\text{DSim}_{\text{Sam},\mathcal{A},\mathcal{S},\mathcal{A},\text{Filter}}^1(1^\lambda)$	
$(C, \text{lk}_{\mathcal{G}_C}, \mathcal{O}) \leftarrow_{\$} \text{Sam}(1^\lambda)$	$(C, \text{lk}_{\mathcal{G}_C}, \mathcal{O}) \leftarrow_{\$} \text{Sam}(1^\lambda)$	
$(\tilde{C}, K, d) \leftarrow_{\$} \text{GCircuit}(C)$	$(\tilde{C}, d, \text{st}_S) \leftarrow_{\$} \mathcal{S}^\mathcal{O}(\text{Filter}(\text{lk}_{\mathcal{G}_C}))$	
$(x, \text{st}_{\mathcal{A}}) \leftarrow_{\$} \mathcal{A}^\mathcal{O}(\tilde{C}, d, \text{lk}_{\mathcal{G}_C})$	$(x, \text{st}_{\mathcal{A}}) \leftarrow_{\$} \mathcal{A}^\mathcal{O}(\tilde{C}, d, \text{lk}_{\mathcal{G}_C})$	
$\tilde{x} \leftarrow_{\$} \text{GInput}(K, x)$	$\tilde{x} \leftarrow_{\$} \mathcal{S}^\mathcal{O}(\mathcal{A}(\text{st}_{\mathcal{A}}), \text{st}_S)$	
$b' \leftarrow_{\$} \mathcal{A}^{\mathcal{O}C(x)}(\tilde{x}, \text{st}_{\mathcal{A}})$	$b' \leftarrow_{\$} \mathcal{A}^{\mathcal{O}\tilde{C}(\tilde{x})}(\tilde{x}, \text{st}_{\mathcal{A}})$	
return b'	return b'	

Fig. 5: Experiments $\text{Ind}_{\mathcal{A},\text{Gb}}^b$, $\text{Sim}_{\mathcal{A},\text{Gb}}^0$, $\text{Sim}_{\mathcal{A},\mathcal{S}}^1$, $\text{DSim}_{\text{Sam},\mathcal{A},\text{Gb},\mathcal{A},\text{Filter}}^0(1^\lambda)$ and $\text{DSim}_{\text{Sam},\mathcal{A},\mathcal{S},\mathcal{A},\text{Filter}}^1(1^\lambda)$.

Definition 4 (Adaptive SIM security). *Let Φ be a leakage function. We say that Gb is adaptively SIM_Φ -secure if for any PPT adversary \mathcal{A} there exists a PPT simulator \mathcal{S} such that*

$$\text{Adv}_{\mathcal{A},\mathcal{S},\text{Gb}}^{\text{Sim}}(1^\lambda) := |\Pr[1 = \text{Sim}_{\mathcal{A},\mathcal{S},\text{Gb}}^0] - \Pr[1 = \text{Sim}_{\mathcal{A},\mathcal{S}}^1]| = \text{negl}(\lambda),$$

where experiments $\text{Sim}_{\mathcal{A},\text{Gb}}^0$ and $\text{Sim}_{\mathcal{A},\mathcal{S}}^1$ are specified in Figure 5.

Definition 5 (Adaptive IND security). *We say that Gb is **adaptively IND-secure** if for any PPT adversary \mathcal{A} which queries circuits C_0 and C_1 with equal input length and inputs x_0 and x_1 such that $C_0(x_0) = C_1(x_1)$ and $\Phi(C_0) = \Phi(C_1)$, there exists a negligible function $\text{negl}(\lambda)$ such that:*

$$\text{Adv}_{\mathcal{A},\text{Gb}}^{\text{Ind}}(1^\lambda) := |\Pr[1 = \text{Ind}_{\mathcal{A},\text{Gb}}^0] - \Pr[1 = \text{Ind}_{\mathcal{A},\text{Gb}}^1]| = \text{negl}(\lambda),$$

where the experiment $\text{Ind}_{\mathcal{A},\text{Gb}}^b$ is specified in Figure 5.

Definition 6 (Sampler classes). A class of samplers \mathcal{C} is a set of PPT adversaries Sam such that $(C, \text{lkg}_C, \mathcal{O}) \leftarrow \$ \text{Sam}(1^\lambda)$.

Definition 7 (Adaptive distributional SIM security (DSIM)). Let \mathcal{C} be a sampler class and Λ be a leakage-function. Garbling scheme Gb is $\text{DSim}_{\Lambda, \text{Filter}}[\mathcal{C}]$ -secure if for any PPT $\text{Sam} \in \mathcal{C}$ and any PPT \mathcal{A} , there exists a PPT simulator \mathcal{S} such that $\text{Adv}_{\text{Sam}, \mathcal{A}, \mathcal{S}, \text{Gb}}^{\text{DSim}}(1^\lambda) :=$

$$|\Pr[1 = \text{DSim}_{\text{Sam}, \mathcal{A}, \text{Gb}, \Lambda, \text{Filter}}^0(1^\lambda)] - \Pr[1 = \text{DSim}_{\text{Sam}, \mathcal{A}, \mathcal{S}, \Lambda, \text{Filter}}^1(1^\lambda)]|$$

is negligible, where Figure 5 defines the experiments $\text{DSim}_{\text{Sam}, \mathcal{A}, \text{Gb}, \Lambda, \text{Filter}}^0(1^\lambda)$ and $\text{DSim}_{\text{Sam}, \mathcal{A}, \mathcal{S}, \Lambda, \text{Filter}}^1(1^\lambda)$.

In the above DSIM definition lkg_C could be e.g. the topology of the circuit C (to match the SIM definition). Alternatively, if we are garbling e.g. an encryption scheme with the key hardcoded in the circuit C , then lkg_C could be the circuit C without the hardcoded random key. \mathcal{O} is a (possibly stateful) PPT oracle which might be, for example, an encryption oracle which depends on the key k embedded into C , $\mathcal{O}_{C(x)}$ might additionally depend on $C(x)$, e.g., when forbidding to send $C(x)$ to a decryption oracle, see Section 5 for an example. See Section 2.1 for more intuition behind the different parameters.

In the DSIM definition, the grey parameters (Λ, Filter) are optional and can be ignored (i.e. Filter can be thought of as identity and Λ as empty) when reading this paper, that is, they are not needed for understanding the bootstrapping proof or DENC example. They are included just as an example on how to extend the definition and to draw connection to SIM definition.

Remark on the optional parameters. W.l.o.g. we consider that $\text{st}_{\mathcal{A}}$ contains variables for lkg_C and x , and, if lkg_C allows to compute C , then we also assume that w.l.o.g., lkg_C and $\text{st}_{\mathcal{A}}$ contain a variable for C which contains the correct value. This will later allow us to discuss specific leakage functions.

In the special case when the sampler leaks $\text{lkg}_C := C$ and $\Lambda(\text{st}_{\mathcal{A}}) = C(x)$, SIM implies DSIM, because the adversary has slightly less information in DSIM than in SIM (since it does not know the randomness used for sampling), and additionally, the simulator \mathcal{S} is stronger in DSIM.

Theorem 2 (SIM implies DSIM). Let Φ be a polynomial-time computable leakage function. If a garbling scheme Gb is SIM_{Φ} -secure, then Gb is $\text{DSim}_{\Lambda, \text{Filter}}[\mathcal{C}]$ -secure, where $\mathcal{C} = \{\text{PPT Sam} : (C, \text{lkg}_C := C) \leftarrow \$ \text{Sam}\}$, $\Lambda(\text{st}_{\mathcal{A}}) = C(x)$ and $\text{Filter}(\text{lkg}_C)$ is such that $\Phi(\text{lkg}_C)$ can be computed given $\text{Filter}(\text{lkg}_C)$.

Proof. Let Gb be a SIM_{Φ} -secure garbling scheme. Assume towards contradiction that there is $\text{Sam} \in \mathcal{C}$ and PPT \mathcal{A} s.t. for all PPT \mathcal{S} $\text{Adv}_{\text{Sam}, \mathcal{A}, \mathcal{S}, \text{Gb}}^{\text{DSim}}(1^\lambda)$ is non-negligible. Define the first stage of the SIM-adversary as $\mathcal{A}'(1^\lambda) := \text{Sam}(1^\lambda)$ and \mathcal{A}' is \mathcal{A} for all further stages. Let \mathcal{S}' be the simulator ensured by SIM-security such that $\text{Adv}_{\mathcal{A}', \mathcal{S}', \text{Gb}}^{\text{Sim}}(1^\lambda)$ is negligible.

Now, define the first stage simulator as $\mathcal{S}(\text{Filter}(\text{lkg}_C)) = \mathcal{S}'(\Phi(\text{lkg}_C))$ and the second stage simulator as $\mathcal{S}(\cdot, \cdot) := \mathcal{S}'(\cdot, \cdot)$ and \mathcal{A}' is \mathcal{A} for all other inputs. Now $\text{Adv}_{\mathcal{A}', \mathcal{S}', \text{Gb}}^{\text{Sim}}(1^\lambda) = \text{Adv}_{\text{Sam}, \mathcal{A}, \mathcal{S}, \text{Gb}}^{\text{DSim}}(1^\lambda)$, and we reached a contradiction. \square

We can also recover the other direction, DSIM implies SIM, if we choose lkg_C to be the sampler's full state and choose $\text{Filter}(\text{lkg}_C) = \Phi(C)$ and $A(\text{st}_A) = C(x)$.

Note however, that the core difference between SIM and DSIM is the idea of *sampler classes* which *restrict* the information lkg which is passed from Sam to \mathcal{A} ⁵. Recall that this restriction models that in garbling scheme application such as private function evaluation, the party choosing the input might have *some* information about the function to be evaluated, but not *all* information—otherwise one party could simply send its function to the other party. Of particular interest to us are circuits with embedded *cryptographic keys* which make the output of the evaluated circuit *indistinguishable* from a distribution which could have been chosen *independently* of the input to the circuit, as, e.g., in IND-CPA-secure encryption where an encryption of a message m is indistinguishable from an encryption of $0^{|m|}$. To

be useful for garbling security, we need the indistinguishable distribution to be generated by a circuit of the same size. Therefore, we demand that the function distribution consists of a fixed circuit C for which only the *randomness* is sampled, i.e., $C := C(\cdot; r)$ is a circuit with randomness r hardcoded into it.

$\begin{array}{l} \text{Sam}(1^\lambda) \\ \hline r \leftarrow_{\$} \{0, 1\}^{p(\lambda)} \\ C_r \leftarrow C_\lambda(\cdot, r) \\ (\text{lkg}_{C_r}, \mathcal{O}) \leftarrow \text{Sam}_{\text{leak}}(r) \\ \mathbf{return} C_r, \text{lkg}_{C_r}, \mathcal{O} \end{array}$	$\begin{array}{l} \text{Out}_{\text{Sam}, \mathcal{D}}^b(1^\lambda) \\ \hline C_r, \text{lkg}_{C_r}, \mathcal{O} \leftarrow_{\$} \text{Sam}(1^\lambda) \\ (x_0, \text{st}) \leftarrow_{\$} \mathcal{D}^{\mathcal{O}}(1^\lambda, \text{lkg}_{C_r}) \\ x_1 \leftarrow 0^{ x_0 } \\ b^* \leftarrow_{\$} \mathcal{D}^{\mathcal{O}_{C_r(x_b)}}(C_r(x_b), \text{st}) \\ \mathbf{return} b^* \end{array}$
--	--

Fig. 6: Output indistinguishability

Definition 8 (Output Indistinguishable Sampler). We define the class \mathcal{C}^{out} of output indistinguishable samplers as the set of PPT Sam such that

Fixed circuit there exists a circuit $C = (C_\lambda)_\lambda$, polynomial p and PPT Sam_{leak} such that Sam can be written as in Fig. 6 (left), and

Output indistinguishability for all PPT distinguishers \mathcal{D} , the advantage

$$\text{Adv}_{\text{Sam}, \mathcal{D}}^{\text{Out}}(1^\lambda) := |\Pr[1 = \text{Out}_{\text{Sam}, \mathcal{D}}^0(1^\lambda)] - \Pr[1 = \text{Out}_{\text{Sam}, \mathcal{D}}^1(1^\lambda)]| = \text{negl}(\lambda),$$

where Fig. 6 defines $\text{Out}_{\text{Sam}, \mathcal{D}}^0(1^\lambda)$ and $\text{Out}_{\text{Sam}, \mathcal{D}}^1(1^\lambda)$.

Remark. Note that \mathcal{D} does not need to receive any leakage about C since the circuit is fixed. We use the notation lkg_{C_r} just to be consistent with definition of DSIM.

5 Distributed Symmetric Encryption (DSE)

This section shows that DSIM security with respect to admissible samplers implies useful security properties on the example of distributed encryption that

⁵ hence, in most cases where DSIM is a meaningful notion, Filter should just be identity.

was introduced in Section 2. For simplicity of exposition, we focus on the case of two servers and a single message. See Section 2 for further introduction to our example and a discussion on extending the example into a length-expanding case while maintaining the same online complexity.

Remember that we are in a setting with two servers who hold secret shares k_1, k_2 of a symmetric key $k = k_1 \oplus k_2$, and who wish to perform distributed encryptions under k . One way to implement such a protocol is by combining a garbling scheme with an oblivious transfer protocol to obtain a two-party computation protocol for evaluating $\text{enc}(k_1 \oplus k_2, x)$ for message $x \in \{0, 1\}^\lambda$.

Below, we describe the protocol between two stateful servers. We refer to the protocol below as *distributed symmetric encryption protocol (DSE)*, using a symmetric encryption scheme se and a garbling scheme Gb . We keep state implicit for conciseness and abstract away the OT protocol. DSE uses a trusted party in the setup phase who generates and secret-shares a symmetric key k .

<u>Phase I</u>	<u>Phase II</u>	
Server 1	Server 2	Server 2
Input: $k_1, 1^\lambda$	Input: $k_2, \text{msg}, \tilde{C}, d, 1^\lambda$	Input: $\tilde{x}, \tilde{C}, d, 1^\lambda$
$r_1 \leftarrow_{\$} \{0, 1\}^\lambda$	$r_2 \leftarrow_{\$} \{0, 1\}^\lambda$	$x' \leftarrow \text{Gb.GEval}(\tilde{C}, \tilde{x}, d)$
$C \leftarrow \text{se.enc}(k_1 \oplus \cdot, \cdot; r_1 \oplus \cdot)$	$x \leftarrow k_2 \parallel \text{msg} \parallel r_2$	return x'
$(\tilde{C}, K, d) \leftarrow_{\$} \text{Gb.GCircuit}(C)$	return x	
return (\tilde{C}, d)		

Security. The protocol shall provide security in a setting where the adversary obtains the key of one of the two servers. We here focus on corruption of the key of Server 2 and model security in this case. The security notion we consider is $\$$ -AE security in the presence of the above distributed protocol under the same key. In this toy example, we consider only a single corrupted execution of the protocol.

Security is defined as indistinguishability of $\text{Denc}_{\mathcal{A}}^0(1^\lambda)$ and $\text{Denc}_{\mathcal{A}}^1(1^\lambda)$ in Fig. 7. The game starts by sampling key shares k_1 and k_2 and computing a garbled circuit \tilde{C} and d , either honestly or via simulator \mathcal{S} . The adversary \mathcal{A} , acting as corrupted Server 2, is then given \tilde{C} , d and k_2 and outputs a message msg to be encrypted. The game proceeds by garbling this message together with encryption randomness r_2 and providing this \tilde{x} to \mathcal{A} . At every step, \mathcal{A} has access to encryption and decryption oracles, modeling honest executions run concurrently. Finally, \mathcal{A} outputs a guess b' .

Definition 9 (Distributed Encryption Security (DENC)). *Let se be a symmetric encryption scheme where ciphertexts are twice as long as the plaintexts. Let Gb be a garbling scheme. Then, the DSE protocol using se and Gb is DENC-secure if for all PPT adversaries \mathcal{A} , there exists a PPT simulator \mathcal{S} such that*

$$|\Pr[1 = \text{Denc}_{\mathcal{A}}^0(1^\lambda)] - \Pr[1 = \text{Denc}_{\mathcal{A}, \mathcal{S}}^1(1^\lambda)]|$$

is negligible, where Fig. 7 defines $\text{Denc}_{\mathcal{A}}^0(1^\lambda)$ and $\text{Denc}_{\mathcal{A}}^1(1^\lambda)$.

$\text{Denc}_{\mathcal{A}}^0(1^\lambda)$	$\text{Denc}_{\mathcal{A},\mathcal{S}}^1(1^\lambda)$	$\text{Hybrid}_{\mathcal{A}}(1^\lambda)$
$k_1 \leftarrow_{\$} \{0, 1\}^\lambda$	$k_1 \leftarrow_{\$} \{0, 1\}^\lambda$	$k_1 \leftarrow_{\$} \{0, 1\}^\lambda$
$k_2 \leftarrow_{\$} \{0, 1\}^\lambda$	$k_2 \leftarrow_{\$} \{0, 1\}^\lambda$	$k_2 \leftarrow_{\$} \{0, 1\}^\lambda$
$k \leftarrow k_1 \oplus k_2$	$k \leftarrow k_1 \oplus k_2$	$k \leftarrow k_1 \oplus k_2$
$r_1 \leftarrow_{\$} \{0, 1\}^\lambda$		
$C \leftarrow \text{enc}(k_1 \oplus \cdot, \cdot; r_1 \oplus \cdot)$		
$(\tilde{C}, K, d) \leftarrow_{\$} \text{GCircuit}(C)$	$(\tilde{C}, d) \leftarrow_{\$} \mathcal{S}(1^\lambda)$	$(\tilde{C}, d) \leftarrow_{\$} \mathcal{S}_{\text{hybrid}}^{\text{ENC,DEC}}(1^\lambda)$
$\text{msg} \leftarrow_{\$} \mathcal{A}_{\text{DEC}}^{\text{ENC}}(1^\lambda, k_2, \tilde{C}, d)$	$\text{msg} \leftarrow_{\$} \mathcal{A}_{\text{DEC}}^{\text{ENC}}(1^\lambda, k_2, \tilde{C}, d)$	$\text{msg} \leftarrow_{\$} \mathcal{A}_{\text{DEC}}^{\text{ENC}}(1^\lambda, k_2, \tilde{C}, d)$
$r_2 \leftarrow_{\$} \{0, 1\}^\lambda$		
$x \leftarrow k_2 \text{msg} r_2$		
$\tilde{x} \leftarrow_{\$} \text{GInput}(K, x)$	$\tilde{x} \leftarrow_{\$} \mathcal{S}(1^\lambda)$	$\tilde{x} \leftarrow_{\$} \mathcal{S}_{\text{hybrid}}^{\text{ENC,DEC}}(1^\lambda)$
$S \leftarrow S \cup \{\tilde{C}(\tilde{x})\}$		$S \leftarrow S \cup \{\tilde{C}(\tilde{x})\}$
$b' \leftarrow_{\$} \mathcal{A}_{\text{DEC}}^{\text{ENC}}(1^\lambda, \tilde{x})$	$b' \leftarrow_{\$} \mathcal{A}_{\text{DEC}}^{\text{ENC}}(1^\lambda, \tilde{x})$	$b' \leftarrow_{\$} \mathcal{A}_{\text{DEC}}^{\text{ENC}}(1^\lambda, \tilde{x})$
return b'	return b'	return b'
<hr/>		
$\text{ENC}(m)$	$\text{ENC}(m)$	$\text{ENC}(m)$
assert $ m = \lambda$	assert $ m = \lambda$	assert $ m = \lambda$
$c \leftarrow_{\$} \text{enc}(k, m)$	$c \leftarrow_{\$} \{0, 1\}^{2\lambda}$	$c \leftarrow_{\$} \text{enc}(k, m)$
$S \leftarrow S \cup \{c\}$		$S \leftarrow S \cup \{c\}$
return c	return c	return c
<hr/>		
$\text{DEC}(c)$	$\text{DEC}(c)$	$\text{DEC}(c)$
assert $c \notin S$		assert $c \notin S$
$m \leftarrow \text{dec}(k, c)$		$m \leftarrow \text{dec}(k, c)$
return m	return \perp	return m

Fig. 7: Security games (left and middle) for DENC-security and hybrid game (right) for Theorem 3. The adversary \mathcal{A} and simulators \mathcal{S} and $\mathcal{S}_{\text{hybrid}}$ are *stateful*. We leave their state implicit for conciseness of notation.

Remark. The purpose of the ENC oracle is to model a weaker form of concurrent security. To closely model the real-life scenario of distributed encryption, it might seem useful to replace the ENC oracle also by a circuit garbling oracle. We remark that this would require studying concurrent composition of DSIM which is a non-trivial question. Concretely, when composing the same garbled circuit

$\text{Sam}(1^\lambda)$ <hr style="border: 0.5px solid black;"/> $k_1 \leftarrow_{\$} \{0, 1\}^\lambda$ $k_2 \leftarrow_{\$} \{0, 1\}^\lambda$ $k \leftarrow k_1 \oplus k_2$ $r_1 \leftarrow_{\$} \{0, 1\}^\lambda$ $C \leftarrow \text{enc}(k_1 \oplus \cdot, \cdot; r_1 \oplus \cdot)$ $\text{lkg}_C \leftarrow k_2$ $\mathcal{O} \leftarrow \text{ENC}, \text{DEC}$ return $(C, \text{lkg}_C, \mathcal{O})$	$\mathcal{O}_{C(x)}$ <hr style="border: 0.5px solid black;"/> $\text{ENC}(m)$ assert $ m = \lambda$ $c \leftarrow_{\$} \text{enc}(k, m)$ $S \leftarrow S \cup \{c\}$ return c <hr style="border: 0.5px solid black;"/> $\text{DEC}(c)$ assert $ m = \lambda$ assert $c \notin S \cup \{C(x)\}$ $m \leftarrow_{\$} \text{dec}(k, c)$ return m	$\mathcal{C}^\mathcal{O}(\tilde{C}, d, \text{lkg}_C)$ <hr style="border: 0.5px solid black;"/> $k_2 \leftarrow \text{lkg}_C$ $\text{msg}, \text{st} \leftarrow_{\$} \mathcal{A}^\mathcal{O}(1^\lambda, k_2, \tilde{C}, d)$ $x \leftarrow \text{msg}$ return (x, st) <hr style="border: 0.5px solid black;"/> $\mathcal{C}^{\mathcal{O}_{C(x)}}(\tilde{x}, \text{st})$ <hr style="border: 0.5px solid black;"/> $b' \leftarrow_{\$} \mathcal{A}^{\mathcal{O}_{C(x)}}(\tilde{x}, \text{st})$ return b'
---	---	--

Fig. 8: Sampler Sam that emulates the sampling of $\text{Denc}^0(1^\lambda)$ and a distinguisher \mathcal{C} which just runs \mathcal{A} .

multiple times in parallel (with different, adversarially chosen inputs), we need to prove a hybrid argument which replaces real garbled circuits by simulated garbled circuits successively. When reducing one step in the hybrid argument to DSIM security, the other garbled circuits become leakage in the DSIM game. However, we need to prove that this additional leakage does not jeopardize security. Unfortunately, proving this property about the leakage requires another reduction to DSIM security. In order to avoid circular reasoning, new techniques are necessary to establish self-composability results for DSIM.

Theorem 3. *If se is an authenticated encryption scheme in the presence of linear related-key attacks ($\$$ -RK-AE-secure) and Gb is $\text{DSim}[\mathcal{C}^{\text{Out}}]$ -secure, then DSE is DENC -secure.*

Remark. We write $\text{DSim}[\mathcal{C}^{\text{Out}}]$ to denote $\text{DSim}_{\Lambda, \text{Filter}}[\mathcal{C}^{\text{Out}}]$ where Λ is empty and Filter is identity. That is, the parameters Λ and Filter are not needed for this section.

Proof. The proof of Theorem 3 proceeds via two high-level game-hops, see Fig. 7 (right) for the hybrid game between $\text{DENC}_{\mathcal{A}}^0(1^\lambda)$ and $\text{DENC}_{\mathcal{A}, \mathcal{S}}^1(1^\lambda)$. The first game hop from $\text{DENC}_{\mathcal{A}}^0(1^\lambda)$ to $\text{Hybrid}_{\mathcal{A}}(1^\lambda)$ reduces t to $\text{DSim}[\mathcal{C}^{\text{Out}}]$ -security, and the second game-hop from $\text{Hybrid}_{\mathcal{A}}(1^\lambda)$ to $\text{DENC}_{\mathcal{A}, \mathcal{S}}^1(1^\lambda)$ reduces to RK-AE security. We first provide the reduction for the 2nd game-hop, since it is easier than the first game-hop.

$\text{Hybrid}_{\mathcal{A}}(1^\lambda)$ to $\text{DENC}_{\mathcal{A}, \mathcal{S}}^1(1^\lambda)$: For any PPT simulator $\mathcal{S}_{\text{hybrid}}$, we define a PPT simulator \mathcal{S} as follows: The simulator \mathcal{S} runs $\mathcal{S}_{\text{hybrid}}$, but answers its ENC queries with random strings of length 2λ and its DEC queries with \perp . Now, to reduce the

indistinguishability of $\text{Hybrid}_{\mathcal{A}}(1^\lambda)$ and $\text{DENC}_{\mathcal{A},\mathcal{S}}^1(1^\lambda)$ to $\$$ -RK-AE, observe that k_1 is not used and thus, k_1 is perfectly random as required. Now, assume toward contradiction that there exists a pair of PPT algorithms $(\mathcal{A}, \mathcal{S}_{\text{hybrid}})$ such that the difference between $\Pr[1 = \text{Hybrid}_{\mathcal{A}}(1^\lambda)]$ and $\Pr[1 = \text{DENC}_{\mathcal{A},\mathcal{S}}^1(1^\lambda)]$ is non-negligible, where \mathcal{S} is derived from $\mathcal{S}_{\text{hybrid}}$ as previously described. Then, we can construct the following adversary \mathcal{B} against $\$$ -RK-AE: \mathcal{B} emulates $\text{Hybrid}_{\mathcal{A}}(1^\lambda)$, except for all ENC and DEC queries which it forwards to its $\$$ -RK-AE-game (with Δ being the all-zeroes string). \mathcal{B} outputs whatever \mathcal{A} outputs. By construction, we have that

$$\Pr[1 = \mathcal{A}(1^\lambda) \xrightarrow{\text{ENC,DEC}} \$\text{-RK-AE}^0(1^\lambda)] = \Pr[1 = \text{Hybrid}_{\mathcal{A}}(1^\lambda)].$$

Additionally, we claim that

$$\Pr[1 = \mathcal{A}(1^\lambda) \xrightarrow{\text{ENC,DEC}} \$\text{-RK-AE}^1(1^\lambda)] = \Pr[1 = \text{DENC}_{\mathcal{A},\mathcal{S}}^1(1^\lambda)].$$

Namely, in both cases, the ENC queries of \mathcal{A} are answered by random strings of length 2λ , and the DEC queries are answered by \perp , and the game is *stateless*, as there is no set S of previously obtained ciphertexts. Therefore, it does not matter whether queries to ENC and DEC are forwarded to the $\$$ -RK-AE¹ game (as done by \mathcal{B}) or simulated as random answers and \perp answers, respectively, without forwarding (as done by \mathcal{S}). Thus, we can conclude $\text{Adv}_{\text{se},\mathcal{A}}^{\$ \text{-RK-AE}}(1^\lambda)$ is non-negligible and reach a contradiction.

$\text{DENC}_{\mathcal{A}}^0(1^\lambda)$ to $\text{Hybrid}_{\mathcal{A}}(1^\lambda)$: Assume towards contradiction that \mathcal{A} is a PPT adversary which has non-negligible advantage in distinguishing between $\text{DENC}_{\mathcal{A}}^0(1^\lambda)$ and $\text{Hybrid}_{\mathcal{A},\mathcal{S}_{\text{hybrid}}}(1^\lambda)$, regardless of how we instantiate $\mathcal{S}_{\text{hybrid}}$. Now, we construct an adversary $(\text{Sam}(1^\lambda), \mathcal{C})$ against $\text{DSim}[\mathcal{C}^{\text{Out}}]$ -security. $(\text{Sam}(1^\lambda), \mathcal{C})$ is shown in Fig. 8. We first prove that Sam is output indistinguishable.

Claim 1. *If se is an $\$$ -RK-AE, then $\text{Sam} \in \mathcal{C}^{\text{Out}}$.*

Written with explicit randomness $r := k_1 || k_2 || r_1$ of length $p(\lambda) := 3\lambda$, circuit $C_\lambda(\cdot, r) := \text{enc}(k_1 \oplus \cdot, \cdot; r_1 \oplus \cdot)$ and $\text{Sam}_{\text{leak}}(r)$ which returns $\text{lk}_{C_r} \leftarrow k_2$ and $\mathcal{O} \leftarrow \text{ENC, DEC}$, we can re-write Sam as in Fig. 9, as required to prove output indistinguishability. We now prove that for all PPT \mathcal{D} ,

$\text{Sam}(1^\lambda)$	$\text{Out}_{\text{Sam},\mathcal{D}}^b(1^\lambda)$
$r \leftarrow_{\$} \{0,1\}^{p(\lambda)}$	$C_r, \text{lk}_{C_r}, \mathcal{O} \leftarrow_{\$} \text{Sam}(1^\lambda)$
$C_r \leftarrow C_\lambda(\cdot, r)$	$(x_0, \text{st}) \leftarrow_{\$} \mathcal{D}^{\mathcal{O}}(1^\lambda, \text{lk}_{C_r})$
$(\text{lk}_{C_r}, \mathcal{O}) \leftarrow \text{Sam}_{\text{leak}}(r)$	$x_1 \leftarrow \mathcal{O}^{ x_0 }$
return $C_r, \text{lk}_{C_r}, \mathcal{O}$	$b^* \leftarrow_{\$} \mathcal{D}^{\mathcal{O}_{C_r(x_b)}}(C_r(x_b), \text{st})$
	return b^*

Fig. 9: Rewritten version of Sam .

$$|\Pr[1 = \text{Out}_{\text{Sam},\mathcal{D}}^0(1^\lambda)] - \Pr[1 = \text{Out}_{\text{Sam},\mathcal{D}}^1(1^\lambda)]|$$

is negligible, via a sequence of game hops shown in Fig. 10.

We inline the code of Sam , move computations downwards when variables are not used before, and inline the code of the circuit C_r . We also write $S \leftarrow S \cup \{c^*\}$

$\text{Out}_{\text{Sam}, \mathcal{D}}^b(1^\lambda)$	$\text{Out}_{\text{Sam}, \mathcal{D}}^b(1^\lambda)$	$\text{ENC}(m)$
$k_1 \leftarrow_{\$} \{0, 1\}^\lambda$	$k_1 \leftarrow_{\$} \{0, 1\}^\lambda$	assert $ m = \lambda$
$k_2 \leftarrow_{\$} \{0, 1\}^\lambda$	$k_2 \leftarrow_{\$} \{0, 1\}^\lambda$	$c \leftarrow_{\$} \text{enc}(k, m)$
$k \leftarrow k_1 \oplus k_2$	$k \leftarrow_{\$} \{0, 1\}^\lambda$	$S \leftarrow S \cup \{c\}$
$r_1 \leftarrow_{\$} \{0, 1\}^\lambda$		return c
$C \leftarrow \text{enc}(k_1 \oplus \cdot, \cdot; r_1 \oplus \cdot)$		
$(x_0, \text{st}) \leftarrow_{\$} \mathcal{D}^{\text{ENC}, \text{DEC}}(1^\lambda, k_2)$	$(x_0, \text{st}) \leftarrow_{\$} \mathcal{D}^{\text{ENC}, \text{DEC}}(1^\lambda, k_2)$	$\text{DEC}(c)$
$x_1 \leftarrow 0^{ x_0 }$	$x_1 \leftarrow 0^{ x_0 }$	assert $ m = \lambda$
	$\text{msg} \ \Delta\ r' \leftarrow x_b$	assert $c \notin S$
$c^* \leftarrow C(x_b)$	$c^* \leftarrow_{\$} \text{enc}(k_1 \oplus \Delta, \text{msg})$	$m \leftarrow_{\$} \text{dec}(k, c)$
$S \leftarrow S \cup \{c^*\}$	$S \leftarrow S \cup \{c^*\}$	return m
$b^* \leftarrow_{\$} \mathcal{D}^{\text{ENC}, \text{DEC}}(c, \text{st})$	$b^* \leftarrow_{\$} \mathcal{D}^{\text{ENC}, \text{DEC}}(c, \text{st})$	
return b^*	return b^*	

Fig. 10: Game hops to show Claim 1.

as an explicit state update instead of writing $c^* = C(x_b)$ as oracle subscript. The encryption process in the grey line uses fresh and uniform randomness since r_1 is uniformly random and so is r_1 xored with an adversarially chosen value.

We now reduce to $\$$ -RK-AE security (Definition 1) in order to replace the ENC oracle by one that returns random strings of length 2λ , the string c^* by a random string of length 2λ and the DEC oracle by an oracle that always returns \perp . After the reduction to $\$$ -RK-AE security, since c^* does not depend on b anymore, we have that $\Pr[b = b^*] = \frac{1}{2}$.

The reduction \mathcal{B} to $\$$ -RK-AE security can answer all ENC queries m by a query $(m, \Delta = 0^\lambda)$ to its own encryption oracle, and it can compute the ciphertext c by choosing $\Delta' := \Delta \oplus k_2$, since then,

$$k \oplus \Delta' = (k_1 \oplus k_2) \oplus (\Delta \oplus k_2) = k_1 \oplus \Delta,$$

which \mathcal{B} expects. Analogously, the reduction proceeds with decryption queries. This concludes the proof of Claim 1 that $\text{Sam} \in \mathcal{C}^{\text{Out}}$.

Hence, by $\text{DSim}[\mathcal{C}^{\text{Out}}]$ -security, there exists a simulator $\mathcal{S}_{\text{DSIM}}$ for $(\text{Sam}, \mathcal{B})$ such that

$$|\Pr[1 = \text{DSim}_{\text{Sam}, \mathcal{B}, \text{Gb}, \mathcal{A}}^0(1^\lambda)] - \Pr[1 = \text{DSim}_{\text{Sam}, \mathcal{B}, \mathcal{S}_{\text{DSIM}}, \mathcal{A}}^1(1^\lambda)]|$$

is negligible. Given $\mathcal{S}_{\text{DSIM}}$, the simulator $\mathcal{S}_{\text{hybrid}}$ with oracle access to ENC and DEC runs

$$(\tilde{C}, d, \text{st}_S) \leftarrow_{\$} \mathcal{S}_{\text{DSIM}}^{\text{ENC}, \text{DEC}}(\text{lk}_{G_C}) \text{ and } \tilde{x} \leftarrow_{\$} \mathcal{S}_{\text{DSIM}}^{\text{ENC}, \text{DEC}}(\tilde{\cdot}, \text{st}_S)$$

and returns $(\tilde{C}, d, \tilde{x})$.

$$\begin{aligned} & \text{Since } \left| \Pr[1 = \text{DSim}_{\text{Sam}, \mathcal{C}, \text{Gb}}^0(1^\lambda)] - \Pr[1 = \text{DSim}_{\text{Sam}, \mathcal{C}, \mathcal{S}_{\text{DSIM}}}^1(1^\lambda)] \right| \\ & = \left| \Pr[1 = \text{DENC}_{\mathcal{A}}^0(1^\lambda)] - \Pr[1 = \text{Hybrid}_{\mathcal{A}, \mathcal{S}_{\text{hybrid}}}(1^\lambda)] \right| \end{aligned}$$

and the former is negligible and the latter non-negligible, we reached a contradiction. \square

6 Bootstrapping for output indistinguishable samplers

In this section, we bootstrap DSIM security for output indistinguishable samplers returning NC^0 circuits to output indistinguishable samplers returning arbitrary polynomial-size circuits. Let us denote the class of output indistinguishable samplers by \mathcal{C}^{Out} and output indistinguishable samplers which only return circuits in NC^0 by $\mathcal{C}^{\text{Out}, \text{NC}^0} \subseteq \mathcal{C}^{\text{Out}}$. Based on a $\text{DSim}_{\mathcal{A}, \text{Filter}}[\mathcal{C}^{\text{Out}, \text{NC}^0}]$ -secure garbling scheme (for NC^0 circuits), we construct a $\text{DSim}_{\mathcal{A}, \text{Filter}}[\mathcal{C}^{\text{Out}}]$ -secure garbling scheme (for arbitrary poly-size circuits): The new garbling scheme $\text{Gb}_{\text{comb}} = (\text{GCircuit}_{\text{comb}}, \text{GInput}_{\text{comb}}, \text{GEval}_{\text{comb}})$ which we construct (cf. Fig. 11) combines two garbling schemes, the *inner* SelSim-secure garbling scheme $\text{Gb}_{\text{in}} = (\text{GCircuit}_{\text{in}}, \text{GInput}_{\text{in}}, \text{GEval}_{\text{in}})$ for arbitrary polynomial-size circuits and an *outer* garbling scheme $\text{Gb}_{\text{outer}} = (\text{GCircuit}_{\text{outer}}, \text{GInput}_{\text{outer}}, \text{GEval}_{\text{outer}})$ for NC^0 circuits which is $\text{DSim}_{\mathcal{A}, \text{Filter}}[\mathcal{C}^{\text{Out}, \text{NC}^0}]$ -secure, to obtain a *combined* garbling scheme Gb_{comb} which is $\text{DSim}_{\mathcal{A}, \text{Filter}}[\mathcal{C}^{\text{Out}}]$ -secure. Gb_{comb} garbles a circuit $C(\cdot)$ and input x as depicted in Figure 11, where C_{in} is a circuit that takes as input x and produces as output the (selectively secure) garbling of C , i.e., \tilde{C}_{in} , a garbling of x and the decoding information d_{in} . Gb_{in} could be any SelSim-secure, projective⁶ garbling scheme, e.g., $\text{Gb}_{\text{in}} = (\text{GCircuit}_{\text{in}}, \text{GInput}_{\text{in}}, \text{GEval}_{\text{in}}) := (\text{GCircuit}_{\text{yao}}, \text{GInput}_{\text{yao}}, \text{GEval}_{\text{yao}})$ is a valid choice (Yao's garbling is provably SelSim secure, assuming only the existence of $\$$ -IND-CPA secure symmetric encryption scheme). We need to show that the circuit C_{in} is indeed in NC^0 and that Gb_{comb} is $\text{DSim}_{\mathcal{A}, \text{Filter}}[\mathcal{C}^{\text{Out}}]$ -secure.

Section 6.1 shows the (easy) statement that the circuit C_{in} is in NC^0 . In a nutshell, this follows from Gb_{in} being a projective garbling scheme and \tilde{C}_{in} and the decoding information d_{in} being just some constant bitstring in C_{in} .

Section 6.2 proves that if the sampler for circuit $C(x)$ is in \mathcal{C}^{Out} , then the sampler for circuit $\tilde{C}(x; r)$ is in $\mathcal{C}^{\text{Out}, \text{NC}^0}$. Section 6.3 then states and proves our main bootstrapping theorem.

6.1 C_{in} is low-depth

Recall that C_{in} is defined as $C_{\text{in}}(\cdot) := [\tilde{C}_{\text{in}}, \text{GInput}_{\text{in}}(\cdot, K_{\text{in}}), d_{\text{in}}]$ where \tilde{C}_{in} , d_{in} and K_{in} are constant values (constant bitstrings). Hence, in order to show that

⁶ A garbling scheme is projective, if for each input bit x_i , the input garbling is one out of two possible strings $K_0(i)$ and $K_1(i)$. For example, Yao's garbling scheme is projective.

$\mathbf{GCircuit}_{\text{comb}}(C)$	$\mathbf{GInput}_{\text{comb}}(K, x)$	$\mathbf{GEval}_{\text{comb}}(\tilde{C}, \tilde{x}, d)$
$r_{\text{in}} \leftarrow_{\$} \{0, 1\}^{12\lambda C }$	$\tilde{x} \leftarrow \mathbf{GInput}_{\text{outer}}(K, x)$	$\tilde{C}_{\text{in}}, \tilde{x}_{\text{in}}, d_{\text{in}} \leftarrow \mathbf{GEval}_{\text{outer}}(\tilde{C}, \tilde{x}, d)$
$\tilde{C}_{\text{in}}, K_{\text{in}}, d_{\text{in}} \leftarrow \mathbf{GCircuit}_{\text{in}}(C(\cdot); r_{\text{in}})$	return \tilde{x}	$y \leftarrow \mathbf{GEval}_{\text{in}}(\tilde{C}_{\text{in}}, \tilde{x}_{\text{in}}, d_{\text{in}})$
$C_{\text{in}}(\cdot) \leftarrow [\tilde{C}_{\text{in}}, \mathbf{GInput}_{\text{in}}(\cdot, K_{\text{in}}), d_{\text{in}}]$		return y
$\tilde{C}, K, d \leftarrow_{\$} \mathbf{GCircuit}_{\text{outer}}(C_{\text{in}})$		
return \tilde{C}, K, d		

Fig. 11: Garbling scheme $\mathbf{Gb}_{\text{comb}}$. Length of r_{in} is chosen s.t. it is compatible with Yao's garbling scheme, however, the same results apply to any projective garbling scheme, just $|r_{\text{in}}|$ might need to be adjusted.

the function C_{in} can be implemented as a constant depth circuit, it is enough to show the following lemma.

Lemma 1 (Low-depth Projective Input Garbling). *If \mathbf{Gb}_{in} is a projective garbling scheme, then $\mathbf{GInput}_{\text{in}}(\cdot; K_{\text{in}}) : x \mapsto \mathbf{GInput}_{\text{in}}(x; K_{\text{in}})$ can be described by a constant-depth circuit.*

Proof. Denote by x_i the i th bit of the input x . Now K_{in} consists of key pairs k_0^i, k_1^i for each bit of x . W.l.o.g., we can assume that K_{in} is a concatenation of all the key pairs in order. The input garbling of x , i.e. $\mathbf{GInput}_{\text{in}}(x; K_{\text{in}})$, outputs $k_{x_1}^1 || \dots || k_{x_\lambda}^\lambda$. This can clearly be done in constant depth by a circuit that just checks each bit of x one by one (in parallel) and outputs the corresponding key $k_{x_i}^i$ for each bit. \square

6.2 Output Indistinguishable Sampling

We now prove that the above circuit transformation, when applied to a circuit sampler $\mathbf{Sam}_{\text{comb}}(1^\lambda)$ in \mathcal{C}^{Out} yields a circuit sampler $\mathbf{Sam}_{\text{outer}}(1^\lambda)$ in $\mathcal{C}^{\text{Out}, \text{NC}^0}$. Since $\mathbf{Sam}_{\text{comb}}$ is output indistinguishable, there exists a polynomial $p(\lambda)$ and a circuit $C = (C_\lambda)_{\lambda \in \mathbb{N}}$ such that $\mathbf{Sam}_{\text{comb}}(1^\lambda)$ can be written as below (left). Then, we define the circuit sampler $\mathbf{Sam}_{\text{outer}}(1^\lambda)$ (with randomness $r || r_{\text{in}}$) as follows (right):

$\mathbf{Sam}_{\text{comb}}(1^\lambda)$	$\mathbf{Sam}_{\text{outer}}(1^\lambda)$
$r \leftarrow_{\$} \{0, 1\}^{p(\lambda)}$	$r \leftarrow_{\$} \{0, 1\}^{p(\lambda)}$
$C_r(\cdot) \leftarrow C(\cdot, r)$	$r_{\text{in}} \leftarrow_{\$} \{0, 1\}^{12\lambda C_r }$
return $C_r, \text{lk}_{C_r}, \mathcal{O}$	$\tilde{C}_{\text{in}}, K_{\text{in}}, d_{\text{in}} \leftarrow \mathbf{GCircuit}_{\text{in}}(C_r(\cdot); r_{\text{in}})$
	$C_{\text{in}}(\cdot) \leftarrow [\tilde{C}_{\text{in}}, \mathbf{GInput}_{\text{in}}(\cdot, K_{\text{in}}), d_{\text{in}}]$
	return $C_{\text{in}}, \text{lk}_{C_r}, \mathcal{O}$

Lemma 2 (Output indistinguishability). *Let $\mathbf{Sam}_{\text{comb}} \in \mathcal{C}^{\text{Out}}$. If Yao's garbling scheme is selectively secure, then $\mathbf{Sam}_{\text{outer}}(1^\lambda) \in \mathcal{C}^{\text{Out}, \text{NC}^0}$.*

Proof. Firstly, by Lemma 1, $\text{Sam}_{\text{outer}}$ produces circuits in NC^0 . Hence remains to show that $\text{Sam}_{\text{outer}}$ is output indistinguishable, i.e., we prove that for all PPT adversaries $\mathcal{D}_{\text{outer}}$,

$$|\Pr[1 = \text{Out}_{\text{Sam}_{\text{outer}}, \mathcal{D}_{\text{outer}}}^0(1^\lambda)] - \Pr[1 = \text{Out}_{\text{Sam}_{\text{outer}}, \mathcal{D}_{\text{outer}}}^1(1^\lambda)]| = \text{negl}(\lambda).$$

We prove indistinguishability of $\text{Out}_{\mathcal{D}_{\text{outer}}}^0(1^\lambda)$ and $\text{Out}_{\mathcal{D}_{\text{outer}}}^1(1^\lambda)$ via several hybrids. The games $\text{Out}_{\mathcal{D}_{\text{outer}}}^b(1^\lambda)$ and H^b are perfectly equivalent; we only inline the definition of circuit $\text{Sam}_{\text{outer}}$ and reorder lines whose order does not affect the functionality.

From $H^b(1^\lambda)$ to H_S^{b+2} , we reduce to the selective simulation-based security of Gb_{in} . The highlighted lines in H^b refer to the SelSim adversary (who first chooses the circuit, the input and its own state C_r, x_b, st , then receives a garbled circuit and garbled input and returns its guess). Note that the SelSim adversary knows r and can hence compute the algorithm \mathcal{O} and pass it in st , so $\mathcal{D}_{\text{outer}}^{\mathcal{O}}$ can be replaced by a code equivalent and efficient algorithm that does not make oracle queries to \mathcal{O} (but simply runs the algorithm \mathcal{O}).

From H_S^{0+2} to H_S^{1+2} , we reduce to the output indistinguishability of Sam_{comb} , note that the two first lines of H_S^{b+2} are code equivalent to Sam_{comb} and the highlighted lines refer to the output indistinguishability adversary (note that the topology $\Phi(C_r)$ can be computed when you know C and C is constant). For this game hop we switch back to \mathcal{O} being an oracle algorithm (as opposed to emulating $\mathcal{D}_{\text{outer}}^{\mathcal{O}}$ by running the actual algorithm \mathcal{O}).

H^b	H_S^{b+2}
$r \leftarrow_{\$} \{0, 1\}^{p(\lambda)}$	$r \leftarrow_{\$} \{0, 1\}^{p(\lambda)}$
$C_r, \text{lk}_{C_r} \leftarrow C(\cdot, r), \text{lk}_{C_r}$	$C_r, \text{lk}_{C_r} \leftarrow C(\cdot, r), \text{lk}_{C_r}$
$(x_0, \text{st}) \leftarrow_{\$} \mathcal{D}_{\text{outer}}^{\mathcal{O}}(1^\lambda, \text{lk}_{C_r})$	$(x_0, \text{st}) \leftarrow_{\$} \mathcal{D}_{\text{outer}}^{\mathcal{O}}(1^\lambda, \text{lk}_{C_r})$
$x_1 \leftarrow 0^{ x_0 }$	$x_1 \leftarrow 0^{ x_0 }$
$r_{\text{in}} \leftarrow_{\$} \{0, 1\}^{12\lambda C_r }$	
$\tilde{C}_{\text{in}}, K_{\text{in}}, d_{\text{in}} \leftarrow \text{GCircuit}_{\text{in}}(C_r(\cdot); r_{\text{in}})$	
$C_{\text{in}}(\cdot) \leftarrow [\tilde{C}_{\text{in}}, \text{GInput}_{\text{in}}(\cdot, K_{\text{in}}), d_{\text{in}}]$	
$(\tilde{C}, \tilde{x}, d) \leftarrow C_{\text{in}}(x_b)$	$(\tilde{C}, \tilde{x}, d) \leftarrow \mathcal{S}(\Phi(C_r), C_r(x_b))$
$b^* \leftarrow_{\$} \mathcal{D}_{\text{outer}}^{\mathcal{O}_{C_r(x_b)}}((\tilde{C}, \tilde{x}, d), \text{st})$	$b^* \leftarrow_{\$} \mathcal{D}_{\text{outer}}^{\mathcal{O}_{C_r(x_b)}}((\tilde{C}, \tilde{x}, d), \text{st})$
return b^*	return b^*

□

6.3 Main theorem

Our main bootstrapping theorem now follows from the above lemmas.

Theorem 4 (Bootstrapping RIND from NC^0 to poly). *If the garbling scheme $\text{Gb}_{\text{outer}} = (\text{GCircuit}_{\text{outer}}, \text{GInput}_{\text{outer}}, \text{GEval}_{\text{outer}})$ is $\text{DSim}_{\Lambda, \text{Filter}}[\mathcal{C}^{\text{Out}}, \text{NC}^0]$ -secure and if $\text{Gb}_{\text{in}} = (\text{GCircuit}_{\text{in}}, \text{GInput}_{\text{in}}, \text{GEval}_{\text{in}})$ is a projective SelSim -secure garbling scheme, then the garbling scheme $\text{Gb}_{\text{comb}} = (\text{GCircuit}_{\text{comb}}, \text{GInput}_{\text{comb}}, \text{GEval}_{\text{comb}})$ achieves $\text{DSim}_{\Lambda, \text{Filter}}[\mathcal{C}^{\text{Out}}]$ security.*

Proof. For all PPT $\text{Sam}_{\text{comb}} \in \text{DSim}_{\Lambda, \text{Filter}}[\mathcal{C}^{\text{Out}}]$ and for all PPT \mathcal{A} , we want to construct a simulator $\mathcal{S}_{\text{comb}}$ such that the following advantage

$$|\Pr[1 = \text{DSim}_{\text{Sam}_{\text{comb}}, \mathcal{A}, \text{Gb}_{\text{comb}}, \Lambda}^0(1^\lambda)] - \Pr[1 = \text{DSim}_{\text{Sam}_{\text{comb}}, \mathcal{A}, \mathcal{S}_{\text{comb}}, \Lambda}^1(1^\lambda)]| \quad (1)$$

is negligible.

$\text{DSim}_{\text{Sam}_{\text{comb}}, \mathcal{A}, \text{Gb}_{\text{comb}}, \Lambda}^0(1^\lambda)$	H^0
$C_r, \text{lk}_{C_r}, \mathcal{O} \leftarrow_{\$} \text{Sam}_{\text{comb}}(1^\lambda)$	$r \leftarrow_{\$} \{0, 1\}^{p(\lambda)}$
$(\tilde{C}, K, d) \leftarrow_{\$} \text{GCircuit}_{\text{comb}}(C_r)$	$C_r(\cdot) \leftarrow C(\cdot, r)$
$(x, \text{st}_{\mathcal{A}}) \leftarrow_{\$} \mathcal{A}^{\mathcal{O}}(\tilde{C}, d, \text{lk}_{C_r})$	$r_{\text{in}} \leftarrow_{\$} \{0, 1\}^{12\lambda \mathcal{C} }$
$\tilde{x} \leftarrow_{\$} \text{GInput}_{\text{comb}}(K, x)$	$\tilde{C}_{\text{in}}, K_{\text{in}}, d_{\text{in}} \leftarrow \text{GCircuit}_{\text{in}}(C_r(\cdot); r_{\text{in}})$
$b' \leftarrow_{\$} \mathcal{A}^{\mathcal{O}_{C_r(x)}}(\tilde{x}, \text{st}_{\mathcal{A}})$	$C_{\text{in}}(\cdot) \leftarrow [\tilde{C}_{\text{in}}, \text{GInput}_{\text{in}}(\cdot, K_{\text{in}}), d_{\text{in}}]$
return b'	$\tilde{C}, K, d \leftarrow_{\$} \text{GCircuit}_{\text{outer}}(C_{\text{in}})$
	$(x, \text{st}_{\mathcal{A}}) \leftarrow_{\$} \mathcal{A}^{\mathcal{O}}(\tilde{C}, d, \text{lk}_{C_r})$
	$\tilde{x} \leftarrow_{\$} \text{GInput}_{\text{outer}}(K, x)$
	$b' \leftarrow_{\$} \mathcal{A}^{\mathcal{O}_{\tilde{C}_r(\tilde{x})}}(\tilde{x}, \text{st}_{\mathcal{A}})$
	return b'
$\text{DSim}_{\text{Sam}_{\text{comb}}, \mathcal{A}, \mathcal{S}_{\text{comb}}, \Lambda}^1(1^\lambda)$	H^1
$C_r, \text{lk}_{C_r}, \mathcal{O} \leftarrow_{\$} \text{Sam}_{\text{comb}}(1^\lambda)$	$r \leftarrow_{\$} \{0, 1\}^{p(\lambda)}$
$(\tilde{C}, d, \text{st}_S) \leftarrow_{\$} \mathcal{S}^{\mathcal{O}}(\text{Filter}(\text{lk}_{C_r}))$	$C_r(\cdot) \leftarrow C(\cdot, r)$
$(x, \text{st}_{\mathcal{A}}) \leftarrow_{\$} \mathcal{A}^{\mathcal{O}}(\tilde{C}, d, \text{lk}_{C_r})$	$r_{\text{in}} \leftarrow_{\$} \{0, 1\}^{12\lambda \mathcal{C} }$
$\tilde{x} \leftarrow_{\$} \mathcal{S}^{\mathcal{O}}(\Lambda(\text{st}_{\mathcal{A}}), \text{st}_S)$	$\tilde{C}_{\text{in}}, K_{\text{in}}, d_{\text{in}} \leftarrow \text{GCircuit}_{\text{in}}(C_r(\cdot); r_{\text{in}})$
$b' \leftarrow_{\$} \mathcal{A}^{\mathcal{O}_{\tilde{C}_r(\tilde{x})}}(\tilde{x}, \text{st}_{\mathcal{A}})$	$C_{\text{in}}(\cdot) \leftarrow [\tilde{C}_{\text{in}}, \text{GInput}_{\text{in}}(\cdot, K_{\text{in}}), d_{\text{in}}]$
return b'	$\tilde{C}, d, \text{st}_S \leftarrow_{\$} \mathcal{S}_{\text{outer}}^{\mathcal{O}}(\text{Filter}(\text{lk}_{C_r}))$
	$(x, \text{st}_{\mathcal{A}}) \leftarrow_{\$} \mathcal{A}^{\mathcal{O}}(\tilde{C}, d, \text{lk}_{C_r})$
	$\tilde{x} \leftarrow_{\$} \mathcal{S}_{\text{outer}}^{\mathcal{O}}(\Lambda(\text{st}_{\mathcal{A}}), \text{st}_S)$
	$b' \leftarrow_{\$} \mathcal{A}^{\mathcal{O}_{\tilde{C}_r(\tilde{x})}}(\tilde{x}, \text{st}_{\mathcal{A}})$
	return b'

The experiments $\text{DSim}_{\text{Sam}_{\text{comb}}, \mathcal{A}, \text{Gb}_{\text{comb}}, \Lambda}^0(1^\lambda)$ and the hybrid game H^0 are code equivalent by inlining the code of Sam_{comb} , $\text{GCircuit}_{\text{comb}}$ and $\text{GInput}_{\text{comb}}$. Now, observe that the lines highlighted in grey in H^0 describe $\text{Sam}_{\text{outer}}(1^\lambda)$ and thus, H^0 is also equivalent to $\text{DSim}_{\text{Sam}_{\text{outer}}, \mathcal{A}, \text{Gb}_{\text{outer}}, \Lambda}^0(1^\lambda)$, as required for the reduction.

Analogously, the hybrid game H^1 is equivalent to $\text{DSim}_{\text{Sam}_{\text{outer}}, \mathcal{A}, \text{Gb}_{\text{outer}}, \Lambda}^1(1^\lambda)$ with $\text{Sam}_{\text{outer}} \in \mathcal{C}^{\text{Out}, \text{NC}^0}$ (Lemma 2) and thus, by $\text{DSim}_{\Lambda, \text{Filter}}[\mathcal{C}^{\text{Out}, \text{NC}^0}]$ -security of Gb_{outer} (assumption), there is a simulator $\mathcal{S}_{\text{outer}}$ s.t. H^0 is indistinguishable from H^1 . If we define simulator \mathcal{S} according to grey lines in H^1 , we notice that H^1 is also code equivalent to $\text{DSim}_{\text{Sam}_{\text{comb}}, \mathcal{A}, \text{Gb}_{\text{comb}}, \Lambda}^1(1^\lambda)$, which concludes the proof. \square

7 Large Yao-Incompressibility Entropy Implies Large Online Complexity

Conceptually, the AIKW lower bounds uses that $f(x)$ has real entropy (at most $|x|$), but if f is a PRG, then $f(x)$ has (HILL) *pseudo-entropy* $|f(x)|$, since $f(x)$ is indistinguishable from a uniformly random string of length $|f(x)|$. AIKW conclude that the input garbling size must be proportional to $|f(x)|$ rather than to $|x|$. Using an analogous argument, HW directly formulate their lower bound for online complexity of maliciously secure MPC in terms of pseudo-entropy, but rely on (Yao) incompressibility instead which is measured by the shortest efficient encoding of $f(x)$. High HILL pseudoentropy of $f(x)$ implies high incompressibility entropy of $f(x)$, since real entropy distribution \mathcal{D} cannot be compressed, but the converse is not necessarily true [HLR07,HMS23], so the use of incompressibility, conceptually, yields a more general result. HW imply lower bounds also on SIM-secure garbling. In this section, we use the HW incompressibility entropy approach to bound the online complexity of a SIM-secure garbling scheme directly to obtain a self-contained proof and slightly improve HW's bound on garbling, cf. Section 1.1.

Remark. Yao incompressibility has been defined in different variations, e.g., whether the definition considers *expected* ([TVZ05]) or *worst-case* (HW, [BSW03]) compression length and whether it uses *perfect* correctness ([TVZ05]) or *imperfect* correctness (HW, [BSW03]). We now state the Yao incompressibility definitions of [TVZ05] and HW, the latter of which simplifies the definition by [BSW03] as is enough in this context.

Definition 10 (Yao Incompressibility Entropy by TVZ [TVZ05]). *Let Φ be a leakage function and $(C, x) \leftarrow_{\mathcal{S}} \mathcal{D}(1^\lambda)$ be an efficiently sampleable distribution. The distribution \mathcal{D} is k_λ -incompressible, if for every (non-uniform) polynomial-size circuit family pair of a compression algorithm $\text{Cmpr}_\lambda(\cdot, \cdot)$ and a decompression function $\text{Decmpr}_\lambda(\cdot, \cdot)$ it holds that*

$$\mathbb{E}_{(C, x) \leftarrow_{\mathcal{S}} \mathcal{D}(1^\lambda)}[|\text{Cmpr}(C(x), \Phi(C))|] \geq k_\lambda,$$

where $|\cdot|$ is the output length of Cmpr .

Definition 11 (Yao Incompressibility Entropy by HW [HW15]). Let Φ be a leakage function and $(C, x) \leftarrow_{\mathcal{S}} \mathcal{D}(1^\lambda)$ be an efficiently sampleable distribution. The distribution \mathcal{D} is k_λ -incompressible, if for every (non-uniform) polynomial-size circuit family pair of a compression algorithm $\text{Cmpr}_\lambda(\cdot, \cdot)$ and a decompression function $\text{Decmpr}_\lambda(\cdot, \cdot)$ s.t. for all x : $\text{Decmpr}(\text{Cmpr}(C(x), \Phi(C)), \Phi(C)) = C(x)$ it holds that if output length $l(\lambda)$ of Cmpr is $< k_\lambda$ then there exists a negligible function ϵ s.t.

$$\Pr_{(C,x) \leftarrow_{\mathcal{S}} \mathcal{D}} [\text{Decmpr}(\text{Cmpr}(C(x), \Phi(C)), \Phi(C)) = C(x)] \leq 1/2 + \epsilon(\lambda)$$

For completeness, we prove our theorems both for Yao incompressibility entropy as defined by HW and by Trevisan, Vadhan and Zuckermann (TVZ [TVZ05]) and refer to Karanko [Kar23] for a more comprehensive overview over different notions of incompressibility entropy and pseudo-entropy and clarifying of their relations—especially, high *expected* incompressibility entropy is implied by high pseudo *Shannon* entropy, while high *worst-case* incompressibility entropy is only implied by high pseudo *min* entropy.

Outline. For concreteness, we now revisit the AIKW lower bound which considers $C := \text{PRG}$. Correctness of the garbling scheme implies that $\tilde{C}(\tilde{x}) = \text{PRG}(x)$. Therefore, the simulator who first creates a garbled circuit \tilde{C} and then gets $y = \text{PRG}(x)$ also needs to create \tilde{x} such that $\tilde{C}(\tilde{x}) = y$. Now, the check that $\tilde{C}(\tilde{x})$ evaluates to y can be performed only knowing y , and since the simulator gets only y , we can run the simulator on a *random* y instead of a PRG output—by PRG security, the simulator should be as successful in creating a simulated input \tilde{x} such that $\tilde{C}(\tilde{x}) = y$. However, if $|\tilde{x}| < |y|$ the simulator has an impossible task, because y has more entropy than \tilde{x} . Therefore, $|\tilde{x}|$ should be proportional to the *computational entropy* of $C(x)$. We now prove that both high TVZ and HW incompressibility entropy of $C(x)$ imply high online complexity of a SIM-secure scheme Gb which garbles C and x .

Theorem 5 (TVZ Incompressibility \Rightarrow High Online Complexity). Let Gb be a SIM-secure garbling scheme with leakage function Φ , and let $(C, x) \leftarrow_{\mathcal{S}} \mathcal{D}(1^\lambda)$ be an efficiently sampleable distribution with incompressibility-entropy $\geq k_\lambda$. with incompressibility-entropy $\geq k_\lambda$. If the online complexity $|\tilde{x}|$ is uniquely determined by $\Phi(C)$, then the expected online complexity is

$$\mathbb{E}_{(C,x) \leftarrow_{\mathcal{S}} \mathcal{D}(1^\lambda), \tilde{x} \leftarrow_{\mathcal{S}} \text{Gb.GInput}} [|\tilde{x}|] \geq k_\lambda - 2.$$

Proof. Assume towards contradiction that

$$\mathbb{E}_{(C,x) \leftarrow_{\mathcal{S}} \mathcal{D}(1^\lambda), \tilde{x} \leftarrow_{\mathcal{S}} \text{Gb.GInput}} [|\tilde{x}|] < k_\lambda - 2, \quad (2)$$

Since the garbling scheme Gb is SIM-secure, there exists a PPT simulator \mathcal{S} , s.t.

$$\Pr \left[\begin{array}{l} \text{Gb.GEval}(\tilde{C}, \tilde{x}, d) = C(x) \\ |\tilde{x}| = |\text{Gb.GInput}(x)| \end{array} \middle| \begin{array}{l} (C, x) \leftarrow_{\mathcal{S}} \mathcal{D}(1^\lambda) \\ \tilde{C}, d, \text{st} \leftarrow_{\mathcal{S}} \mathcal{S}(1^\lambda, \Phi(C)) \\ \tilde{x} \leftarrow_{\mathcal{S}} \mathcal{S}(1^\lambda, C(x), \text{st}) \end{array} \right] \geq 1 - \mu(\lambda),$$

where μ is negligible. The garbled input that the simulator returns must (almost always) have the correct online complexity, since otherwise we can use the length of the garbled input to distinguish real garbling from simulated garbling. We now use the simulator \mathcal{S} to build a pair of efficient algorithms compressor **Cmpr** and decompressor **Decmpr**. Let $\mathcal{S}_\lambda^{\text{det}}$ be a *deterministic* version of the simulator, i.e., $\mathcal{S}_\lambda^{\text{det}}$ is equal to the simulator $\mathcal{S}(1^\lambda, \cdot; r_\lambda)$, where r_λ is the internal randomness that maximizes the above probability. We here use that [TVZ05] allow **Cmpr** and decompressor **Decmpr** to be a non-uniform circuit family. Then, we define compressor **Cmpr** and decompressor **Decmpr** as follows, writing $\text{lk}_{\mathbf{g}_\Phi}$ for a supposed output of $\Phi(C)$.

Cmpr ($y, \text{lk}_{\mathbf{g}_\Phi}$)	Decmpr ($b z, \text{lk}_{\mathbf{g}_\Phi}$)
$\tilde{C}, d, \text{st} \leftarrow \mathcal{S}_\lambda^{\text{det}}(\text{lk}_{\mathbf{g}_\Phi})$ // 1st stage of simulator	if $b = 1$
$\tilde{x} \leftarrow \mathcal{S}_\lambda^{\text{det}}(y, \text{st})$ // 2nd stage of simulator	return z
if $\text{GEval}(\tilde{C}, \tilde{x}, d) = y$ AND $ \tilde{x} = \text{Gb.GInput}(x) $	$\tilde{C}, d, \text{st} \leftarrow \mathcal{S}_\lambda^{\text{det}}(\text{lk}_{\mathbf{g}_\Phi})$
return $0 \tilde{x}$	// 1st stage of simulator
return $1 y$	return $\text{GEval}(\tilde{C}, z, d)$
return $\tilde{x}_{1, \dots, k-2}$	

The if-clause in the compressor is false only with negligible probability, because by averaging argument, we can show that there is simulator's randomness r_λ that achieves at least success probability $1 - 2\mu(\lambda)$. Hence,

$$\begin{aligned}
\mathbb{E}_{\mathcal{D}}[|\text{Cmpr}(C(x), \Phi(C))|] &= \sum_{C, x \in \mathcal{D}, \text{if true}} \Pr_{\mathcal{D}}[C, x] \underbrace{|\text{Cmpr}(C(x), \Phi(C))|}_{=|\text{Gb.GInput}(x)|+1} \\
&+ \underbrace{\sum_{C, x \in \mathcal{D}, \text{if false}} \Pr_{\mathcal{D}}[C, x] \underbrace{|\text{Cmpr}(C(x), \Phi(C))|}_{\text{polynomial (WLOG)}}}_{\text{negl}(\lambda)} \\
&\leq \mathbb{E}_{\mathcal{D}}[|\text{Gb.GInput}(x)| + 1] + \text{negl}(\lambda) \\
&= \underbrace{\mathbb{E}_{\mathcal{D}}[|\text{Gb.GInput}(x)|]}_{\leq k-2 \text{ by (2)}} + 1 + \text{negl}(\lambda) < k - 1 + \text{negl}(\lambda)
\end{aligned}$$

which is a contradiction with incompressibility. □

Remark. Note that as opposed to Theorem 5, in Theorem 6, the online complexity is k instead of $k - 2$, but, in turn, we only obtain this lower bound on the online complexity for *some* (C, x) pair rather than *in expectation*.

Theorem 6 (HW Incompressibility \Rightarrow High Online Complexity). *Let Gb be a SIM-secure garbling scheme with leakage function Φ , and let $(C, x) \leftarrow_{\mathcal{S}} \mathcal{D}(1^\lambda)$ be an efficiently samplable distribution with HW incompressibility-entropy*

$\geq k_\lambda$. Then, there exists (at least) one pair in the support of \mathcal{D} whose online complexity is $|\tilde{x}| \geq k_\lambda$. Again, we assume that the online complexity $|\tilde{x}|$ is uniquely determined by $\Phi(C)$.

Proof. The part in gray here is copied in verbatim from proof of Theorem 5. Assume towards contradiction that for all (x, C) pairs in \mathcal{D}

$$|\tilde{x}| < k_\lambda, \quad (3)$$

Since the garbling scheme Gb is SIM-secure, there exists a PPT simulator \mathcal{S} , s.t.

$$\Pr \left[\begin{array}{l} \text{Gb.GEval}(\tilde{C}, \tilde{x}, d) = C(x) \\ |\tilde{x}| = |\text{Gb.GInput}(x)| \end{array} \middle| \begin{array}{l} (C, x) \leftarrow_{\mathfrak{S}} \mathcal{D}(1^\lambda) \\ \tilde{C}, d, \text{st} \leftarrow_{\mathfrak{S}} \mathcal{S}(1^\lambda, \Phi(C)) \\ \tilde{x} \leftarrow_{\mathfrak{S}} \mathcal{S}(1^\lambda, C(x), \text{st}) \end{array} \right] \geq 1 - \mu(\lambda),$$

where μ is negligible. The garbled input that the simulator returns must (almost always) have the correct online complexity, since otherwise we can use the length of the garbled input to distinguish real garbling from simulated garbling. We now use the simulator \mathcal{S} to build a pair of efficient algorithms compressor Cmpr and decompressor Decmpr. Let $\mathcal{S}_\lambda^{\text{det}}$ be a *deterministic* version of the simulator, i.e., $\mathcal{S}_\lambda^{\text{det}}$ is equal to the simulator $\mathcal{S}(1^\lambda, \cdot; r_\lambda)$, where r_λ is the internal randomness that maximizes the above probability. We here use that HW allow Cmpr and decompressor Decmpr to be a non-uniform circuit family. Then, we define compressor Cmpr and decompressor Decmpr as follows, writing lkg_Φ for a supposed output of $\Phi(C)$.

Cmpr (y, lkg)	Decmpr (z, lkg)
$\tilde{C}, d, \text{st} \leftarrow \mathcal{S}_\lambda^{\text{det}}(\text{lkg})$	$\tilde{C}, d, \text{st} \leftarrow \mathcal{S}_\lambda^{\text{det}}(\text{lkg})$
$\tilde{x} \leftarrow \mathcal{S}_\lambda^{\text{det}}(y, \text{st})$	return GEval(\tilde{C}, z, d)
return $\tilde{x}_{1, \dots, k-1}$	

Now

$$\begin{aligned} & \Pr_{C, x \leftarrow_{\mathfrak{S}} \mathcal{D}} [\text{Decmpr}(\text{Cmpr}(C(x), \Phi(C)), \Phi(C)) = C(x)] \\ &= \underbrace{\Pr[\mathcal{S}_\lambda^{\text{det}} \text{ correct}]}_{\geq 1 - 2\mu \text{ by avg. arg.}} \underbrace{\Pr[|\tilde{x}| \leq k - 1 \mid \mathcal{S}_\lambda^{\text{det}} \text{ correct}]}_{=1 \text{ by (3)}} \\ &\geq 1 - \text{negl}(\lambda) \end{aligned}$$

which is a contradiction with HW incompressibility. □

GCircuit_{yao}(C)	GInput_{yao}(K, x)
for $g \in \text{Gates}(C)$ $K_g(0) \leftarrow \{0, 1\}^\lambda$ $K_g(1) \leftarrow \{0, 1\}^\lambda$ for $g \in \text{Gates}(C)$: $K_{\text{left}} \leftarrow K_{\text{LeftPred}}(g, C)$ $K_{\text{right}} \leftarrow K_{\text{RightPred}}(g, C)$ $K_{\text{output}} \leftarrow K_g$ for $(b_{\text{left}}, b_{\text{right}}) \in \{0, 1\}^2$: $d \leftarrow (\text{op}(g))(b_{\text{left}}, b_{\text{right}})$ $c_{\text{in}} \leftarrow \text{enc}(K_{\text{left}}(b_{\text{left}}), K_{\text{out}}(d))$ $c_{\text{out}} \leftarrow \text{enc}(K_{\text{right}}(b_{\text{right}}), c_{\text{in}})$ $\tilde{g} \leftarrow \tilde{g} \cup \{c_{\text{out}}\}$ $\tilde{C}[g] \leftarrow \tilde{g}$ $K \leftarrow \{(K_g, g), g \in \text{InputGates}(C)\}$ $K_{\text{out}} \leftarrow \{(\text{sort}(K_g(0), K_g(1)), g),$ $\quad g \in \text{OutputGates}(C)\}$ for $g \in \text{OutputGates}(C)$: if $K_{\text{out}}(g) = (K_g(0), K_g(1))$ then $\text{out}_g \leftarrow 0$ else $\text{out}_g \leftarrow 1$ $d \leftarrow \{(\text{out}_g, g) : g \in \text{OutputGates}(C)\}$ return $((\tilde{C}, K_{\text{out}}), K, d)$	for $g \in \text{InputGates}(C)$: $d_g \leftarrow g\text{-th pos. of } x$ $\tilde{x}[g] \leftarrow K_g(d_g)$ return \tilde{x} <hr style="border: 0.5px solid black; margin: 5px 0;"/> GEval_{yao}(\tilde{C}', \tilde{x}, d) $(\tilde{C}, K_{\text{out}}) \leftarrow \tilde{C}'$ for $g \in \text{InputGates}(C)$: $k_g \leftarrow \tilde{x}[g]$ for $g \in \text{Gates}(C) \setminus \text{InputGates}(C)$: $k_{\text{left}} \leftarrow k_{\text{LeftPred}}(g, \Phi(C))$ $k_{\text{right}} \leftarrow k_{\text{RightPred}}(g, \Phi(C))$ for $c \in \tilde{C}[g]$: $m_{\text{out}} \leftarrow \text{dec}(k_{\text{right}}, c)$ $m_{\text{in}} \leftarrow \text{dec}(k_{\text{left}}, m_{\text{out}})$ if $m_{\text{in}} \neq \perp$ then $k_g \leftarrow m_{\text{in}}$ for $g \in \text{OutputGates}(C)$: $\text{map}_g \leftarrow d(g)$ if $k_g = K_{\text{out}}(g)(\text{map}_g)$ then $y[g] = 0$ if $k_g = K_{\text{out}}(g)(1 - \text{map}_g)$ then $y[g] = 1$ return y

Fig. 12: Yao's garbling scheme Gb_{yao} with split output map.

8 Garbling NC^1 Circuits with Online Complexity $\mathcal{O}(|x|)$

In this section we construct a garbling scheme with online complexity $2n\lambda$ and prove that it is adaptively indistinguishable for NC^1 circuits. Our construction can thus be seen as JSW [JSW17] instantiated just with Yao's garbling scheme and without the somewhere equivocal encryption layer. In a nutshell, the JSW construction garbles a circuit C by first constructing a new circuit C' consisting of two copies of C connected through selector gates that each forward the output of one of the copies of C . The circuit C' is then garbled using Yao's garbling scheme and the resulting garbled circuit is encrypted with a somewhere equivocal encryption scheme [HJO⁺16]. The somewhere equivocal encryption scheme contributes to the online complexity as decryption keys need to be transmitted in the online phase. We show through a new security analysis that the final encryption step can actually be omitted, thus reducing the online complexity

to $2\lambda|x|$. We now present the construction and then discuss how our proof (cf. Appendix B) differs from JSW.

8.1 Yao’s garbling scheme

Let $\text{se} = (\text{enc}, \text{dec})$ be a symmetric encryption scheme with key space $\{0, 1\}^\lambda$. Yao’s garbling scheme Gb_{yao} [Yao82a, Yao86] consists of algorithms $\text{GCircuit}_{\text{yao}}$, $\text{GInput}_{\text{yao}}$ and $\text{GEval}_{\text{yao}}$ that are shown in Fig. 12. For each wire, two uniformly random keys are sampled and assigned values 0 and 1, respectively. To garble a gate g , its left and right predecessor and the corresponding wire keys are determined, and four ciphertexts are computed according to gate operation $op(g)$.

Our description differs slightly from the literature in the split between garbled circuit and output decoding information: Instead of the more traditional view that treats the whole output map (consisting of wire keys and their mapping to bits) as part of the output decoding information, we separate the wire keys K_{out} (sorted lexicographically to hide their association with bits) from the actual map d (pointing now to entries in the key list). The wire keys then become part of the garbled circuit. This is in line with what JSW called *output-key security* in their construction:

```

( $\text{Sel}_0 \circ C || C$ ) ( $x$ )
-----
 $y_0 \leftarrow C(x)$ 
 $y_1 \leftarrow C(x)$ 
 $y \leftarrow y_0$ 
return  $y$ 

```

Fig. 13: Circuit C' .

Yao’s garbling scheme is actually adaptively simulatable even when the output wire keys (without the mapping to bits) are sent with the garbled gates in the offline phase. We refer to this version as *weak online Yao*. Interestingly, this is reminiscent of the point-and-permute technique [Rog91] applied only to the output wires, and reduces the online complexity to $n\lambda + m$ (assuming a suitable encoding of gate indices). Note that this change does not affect the combined information contained in the garbled circuit and output decoding information, we simply redistribute it.

8.2 Our construction

For $b \in \{0, 1\}$, a *selector gate* Sel_b takes as input two bits u_0 and u_1 and outputs u_b . We define our garbling scheme $\text{Gb} = (\text{GCircuit}, \text{GInput}, \text{GEval})$ as

- $\text{GCircuit}(C) := \text{GCircuit}_{\text{yao}}(C')$, where $C' := \text{Sel}_0 \circ C || C$ (Fig. 13) is a circuit consisting of two copies of C connected by m selector gates Sel_0 .
- $\text{GInput}(K, x) := \text{GInput}_{\text{yao}}(K, x || x)$.
- $\text{GEval}(\tilde{C}, \tilde{x}, d) := \text{GEval}_{\text{yao}}(\tilde{C}, \tilde{x}, d)$

The input size of C' is twice that of C , and hence the input encoding size is twice that of Yao’s garbling scheme.

Comparison to JSW Construction. JSW additionally encrypt the garbled circuit with a layer of somewhere equivocal encryption (SEE) where the simulator can revoke at pebbling complexity many ciphertexts. The overall construction has

an online complexity independent of the output size of the circuit. However, the online complexity of JSW is affected by the SEE decryption key, which is sent together with the garbled input. The decryption key removes the layer of somewhere equivocal encryption and then the garbled circuit can be evaluated on the garbled input. By the formula in [HJO⁺16] (Table 1), the size of the decryption key is $t \cdot s \cdot \lambda \cdot \log n$, where for JSW, t is the number of gates to be revoked and s is the size of a garbled gate, which is effectively linear in λ . Since the number of gates to be revoked is either the *width* w or the *depth* d of the circuit, the size of the decryption key is either $\mathcal{O}(d\lambda^2)$ or $\mathcal{O}(w\lambda^2)$.

8.3 Security

Theorem 7. *Assuming IND-CPA security of symmetric encryption scheme se , the garbling scheme Gb is IND-secure (Def. 5) for NC^1 circuits and has an online complexity of $2n\lambda$.*

Our proof of Theorem 7 follows JSW and HJOSW [HJO⁺16], except for the game-hop moving from left selectors to right selectors. Here, we show that $C_0(x_0) = C_1(x_1)$ implies that selector gates do not need to be garbled in an input-dependent way. Namely, recall that $\text{Sel}_b(u_0, u_1)$ returns u_b , but $C_0(x_0) = C_1(x_1)$ implies that $u_0 = u_1$ and so the output of $\text{Sel}_b(u_0, u_1)$ is *independent* of bit b . Moreover, $u_0 = u_1$ implies that both input bits to Sel_b and the output bit of Sel_b are all equal, so regardless of b , the output 0-key is encrypted under the input left and right 0-keys and the output 1-key is encrypted under the input left and right 1-keys. The “mixed” ciphertexts (encryptions under the left 1-key and the right 0-key as well as encryptions under the left 0-key and the right 1-key) will never be decrypted, since $C_0(x_0) = C_1(x_1)$ implies that the input bit to the selector gate from $C_0(x_0)$ is equal to the input bit to the selector gate from $C_1(x_1)$, hence their content does not matter. We provide the proof of Theorem 7 in Appendix B.

Acknowledgments

We thank Christoph Egger, Pierre Meyer, the cryptography group at ENS Paris and the anonymous reviewers of Asiacrypt 2023 for the interesting discussion, and for pointing us towards the work of Hubáček and Wichs [HW15].

This work was supported by the Research Council of Finland, Blockchain Technology Laboratory at the University of Edinburgh and Input Output Global.

References

- ADL⁺22. Shashank Agrawal, Wei Dai, Atul Luykx, Pratyay Mukherjee, and Peter Rindal. Paradise: Efficient threshold authenticated encryption in fully malicious model. In Takanori Isobe and Santanu Sarkar, editors, *Progress in Cryptology – INDOCRYPT 2022*, pages 26–51, Cham, 2022. Springer International Publishing.

- Agr19. Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 191–225. Springer, Heidelberg, May 2019.
- AIK04. Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0 . In *45th Annual Symposium on Foundations of Computer Science*, pages 166–175. IEEE Computer Society Press, October 2004.
- AIKW13. Benny Applebaum, Yuval Ishai, Eyal Kushilevitz, and Brent Waters. Encoding functions with constant online rate or how to compress garbled circuits keys. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 166–184. Springer, Heidelberg, August 2013.
- AMMR18. Shashank Agrawal, Payman Mohassel, Pratyay Mukherjee, and Peter Rindal. DiSE: Distributed symmetric-key encryption. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018: 25th Conference on Computer and Communications Security*, pages 1993–2010. ACM Press, October 2018.
- App14. Benny Applebaum. *Cryptography in Constant Parallel Time*. Information Security and Cryptography. Springer, Heidelberg, 2014.
- AS16. Prabhanjan Vijendra Ananth and Amit Sahai. Functional encryption for turing machines. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A: 13th Theory of Cryptography Conference, Part I*, volume 9562 of *Lecture Notes in Computer Science*, pages 125–153. Springer, Heidelberg, January 2016.
- BBO07. Mihir Bellare, Alexandra Boldyreva, and Adam O’Neill. Deterministic and efficiently searchable encryption. In Alfred Menezes, editor, *Advances in Cryptology – CRYPTO 2007*, volume 4622 of *Lecture Notes in Computer Science*, pages 535–552. Springer, Heidelberg, August 2007.
- BDF⁺18. Chris Brzuska, Antoine Delignat-Lavaud, Cédric Fournet, Konrad Kohbrok, and Markulf Kohlweiss. State separation for code-based game-playing proofs. In Thomas Peyrin and Steven Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018, Part III*, volume 11274 of *Lecture Notes in Computer Science*, pages 222–249. Springer, Heidelberg, December 2018.
- BHK13. Mihir Bellare, Viet Tung Hoang, and Sriram Keelveedhi. Instantiating random oracles via UCEs. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 398–415. Springer, Heidelberg, August 2013.
- BHR12a. Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology – ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 134–153. Springer, Heidelberg, December 2012.
- BHR12b. Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012: 19th Conference on Computer and Communications Security*, pages 784–796. ACM Press, October 2012.
- BK03. Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In Eli Biham,

- editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 491–506. Springer, Heidelberg, May 2003.
- BK08. Hans L Bodlaender and Arie MCA Koster. Combinatorial optimization on graphs of bounded treewidth. *The Computer Journal*, 51(3):255–269, 2008.
- BLMR13. Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 410–428. Springer, Heidelberg, August 2013.
- BO23. Chris Brzuska and Sabine Oechsner. A state-separating proof for yao’s garbling scheme. In *36th IEEE Computer Security Foundations Symposium - CSF 2023*, pages 137–152. IEEE, 2023.
- Bod88. Hans L Bodlaender. Dynamic programming on graphs with bounded treewidth. In *Automata, Languages and Programming: 15th International Colloquium*, pages 105–118. Springer, 1988.
- Bod98. Hans L. Bodlaender. A partial k-ary tree decomposition of graphs with bounded treewidth. *Theoretical Computer Science*, 209(1):1–45, 1998.
- BSW03. Boaz Barak, Ronen Shaltiel, and Avi Wigderson. Computational analogues of entropy. In *APPROX 2003 and 7th International Workshop on Randomization and Approximation Techniques in Computer Science, RANDOM 2003*, volume 2764 of *Lecture Notes in Computer Science*, pages 200–215. Springer, 2003.
- Can01. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science*, pages 136–145. IEEE Computer Society Press, October 2001.
- CLTV15. Ran Canetti, Huijia Lin, Stefano Tessaro, and Vinod Vaikuntanathan. Obfuscation of probabilistic circuits and applications. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015: 12th Theory of Cryptography Conference, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 468–497. Springer, Heidelberg, March 2015.
- DNRS99. Cynthia Dwork, Moni Naor, Omer Reingold, and Larry J. Stockmeyer. Magic functions. In *40th Annual Symposium on Foundations of Computer Science*, pages 523–534. IEEE Computer Society Press, October 1999.
- Fre90. Eugene C. Freuder. Complexity of k-tree structured constraint satisfaction problems. In *Proceedings of the Eighth National Conference on Artificial Intelligence - Volume 1, AAAI’90*, pages 4–9. AAAI Press, 1990.
- GGH⁺13. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual Symposium on Foundations of Computer Science*, pages 40–49. IEEE Computer Society Press, October 2013.
- GM09. Martin Grohe and Dániel Marx. On tree width, bramble size, and expansion. *Journal of Combinatorial Theory, Series B*, 99(1):218–228, 2009.
- Gol93. Oded Goldreich. A uniform-complexity treatment of encryption and zero-knowledge. *Journal of Cryptology*, 6(1):21–53, March 1993.
- Gol11. Oded Goldreich. Candidate one-way functions based on expander graphs. In Oded Goldreich, editor, *Studies in Complexity and Cryptography. Miscellaneous on the Interplay between Randomness and Computation*, volume 6650 of *Lecture Notes in Computer Science*, pages 76–87. Springer, 2011.

- GS18. Sanjam Garg and Akshayaram Srinivasan. Adaptively secure garbling with near optimal online complexity. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 535–565. Springer, Heidelberg, April / May 2018.
- HJO⁺16. Brett Hemenway, Zahra Jafargholi, Rafail Ostrovsky, Alessandra Scafuro, and Daniel Wichs. Adaptively secure garbled circuits from one-way functions. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part III*, volume 9816 of *Lecture Notes in Computer Science*, pages 149–178. Springer, Heidelberg, August 2016.
- HLR07. Chun-Yuan Hsiao, Chi-Jen Lu, and Leonid Reyzin. Conditional computational entropy, or toward separating pseudoentropy from compressibility. In Moni Naor, editor, *Advances in Cryptology – EUROCRYPT 2007*, volume 4515 of *Lecture Notes in Computer Science*, pages 169–186. Springer, Heidelberg, May 2007.
- HMS23. Iftach Haitner, Noam Mazor, and Jad Silbak. Incompressibility and next-block pseudoentropy. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023*, volume 251 of *LIPICs*, pages 66:1–66:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- HW15. Pavel Hubacek and Daniel Wichs. On the communication complexity of secure function evaluation with long output. In Tim Roughgarden, editor, *ITCS 2015: 6th Conference on Innovations in Theoretical Computer Science*, pages 163–172. Association for Computing Machinery, January 2015.
- JKKR17. Abhishek Jain, Yael Tauman Kalai, Dakshita Khurana, and Ron Rothblum. Distinguisher-dependent simulation in two rounds and its applications. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part II*, volume 10402 of *Lecture Notes in Computer Science*, pages 158–189. Springer, Heidelberg, August 2017.
- JO20. Zahra Jafargholi and Sabine Oechsner. Adaptive security of practical garbling schemes. In Karthikeyan Bhargavan, Elisabeth Oswald, and Manoj Prabhakaran, editors, *Progress in Cryptology - INDOCRYPT 2020: 21st International Conference in Cryptology in India*, volume 12578 of *Lecture Notes in Computer Science*, pages 741–762. Springer, Heidelberg, December 2020.
- JS14. Maurice J. Jansen and Jayalal Sarma. Balancing bounded treewidth circuits. *Theory Comput. Syst.*, 54(2):318–336, 2014.
- JSW17. Zahra Jafargholi, Alessandra Scafuro, and Daniel Wichs. Adaptively indistinguishable garbled circuits. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017: 15th Theory of Cryptography Conference, Part II*, volume 10678 of *Lecture Notes in Computer Science*, pages 40–71. Springer, Heidelberg, November 2017.
- JW16. Zahra Jafargholi and Daniel Wichs. Adaptive security of Yao’s garbled circuits. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B: 14th Theory of Cryptography Conference, Part I*, volume 9985 of *Lecture Notes in Computer Science*, pages 433–458. Springer, Heidelberg, October / November 2016.
- Kar23. Pihla Karanko. Different flavours of hill pseudoentropy and yao incompressibility entropy. Cryptology ePrint Archive, Paper 2023/1867, 2023. <https://eprint.iacr.org/2023/1867>.

- Khu21. Dakshita Khurana. Non-interactive distributional indistinguishability (NIDI) and non-malleable commitments. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part III*, volume 12698 of *Lecture Notes in Computer Science*, pages 186–215. Springer, Heidelberg, October 2021.
- KKP21. Chethan Kamath, Karen Klein, and Krzysztof Pietrzak. On treewidth, separators and Yao’s garbling. In Kobbi Nissim and Brent Waters, editors, *Theory of Cryptography - 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8-11, 2021, Proceedings, Part II*, volume 13043 of *Lecture Notes in Computer Science*, pages 486–517. Springer, 2021.
- LP09. Yehuda Lindell and Benny Pinkas. A proof of security of Yao’s protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, April 2009.
- Muk20. Pratyay Mukherjee. Adaptively secure threshold symmetric-key encryption. In Karthikeyan Bhargavan, Elisabeth Oswald, and Manoj Prabhakaran, editors, *Progress in Cryptology - INDOCRYPT 2020: 21st International Conference in Cryptology in India*, volume 12578 of *Lecture Notes in Computer Science*, pages 465–487. Springer, Heidelberg, December 2020.
- Rog91. Phillip Rogaway. *The round complexity of secure protocols*. PhD thesis, MIT, 1991.
- RS86. Neil Robertson and P.D Seymour. Graph minors. ii. algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986.
- TVZ05. Luca Trevisan, Salil Vadhan, and David Zuckerman. Compression of samplable sources. *Comput. Complex.*, 14(3):186–227, December 2005.
- Yao82a. Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 160–164. IEEE Computer Society Press, November 1982.
- Yao82b. Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 80–91. IEEE Computer Society Press, November 1982.
- Yao86. Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th Annual Symposium on Foundations of Computer Science*, pages 162–167. IEEE Computer Society Press, October 1986.

A Invertibility of low treewidth functions

Recall that Kamath, Klein and Pietrzak (KKP [KKP21]) show indistinguishability-based security of Yao’s garbling scheme for garbling circuits of constant treewidth. Existing results show that many NP-complete problems become efficiently solvable when restricted to instances with bounded treewidth [Bod88, Fre90, BK08] which implies that circuits of low treewidth are efficiently invertible. On the one hand, invertibility of low-treewidth circuits means that the KPP result is significantly stronger than originally stated and shows that low-treewidth circuits can be garbled by Yao’s garbling scheme with *simulation-based* security. On the other hand, invertibility of low-treewidth circuits also means that the KPP result is not useful for garbling cryptographic functions and especially not for garbling PRGs and circumventing the AIKW impossibility result.

For self-containedness, we now provide an explicit algorithm which inverts any circuit with constant treewidth in polynomial time, and any circuit with polylogarithmic treewidth in quasipolynomial time.

A.1 The inversion algorithm

The inversion algorithm $\text{Inverter}_{\text{Sep}}$ we describe below is an adaptation of Goldreich’s inverter for circuits with low *expansion* [Gol11]. Expansion and treewidth of a graph both measure how well-connected a graph is, see Section A.4 for their formal relation. Goldreich’s strategy is a divide-and-conquer approach: Split the graph of the circuit into two roughly equally sized halves, guess the values at the *merging points*, and then invert both halves individually, using the strategy recursively.

The algorithm $\text{Inverter}_{\text{Sep}}$ uses the same strategy and relies on separators as the merging points. A separator is a subset of vertices that, when removed, separates the graph into at least two components of roughly equal size, and we take advantage of the existence of separators whose size is at most the treewidth [RS86, Bod98]. The recursive inversion subroutine $\text{Rec-Inverter}_{\text{Sep}}$ of $\text{Inverter}_{\text{Sep}}$ takes as input a *set* Y of sinks together with their values and tries to invert these. This is useful because the output nodes of the main circuit become sinks of the sub-circuit and thus outputs of the sub-circuit which need to be inverted.

Both algorithms are parametrized with a partial function Sep which, given a circuit, outputs a smallest balanced separator for it. We assume that Sep works on C and all of its sub-components defined recursively by the separators provided by Sep . Sep can either be considered non-uniform information, but can also be computed without too much effort in some cases. Namely, for a circuit of size $|C|$ and treewidth tw , computing Sep takes time at most $|C|^{tw}$ which is polynomial for constant treewidth and quasipolynomial for logarithmic treewidth.

A.2 Definitions and notation

For a circuit C , we use the notation C both for the circuit as well as the underlying directed acyclic graph (DAG). We denote the vertex set and directed edge

set of C by \mathcal{V}^C and \mathcal{E}^G , respectively. For any $A, B \subseteq \mathcal{V}^C$, the edges that have sources in A and are incident on B are denoted by $\mathcal{E}^C(A, B)$. For a subset of the nodes $S \subseteq \mathcal{V}^C$, we call a map $\mathbf{a} : S \rightarrow \{0, 1\}$ an assignment.

Definition 12 (Tree-decomposition [RS86, Bod98]). A tree decomposition of a directed graph $G = (\mathcal{V}, \mathcal{E})$ is a tree, T , with nodes $\mathcal{X}_1, \dots, \mathcal{X}_p$, where each $\mathcal{X}_i \subseteq \mathcal{V}$, satisfies the following properties:

1. Each graph vertex is contained in at least one tree node (i.e., $\cup_{i \in [1, p]} \mathcal{X}_i = \mathcal{V}$).
2. For every edge $(v, w) \in \mathcal{E}$, there exists a node \mathcal{X}_i that contains both v and w .
3. The tree nodes containing a vertex v form a connected subtree of T .

The width of a tree decomposition is the size of its largest node \mathcal{X}_i minus one. Its treewidth tw is the minimum width amongst all tree decompositions of G .

Definition 13 (Separators [RS86]). For a graph $G = (\mathcal{V}, \mathcal{E})$, a set $S \subseteq \mathcal{V}$ is said to be a (balanced) separator if the graph restricted to $\mathcal{V} \setminus S$ has at least two components, and each of these components has size at most $\frac{2}{3}|\mathcal{V}|$.

Theorem 8 ([RS86, Bod98]). A graph G with treewidth tw has a balanced separator of size at most tw .

A.3 Inverting circuits with bounded treewidth

Theorem 9. Let C be a circuit of treewidth tw . Then, there exists a non-uniform algorithm $\text{Inverter}_{\text{Sep}}$ with advice Sep that runs in time $\mathcal{O}(2^{3tw \cdot \log |C|})$ and inverts C with probability 1, i.e.,

$$\Pr_{x \leftarrow \{0,1\}^n, y \leftarrow C(x)} [\text{Inverter}_{\text{Sep}}(y, C) \in C^{-1}(x)] = 1$$

where n is the input length of C . In particular, if C is polynomial in the security parameter, then $\text{Inverter}_{\text{Sep}}$ runs in polynomial time for constant treewidth and in quasipolynomial time for polylogarithmic treewidth.

The inversion algorithm $\text{Inverter}_{\text{Sep}}$ for advice Sep is shown in Fig. 14.

Proof. We start by introducing some notation. Given a set of separators S for a circuit C and an assignment $\mathbf{a} : S \rightarrow \{0, 1\}$, we denote by $\text{Partition}(C, S, \mathbf{a})$ the set of connected components CC emerging from splitting C at the separators. We consider the separators neighboring a connected component CC to be part of it. To make the splitting meaningful, we need to make the following modifications to C before splitting:

A separator can only be placed on a gate with a *single* input. If a separator were to be placed on a *two-input* node, then the two input wires and the single output wire⁷ is replaced by an identity node, and instead of one separator, three separators are placed on each of the three identity gates.

⁷ High fan-out can be placed below the new identity node.

$\text{Inverter}_{\text{Sep}}(C, y)$	$\text{Rec-Inverter}_{\text{Sep}}(C, X, Y)$
$\text{bound} \leftarrow \log C $ $Y \leftarrow \{(i, y_i) \mid i \in \text{sinks}(C)\}$ $X \leftarrow \emptyset$ $X' \leftarrow \text{Rec-Inverter}(C, X, Y)$ $x' \leftarrow \text{encodeAsg}(X')$ return x'	if $ C \leq \text{bound}$ then : $X' \leftarrow \text{bruteforce}(C, X, Y)$ return X' else : $S \leftarrow \text{Sep}(C)$ <i>// iterate over all assignments to S</i> for $\mathbf{a} : S \rightarrow \{0, 1\}$: $X' \leftarrow \emptyset$ for $(CC, X_{CC}, Y_{CC}) \in \text{Partition}(C, S, \mathbf{a})$: $X'' \leftarrow \text{Rec-Inverter}_{\text{Sep}}(CC, X_{CC}, Y_{CC})$ $X' \leftarrow X' \cup X''$ if $X \setminus S = \{x : (x, 0) \in X' \vee (x, 1) \in X'\}$ then $X' \leftarrow X' \cup \{(x, \mathbf{a}(x)) : x \in S \cap X\}$ return X' $X' \leftarrow \{(x, \perp) : x \in X'\}$ return X'
$\text{bruteforce}(C, X, Y)$ <i>// iterate over all assignments to X</i> for $\mathbf{a} : X \rightarrow \{0, 1\}$: if $C(\mathbf{a}(X)) = Y$ then $X' \leftarrow \{(x, \mathbf{a}(x)) : x \in X\}$ return X' $X' \leftarrow \{(x, \perp) : x \in X\}$ return X'	

Fig. 14: Inversion algorithm for low treewidth circuits.

Now, a separator node is either a *constant value* source or a sink of the connected component CC (or disconnected from it in which case it is not included). If it is a sink, then it provides an additional output value which needs to be respected in the process of inversion.

Since the set of input values becomes smaller and the set of output nodes changes in their process, we actually consider the triple

$$(CC, X_{CC}, Y_{CC}) \in \text{Partition}(C, S, \mathbf{a})$$

to be specified by $\text{Partition}(C, S, \mathbf{a})$, i.e., the set X_{CC} of input nodes as well as the set Y_{CC} of pairs (g, y_g) , where g is the index of the output node and y_g is its value.

If X is a set of input nodes and $\mathbf{a} : X \rightarrow \{0, 1\}$ and Y is a set of pairs (g, y_g) , then we use the notation

$$C(\mathbf{a}(X)) = Y$$

to express that circuit C when evaluated on inputs X carrying values according to \mathbf{a} , the output gates g all carry the values y_g as specified by Y .

Finally, note that when using $|C|$, we refer to the number of *non-constant* gates in C .

Inverter runtime. We first turn to the runtime of inverting a circuit C when $|C| \leq \text{bound} = \log |C|$ (base case). In particular, the input size is at a subset of $|C|$ and hence is upper bounded by $|C|$. Note that **bruteforce** runs in time

exponential in the input size, since it enumerates over all possible assignments to the input. Hence, inversion at the base case runs in time $\mathcal{O}(2^{\log |C|}) = \mathcal{O}(|C|)$.

For the recursion, observe that by definition of a balanced separator ([Definition 13](#)), each connected component is at most of size $\frac{2}{3}$ of the size of the previously largest connected component. Thus, the depth of the recursion is at most logarithmic in $|C|$. Each recursion involves 2^{3tw} possible branches. Hence, the overall runtime is $(2^{3tw})^{\log |C|}$, which is polynomial in the security parameter if tw is constant and quasipolynomial if tw is logarithmic.

Inverter correctness. We now prove that if an inverse exists, then the inverter indeed finds an inverse.

Base Case. `bruteforce`(C, X, Y) tries out all possible inputs. Therefore, if a pre-image exists, then `bruteforce` finds it.

Recursion. Given input (C, X, Y) which has (at least) solution X^* and assuming that all recursion subroutines find pre-images if they exist are sound, we need to show that `ReInverter`(C, X, Y) finds a pre-image of Y under C . To show this, it suffices to argue that for *some* assignment $\mathbf{a} : S \rightarrow \{0, 1\}$, all subroutines return a pre-image. Since we know that solution X^* exists, we can run C on solution X^* and see which values the separators S take to obtain assignment $\mathbf{a} : S \rightarrow \{0, 1\}$. Since X^* is a solution for Y under C and is consistent with \mathbf{a} , the corresponding restrictions of X^* are solutions for the subroutines. By the assumption that the subroutines find a pre-image if it exists, they all return a pre-image. As the connected components form a partition, each input is contained in exactly one partition or is a separator. Thus, indeed, all of X is covered by the union of the separators and the connected components and `ReInverter`(C, X, Y) returns a complete pre-image. □

A.4 Connection to circuit expansion

Treewidth and expansion both capture *how connected* a graph is, and Grohe and Marx [[GM09](#)] show that treewidth and expansion are indeed closely related, as the treewidth of a graph is at least as high as its (vertex) expansion. We recall (parameterized) vertex expansion and Proposition 1 of [[GM09](#)]:

Definition 14 (Vertex Expansion [[GM09](#)]). *Let G be a graph. For a set X , let $S(X)$ denote the set of vertices in $V(G) \setminus X$ that are adjacent to X . For every $\alpha \in [0, 1]$, the vertex expansion of G with parameter α is defined as*

$$\min_{\substack{X \subseteq V(G) \\ 0 < |X| \leq \alpha \cdot |V(G)|}} \frac{|S(X)|}{|X|}.$$

Theorem 10 (Grohe-Marx [[GM09](#)]). *Let $n \geq 1$ and $0 \leq \alpha \leq 1$. Then for every n -vertex graph G we have*

$$tw(G) \geq \lfloor vx_\alpha(G) \cdot (\alpha/2) \cdot n \rfloor.$$

Note that when $\alpha = \frac{1}{2}$, we recover the usual (un-parameterized) notion of vertex expansion. We then obtain from Theorem 10 that

$$tw(G) \geq \left\lfloor \frac{n}{4} \cdot vx(G) \right\rfloor.$$

Hence, as soon as $n > 3$ the treewidth of any n -vertex graph is at least that of its vertex expansion.

B Proof of Theorem 7

In order to prove adaptive indistinguishability of our construction, we define a number of hybrid games that differ in which circuits the circuit C' consists of. Each hybrid game $\text{Hyb}_{\mathcal{A}}^{b_L, b_R, \text{sel}}$ in Fig. 15 garbles circuit $C' = \text{Sel}_{\text{sel}} \circ C_{b_L} || C_{b_R}$. It is easy to see that $\text{Hyb}_{\mathcal{A}}^{0,0,0}$ is equivalent to game $\text{Ind}_{\mathcal{A}, \text{Gb}}^0$, and $\text{Hyb}_{\mathcal{A}}^{1,1,0}$ is equivalent to game $\text{Ind}_{\mathcal{A}, \text{Gb}}^1$. We then establish two statements about the relation between hybrids that will allow us to modify the circuit C' that is garbled by switching the internal circuit and the selector gates.

Circuit switch. The following lemma allows us to switch the circuit that is not selected by a circuit of equivalent topology as long as both circuit/input pairs lead to the same output:

Lemma 3 (Circuit switch). *Assuming the IND-CPA security of the underlying encryption scheme, for any PPT adversary \mathcal{A} which queries circuits C_0, C_1 with equal input length and inputs x_0 and x_1 such that $C_0(x_0) = C_1(x_1)$ and $\Phi(C_0) = \Phi(C_1)$,*

$$\left| \Pr[1 = \text{Hyb}_{\mathcal{A}}^{0,0,0}] - \Pr[1 = \text{Hyb}_{\mathcal{A}}^{0,1,0}] \right| = \text{negl}(\lambda) \quad (4)$$

and

$$\left| \Pr[1 = \text{Hyb}_{\mathcal{A}}^{0,1,1}] - \Pr[1 = \text{Hyb}_{\mathcal{A}}^{1,1,1}] \right| = \text{negl}(\lambda). \quad (5)$$

Lemma 3 will be proved in Appendix C through a reduction to the adaptive simulatability of Yao's garbling scheme: Intuitively, simulatability guarantees garbling circuit C_0 on input x_0 is indistinguishable from a simulation given topology $\Phi(C_0)$ and output $C_0(x_0)$, which in turn is indistinguishable from garbling C_1 on input x_1 as both circuits have the same topology and output.

```

HybℳbL, bR, sel(1λ)
-----
(C0, C1, st) ←ℳ ℳ(1λ)
n ← InputLength(C0)
K ←ℳ Ksample(1n)
C' ← Selsel ∘ CbL || CbR
(Ĉ, d) ←ℳ GCircuityao(C')
(x0, x1, st') ←ℳ ℳ(Ĉ, d, st)
x ← xbL || xbR
x̃ ←ℳ GInputyao(K, x)
b' ←ℳ ℳ(x̃, st')
return b'

```

Fig. 15: Hybrid $\text{Hyb}_{\mathcal{A}}^{b_L, b_R, \text{sel}}$.




Fig. 16: Overview of game hops in proof of [Theorem 7](#).

Selector switch. The second lemma allows to switch the selector gates:

Lemma 4 (Selector switch). *Assuming the IND-CPA security of the underlying encryption scheme, we have that for any PPT adversary \mathcal{A} which queries NC^1 circuits C_0, C_1 with equal input length and inputs x_0 and x_1 such that $C_0(x_0) = C_1(x_1)$ and $\Phi(C_0) = \Phi(C_1)$,*

$$\left| \Pr \left[1 = \text{Hyb}_{\mathcal{A}}^{0,1,0} \right] - \Pr \left[1 = \text{Hyb}_{\mathcal{A}}^{0,1,1} \right] \right| = \text{negl}(\lambda) \quad (6)$$

and

$$\left| \Pr \left[1 = \text{Hyb}_{\mathcal{A}}^{1,1,1} \right] - \Pr \left[1 = \text{Hyb}_{\mathcal{A}}^{1,1,0} \right] \right| = \text{negl}(\lambda). \quad (7)$$

Proof of Theorem 7. The proof proceeds as a sequence of game hops, where indistinguishability of games for each hop was shown in [Lemma 3](#) and [4](#), respectively. [Fig. 16](#) shows the circuit to be garbled in each game. Observing that $\text{Hyb}_{\mathcal{A}}^{0,0,0}$ is equivalent to game $\text{Ind}_{\mathcal{A}, \text{Gb}}^0$, and $\text{Hyb}_{\mathcal{A}}^{1,1,0}$ is equivalent to game $\text{Ind}_{\mathcal{A}, \text{Gb}}^1$ then concludes our proof.

B.1 Pebbling Strategies

To prove [Lemma 10](#), we now describe a pebbling strategy `PebbleSelectors` which moves from the left to the right configuration for a selector circuit using G moves and L-R moves. We analyze the complexity of the `PebbleSelectors` afterwards.

Pebbling Strategy. For every gate i in circuit C , we denote by $\text{LeftPred}(C, i)$ and $\text{RightPred}(C, i)$ the two predecessor gates.

- $\text{PebbleSelectors}(C)$:
 1. For each selector gate i in C :
 - Run $\text{RecPutGrey}(C, \text{LeftPred}(C, i))$ and $\text{RecPutGrey}(C, \text{RightPred}(C, i))$
 - Replace the L pebble on gate i with a R pebble.
 - Run $\text{RecRemoveGrey}(C, \text{LeftPred}(C, i))$ and $\text{RecRemoveGrey}(C, \text{RightPred}(C, i))$
- $\text{RecPutGrey}(C, i)$:
 1. If gate i an input gate, put a grey pebble on i and **return**.
 2. Run $\text{RecPutGrey}(C, \text{LeftPred}(C, i))$, $\text{RecPutGrey}(C, \text{RightPred}(C, i))$.
 3. Put a grey pebble on gate i .
 - Run $\text{RecRemoveGrey}(C, \text{LeftPred}(C, i))$ and $\text{RecRemoveGrey}(C, \text{RightPred}(C, i))$
- $\text{RecRemoveGrey}(C, i)$: This is the same as RecPutGrey , except instead of putting a grey pebble on gate i , in steps 1 and 3 we remove it.

To prove Lemma 10, we establish the following sublemma.

Lemma 5 (Single Output Gate Pebbling Complexity). *Let C be a selector circuit of depth d , with selector gates at the very bottom (on the outputs) and m outputs. If we start with an arbitrary configurations of left and right pebbles on the selector gates and white pebbles everywhere on C , then for every output gate i , there is a (G, LR) -pebbling strategy using at most $2d - 1$ grey pebbles and t steps such that in the end, the i -th output (selector) gate is pebbled right instead of left (or vice versa), all other selector gates are in the same configuration as in the beginning, and all other gates are pebbled white.*

Here, $t = 4^{d-1}$.

Notice that our pebbling strategy is simply the JW pebbling strategy with the addition of (nearly identical) rules of pebbling selector gates, and hence the analysis is essentially the same, and is included here for completeness.

More specifically, we need to analyze both the runtime and space complexity (number of grey pebbles) of RecPutGrey .

Runtime. We want to find a closed expression for $t : \mathbb{N} \rightarrow \mathbb{N}$, which is the number of steps made to place a pebble on layer d (and nowhere else). Note that the number of steps needed to remove a pebble and the number of steps to place a pebble are identical. Thus, we know that $t(1) = 1$ and $t(d) = 4t(d-1)$. Thus, we obtain $t(2) = 4$, $t(3) = 4 \cdot 4$ and $t(d) = 4^{d-1}$.

Pebbling complexity. Similarly, we want to find a closed expression for $s : \mathbb{N} \rightarrow \mathbb{N}$ which is the number of grey simultaneous pebbles needed to place a pebble on layer d . To count the number of grey pebbles needed, first observe that, in order to place a pebble on layer d , the algorithm first places a pebble on layer $d - 1$ and then another pebble on layer $d - 1$ and then a pebble on layer d . Then, to remove the first of the two pebbles from layer $d - 1$, one needs $s(d - 1)$ pebbles, while there are two more pebbles lying on layer d and lying on the co-parent. Thus, we have that $s(d) = s(d - 1) + 2$. Moreover, $s(1) = 1$. Hence, we obtain $s(d) = 1 + 2 \cdot (d - 1) = 2d - 1$. \square

C Proof of Circuit Switch Lemma

We first establish the adaptive simulatability of weak online Yao (cf. Section 8.1) before returning to the proof of Lemma 3.

Lemma 6. *Assuming the IND-CPA security of the underlying encryption scheme, weak online Yao is adaptively SIM-secure for NC^1 circuits.*

Proof sketch. The statement follows directly from a modification to the adaptive simulatability proof due to Jafargholi and Wichs [JW16]: Observe that when the output map is computed in the online phase, the map actually only assigns bit values to existing output wire keys that were sampled in the offline phase. The output wire keys are moreover not used as encryption keys anywhere else in the circuit. Since the encryption scheme used to garble gates is further assumed to be IND-CPA secure, revealing the output wire keys (used as messages in the IND-CPA game) does not affect security. Wire keys are moreover chosen completely independently of the rest of the garbling procedure and hence we can use the messages provided by an IND-CPA adversary. \square

Lemma 3. We are going to prove the first statement, the second statement follows from an analogous argument. Let \mathcal{S}_{yao} be the simulator for Gb_{yao} obtained from Lemma 6 when garbling C_0 . Games $\text{Hyb}_{\mathcal{A}}^{0,0,0}$ and $\text{Hyb}_{\mathcal{A}}^{0,1,0}$ differ only in the circuit whose output is not selected. Indistinguishability of the two games is shown in two steps. We define an intermediate game $\text{Hyb}_{\mathcal{S}_{\text{yao}}}$ that *simulates* the garbling of C_0 on input x_0 using \mathcal{S}_{yao} , given only the topology of C_0 and output $y = C_0(x_0) = C_1(x_1)$. This game is shown to be indistinguishable from both $\text{Hyb}_{\mathcal{A}}^{0,0,0}$ and $\text{Hyb}_{\mathcal{A}}^{0,1,0}$ via reductions to adaptive simulatability of garbling C_0 and C_1 respectively. Since $C_0(x_0) = C_1(x_1)$ and both circuits have the same topology, and hence both circuits share the same simulator, game $\text{Hyb}_{\mathcal{S}_{\text{yao}}}$ can be shown to be indistinguishable from both $\text{Hyb}_{\mathcal{A}}^{0,0,0}$ and $\text{Hyb}_{\mathcal{A}}^{0,1,0}$.

For the game hop from $\text{Hyb}_{\mathcal{A}}^{0,0,0}$ to $\text{Hyb}_{\mathcal{S}_{\text{yao}}}$, we construct adaptive simulatability distinguisher \mathcal{B} from adaptive indistinguishability distinguisher \mathcal{A} as shown in Fig. 17. The reduction obtains two circuits C_0, C_1 from \mathcal{A} and forwards C_0 to the challenger. Upon receiving garbled circuit \tilde{C}_0 and output wire keys L_{out} , \mathcal{B} computes a random output map K_{out} using the keys L_{out} . With the help of K_{out} , the reduction then garbles $\text{Sel}_0 \circ C_0$ that can be combined with \tilde{C}_0 to

Reduction $\mathcal{B}(1^\lambda)$:

$(C_0, C_1) \leftarrow_{\$} \mathcal{A}(1^\lambda)$

Send C_0 to the challenger and receive $(\tilde{C}_0, L_{\text{out}})$

for $g \in \text{OutputGates}(C_0)$:

$y_g \leftarrow_{\$} \{0, 1\}$

$(k_L, k_R) \leftarrow L_{\text{out}}(g)$

$K_{\text{out}}(y_g) \leftarrow k_L$

$K_{\text{out}}(1 - y_g) \leftarrow k_R$

Garble $\text{Sel}_0 \circ C_0$ according to Gb_{yao} and obtain (\tilde{C}, K, d^*) ,

using K_{out} as $K_{\text{RightPred}}$ for selector gates

$\tilde{C}^* \leftarrow \tilde{C} \cup \tilde{C}_0$

$(x_0, x_1) \leftarrow_{\$} \mathcal{A}(\tilde{C}^*, d^*)$

Send x_0 to the challenger and receive (\tilde{x}^*, d)

$b' \leftarrow_{\$} \mathcal{A}(\tilde{x}^*)$

return b'

Fig. 17: Reduction \mathcal{B} .

obtain d^* as well as either the garbled circuit $\tilde{C}^* = \text{Sel}_0 \circ C_0 || C_0$ or a similar garbled circuit \tilde{C}^* where C_0 is simulated. Finally, \mathcal{B} queries \mathcal{A} with input \tilde{C}^* and d^* to obtain inputs x_0, x_1 , sends x_0 to the challenger, forwards the response to \mathcal{A} , and returns \mathcal{A} 's output bit b' .

Observe that the output map K_{out} computed by \mathcal{B} does not necessarily match the one of the challenger. This does however not influence the distributions of ciphertext in a garbled selector gate since $\text{Sel}_0(b_L, b_R) = \text{Sel}_0(b_L, 1 - b_R)$. Hence \tilde{C}^* is either $\text{Sel}_0 \circ C_0 || C_0$ garbled according to Yao, or a simulated garbled circuit. In the first case, \mathcal{B} is equivalent to $\text{Hyb}_{\mathcal{A}}^{0,0,0}$. In the second case, \mathcal{B} is equivalent to $\text{Hyb}_{\mathcal{S}_{\text{yao}}}$. The game hop from $\text{Hyb}_{\mathcal{S}_{\text{yao}}}$ to $\text{Hyb}_{\mathcal{A}}^{0,1,0}$ follows via a similar reduction that sends C_1 and x_1 to the challenger instead. \square

D Selector Switch

We use the language of pebbling games to express our result. We proceed as follows: To prove the indistinguishability of $\text{Hyb}_{\mathcal{A}}^{0,1,0}$ and $\text{Hyb}_{\mathcal{A}}^{0,1,1}$, we define different *pebbles* which we can place on gates and which describe how different gates are garbled. Given such a *pebbling configuration*, we describe a hybrid game which garbles the circuit according to the pebbling configuration. We then define *legal moves* between pebbling configurations (also knowns as *pebbling rules*), and we show that

- (1) each *legal move* reduces to IND-CPA security, and that
- (2) we do not make too many moves.

Additionally, we also prove that we do not need more than a certain number of so-called grey pebbles, but let's return to this later.

D.1 Pebbling Colors and Rules

We now first introduce our circuit class, then our pebbling colors, and then our pebbling rules.

Definition 15 (Selector Circuit). *We call a circuit a selector circuit if it consists of two (potentially identical) circuits C_0 and C_1 , composed in parallel and a layer of selector gates at the bottom determining whether the output of the left or the right circuit is returned.*

Colors Selector nodes either have a *left* or a *right* pebble on them. Other nodes have either a *white* or a *grey* pebble on them, leading to four colors: *left* (L), *right* (R), *white* (W) and *grey* (G).

Operations associated with colors. A gate with a *white* pebbled is garbled according to the operation specified by this gate. For a gate with a *grey* pebble on it, all ciphertexts contain the *active* output key, i.e., the operation computed by a grey bit depends on which key is active thus also called *input-dependent* garbling. For a selector gate with left pebble, we have as operation $\text{Sel}_0(x_{\text{left}}, x_{\text{right}}) = x_{\text{left}}$, and for a selector gate with a right pebble on it, we have as operation $\text{Sel}_1(x_{\text{left}}, x_{\text{right}}) = x_{\text{right}}$. Note that selector garblings are not input-dependent and that $\text{Sel}_0(0, 0) = \text{Sel}_1(0, 0) = 0$ and $\text{Sel}_0(1, 1) = \text{Sel}_1(1, 1) = 1$.

Configurations. In Section 8.3, different hybrid games used different variants of the circuit $C' = \text{Sel}_{\text{sel}} \circ C_{b_L} || C_{b_R}$. In this section, we are interested in $C' = \text{Sel}_{\text{sel}} \circ C_0 || C_1$ for $\text{sel} \in \{0, 1\}$ and $C' = \text{Sel}_{\text{sel}} \circ C_1 || C_0$ for $\text{sel} \in \{0, 1\}$. Since we want to swap a left selector by a right selector (or vice versa), in this section, the left or right *pebble* on the selector gate determines whether it is a left or right selector. The hybrids $\text{Hyb}_{\mathcal{A}}^{0,1,0}$ and $\text{Hyb}_{\mathcal{A}}^{0,1,1}$ all garble the entire circuit as in Yao's garbling scheme and use either left or right selectors. We give names to the corresponding pebbling configurations.

Definition 16 (Left/Right Configuration). *If C' is a selector circuit, we call a pebbling configuration $\text{config}_{\text{left}}$ of C' the left configuration if all selector gates are pebbled left and all other gates are pebbled white. We call a pebbling configuration $\text{config}_{\text{right}}$ of C' the right configuration if all selector gates are pebbled right and all other gates are pebbled white.*

Pebbling Rules. Our pebbling games only allow the use of the following two pebbling rules, also illustrated in Fig. 18.

grey (G) If both parents of a node N are grey, then a white pebble on N can be replaced by a grey pebble and a grey pebble on N can be replaced by a white pebble.




Fig. 18: Pebbling rules for input dependent gates (left) and left to right selector gate pebbling rules (right).

left-right (LR) If both parents of a selector node N are grey, then a left pebble on N can be replaced by a right pebble and a right pebble on N can be replaced by a left pebble.

Definition 17 (G-LR pebbling strategy). We call a pebbling strategy G -LR, if it only uses the G rule and the LR rule.

Garbling subroutine. We now define a new garbling subroutine $2\text{Garble}_{\text{config}, \text{bit}}$ in Fig. 19. It is parameterized by a configuration config that indicates how to garble each gate, and partial function bit that contains the current guessed values of the output of a gate. $\text{Gb}_{\text{config}, \text{bit}}$ takes as input two circuits C_0 and C_1 , constructs the circuit $\text{Sel}_{\text{config}} \circ C_0 || C_1$ and then garbles $\text{Sel}_{\text{config}} \circ C_0 || C_1$ as determined by the colors of the pebbling configuration. Note that

$$\begin{aligned}
 \text{GCircuit}(\text{Sel}_0 \circ C_0 || C_1) &= 2\text{Garble}_{\text{config}_{\text{left}}, \perp}(C_0, C_1) & (8) \\
 \text{GCircuit}(\text{Sel}_1 \circ C_0 || C_1) &= 2\text{Garble}_{\text{config}_{\text{right}}, \perp}(C_0, C_1) \\
 \text{GCircuit}(\text{Sel}_1 \circ C_1 || C_1) &= 2\text{Garble}_{\text{config}_{\text{right}}, \perp}(C_1, C_1) \\
 \text{GCircuit}(\text{Sel}_0 \circ C_1 || C_1) &= 2\text{Garble}_{\text{config}_{\text{left}}, \perp}(C_1, C_1) \\
 \text{GInput}(K, x_0 || x_1) &= 2\text{Input}(K, x_0, x_1) \\
 \text{GEval}(\tilde{C}, \tilde{x}, d) &= 2\text{Eval}(\tilde{C}, \tilde{x}, d)
 \end{aligned}$$

D.2 Main lemmas for the proof of Lemma 4

For the proof of Lemma 4, we prove Equation 6, and the proof of Equation 7 is analogous. The proof of Equation 6 consists of four lemmas.

The first lemma introduces the guessing of ℓ many bits into the game and shows that we loose (exactly) a factor of $2^{-\ell}$ in security in this step. The next two lemmas bound the advantage of an adversary between two hybrids that differ exactly by one pebbling rule step. And the fourth lemma bounds the number of these steps by a polynomial. We now first state the lemmas, then show that Equation 6 follows from them, and then provide the proof of the lemmas in Appendix D.3, Appendix D.4, Appendix D.5 and Appendix B.1.

<pre> 2Garble_{config,bit}(C₀, C₁) <hr/> for g ∈ Gates(Sel ∘ C₀ C₁) : if config(g) = (reduction, op₀, op₁) then p^L ← LeftPred(g, Sel ∘ C₀ C₁) p^R ← RightPred(g, Sel ∘ C₀ C₁) for g ∈ Gates(Sel ∘ C₀ C₁) \ {p^L, p^R} : K_g(0) ←_{\$} {0, 1}^λ, K_g(1) ←_{\$} {0, 1}^λ b^L ← bit(p^L); b^R ← bit^R(p^R) K_{p^L}(1 - b_L) ← ⊥; K_{p^L}(1) ← GETA_L(b_L) K_{p^R}(1 - b_R) ← ⊥; K_{p^R}(1) ← GETA_R(b_R) for g ∈ Gates(Sel ∘ C₀ C₁) : K_{left} ← K_{LeftPred}(g, Sel ∘ C₀ C₁) K_{right} ← K_{RightPred}(g, Sel ∘ C₀ C₁) K_{output} ← K_g if config(\tilde{g}) = W : \tilde{g} ←_{\$} garble^{white}(K_{left}, K_{right}, K_{out}, op(g)) if config(\tilde{g}) = L : \tilde{g} ←_{\$} garble^{white}(K_{left}, K_{right}, K_{out}, Sel₀) if config(\tilde{g}) = R : \tilde{g} ←_{\$} garble^{white}(K_{left}, K_{right}, K_{out}, Sel₁) if config(\tilde{g}) = G : d ← bit(g) \tilde{g} ←_{\$} garble^{grey}(K_{left}, K_{right}, K_{out}, d) if config(\tilde{g}) = (reduction, op₀, op₁) : \tilde{g} ←_{\$} garble^{reduction}(K_{out}, op₀, op₁) $\tilde{C}[g] \leftarrow \tilde{g}$ K ← {(K_g, g), g ∈ InputGates(Sel ∘ C₀ C₀)} for g ∈ OutputGates(C) : if K_{out}(g) = (K_g(0), K_g(1)) then out_g ← 0 else out_g ← 1 d ← {(out_g, g) : g ∈ OutputGates(C)} return ((\tilde{C}, K_{out}), K, d) <hr/> 2Input(K, x₀, x₁) <hr/> for g ∈ InputGates(Sel₀ ∘ C₀ C₀) : d_g ← Value of x₀ x₁ at position g $\tilde{x}[g] \leftarrow K_g[d_g]$ return \tilde{x} </pre>	<pre> garble^{white}(K_{left}, K_{right}, K_{out}, op) <hr/> $\tilde{g} \leftarrow \emptyset$ for (b_{left}, b_{right}) ∈ {0, 1}² : d ← op(b_{left}, b_{right}) if LeftPred(g, Sel ∘ C₀ C₁) = p^{L/R} then c_{in} ← ENC_{L/R}(b_{left}, K_{out}(d), K_{out}(d)) else c_{in} ←_{\$} enc(K_{left}(b_{left}), K_{out}(d)) if RightPred(g, Sel ∘ C₀ C₁) = p^{L/R} then c_{out} ← ENC_{L/R}(b_{right}, c_{in}, c_{in}) else c_{out} ←_{\$} enc(K_{right}(b_{right}), c_{in}) $\tilde{g} \leftarrow \tilde{g} \cup \{c_{out}\}$ return \tilde{g} <hr/> garble^{grey}(K_{left}, K_{right}, K_{out}, d) <hr/> $\tilde{g} \leftarrow \emptyset$ for (b_{left}, b_{right}) ∈ {0, 1}² : if LeftPred(g, Sel ∘ C₀ C₁) = p^{L/R} then c_{in} ← ENC_{L/R}(b_{left}, K_{out}(d), K_{out}(d)) else c_{in} ←_{\$} enc(K_{left}(b_{left}), K_{out}(d)) if RightPred(g, Sel ∘ C₀ C₁) = p^{L/R} then c_{out} ← ENC_{L/R}(b_{right}, c_{in}, c_{in}) else c_{out} ←_{\$} enc(K_{right}(b_{right}), c_{in}) $\tilde{g} \leftarrow \tilde{g} \cup \{c_{out}\}$ return \tilde{g} <hr/> garble^{reduction}(K_{out}, d) <hr/> $\tilde{g} \leftarrow \emptyset$ for (b_{left}, b_{right}) ∈ {0, 1}² : d₀ ← op₀(b_{left}, b_{right}), d₁ ← op₁(b_{left}, b_{right}) c_{in}⁰ ← ENC_L(b_{left}, K_{out}(d₀), K_{out}(d₀)) c_{in}¹ ← ENC_L(b_{left}, K_{out}(d₀), K_{out}(d₁)) c_{out} ← ENC_R(b_{right}, c_{in}⁰, c_{in}¹) return \tilde{g} <hr/> 2Eval(\tilde{C}, \tilde{x}, d) <hr/> Same as GEval_{yao} See Fig. 12. </pre>
--	---

Fig. 19: 2Garble subroutine

$\text{PebHyb}_{\mathcal{A}}^{b_L, b_R, \text{config}, \ell}(1^\lambda)$	$\mathcal{R}_{\text{config}_{\text{new}}, \ell}^{\text{config}_{\text{old}}, \mathcal{A}}(1^\lambda)$
$(C_0, C_1, \text{st}) \leftarrow_{\$} \mathcal{A}(1^\lambda)$ $n \leftarrow \text{InputLength}(C_0)$ $K \leftarrow_{\$} \text{Ksample}(1^n)$ $C' \leftarrow \text{Sel} \circ C_{b_L} C_{b_R}$ for every g s.t. $\text{config}(g) = G$: $\text{bit}(g) \leftarrow_{\$} \{0, 1\}$ $r \leftarrow \ell - \text{Dom}(\text{bit}) $ $s \leftarrow_{\$} \{0, 1\}^r$ $(\tilde{C}, d) \leftarrow_{\$} 2\text{Garble}_{\text{config}, \text{bit}}(C_{b_L}, C_{b_R})$ $(x_0, x_1, \text{st}') \leftarrow_{\$} \mathcal{A}(\tilde{C}, d, \text{st})$ $x \leftarrow x_{b_L} x_{b_R}$ for every g s.t. $\text{config}(g) = G$: if value of g in $C'(x) \neq \text{bit}(g)$ then return 0 if $s \neq 0^r$ then return 0 $\tilde{x} \leftarrow 2\text{Input}(K, x_{b_L} x_{b_R})$ $b' \leftarrow_{\$} \mathcal{A}(\tilde{x}, \text{st}')$ return b'	$(C_0, C_1, \text{st}) \leftarrow_{\$} \mathcal{A}(1^\lambda)$ $n \leftarrow \text{InputLength}(C_0)$ $K \leftarrow_{\$} \text{Ksample}(1^n)$ $C' \leftarrow \text{Sel} \circ C_{b_L} C_{b_R}$ for every g s.t. $\text{config}_{\text{old}}(g) = G$: $\text{bit}(g) \leftarrow_{\$} \{0, 1\}$ $r \leftarrow \ell - \text{Dom}(\text{bit}) $ $s \leftarrow_{\$} \{0, 1\}^r$ for every g : if $\text{config}_{\text{old}}(g) = \text{config}_{\text{new}}(g)$ then : $\text{config}(g) \leftarrow \text{config}_{\text{old}}(g)$ else $op_0 \leftarrow \text{op}(g, \text{config}_{\text{old}}(g))$ $op_1 \leftarrow \text{op}(g, \text{config}_{\text{new}}(g))$ $(\tilde{C}, d) \leftarrow_{\$} 2\text{Garble}_{\text{config}, \text{bit}}(C_{b_L}, C_{b_R})$ $(x_0, x_1, \text{st}') \leftarrow_{\$} \mathcal{A}(\tilde{C}, d, \text{st})$ $x \leftarrow x_{b_L} x_{b_R}$ for every g s.t. $\text{config}_{\text{old}}(g) = G$: if value of g in $C'(x) \neq \text{bit}(g)$ then return 0 if $s \neq 0^r$ then return 0 $\tilde{x} \leftarrow 2\text{Input}(K, x_{b_L} x_{b_R})$ $b' \leftarrow_{\$} \mathcal{A}(\tilde{x}, \text{st}')$ return b'

Fig. 20: Hybrid pebbling games and reduction

Lemma 7 (Guessing). *Assuming the IND-CPA security of the underlying encryption scheme, we have that for any PPT adversary \mathcal{A} which queries circuits C_0, C_1 with equal input length and inputs x_0 and x_1 such that $C_0(x_0) = C_1(x_1)$, and for all $\ell \in \mathbb{N}$,*

$$\begin{aligned} & \left| \Pr \left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}_{\text{left}},\ell} \right] - \Pr \left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}_{\text{right}},\ell} \right] \right| \\ &= 2^{-\ell} \left| \Pr \left[1 = \text{Hyb}_{\mathcal{A}}^{0,1,0} \right] - \Pr \left[1 = \text{Hyb}_{\mathcal{A}}^{0,1,1} \right] \right| \end{aligned} \quad (9)$$

Lemma 8 (G Pebbling move). *Let $\text{config}^{\text{old}}$ and $\text{config}^{\text{new}}$ be two pebbling configurations which are one G pebbling move away from each other, and assume w.l.o.g., that a grey pebble is removed in this step. Then, for all PPT adversaries \mathcal{A} ,*

$$\begin{aligned} & \left| \Pr \left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}^{\text{old}},\ell} \right] - \Pr \left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}^{\text{new}},\ell} \right] \right| \\ & \leq \left| \Pr \left[1 = \mathcal{R}_{\text{config}_{\text{new}},\ell}^{\text{config}_{\text{old}},\mathcal{A}}, \text{IND-CPA}_{L,R}^0 \right] - \Pr \left[1 = \mathcal{R}_{\text{config}_{\text{new}},\ell}^{\text{config}_{\text{old}},\mathcal{A}}, \text{IND-CPA}_{L,R}^1 \right] \right| \end{aligned}$$

where Fig. 20 defines $\mathcal{R}_{\text{config}_{\text{new}},\ell}^{\text{config}_{\text{old}},\mathcal{A}}$.

Lemma 9 (Bounding a L-R Pebbling move). *Let $\text{config}^{\text{old}}$ and $\text{config}^{\text{new}}$ be two pebbling configurations which are one L - R pebbling move away from each other, and assume w.l.o.g., that a left pebble is replaced by a right pebble in this step. Then, for all PPT adversaries \mathcal{A} ,*

$$\begin{aligned} & \left| \Pr \left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}^{\text{old}},\ell} \right] - \Pr \left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}^{\text{new}},\ell} \right] \right| \\ & \leq \left| \Pr \left[1 = \mathcal{R}_{\text{config}_{\text{new}},\ell}^{\text{config}_{\text{old}},\mathcal{A}}, \text{IND-CPA}_{L,R}^0 \right] - \Pr \left[1 = \mathcal{R}_{\text{config}_{\text{new}},\ell}^{\text{config}_{\text{old}},\mathcal{A}}, \text{IND-CPA}_{L,R}^1 \right] \right| \end{aligned}$$

where Fig. 20 defines $\mathcal{R}_{\text{config}_{\text{new}},\ell}^{\text{config}_{\text{old}},\mathcal{A}}$.

Lemma 10 (Pebbling Complexity). *Let C be a selector circuit of depth d , with selector gates at the very bottom (on the outputs) and m outputs. The algorithm $\text{PebbleSelectors}(C)$ is a (G,LR) pebbling strategy and requires $2d - 1$ grey pebbles and $m \cdot t$ steps to move from the left to the right configuration. Here, $t = m \cdot 4^d$*

To see that Equation 6 follows, observe that Lemma 10 says that there is a pebbling strategy which needs at most $2d - 1$ grey pebbles, thus, we can choose $\ell = 2d - 1$. Since the depth d is logarithmic in the security parameter, the factor $2^{-\ell}$ which we lose in Lemma 7 is polynomial in the security parameter.

Moreover, Lemma 8 and Lemma 9 ensure that the distinguishing advantage between two adjacent hybrids is negligible, and Lemma 10 ensures that we need only $m \cdot 4^{d-1}$ hybrid games. $m \cdot 4^{d-1}$ is also a polynomial in the security parameter, since d is logarithmic in the security parameter and since the number of inputs m is determined by a circuit written by a PPT adversary. Thus, indeed, if the underlying symmetric encryption scheme is IND-CPA secure, Equation 6 holds. \square

D.3 Proof of Lemma 7 (Guessing)

Due to Equation 8, we have that for all $\ell \in \mathbb{N}$,

$$\Pr\left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}_{\text{left}},\ell} \mid s = 0^{|\ell|}\right] = \Pr\left[1 = \text{Hyb}_{\mathcal{A}}^{0,1,0}\right]$$

and

$$\Pr\left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}_{\text{left}},\ell} \mid s \neq 0^{|\ell|}\right] = 1.$$

Moreover, $\Pr[s = 0^\ell] = 2^{-\ell}$ and thus,

$$\Pr\left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}_{\text{left}},\ell}\right] = 1 \cdot (1 - 2^{-\ell}) + 2^{-\ell} \Pr\left[1 = \text{Hyb}_{\mathcal{A}}^{0,1,0}\right] \quad (10)$$

Analogously,

$$\Pr\left[1 = \text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}_{\text{right}},\ell}\right] = 1 \cdot (1 - 2^{-\ell}) + 2^{-\ell} \Pr\left[1 = \text{Hyb}_{\mathcal{A}}^{0,1,1}\right] \quad (11)$$

Putting equation 10 and 11 together directly yields Equation 9. \square

D.4 Proof of Lemma 8 (G move)

The reduction $\mathcal{R}_{\text{config}_{\text{new}},\ell}^{\text{config}_{\text{old}},\mathcal{A}}$ (cf. Fig. 20) proceeds exactly as $\text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}^{\text{old}},\ell}$ except for the gate g where $\text{config}^{\text{old}}$ and $\text{config}^{\text{new}}$ differ. Here, the operation induced by $\text{config}^{\text{old}}$ always encrypts the active key (denoted op_0), and the operation induced by $\text{config}^{\text{new}}$ (denoted op_1) encrypts whichever key should be encrypted according to the circuit $C' = \text{Sel} \circ C_0 \parallel C_1$. Here, the reduction uses its left-or-right oracles for both parents of g .

In order to simulate the rest of the garbling properly, the reduction also queries the encryption oracles (twice with the same message) whenever the keys of the parents of g are used for encrypting a key. Since the parents of g are grey-pebbled, we know that they only encrypt the active key. The active key is accessible to the reduction because it can retrieve it via the GETA query from the IND-CPA game and does so in the very beginning.

To see the soundness of the reduction, we can inline the code of the IND-CPA game into the reduction and observe that indeed, the emulation is perfect.

To see this from the code of the reduction, first observe that the code marked in pink (cf. Fig. 20) is not used by $\text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}^{\text{old/new}},\ell}$. The reduction computes a new configuration config based on $\text{config}^{\text{old}}$ and $\text{config}^{\text{new}}$ which marks g as a special reduction gate. For the encryptions of all gates which are not g , the reduction perfectly emulates $\text{PebHyb}_{\mathcal{A}}^{0,1,\text{config}^{\text{old/new}},\ell}$. For g itself, if the IND-CPA $_{L/R}^b$ game has bit 0, the encryption always encrypts according to op_0 , i.e., $\text{config}^{\text{old}}$. If the bit is 1, then the encryption always encrypts according to op_1 , except for

the active keys. For active keys, however, op_0 and op_1 are identical, they both encrypt the active key. Thus, we have that

$$\begin{aligned} \mathcal{R}_{\text{config}_{\text{new}}, \ell}^{\text{config}_{\text{old}}, \mathcal{A}} &\rightarrow \text{IND-CPA}_{L/R}^0 = \text{PebHyb}_{\mathcal{A}}^{0,1, \text{config}^{\text{old}}, \ell} \\ \mathcal{R}_{\text{config}_{\text{new}}, \ell}^{\text{config}_{\text{old}}, \mathcal{A}} &\rightarrow \text{IND-CPA}_{L/R}^1 = \text{PebHyb}_{\mathcal{A}}^{0,1, \text{config}^{\text{new}}, \ell} \end{aligned}$$

and Lemma 8 follows. \square

D.5 Proof of Lemma 9 (L-R move)

The reduction and essentially all arguments in this proof are analogous to the proof of Lemma 8 (G move). The core difference is the argument for why, in this game hop, changing from a left to a right pebble on a selector gate does not change which output key is encrypted under the two active input keys when garbling the selector gate.

Since $C_0(x_0) = C_1(x_1)$, we have two cases. In case 0, both the 0-keys of the parent nodes of the selector gate are active and in this case $\text{Sel}_0(0,0) = \text{Sel}_1(0,0) = 0$ as required. In case 1, both the 1-keys of the parent nodes of the selector gate are active and in this case, we also have $\text{Sel}_0(1,1) = \text{Sel}_1(1,1) = 1$ as needed. \square