

Feasibility and Infeasibility of Secure Computation with Malicious PUFs

DANA DACHMAN-SOLED* NILS FLEISCHHACKER† JONATHAN KATZ*
ANNA LYSYANSKAYA‡ DOMINIQUE SCHRÖDER§

Abstract

A recent line of work has explored the use of *physically uncloneable functions (PUFs)* for secure computation, with the goals of (1) achieving universal composability without additional setup, and/or (2) obtaining unconditional security (i.e., avoiding complexity-theoretic assumptions). Initial work assumed that all PUFs, even those created by an attacker, are honestly generated. Subsequently, researchers have investigated models in which an adversary can create *malicious* PUFs with arbitrary behavior. Researchers have considered both malicious PUFs that might be stateful, as well as malicious PUFs that can have arbitrary behavior but are guaranteed to be stateless.

We settle the main open questions regarding secure computation in the malicious-PUF model:

- We prove that unconditionally secure oblivious transfer is impossible, even in the stand-alone setting, if the adversary can construct (malicious) *stateful* PUFs.
- If the attacker is limited to creating (malicious) *stateless* PUFs, then universally composable two-party computation is possible without computational assumptions.

1 Introduction

A *physically uncloneable function (PUF)* [17, 18, 1, 14] is a physical object generated via a process that is intended to create “unique” objects with “random” behavior. PUFs can be probed and their response measured, and a PUF thus defines a function. (We ignore for now the possibility of slight variability in the response, which can be corrected using standard techniques.) At an abstract level, this function has two important properties: it is *random*, and it *cannot be copied* even by the entity who created the PUF. The latter implies that the PUF can only be queried by the party currently holding the PUF, something which distinguishes a PUF from a random oracle.

Since their introduction, several cryptographic applications of PUFs have been suggested, in particular in the area of secure computation. PUFs are especially interesting in this setting because they can potentially be used (1) to obtain *universally composable (UC)* protocols [5] without additional setup, thus bypassing known impossibility results that hold for universal composition in the “plain” model [6, 7], and (2) to construct protocols with *unconditional* security, i.e., without relying on any cryptographic assumptions.

*University of Maryland. Email: danadach@ece.umd.edu, jkatz@cs.umd.edu.

†Johns Hopkins University. Email: mail@nilsfleischhacker.de.

‡Brown University. Email: anna.lysyanskaya@brown.edu.

§Friedrich-Alexander-University Erlangen-Nürnberg. Email: dominique.schroeder@fau.de.

Initial results in this setting [19, 20] showed constructions of oblivious transfer with stand-alone security based on PUFs. Brzuska et al. [4] later formalized PUFs within the UC framework, and showed UC constructions of bit commitment, key agreement, and oblivious transfer (and hence secure computation of arbitrary functionalities) with *unconditional* security. The basic feasibility questions related to PUFs thus seemed to have been resolved.

Ostrovsky et al. [16], however, observed that the previous results implicitly assume that all PUFs, including those created by the attacker, are honestly generated. They point out that this may not be a reasonable assumption: nothing forces the attacker to use the recommended process for manufacturing PUFs and it is not clear, in general, how to “test” whether a PUF sent by some party was generated correctly or not. (Assuming a trusted entity who creates the PUFs is not a panacea, as one of the goals of using PUFs is to avoid reliance on trusted parties.) Addressing this limitation, Ostrovsky et al. defined a model in which an attacker can create *malicious* PUFs with arbitrary, adversary-specified behavior. The previous protocols can be easily attacked in this new adversarial setting, but Ostrovsky et al. showed that it is possible to construct universally composable protocols for secure computation in the malicious-PUF model under additional, number-theoretic assumptions. They explicitly left open the question of whether *unconditional* security is possible in the malicious-PUF model. Recently, Damgård and Scafuro [8] have made partial progress on this question by presenting a commitment scheme with unconditional security in the malicious-PUF model.

Stateful vs. stateless (malicious) PUFs. Honestly generated PUFs are stateless; that is, the output of an honestly generated PUF on some input is independent of previous inputs to the PUF. Ostrovsky et al. noted that maliciously generated PUFs might be stateful or stateless. (The positive results mentioned earlier remain secure even against attackers who can create malicious, stateful PUFs.) Allowing the adversary to create stateful PUFs is obviously more general. Nevertheless, assuming that the adversary is limited to producing stateless PUFs may be reasonable; indeed, depending on the physical technology used to implement the PUFs, incorporating dynamic state in a PUF may be infeasible.

1.1 Our Results

Spurred by the work of Ostrovsky et al. and Damgård and Scafuro, we reconsider the possibility of unconditionally secure computation based on malicious PUFs and resolve the main open questions in this setting. Specifically, we show:

1. Unconditionally secure oblivious transfer (and thus unconditionally secure computation of general functions) is impossible when the attacker can create malicious *stateful* PUFs. Our result holds even with regard to stand-alone security, and even for indistinguishability-based (as opposed to simulation-based) security notions.
2. If the attacker is limited to creating malicious, but *stateless*, PUFs, then universally composable oblivious transfer (OT) and two-party computation of general functionalities are possible. Our oblivious-transfer protocol is efficient and requires each party to create only a single PUF for any bounded number of OT executions. The protocol is also conceptually simple, which we view as positive in light of the heavy machinery used in [16].

Comparison to the previous version of this work. We improve on the proceedings version of this work in several respects. First, we provide a simpler and corrected version of an ideal

functionality corresponding to PUFs. We also provide a full proof for our impossibility result. Finally, we show a simpler protocol for our positive result; specifically, we show how to realize OT using a single PUF rather than two PUFs.

1.2 Other Related Work

Hardware tokens have also been proposed as a physical assumption on which to base secure computation [13]. PUFs are incomparable to hardware tokens: on the one hand, hardware tokens can implement arbitrary code whereas PUFs can only implement a “random function”; on the other hand, since hardware tokens must be efficient, they cannot implement the functionality of a PUF. For this reason, known results (in particular the fact that UC oblivious transfer is impossible with stateless tokens [10]) do not directly translate from one model to the other.

Impossibility results for (malicious) PUFs are also not implied by impossibility results in the random-oracle model (e.g., [2]). A random oracle can be queried by any party at any time, whereas an honestly generated PUF can only be queried by the party who currently holds it. Indeed, we show that oblivious transfer *is* possible when malicious PUFs are assumed to be stateless; in contrast, oblivious transfer is impossible in the random-oracle model [11, 2].

Ostrovsky et al. [16] consider a second malicious model where the attacker can *query* honestly generated PUFs in a non-prescribed manner. They show that secure computation is impossible if both this and maliciously generated PUFs are allowed. We do not consider the possibility of malicious queries in this work.

Rührmair and van Dijk [21] show impossibility results in a malicious-PUF model that differs significantly from (and appears to correspond to a stronger adversary than in) the models considered in prior work [16, 8] and here. In other work [22], van Dijk and Rührmair informally discussed the idea of using the Impagliazzo-Rudich technique [11] to prove impossibility in the context of PUFs, but did not give any formal proofs.

2 Formalizing Physically Uncloneable Functions

Our goal in this section is to provide ideal functionalities for PUFs that correspond to the models used in most of the recent proofs of security for PUF-based protocols. We provide some high-level intuition as well as justification for our specific formal model. Nevertheless, we believe that both our impossibility result (in Section 3) and our feasibility results (in Section 4) are understandable based only on the informal description of our model given in this section.

We begin in Figure 1 with an ideal functionality corresponding to honest PUFs, i.e., where it is assumed that all parties are only able to generate PUFs according to some mandated specification. An honest PUF essentially behaves as a random oracle; in particular, the first time any input c is queried to the PUF, a random value y is returned; if c is queried subsequently, the same value y is returned. The key difference between a PUF and a random oracle, however, is that *only one party can access the PUF at any given time*; this is meant to model the fact that only one party can be in possession of the physical object corresponding to the PUF at any time. The functionality handles this issue by keeping track of the current holder of the PUF. The party in possession of the PUF has the ability to transfer the PUF to another party at any time; this is modeled in the ideal functionality via the *handover* query.

Our ideal functionality for honest PUFs is inspired by the original specification given by Brzuska

$\mathcal{F}_{\text{HPUF}}$ is parameterized by security parameter 1^λ and runs with parties P_1, P_2 , and adversary Sim. It begins with a list \mathcal{L} initialized to \emptyset . It supports the following interface:

- If a party P sends `init`, choose uniform identifier $\text{id} \in \{0, 1\}^\lambda$ and set $\mathcal{L}_{\text{id}} = \emptyset$. Return `id` to P , and store $(\text{id}, P, \mathcal{L}_{\text{id}})$ in \mathcal{L} .
- If a party P sends `(eval, id, c)` then check whether there is an entry $(\text{id}, P, \mathcal{L}_{\text{id}}) \in \mathcal{L}$ and return \perp if not. Otherwise, if there is an entry $(c, y) \in \mathcal{L}_{\text{id}}$, return y to P ; if not, choose uniform $y \in \{0, 1\}^\lambda$, store (c, y) in \mathcal{L}_{id} , and return y to P .
- If a party P sends `(handover, id, P')` then check whether there is an entry $(\text{id}, P, \mathcal{L}_{\text{id}}) \in \mathcal{L}$ and return \perp if not. Otherwise, replace the tuple $(\text{id}, P, \mathcal{L}_{\text{id}})$ in \mathcal{L} with the tuple $(\text{id}, P', \mathcal{L}_{\text{id}})$, and send $(P, \text{handover}, \text{id}, P')$ to P' and Sim.

Figure 1: The ideal functionality $\mathcal{F}_{\text{HPUF}}$ for honestly generated PUFs.

et al. [4], though we have chosen to simplify the ideal functionality without sacrificing anything significant. For completeness, we briefly justify the simplifications we have made:

- Physical PUFs are typically *noisy*; that is, when queried on the same input twice, they will return close—but not identical—outputs. Moreover, the output of a PUF may not be uniform, but may instead only have high entropy. Some prior work has explicitly modeled these features as part of the ideal functionality defining PUFs.

Brzuska et al. [4] have already shown that both these issues can be handled using *fuzzy extractors* [9], and therefore one can assume non-noisy, uniform outputs without loss of generality. Looking ahead, we note that this holds also in the malicious-PUF setting.

- Prior work has allowed the adversary to directly query the PUF during the point in time when the PUF is transferred from one party to another. In our setting, where at least one party is corrupted by the adversary, this does not give the adversary any additional power. For our positive results, one can verify that if *no* party is corrupted then our protocols remain secure against an eavesdropping attacker even if it can query the PUF when it is transferred between the honest parties.

Ostrovsky et al. [16] initiated the formal study of malicious PUFs. They considered two different types of malicious behavior, the first modeling the case where the attacker can produce a PUF that behaves arbitrarily, and the second addressing the case where the attacker can query an honest PUF in a disallowed manner. In this work we only consider the first type of malicious behavior. Figure 2 specifies the ideal functionality in that case. Note that honest users will not query the functionality with the initial query.

Our ideal functionality allows corrupted parties to create PUFs having behavior specified by an arbitrary circuit C . (Of course, if the attacker is limited to running in polynomial time then C must have polynomial size; we return to this point below.) Although not explicitly included in the description of the functionality, the circuit is allowed to have *oracle gates* that enable queries to a (freshly created) PUF that cannot be directly accessed by any party; this ensures that malicious PUFs are at least as powerful as honest PUFs. The functionality encompasses both the case where the attacker can create *stateful* PUFs maintaining arbitrary state (by setting $\ell = *$), as well as the case where the attacker is limited to creating *stateless* PUFs (by setting $\ell = 0$).

We remark for completeness that we have corrected what appears to be a technical issue in the

$\mathcal{F}_{\text{PUF}}^\ell$ is parameterized by security parameter 1^λ and length parameter ℓ , and runs with parties P_1, P_2 , and adversary Sim . It begins with a list \mathcal{L} initialized to \emptyset . It supports the following interface:

- If a party P sends `init`, choose uniform identifier $\text{id} \in \{0, 1\}^\lambda$ and set $\mathcal{L}_{\text{id}} = \emptyset$. Return id to P , and store $(\text{id}, 0, P, \mathcal{L}_{\text{id}})$ in \mathcal{L} .
- If a party P sends `(initval, C)`, where C is a circuit, choose uniform identifier $\text{id} \in \{0, 1\}^\lambda$ and set $\text{state} = \epsilon$. Return id to P , and store $(\text{id}, 1, P, (C, \text{state}))$ in \mathcal{L} . (The circuit C may have special *oracle gates*; see the text for further discussion.)
- If a party P sends `(eval, id, c)` then check whether there is an entry $(\text{id}, b, P, \star) \in \mathcal{L}$ and return \perp if not. Otherwise:
 - $b = 0$: Say there is an entry $(\text{id}, 0, P, \mathcal{L}_{\text{id}}) \in \mathcal{L}$. If there is an entry $(c, y) \in \mathcal{L}_{\text{id}}$, return y to P ; if not, choose uniform $y \in \{0, 1\}^\lambda$, store (c, y) in \mathcal{L}_{id} , and return y to P .
 - $b = 1$: Say there is an entry $(\text{id}, 1, P, (C, \text{state})) \in \mathcal{L}$. Compute $(y, \text{state}') := C(c, \text{state})$ and return y to P . If $\text{state}' \in \{0, 1\}^\ell$ then replace the tuple $(\text{id}, 1, P, (C, \text{state}))$ in \mathcal{L} with the tuple $(\text{id}, 1, P, (C, \text{state}'))$.
- If a party P sends `(handover, id, P')` then check whether there is an entry $(\text{id}, b, P, X) \in \mathcal{L}$ and return \perp if not. Otherwise, replace the tuple (id, b, P, X) in \mathcal{L} with the tuple (id, b, P', X) , and send $(P, \text{handover}, \text{id}, P')$ to P' and Sim .

Figure 2: The ideal functionality $\mathcal{F}_{\text{PUF}}^\ell$ that allows for maliciously generated PUFs. The parameter ℓ determines the size of the state; $\ell = 0$ corresponds to stateless PUFs, and $\ell = *$ corresponds to stateful PUFs with unbounded state.

original formulation of malicious PUFs by Ostrovsky et al. In particular, their work allowed the attacker to construct PUFs computing *arbitrary* functions, even those not computable in polynomial time. Leaving aside the question of whether creating such PUFs is feasible, we note that allowing such behavior may cause problems when trying to use cryptographic hardness assumptions along with PUFs to design secure protocols. Although Ostrovsky et al. propose a patch for this issue (namely, restricting PUFs to computing “admissible” functions that do not allow solving certain cryptographic problems), the patch is not satisfactory because the simulator in the security proof will be unable to simulate the behavior of the PUF.

We have also given what we believe is a more natural definition that allows the attacker to exactly specify the behavior of the PUF. In contrast, Ostrovsky et al. assume a fixed *distribution* over PUFs, and have the ideal functionality sample a malicious PUF from this distribution. Since the attacker can simulate any such distribution on its own, the attacker in our formulation is at least as strong as the attacker in theirs.

Finally, we note that the functionality allows parties to identify a PUF via its identifier id . In particular, this ensures that if an honest P sends an honestly generated PUF to a (potentially malicious) party P' , and then at some later point in time P' is supposed to send that same PUF back to P , then P can verify that the PUF sent by P' is indeed the correct one. Although a physical PUF may not obviously have this property (since it may not have an obvious public identifier), a trivial way to achieve the same effect (for honestly generated PUFs) is for P to “mark” a PUF by querying it at a random point and recording the response; when the PUF is later returned, P can verify that the same response is returned for the same challenge. The presence of these identifiers only makes our negative result stronger; for our positive result, we implicitly assume that the honest party generating a PUF uses the above “marking” approach.

3 Impossibility Result for Malicious, Stateful PUFs

In this section we prove the impossibility of constructing unconditionally secure oblivious transfer (OT) when the attacker is able to create malicious, stateful PUFs. Our impossibility result applies even for an indistinguishability-based, stand-alone definition of security for OT. We assume perfect correctness for simplicity, however this is not essential. For the purposes of our proof, we consider bit OT protocols. We show:

Theorem 1. *Let Π be a PUF-based OT protocol where the sender \mathcal{S} and receiver \mathcal{R} each make polynomially (in security parameter λ) many PUF queries. Consider executions where the sender \mathcal{S} is given uniform inputs $s_0, s_1 \in \{0, 1\}$, the receiver \mathcal{R} is given uniform input $b \in \{0, 1\}$, and the receiver learns s_b . Then at least one of the following holds:*

1. *There is an unbounded adversary \mathcal{S}^* that uses malicious, stateful PUFs, makes polynomially many PUF queries, and outputs the bit b of \mathcal{R} with probability at least $1/2 + 1/\text{poly}(\lambda)$.*
2. *There is an unbounded adversary \mathcal{R}^* that uses malicious, stateful PUFs, makes polynomially many PUF queries, and outputs both bits s_0, s_1 of \mathcal{S} with probability at least $1/2 + 1/\text{poly}(\lambda)$.*

We begin with a high-level overview of our proof before giving the details in the following sections. The starting point for our result is the impossibility of constructing OT in the random-oracle model. Such a result is implied by the impossibility of key agreement in the random-oracle model [11, 2], but we sketch a direct proof here.

Consider an OT protocol in the random-oracle model between a sender \mathcal{S} and receiver \mathcal{R} . We show that either \mathcal{S} or \mathcal{R} can attack the protocol in the same sense as in Theorem 1. Imagine that both parties run the protocol honestly and then at the end of the protocol they each run the Eve algorithm from [2] to obtain a set Q of queries to/answers from the random oracle. This set Q contains all “intersection queries” between \mathcal{S} and \mathcal{R} , namely, all queries that were made by both parties to the random oracle. (Note that the setting here is different from the key-agreement setting in which a third party—i.e., an eavesdropper—runs the Eve algorithm; in our setting, finding intersection queries is trivial for \mathcal{S} and \mathcal{R} since all intersection queries are, by definition, already contained in the view of each parties. The point of running the Eve algorithm is for both parties to reconstruct the *same* set Q containing all intersection queries.) As in [2], conditioned on the transcript of the protocol and this set Q , the views of \mathcal{S} and \mathcal{R} are independent. Then, following [2], with high probability the distribution over \mathcal{R} ’s view conditioned on \mathcal{S} ’s view and Q is statistically close to the distribution over \mathcal{R} ’s view conditioned only on the transcript and Q .

To use the above to obtain an attack, we first consider the distribution D over \mathcal{R} ’s view conditioned on \mathcal{S} ’s view and Q . If the protocol is secure against a malicious sender, then the probability that $b = 0$ (resp., $b = 1$) in that view must be roughly $1/2$, and in particular it is possible to sample views of \mathcal{R} consistent with both $b = 0$ and $b = 1$. Next consider the distribution D' over \mathcal{R} ’s view conditioned on the transcript and Q . Since this distribution is statistically close to the previous distribution, it must again be possible to sample views of \mathcal{R} consistent with both $b = 0$ and $b = 1$. But since \mathcal{R} can sample from distribution D' , this means that a malicious \mathcal{R}^* can with high probability sample a view consistent with $b = 0$ and the protocol transcript *and* a view consistent with $b = 1$ and the given transcript. Correctness of the protocol then implies that \mathcal{R} can with high probability discover both of \mathcal{S} ’s inputs.

From random oracles to PUFs. The problem with extending the above to the PUF model is that, unlike a random oracle, a PUF can only be queried by the party who currently holds it. This

means that the above attack, as described, will not work; in particular, since at most one party will hold each PUF at the end of the protocol, it is not clear that both parties can compute Q . In fact, as we show in Section 4, there *does* exist an unconditionally secure OT protocol in the *stateless* malicious-PUF model. This means that any impossibility result must exploit the fact that the malicious party can create *stateful* PUFs.

To illustrate the main ideas, consider a protocol in which four PUFs are used. $\text{PUF}_{\mathcal{S}}$ and $\text{PUF}'_{\mathcal{S}}$ are created by \mathcal{S} , with $\text{PUF}_{\mathcal{S}}$ held by \mathcal{S} at the end of the protocol and $\text{PUF}'_{\mathcal{S}}$ held by \mathcal{R} at the end of the protocol. Similarly, $\text{PUF}_{\mathcal{R}}$, $\text{PUF}'_{\mathcal{R}}$ are created by \mathcal{R} , with $\text{PUF}_{\mathcal{R}}$ held by \mathcal{R} at the end of the protocol and $\text{PUF}'_{\mathcal{R}}$ held by \mathcal{S} at the end of the protocol. Consider the set Q that contains the following “intersection queries”:

1. All queries that *both* parties made to $\text{PUF}'_{\mathcal{S}}$ or $\text{PUF}'_{\mathcal{R}}$ (as in [2]).
2. All queries that \mathcal{R} made to $\text{PUF}_{\mathcal{S}}$, and all queries that \mathcal{S} made to $\text{PUF}_{\mathcal{R}}$.

The set Q of queries described above is a *superset* of all the “intersection queries” made by \mathcal{S} and \mathcal{R} because it contains any query made by both parties to $\text{PUF}'_{\mathcal{S}}$ or $\text{PUF}'_{\mathcal{R}}$, all queries \mathcal{R} made to $\text{PUF}_{\mathcal{S}}$ (which is a superset of the queries made by both parties to $\text{PUF}_{\mathcal{S}}$), and all queries made by \mathcal{S} to $\text{PUF}_{\mathcal{R}}$ (which is a superset of the queries made by both parties to $\text{PUF}_{\mathcal{R}}$).

We now want to provide a way for both parties to obtain Q . For queries of the first type, this can be achieved using the Eve algorithm from [2] if we can provide a way for \mathcal{S} to query $\text{PUF}'_{\mathcal{S}}$ (resp., for \mathcal{R} to query $\text{PUF}'_{\mathcal{R}}$) at the end of the protocol. This, in turn, can be done if \mathcal{S} constructs $\text{PUF}'_{\mathcal{S}}$ with known code, so that \mathcal{S} can effectively “query” $\text{PUF}'_{\mathcal{S}}$ even when $\text{PUF}'_{\mathcal{S}}$ is no longer in its possession (and analogously for $\text{PUF}'_{\mathcal{R}}$). Specifically, we have each party embed a t -wise independent function in the PUF they create, where t is large enough so that the behavior of the PUF is indistinguishable from a random function as far as execution of the protocol (and the attack) is concerned.

For the second type of queries above, we rely on the ability of \mathcal{S} and \mathcal{R} to create *stateful* PUFs. Specifically, we have \mathcal{S} create $\text{PUF}_{\mathcal{S}}$ in such a way that the PUF records (in an undetectable fashion) all the queries that \mathcal{R} makes to $\text{PUF}_{\mathcal{S}}$; this allows \mathcal{S} to later recover those queries once $\text{PUF}_{\mathcal{S}}$ is in its possession. Queries that \mathcal{S} makes to $\text{PUF}_{\mathcal{R}}$ are handled in a similar fashion.

To complete the proof, we then show that the set of intersection queries as defined above is enough for the analysis from [2] to apply. Here it is crucial for the parties to find intersection queries immediately after each message is sent (as opposed to waiting until the end of the protocol) in order to handle the fact that the PUFs are exchanged between the parties.

3.1 Proof Details

Let Π be a PUF-based OT protocol with ℓ rounds, where a round involves each party alternately sending a message, and possibly transferring PUFs, to the other party. (So, we have 2ℓ messages overall.) We assume without loss of generality that \mathcal{S} sends the first message of the protocol. and \mathcal{R} sends the final message. We also assume without loss of generality that Π is in *normal form*, meaning that:

1. \mathcal{S} and \mathcal{R} each ask at most one PUF query in each round, and
2. the party receiving the final message of the protocol does not query the oracle after receiving this message.

Any protocol can be transformed to be in normal form without affecting its security.

Let z be (a bound on) the total number of PUFs used in the protocol, and let N be a bound on the total number of queries made by each party to those PUFs. To simplify notation, we treat the PUFs $\{\text{PUF}_1, \dots, \text{PUF}_z\}$ as being defined by a single random function H ; the query $q = (j, q')$, which corresponds to the query $\text{PUF}_j(q')$ and can only be made by the party who holds PUF_j , is answered with $H(j, q')$.

For $i \in [2\ell]$ we denote by S_{back}^i the set of $j \in [z]$ such that PUF_j is sent back to its creator along with the i th message. We let S_{other}^i be the set of indices j such that PUF_j is not held by its creator after the i th message is sent. (Note that there can be PUFs that are not in either set.)

We let $\mathbf{M} = m_1, \dots, m_{2\ell}$ denote the messages exchanged in a protocol execution, and let $\mathbf{M}^i = m_1, \dots, m_i$ denote the first i messages sent. An *augmented message* \tilde{m}_i consists of m_i along with a set ψ^i that contains all queries made (up to that point in the protocol) to $\{\text{PUF}_j\}_{j \in S_{\text{back}}^i}$ by the party who sent message m_i . We let $\tilde{\mathbf{M}}, \tilde{\mathbf{M}}^i$ denote augmented transcripts, and set $\Psi^i \stackrel{\text{def}}{=} \cup_{j \leq i} \psi^j$.

We let $\mathcal{V}_{\mathcal{S}}^i$ denote the view of \mathcal{S} up to the moment after the i -th message is sent, which includes \mathcal{S} 's randomness $r_{\mathcal{S}}$, the partial transcript \mathbf{M}^i , and all PUF query/answer pairs known to \mathcal{S} so far. $\mathcal{V}_{\mathcal{R}}^i$ is defined analogously for \mathcal{R} . We let $\mathcal{Q}(\cdot)$ be an operator that extracts the set of queries from a set of query/answer pairs or a view.

Executions and distributions. A (full) execution of protocol Π can be described by a tuple $(r_{\mathcal{S}}, r_{\mathcal{R}}, H)$ where $r_{\mathcal{S}}$ denotes \mathcal{S} 's randomness, $r_{\mathcal{R}}$ denotes \mathcal{R} 's randomness, and H is a random function that determines the behavior of the PUFs used throughout the protocol. (We assume the parties' randomness is also used to determine their inputs to the protocol.) We denote by \mathcal{E} the experiment in which uniform $(r_{\mathcal{S}}, r_{\mathcal{R}}, H)$ are chosen and then the protocol is executed.

Fix a sequence of augmented messages $\tilde{\mathbf{M}}^i = [\tilde{m}_1, \dots, \tilde{m}_i]$ and a set of query/answer pairs \mathbf{P} for which $\Pr_{\mathcal{E}}[(\tilde{\mathbf{M}}^i, \mathbf{P})] > 0$. We denote by $\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})$ the joint distribution over $(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i)$ (as generated by \mathcal{E}) conditioned on the augmented transcript of the first i messages being equal to $\tilde{\mathbf{M}}^i$ as well as $H(j, q') = a$ for all $((j, q'), a) \in \mathbf{P}$. The event $\text{Good}(\tilde{\mathbf{M}}^i, \mathbf{P})$ holds in this distribution if and only if $\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \cap \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \subseteq \mathbf{P}^+$, where $\mathbf{P}^+ \stackrel{\text{def}}{=} \mathbf{P} \cup \Psi^i$, and we define $\mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})$ to be the distribution $\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})$ conditioned on $\text{Good}(\tilde{\mathbf{M}}^i, \mathbf{P})$. For complete transcripts $\tilde{\mathbf{M}}$, the distributions $\mathcal{V}(\tilde{\mathbf{M}}, \mathbf{P})$ and $\mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}, \mathbf{P})$ are defined similarly.

The Eve algorithm. We now define a deterministic algorithm Eve. We imagine Eve as running in 2ℓ steps, where in step i it is given the next value \tilde{m}_i in the augmented transcript of a protocol execution, and is assumed to have access to the PUFs in S_{other}^i .

Construction 1. Let $\varepsilon < 1/100$ be a parameter. Eve begins with a set \mathbf{P} of query/answer pairs initialized to \emptyset . In step i , given \tilde{m}_i , do: as long as there is a query $q = (j, q') \notin \mathbf{P}^+$ such that

$$\Pr_{(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i) \leftarrow \mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})} [q \in \mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \wedge j \in S_{\text{other}}^i] \geq \frac{\varepsilon^2}{100m} \quad \text{or} \quad \Pr_{(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i) \leftarrow \mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})} [q \in \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \wedge j \in S_{\text{other}}^i] \geq \frac{\varepsilon^2}{100m},$$

Eve queries the lexicographically first such $q = (j, q')$ to H and adds $(q, H(q))$ to \mathbf{P} .

The output of Eve is the final set \mathbf{P} .

Let $\Delta(A, B)$ denote the statistical difference between distributions A, B . We show:

Lemma 2. Construction 1 satisfies the following:

1. The expected number of PUF queries made by Eve (where the expectation is taken over uniform choice of (r_S, r_R, H)) is $\text{poly}(N/\varepsilon)$.
2. Let \mathbf{P} be the set output by Eve. (This is a random variable that depends on $\tilde{\mathbf{M}}$, which in turn depends on (r_S, r_R, H) .) Then with probability at least $1 - \varepsilon/2$,

$$\Delta(\mathcal{V}_S(\tilde{\mathbf{M}}, \mathbf{P}) \times \mathcal{V}_R(\tilde{\mathbf{M}}, \mathbf{P}), \mathcal{V}(\tilde{\mathbf{M}}, \mathbf{P})) \leq \varepsilon/2,$$

where $\mathcal{V}_S(\tilde{\mathbf{M}}, \mathbf{P})$ (resp., $\mathcal{V}_R(\tilde{\mathbf{M}}, \mathbf{P})$) is the distribution of \mathcal{S} 's view (resp., \mathcal{R} 's view) in $\mathcal{V}(\tilde{\mathbf{M}}, \mathbf{P})$.

The proof of this lemma is very similar to the analogous proof by Barak and Mahmoody [2], and we have based our notation and our description of the Eve algorithm on their work. The main difference—besides the fact that we consider augmented transcripts—is that we prove (near) independence of \mathcal{S} 's and \mathcal{R} 's views, conditioned on the augmented transcript and the queries of Eve, even though Eve can only make queries to a (different) subset of the PUFs (namely, those in S_{other}^i) after each message m_i of the protocol is sent. Briefly, this is sufficient since for any PUF not in S_{other}^i , it must be the case that the augmented transcript contains all queries that one of the two parties made to that PUF up to the point when m_i is sent. Since S_{other}^i can change after each message is sent, it is important that Eve proceeds in an *online* fashion, asking queries each time a message is sent, rather than delaying its queries to the end of the protocol as in [2].

We give the proof of Lemma 2 in Section 3.2; readers willing to take the lemma on faith may skip to Section 3.3, where we use the lemma to prove Theorem 1.

3.2 Analysis of the Eve Algorithm

Here, we prove Lemma 2.

Recall that H is sampled uniformly from the set of all functions from $[z] \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. For any partial function F with domain $D = [z] \times \{0, 1\}^n$, we denote by $\Pr_H[F]$ the probability (over choice of H) that H is consistent with F . We rely on the following lemma [2]:

Lemma 3. For consistent partial functions F_1, F_2 it holds that

$$\Pr_H[F_1 \cup F_2] = \frac{\Pr_H[F_1] \cdot \Pr_H[F_2]}{\Pr_H[F_1 \cap F_2]}.$$

We first present a product characterization of the distribution $\mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})$.

Lemma 4. For any $(\tilde{\mathbf{M}}^i, \mathbf{P})$ there exists a distribution \mathbf{S} (resp., \mathbf{R}) over \mathcal{S} 's (resp., \mathcal{R} 's) views such that the distribution $\mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})$ is identical to the product distribution $(\mathbf{S} \times \mathbf{R})$ conditioned on the event $\text{Good}(\tilde{\mathbf{M}}^i, \mathbf{P})$. Namely,

$$\mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P}) \equiv ((\mathbf{S} \times \mathbf{R}) \mid \mathcal{Q}(\mathbf{S}) \cap \mathcal{Q}(\mathbf{R}) \subseteq \mathbf{P}^+),$$

where $\mathbf{P}^+ \stackrel{\text{def}}{=} \mathbf{P} \cup \Psi^i$ (note that Ψ^i is implicit in $\tilde{\mathbf{M}}^i$).

Proof. Suppose $(\mathcal{V}_S^i, \mathcal{V}_R^i) \leftarrow \mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})$ is such that $\mathcal{Q}(\mathcal{V}_S^i) \cap \mathcal{Q}(\mathcal{V}_R^i) \subseteq \mathbf{P}^+$. For such $(\mathcal{V}_S^i, \mathcal{V}_R^i)$ we show that $\Pr_{\mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})}[(\mathcal{V}_S^i, \mathcal{V}_R^i)] = \alpha(\tilde{\mathbf{M}}^i, \mathbf{P}) \cdot \alpha_S \cdot \alpha_R$ where $\alpha(\tilde{\mathbf{M}}^i, \mathbf{P})$ depends only on $(\tilde{\mathbf{M}}^i, \mathbf{P})$, where α_S depends only on \mathcal{V}_S^i , and where α_R depends only on \mathcal{V}_R^i . This means that if we let \mathbf{S} be the distribution over $\text{Supp}(\mathcal{V}_S^i)$ such that $\Pr_{\mathbf{S}}[\mathcal{V}_S^i]$ is proportional to α_S and let \mathbf{R} be the distribution

over $\text{Supp}(\mathcal{V}_{\mathcal{R}}^i)$ such that $\Pr_{\mathbf{R}}[\mathcal{V}_{\mathcal{R}}^i]$ is proportional to $\alpha_{\mathcal{R}}$, then $\mathcal{GV}(\tilde{\mathbf{M}}^i, \mathbf{P})$ is proportional (and hence equal to) the distribution $((\mathbf{S} \times \mathbf{R}) \mid \mathcal{Q}(\mathbf{S}) \cap \mathcal{Q}(\mathbf{R}) \subseteq \mathbf{P}^+)$.

Since $\mathcal{Q}(\mathbf{S}) \cap \mathcal{Q}(\mathbf{R}) \subseteq \mathbf{P}^+$ we have

$$\Pr_{\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})} [(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i)] = \Pr_{\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})} [(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i) \wedge \text{Good}(\tilde{\mathbf{M}}^i, \mathbf{P})] = \Pr_{\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})} [\text{Good}(\tilde{\mathbf{M}}^i, \mathbf{P})] \cdot \Pr_{\mathcal{GV}(\tilde{\mathbf{M}}^i, \mathbf{P})} [(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i)]. \quad (1)$$

On the other hand, by the definition of conditional probability we have

$$\Pr_{\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})} [(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i)] = \frac{\Pr_{\mathcal{E}}[(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i, \tilde{\mathbf{M}}^i, \mathbf{P})]}{\Pr_{\mathcal{E}}[(\tilde{\mathbf{M}}^i, \mathbf{P})]}. \quad (2)$$

Therefore, by Equations (1) and (2) we have

$$\Pr_{\mathcal{GV}(\tilde{\mathbf{M}}^i, \mathbf{P})} [(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i)] = \frac{\Pr_{\mathcal{E}}[(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i, \tilde{\mathbf{M}}^i, \mathbf{P})]}{\Pr_{\mathcal{E}}[(\tilde{\mathbf{M}}^i, \mathbf{P})] \cdot \Pr_{\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})}[\text{Good}(\tilde{\mathbf{M}}^i, \mathbf{P})]}. \quad (3)$$

The denominator of Equation (2) only depends on $(\tilde{\mathbf{M}}^i, \mathbf{P})$ and so we can take $\beta(\tilde{\mathbf{M}}^i, \mathbf{P}) = \Pr_{\mathcal{E}}[(\tilde{\mathbf{M}}^i, \mathbf{P})] \cdot \Pr_{\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})}[\text{Good}(\tilde{\mathbf{M}}^i, \mathbf{P})]$. In the following we analyze the numerator.

We claim:

$$\Pr_{\mathcal{E}}[(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i, \tilde{\mathbf{M}}^i, \mathbf{P})] = \Pr[\mathbf{r}_{\mathcal{S}} = r_{\mathcal{S}}] \cdot \Pr[\mathbf{r}_{\mathcal{R}} = r_{\mathcal{R}}] \cdot \Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \cup \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \cup \mathbf{P}].$$

The reason is that the necessary and sufficient condition that $(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i, \tilde{\mathbf{M}}^i, \mathbf{P})$ occurs in an execution of the system is that when we sample a uniform $(r_{\mathcal{S}}, r_{\mathcal{R}}, H)$ it holds that $r_{\mathcal{S}}$ equals \mathcal{S} 's randomness, $r_{\mathcal{R}}$ equals \mathcal{R} 's randomness, and H is consistent with $\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \cup \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \cup \mathbf{P}$. These conditions imply that \mathcal{S} and \mathcal{R} will produce transcript $\tilde{\mathbf{M}}^i$ as well.

By Lemma 3, the fact that $\Psi^i \subseteq \mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \cup \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i)$, and the fact that $\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \cap \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \subseteq \mathbf{P} \cup \Psi^i = \mathbf{P}^+$, we have

$$\begin{aligned} \Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \cup \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \cup \mathbf{P}] &= \Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \cup \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \cup \mathbf{P}^+] \\ &= \Pr_H[\mathbf{P}^+] \cdot \Pr_{\mathcal{E}}[(\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \cup \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i)) \setminus (\mathbf{P}^+)] \\ &= \Pr_H[\mathbf{P}^+] \cdot \frac{\Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \setminus \mathbf{P}^+] \cdot \Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \setminus \mathbf{P}^+]}{\Pr_{\mathcal{E}}[(\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \cap \mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i)) \setminus \mathbf{P}^+]} \\ &= \Pr_H[\mathbf{P}^+] \cdot \Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \setminus \mathbf{P}^+] \cdot \Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \setminus \mathbf{P}^+]. \end{aligned}$$

Therefore:

$$\begin{aligned} \Pr_{\mathcal{GV}(\tilde{\mathbf{M}}^i, \mathbf{P})} [(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i)] &= \Pr[\mathbf{r}_{\mathcal{S}} = r_{\mathcal{S}}] \cdot \Pr[\mathbf{r}_{\mathcal{R}} = r_{\mathcal{R}}] \cdot \Pr_H[\mathbf{P}^+] \\ &\quad \times \frac{\Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \setminus \mathbf{P}^+] \cdot \Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \setminus \mathbf{P}^+]}{\beta(\tilde{\mathbf{M}}^i, \mathbf{P})}, \end{aligned}$$

and so we can take

$$\alpha_{\mathcal{S}} = \Pr[\mathbf{r}_{\mathcal{S}} = r_{\mathcal{S}}] \cdot \Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{S}}^i) \setminus \mathbf{P}^+], \quad \alpha_{\mathcal{R}} = \Pr[\mathbf{r}_{\mathcal{R}} = r_{\mathcal{R}}] \cdot \Pr_{\mathcal{E}}[\mathcal{Q}(\mathcal{V}_{\mathcal{R}}^i) \setminus \mathbf{P}^+],$$

$$\text{and } \alpha(\tilde{\mathbf{M}}^i, \mathbf{P}) = \frac{\Pr_H[\mathbf{P}^+]}{\beta(\tilde{\mathbf{M}}^i, \mathbf{P})}. \quad \square$$

In the remainder of this section, let $\varepsilon_1 = \varepsilon^2/100$.

Lemma 5. *Let $(\tilde{\mathbf{M}}^i, \mathbf{P})$ be the augmented partial transcript and the set of oracle query/answer pairs made by Eve when the last message in $\tilde{\mathbf{M}}^i$ is sent, with $\Pr_{\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})}[\text{Good}(\mathbf{M}^i, \mathbf{P})] > 0$. For every such $(\tilde{\mathbf{M}}^i, \mathbf{P})$, there is a bipartite graph G with vertex sets $(\mathcal{U}_{\mathcal{S}}, \mathcal{U}_{\mathcal{R}})$ and edges E such that:*

1. *Every vertex u in $\mathcal{U}_{\mathcal{S}}$ has a corresponding view \mathcal{S}_u for \mathcal{S} and a set $\mathcal{Q}_u = \mathcal{Q}(\mathcal{S}_u) \setminus \mathbf{P}^+$. The same holds for vertices in $\mathcal{U}_{\mathcal{R}}$ with \mathcal{R} in place of \mathcal{S} .*
2. *There is an edge between $u \in \mathcal{U}_{\mathcal{S}}$ and $v \in \mathcal{U}_{\mathcal{R}}$ if and only if $\mathcal{Q}_u \cap \mathcal{Q}_v = \emptyset$.*
3. *Every vertex is connected to at least a $(1 - 2\varepsilon_1)$ -fraction of vertices in the other component.*
4. *The distribution $(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i) \leftarrow \mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})$ is identical to sampling a uniform edge $(u, v) \leftarrow E$ and taking $(\mathcal{S}_u, \mathcal{R}_v)$ (i.e., the views corresponding to u and v).*

Proof. For fixed $(\tilde{\mathbf{M}}^i, \mathbf{P})$, the bipartite graph $G = (\mathcal{U}_{\mathcal{S}}, \mathcal{U}_{\mathcal{R}}, E)$ is defined as follows. Every node $u \in \mathcal{U}_{\mathcal{S}}$ will have a corresponding partial view \mathcal{S}_u of \mathcal{S} that is in the support of the distribution \mathbf{S} from Lemma 4. We let the number of nodes corresponding to a view $\mathcal{V}_{\mathcal{S}}^i$ be proportional to $\Pr_{\mathbf{S}}[\mathbf{S} = \mathcal{V}_{\mathcal{S}}^i]$, meaning that \mathbf{S} corresponds to the uniform distribution over the vertices in $\mathcal{U}_{\mathcal{S}}$. Similarly, every node $v \in \mathcal{U}_{\mathcal{R}}$ will have a corresponding partial view \mathcal{R}_v such that \mathbf{R} corresponds to the uniform distribution over $\mathcal{U}_{\mathcal{R}}$.

For $u \in \mathcal{U}_{\mathcal{S}}$, we define $\mathcal{Q}_u = \mathcal{Q}(\mathcal{S}_u) \setminus (\mathbf{P} \cup \Psi^i) = \mathcal{Q}(\mathcal{S}_u) \setminus \mathbf{P}^+$ to be the set of queries *outside* \mathbf{P}^+ that were asked by \mathcal{S} in the view \mathcal{S}_u . We define $\mathcal{Q}_v = \mathcal{Q}(\mathcal{R}_v) \setminus \mathbf{P}^+$ similarly. We put an edge between nodes u and v in G (denoted by $u \sim v$) if and only if $\mathcal{Q}_u \cap \mathcal{Q}_v = \emptyset$.

It can be seen that the distribution $\mathcal{G}\mathcal{V}(\tilde{\mathbf{M}}^i, \mathbf{P})$ is equal to the distribution obtained by choosing a uniform edge $u \sim v$ of G and then outputting the views $(\mathcal{S}_u, \mathcal{R}_v)$. It is thus immediate that properties 1, 2, and 4 hold. It remains to show property 3. To show this, we will argue that the graph G is *dense* as formalized in the next claim.

Claim 6. *For every $u \in \mathcal{U}_{\mathcal{S}}$, $d(u) \geq |\mathcal{U}_{\mathcal{R}}| \cdot (1 - 2\varepsilon_1)$ and for every $v \in \mathcal{U}_{\mathcal{R}}$, $d(v) \geq |\mathcal{U}_{\mathcal{S}}| \cdot (1 - 2\varepsilon_1)$, where $d(w)$ is the degree of vertex w .*

To prove the claim, we first show that for every $w \in \mathcal{U}_{\mathcal{S}}$, it holds that $\sum_{v \in \mathcal{U}_{\mathcal{R}}, w \not\sim v} d(v) \leq \varepsilon_1 \cdot |E|$. The reason is that the probability of vertex v being chosen when we choose a uniform edge is $d(v)/|E|$, and if $\sum_{v \in \mathcal{U}_{\mathcal{R}}, w \not\sim v} d(v)/|E| > \varepsilon_1$ it means that $\Pr_{(u,v) \leftarrow E}[\mathcal{Q}_w \cap \mathcal{Q}_u \neq \emptyset] \geq \varepsilon_1$. Moreover, note that $\mathcal{Q}_w \cap \mathcal{Q}_u$ can only contain queries of the form $q = (j, q')$ where $j \in S_{\text{back}}^i$ since, for any $j \notin S_{\text{back}}^i$, the set \mathbf{P}^+ contains all queries made by at least one of the parties to PUF_j by the time the i th message was sent. Hence, by the pigeonhole principle (since $|\mathcal{Q}_w| \leq m$), there must exist $q = (j, q')$ where $j \in \text{back}^i$ such that $\Pr_{(u,v) \leftarrow E}[q \in \mathcal{Q}_v] \geq \varepsilon_1/m$. But this is a contradiction, because if that holds then q should have been in \mathbf{P} by definition of Eve (and hence q could not be in \mathcal{Q}_w). The same argument shows that for every $w \in \mathcal{U}_{\mathcal{R}}$, $\sum_{u \in \mathcal{U}_{\mathcal{S}}, w \not\sim u} d(u) \leq \varepsilon_1 \cdot |E|$. Thus, for every vertex $w \in \mathcal{U}_{\mathcal{S}} \cup \mathcal{U}_{\mathcal{R}}$, $|E^{\not\sim}(w)| \leq \varepsilon_1 |E|$ where $E^{\not\sim}(w)$ denotes the set of edges that do not contain any neighbor of w (i.e., $E^{\not\sim}(w) = \{(u, v) \in E \mid u \not\sim w \wedge w \not\sim v\}$).

The following claim was proved in [2]:

Claim 7. *For $\varepsilon_1 \leq 1/2$, let $G = (\mathcal{U}_{\mathcal{S}}, \mathcal{U}_{\mathcal{R}}, E)$ be a nonempty bipartite graph where $|E^{\not\sim}(w)| \leq \varepsilon_1 \cdot |E|$ for all vertices $w \in \mathcal{U}_{\mathcal{S}} \cup \mathcal{U}_{\mathcal{R}}$. Then $d(u) \geq |\mathcal{U}_{\mathcal{R}}| \cdot (1 - 2\varepsilon_1)$ for all $u \in \mathcal{U}_{\mathcal{S}}$ and $d(v) \geq |\mathcal{U}_{\mathcal{S}}| \cdot (1 - 2\varepsilon_1)$ for all $v \in \mathcal{U}_{\mathcal{R}}$.*

This claim completes the proof of Claim 6 and therefore the proof of Lemma 5. \square

We say that event **Fail** holds if and only if for some $i \in [2\ell]$, immediately after the i th message is sent, \mathcal{S} or \mathcal{R} makes a query q that was made already by the other party but is not contained in $\mathsf{P}^+ := \mathsf{P} \cup \Psi^i$. If the first query that makes **Fail** happen is the $(i + 1)$ st query and i is odd, we say event **RFail** $_i$ happened, and if i is even we say event **SFail** $_i$ happened.

We prove the following lemma:

Lemma 8. *For odd $i \in [2\ell]$ and every $(\mathcal{V}_{\mathcal{R}}^i, \tilde{\mathsf{M}}^i, \mathsf{P})$ sampled by executing the system it holds that*

$$\Pr_{\mathcal{G}\mathcal{V}(\tilde{\mathsf{M}}^i, \mathsf{P})} [\mathbf{RFail}_i \mid \mathcal{V}_{\mathcal{R}}^i] \leq \frac{3\varepsilon_1}{2m}.$$

A symmetric statement holds for even $i \in [2\ell]$ and \mathcal{S} .

Proof. Let $q = (j, q')$ be the $(i + 1)$ st query of the protocol, made by \mathcal{R} immediately after the last message \tilde{m}_i in $\tilde{\mathsf{M}}^i$ is sent from \mathcal{S} to \mathcal{R} . By Lemma 5, the distribution $\mathcal{G}\mathcal{V}(\tilde{\mathsf{M}}^i, \mathsf{P})$ conditioned on getting $\mathcal{V}_{\mathcal{R}}^i$ as \mathcal{R} 's view is the same as sampling a uniform edge $u \sim v$ in the graph G conditioned on $\mathcal{R}_v = \mathcal{V}_{\mathcal{R}}^i$. We prove Lemma 8 even conditioned on choosing any vertex v such that $\mathcal{R}_v = \mathcal{V}_{\mathcal{R}}^i$. For such fixed v , the distribution of \mathcal{S} 's view \mathcal{S}_v when we choose a uniform edge $u \sim v'$ conditioned on $v = v'$ is the same as choosing a uniform neighbor $u \leftarrow N(v)$ of the node v and then selecting \mathcal{S} 's view \mathcal{S}_u corresponding to node u . Let $S = \{u \in \mathcal{U}_{\mathcal{S}} \mid q \in \mathcal{Q}_u\}$. Note that if $q = (j, q')$ is such that $j \notin \mathcal{S}_{\text{other}}^i$ then we have

$$\Pr_{u \leftarrow N(v)} [q \in \mathcal{Q}_u] = 0.$$

This is because \mathcal{R} can only query a PUF that it holds at the point right after the i th message is sent. However, if PUF_j is such that \mathcal{R} currently holds it, and $\text{PUF}_j \notin \mathcal{S}_{\text{other}}^i$, then PUF_j must have been created by \mathcal{R} . Thus, by definition of the augmented transcript $\tilde{\mathsf{M}}^i$, all queries made by \mathcal{S} to PUF_j up to this point in the protocol are included in $\Psi^i \subseteq \mathsf{P}^+$ and thus cannot be in $\mathcal{Q}_u = \mathcal{Q}(\mathcal{S}_u) \setminus \mathsf{P}^+$.

We can therefore focus our attention on queries $q = (j, q')$ such that $j \in \mathcal{S}^i$. We have

$$\begin{aligned} \Pr_{u \leftarrow N(v)} [q \in \mathcal{Q}_u] &\leq \frac{|S|}{d(v)} \leq \frac{|S|}{(1 - 2\varepsilon_1) \cdot \mathcal{U}_{\mathcal{S}}} \leq \frac{|S| \cdot |\mathcal{U}_{\mathcal{R}}|}{(1 - 2\varepsilon_1) \cdot |E|} \\ &\leq \frac{\sum_{u \in S} d(u)}{(1 - 2\varepsilon_1)^2 \cdot |E|} \leq \frac{\varepsilon_1}{(1 - 2\varepsilon_1)^2 \cdot m} \leq \frac{3\varepsilon_1}{2m}. \end{aligned}$$

The second and fourth inequalities are by Lemma 5. The third inequality is because $|E| \leq |\mathcal{U}_{\mathcal{S}}| \cdot |\mathcal{U}_{\mathcal{R}}|$, and the sixth inequality is because $\varepsilon_1 < \varepsilon < 1/100$. The fifth inequality is by definition of **Eve**, who asks high-probability queries (j, q') , for $j \in \mathcal{S}^i$, as long as such queries exist. Namely, when we choose a uniform edge $u \sim v$ (which by Lemma 5 is the same as sampling $(\mathcal{V}_{\mathcal{S}}^i, \mathcal{V}_{\mathcal{R}}^i) \leftarrow \mathcal{G}\mathcal{V}(\tilde{\mathsf{M}}^i, \mathsf{P})$), it holds that $u \in S$ with probability $\sum_{u \in S} d(u)/|E|$. But for all $u \in S$ it holds that $q \in \mathcal{Q}_u$, and so if $\sum_{u \in S} d(u)/|E| > \varepsilon_1/m$ the query q would have been made by **Eve** already, and by property 2 of Lemma 5 it cannot be the case that q is in any set \mathcal{Q}_u . \square

Given Lemma 8, Lemma 2 holds via the same analysis as in [2].

3.3 Attacking the Protocol

We now show how to use the Eve algorithm from Construction 1 to derive an attack on protocol Π . First note that although Eve, as defined, makes an *expected* polynomial number of queries to the PUFs, we can modify Eve in the standard way so that it makes at most $t = \text{poly}(N/\epsilon)$ queries to the PUFs and such that with probability at least $1 - \epsilon$ over the augmented transcript \tilde{M} and the output P of Eve it holds that

$$\Delta(\mathcal{V}_S(\tilde{M}, P) \times \mathcal{V}_R(\tilde{M}, P), \mathcal{V}(\tilde{M}, P)) \leq \epsilon. \quad (4)$$

We assume this Eve is used in all that follows.

A second, crucial observation is that *both a malicious \mathcal{S} and a malicious \mathcal{R}* can create PUFs so that they can run Eve during an execution of Π . Here, we rely critically on the fact that the parties can create malicious, stateful PUFs. We describe how a malicious \mathcal{S} can run Eve, but note that the situation is symmetric from the point of view of a malicious \mathcal{R} .

- Let $t^* \stackrel{\text{def}}{=} t + 2N$. All PUFs created by \mathcal{S} will be modified in the following two ways: first, instead of being created (honestly) as a random function, each PUF will be defined by choosing an independent key k for a t^* -wise independent function h . Second, each PUF will use its state to keep track¹ of the queries made to it by \mathcal{R} .
- \mathcal{S} otherwise runs the protocol honestly, running Eve after each message of the protocol is sent. We describe how this is done both when \mathcal{S} sends and when it receives a message:
 - Consider the case after \mathcal{S} sends the i th message m_i . Observe that \mathcal{S} knows ψ^i because it knows S_{back}^i and it certainly knows the queries it made to those PUFs. Moreover, we show that it can access the PUFs in S_{other}^i . The PUFs in S_{other}^i are of two types: those created by \mathcal{R} but currently held by \mathcal{S} , and those created by \mathcal{S} but currently held by \mathcal{R} . PUFs of the first type can be directly accessed by \mathcal{S} . PUFs of the second type can still be computed by \mathcal{S} because they are defined by a t^* -wise independent hash function whose key is known by \mathcal{S} . We conclude that \mathcal{S} is able to run Eve in this case.
 - Consider the case when \mathcal{S} receives the i th message m_i . Along with that message, \mathcal{S} receives a set of PUFs S_{back}^i from \mathcal{R} . Since (by definition) all those PUFs were created by \mathcal{S} , it can extract all the queries made by \mathcal{R} to those PUFs and hence compute ψ^i . Exactly as before, \mathcal{S} can also access all the PUFs in S_{other}^i . We conclude that \mathcal{S} is able to run Eve in this case as well.

The total number of queries made to any PUF throughout the entire experiment, whether by \mathcal{S} , \mathcal{R} as part of the protocol or by Eve, is at most $t + 2N = t^*$. This means that the resulting distribution is identical to the one analyzed in Lemma 2 (with the exception that we bound the number of queries made by Eve, as discussed above).

With the above in place, we show that either \mathcal{S} or \mathcal{R} can carry out a successful attack by running Eve. We let $\tilde{\Pi}$ denote the experiment in which \mathcal{S} and \mathcal{R} run an honest execution of the protocol and Eve is additionally run with $\epsilon < 1/400$. We let $\mathcal{V}_R(\tilde{M}, P, \mathcal{V}_S)$ denote the distribution

¹This is easy to do by having \mathcal{S} choose a random “trapdoor” td , and then create a circuit with the following behavior: on input $x \neq \text{td}$, return $h_k(x)$ and concatenate x to the state; on input td , return the current state. Note that creating PUFs in this way will have only a negligible effect on the output of an honest execution of any PUF-based protocol, since the probability that the PUF is queried with td during execution of the protocol is negligible.

on the view of \mathcal{R} in $\tilde{\Pi}$ conditioned on the augmented transcript of the first i messages being equal to \tilde{M}^i , the view of \mathcal{S} being equal to \mathcal{V}_S , and $H(j, q') = a$ for all $((j, q'), a) \in P$. Given a view $\mathcal{V}_{\mathcal{R}}$ of \mathcal{R} , we let $\text{in}(\mathcal{V}_{\mathcal{R}})$ denote the input bit of \mathcal{R} in that view, and let $\text{out}(\mathcal{V}_{\mathcal{R}})$ be the output bit of \mathcal{R} implicit in that view (i.e., as dictated by Π). We define $\text{in}(\mathcal{V}_{\mathcal{R}})$ analogously for \mathcal{S} .

Let $p(\cdot)$ be a sufficiently large polynomial. Clearly, one of the following must hold:

Case 1: For infinitely many λ , with probability at least $1/p(\lambda)$ over $(\tilde{M}, P, \mathcal{V}_S)$ generated by a run of $\tilde{\Pi}$ either

$$\Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{M}, P, \mathcal{V}_S)}[\text{in}(\mathcal{V}_{\mathcal{R}}) = 0] < 0.45 \quad \text{or} \quad \Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{M}, P, \mathcal{V}_S)}[\text{in}(\mathcal{V}_{\mathcal{R}}) = 1] < 0.45.$$

Case 2: For infinitely many λ , with probability at least $1 - 1/p(\lambda)$ over $(\tilde{M}, P, \mathcal{V}_S)$ generated by a run of $\tilde{\Pi}$

$$\Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{M}, P, \mathcal{V}_S)}[\text{in}(\mathcal{V}_{\mathcal{R}}) = 0] \geq 0.45 \quad \text{and} \quad \Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{M}, P, \mathcal{V}_S)}[\text{in}(\mathcal{V}_{\mathcal{R}}) = 1] \geq 0.45.$$

We show that if case 1 holds then a malicious sender can successfully attack the protocol, whereas if case 2 holds then a malicious receiver can successfully attack the protocol.

3.3.1 Attack by a Malicious Sender

Assume case 1 holds. The attack by a malicious sender \mathcal{S}^* proceeds as follows. It runs the protocol with the honest receiver, additionally running Eve after each message is sent (as described previously). At the end of the execution, it has an augmented transcript \tilde{M} , the set of queries P output by Eve, and its own view \mathcal{V}_S . It then computes the probabilities

$$p_0 \stackrel{\text{def}}{=} \Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{M}, P, \mathcal{V}_S)}[\text{in}(\mathcal{V}_{\mathcal{R}}) = 0] \quad \text{and} \quad p_1 \stackrel{\text{def}}{=} \Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{M}, P, \mathcal{V}_S)}[\text{in}(\mathcal{V}_{\mathcal{R}}) = 1].$$

Finally, if $p_0 < 0.45$ it outputs 1; if $p_1 < 0.45$ it outputs 0; and otherwise it outputs a random bit. (Note that $p_0 + p_1 = 1$. Therefore, if $p_0 < 0.45$ we must have $p_1 \geq 0.55$ and vice versa.)

We now analyze the probability with which the output of \mathcal{S}^* is equal to \mathcal{R} 's input bit. Since case 1 holds, we know that for infinitely many λ the probability that $p_0 < 0.45$ or $p_1 < 0.45$ is at least $1/p(\lambda)$. When that happens, \mathcal{S}^* correctly predicts \mathcal{R} 's output with probability at least 0.55; otherwise, \mathcal{S}^* guesses \mathcal{R} 's output with probability 0.5. Overall, then, \mathcal{S}^* outputs \mathcal{R} 's input with probability at least

$$0.55 \cdot \frac{1}{p(\lambda)} + 0.5 \cdot \left(1 - \frac{1}{p(\lambda)}\right) = \frac{1}{2} + \frac{1}{20 \cdot p(\lambda)}$$

for infinitely many λ .

3.3.2 Attack by a Malicious Receiver

Assume case 2 holds. The attack by a malicious receiver \mathcal{R}^* proceeds as follows. It runs the protocol with the honest sender, additionally running Eve after each message is sent (as described

previously). At the end of the execution, it has an augmented transcript $\tilde{\mathbf{M}}$ and the set of queries \mathbf{P} output by Eve. It then computes the probabilities

$$p'_0 \stackrel{\text{def}}{=} \Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 0] \quad \text{and} \quad p'_1 \stackrel{\text{def}}{=} \Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 1].$$

If $p'_0 = 0$ or $p'_1 = 0$ then it outputs \perp and terminates. Otherwise, it samples views $\mathcal{V}_{\mathcal{R}}^0$ and $\mathcal{V}_{\mathcal{R}}^1$ from $\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}, \text{in}(\mathcal{V}_{\mathcal{R}}) = 0)$ and $\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}, \text{in}(\mathcal{V}_{\mathcal{R}}) = 1)$, respectively (where $\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}, \text{in}(\mathcal{V}_{\mathcal{R}}) = b)$ denotes that the distribution is additionally conditioned on $\text{in}(\mathcal{V}_{\mathcal{R}}) = b$). Finally, it outputs $s'_0 = \text{out}(\mathcal{V}_{\mathcal{R}}^0)$ and $s'_1 = \text{out}(\mathcal{V}_{\mathcal{R}}^1)$.

We are interested in the probability with which s'_0, s'_1 correspond to the inputs bits of \mathcal{S} . Toward this, we first prove the following lemma:

Lemma 9. *With probability at least $1 - \sqrt{\varepsilon} - \varepsilon > 1 - 2\sqrt{\varepsilon}$ over $(\tilde{\mathbf{M}}, \mathbf{P}, \mathcal{V}_{\mathcal{S}})$ generated in $\tilde{\Pi}$:*

$$\Delta \left(\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}), \mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}, \mathcal{V}_{\mathcal{S}}) \right) \leq \sqrt{\varepsilon}.$$

Proof. Note that for any fixed $(\tilde{\mathbf{M}}, \mathbf{P})$,

$$\Delta \left(\mathcal{V}_{\mathcal{S}}(\tilde{\mathbf{M}}, \mathbf{P}) \times \mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}), \mathcal{V}(\tilde{\mathbf{M}}, \mathbf{P}) \right) = \mathbf{Exp}_{\mathcal{V}_{\mathcal{S}} \sim \mathcal{V}_{\mathcal{S}}(\tilde{\mathbf{M}}, \mathbf{P})} \left[\Delta \left(\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}), \mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}, \mathcal{V}_{\mathcal{S}}) \right) \right].$$

Moreover, from Equation (4) we know that with probability at least $1 - \varepsilon$ over $(\tilde{\mathbf{M}}, \mathbf{P})$

$$\Delta \left(\mathcal{V}_{\mathcal{S}}(\tilde{\mathbf{M}}, \mathbf{P}) \times \mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}), \mathcal{V}(\tilde{\mathbf{M}}, \mathbf{P}) \right) \leq \varepsilon.$$

The lemma follows using Markov's inequality. \square

Since case 2 holds, we know that for infinitely many λ the probability that both

$$\Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}, \mathcal{V}_{\mathcal{S}})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 0] \geq 0.45 \quad \text{and} \quad \Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}, \mathcal{V}_{\mathcal{S}})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 1] \geq 0.45$$

is at least $1 - 1/p(\lambda)$. By perfect correctness, we also have

$$\Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}, \mathcal{V}_{\mathcal{S}})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 0 \wedge \text{out}(\mathcal{V}_{\mathcal{R}}) \neq s_0] = 0 \quad \text{and} \quad \Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P}, \mathcal{V}_{\mathcal{S}})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 1 \wedge \text{out}(\mathcal{V}_{\mathcal{R}}) \neq s_1] = 0.$$

Thus, using Lemma 9 and the fact that $\varepsilon \leq 1/400$ (so $2\sqrt{\varepsilon} \leq 0.1$), and taking p large enough so that $1/p(\lambda) \leq 0.1$, we have that (for infinitely many λ) with probability at least $1 - 2\sqrt{\varepsilon} - 1/p(\lambda) \geq 0.8$ over $(\tilde{\mathbf{M}}, \mathbf{P}, \mathcal{V}_{\mathcal{S}})$ generated in $\tilde{\Pi}$ all the following hold:

- $\Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 0] \geq 0.4$ and $\Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 1] \geq 0.4$.
- $\Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 0 \wedge \text{out}(\mathcal{V}_{\mathcal{R}}) \neq s_0] \leq 0.05$.
- $\Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P})} [\text{in}(\mathcal{V}_{\mathcal{R}}) = 1 \wedge \text{out}(\mathcal{V}_{\mathcal{R}}) \neq s_1] \leq 0.05$.

Returning to our analysis of \mathcal{R}^* , we thus see that, for infinitely many λ , with probability at least 0.8 it holds that $p'_0 \neq 0$, $p'_1 \neq 0$, and

$$\Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P})} [\text{out}(\mathcal{V}_{\mathcal{R}}) \neq s_0 \mid \text{in}(\mathcal{V}_{\mathcal{R}}) = 0] \leq 1/8 \quad \text{and} \quad \Pr_{\mathcal{V}_{\mathcal{R}}(\tilde{\mathbf{M}}, \mathbf{P})} [\text{out}(\mathcal{V}_{\mathcal{R}}) \neq s_1 \mid \text{in}(\mathcal{V}_{\mathcal{R}}) = 1] \leq 1/8.$$

So with probability at least $0.8 \cdot (1 - \frac{1}{8} - \frac{1}{8}) = 0.6$, the output of \mathcal{R}^* is equal to the input of \mathcal{S} .

4 Feasibility Results for Malicious, Stateless PUFs

We show that universally composable two-party computation is possible if the adversary is limited to creating *stateless* malicious PUFs. The core of our result is an unconditionally secure construction of a universally composable OT protocol in this model, described in Section 4.1. In Section 4.2 we briefly discuss how this protocol can be used to obtain the claimed result.




Figure 3: Our OT protocol. Following a preprocessing phase, the parties can execute N instances of oblivious transfer.

4.1 Universally Composable Oblivious Transfer

Our OT protocol adapts the protocol of Brzuska et al. [4], which was proven secure against attackers limited to honestly generated PUFs. It is easy to see that their protocol is not secure against attackers who can create malicious stateless PUFs. We show that having the sender create the PUF instead of the receiver is sufficient to obtain security in that case.

The protocol, described in Figure 3, consists of a preprocessing phase run by the sender \mathcal{S} and receiver \mathcal{R} , followed by a pre-determined number N of oblivious transfers. In the preprocessing phase, \mathcal{S} first creates a PUF PUF and sends it to the receiver. The receiver then chooses N uniform

values c_1, \dots, c_N and, for each one, computes $r_i := \text{PUF}(i, c_i)$. It then sends PUF back to \mathcal{S} . (Recall from Section 2 that we assume \mathcal{S} can verify that \mathcal{R} sent back the same PUF that \mathcal{S} created.)

When the parties want to execute the i th oblivious transfer, they proceed as follows. \mathcal{S} begins by sending a uniform value x . Then \mathcal{R} , with choice bit b , computes $v := c_i \oplus (b \cdot x)$ and sends v back to \mathcal{S} . The sender then computes $\hat{r}_0 := \text{PUF}(i, v)$ and $\hat{r}_1 := \text{PUF}(i, v \oplus x)$, and uses these values to “mask” its inputs s_0, s_1 . Since \mathcal{R} knows $\text{PUF}(i, v \oplus (b \cdot x)) = \text{PUF}(i, c_i)$, it can recover s_b .

We prove that this protocol is secure even if a malicious \mathcal{S} can create a malicious (but stateless) PUF. Our proof of security is unconditional, but assumes that the malicious party is limited to querying a PUF on polynomially many points. (As usual, we also assume an authenticated communication channel between \mathcal{S} and \mathcal{R} , but omit explicit mention of it.)

Theorem 10. *The protocol in Figure 3 securely realizes \mathcal{F}_{OT} in the $\mathcal{F}_{\text{PUF}}^0$ -hybrid model.*

Proof. For simplicity we assume $N = 1$, but the proof extends in a straightforward way for $N > 1$. The case where both \mathcal{S} and \mathcal{R} are honest is trivial, and so we focus on the case where one of the parties is corrupted.

Receiver is corrupted. We take the corrupted receiver \mathcal{R}^* to be the dummy adversary who simply forwards messages to/from the environment \mathcal{Z} . We describe an ideal-world simulator Sim that plays the role of the receiver in an interaction with the ideal-world OT functionality \mathcal{F}_{OT} . We then argue that no environment \mathcal{Z} can distinguish an interaction between an honest \mathcal{S} and Sim in the ideal world from an execution of our protocol between \mathcal{S} and \mathcal{R}^* in the $\mathcal{F}_{\text{PUF}}^0$ -hybrid world.

Sim is defined as follows:

- Sim simulates a copy of an honestly generated PUF PUF sent by \mathcal{S} . When \mathcal{Z} requests to make a query $\text{PUF}(i, c)$, a fresh random value is chosen and returned to \mathcal{Z} . At some point, \mathcal{Z} indicates that PUF should be returned to \mathcal{S} . (Recall from Section 2 that we assume \mathcal{S} can verify that its PUF is returned, and the protocol does not proceed until this is done.) Let Q denote the set of queries of the form $\text{PUF}(1, \star)$ made in this step.
- When the second phase of the protocol is initiated, Sim chooses a uniform value x and sends it to \mathcal{Z} (as if sent by \mathcal{S}). Next, \mathcal{Z} specifies a message v to be sent to \mathcal{S} . Then:
 - If $v \in Q$ and $v \oplus x \in Q$ then Sim aborts.
 - If $v \in Q$ but $v \oplus x \notin Q$ then Sim sets $b = 0$ and $\hat{r} = \text{PUF}(1, v)$ (i.e., \hat{r} is the value returned previously in response to the same query).
 - If $v \notin Q$ but $v \oplus x \in Q$ then Sim sets $b = 1$ and $\hat{r} = \text{PUF}(1, v \oplus x)$.
 - If $v \notin Q$ and $v \oplus x \notin Q$ then Sim sets $b = 1$ and chooses uniform \hat{r} .

Sim sends b to \mathcal{F}_{OT} and receives in return a bit s_b . It then sets $S_b := s_b \oplus \hat{r}$ and chooses uniform $S_{\bar{b}}$. Finally, it sends S_0, S_1 to \mathcal{Z} .

It is not hard to see that the simulation is perfect unless $v \in Q$ and $v \oplus x \in Q$. We show that this event, which we call **Bad**, occurs with negligible probability. Let $p = p(\lambda)$ be a bound on the number of queries to PUF made by \mathcal{Z} , so $|Q| \leq p$. Event **Bad** can only possibly occur if there exist $q_1, q_2 \in Q$ such that $q_1 \oplus q_2 = x$. Since there are at most p^2 pairs $q_1, q_2 \in Q$ and x is uniform, this means that the probability of **Bad** is at most $p^2/2^\lambda$, which is negligible.

Sender is corrupted. We take the corrupted sender \mathcal{S}^* to be the dummy adversary who simply forwards messages to/from the environment \mathcal{Z} . We describe an ideal-world simulator Sim that plays the role of the sender in an interaction with the ideal-world OT functionality \mathcal{F}_{OT} . We then argue that no environment \mathcal{Z} can distinguish an interaction between an honest \mathcal{R} and Sim in the ideal world from an execution of our protocol between \mathcal{R} and \mathcal{S}^* in the $\mathcal{F}_{\text{PUF}}^0$ -hybrid world.

Sim is defined as follows:

- We assume without loss of generality that \mathcal{Z} requests creation of a malicious PUF PUF^* , defined via a circuit C . Any queries by \mathcal{Z} to PUF^* , whether in this phase or the next phase, are answered in the obvious way by Sim . (Note that C may require oracle access to an honestly generated PUF, but that can be handled by Sim in the natural way.)

Sim simulates the sending of PUF^* to \mathcal{R} as well as its return.

- \mathcal{Z} specifies a message x (to be sent to \mathcal{R}), and in response Sim chooses a uniform v and sends it to \mathcal{Z} . Next, \mathcal{Z} specifies messages S_0, S_1 (to be sent to \mathcal{R}). At this point, Sim locally computes $\hat{r}_0 := \text{PUF}^*(1, v)$ and $\hat{r}_1 := \text{PUF}^*(1, v \oplus x)$, sets $s_b = S_b \oplus \hat{r}_b$ for $b \in \{0, 1\}$, and sends (s_0, s_1) to \mathcal{F}_{OT} .

It is immediate that the simulation is perfect. (Here we crucially rely on the fact that PUF^* is stateless, so it always returns the same response to the same challenge, and contains no information about \mathcal{R} 's queries.) \square

4.2 From UC Oblivious Transfer to UC Two-Party Computation

We observe that our UC oblivious-transfer protocol can be used to obtain UC two-party computation of any functionality. The main idea is to first construct a protocol with semi-honest security based on Yao's garbled-circuit protocol (using a PUF to implement a pseudorandom function), and then to apply the compiler of Ishai, Prabhakaran, and Sahai [12].

Semi-honest secure two-party computation. Lindell and Pinkas presented a proof for Yao's two-party secure-computation protocol [15]. They show how to instantiate the garbling part of the protocol with a private-key encryption scheme having certain properties. In addition, they show that any pseudorandom function is sufficient to instantiate such a private-key encryption scheme.

Our main observation is that we can replace the pseudorandom function with a PUF. Specifically, the circuit-generator in Yao's protocol will create a PUF PUF and define $F_k(x) = \text{PUF}(k, x)$. If the circuit generator is honest, then this defines a good pseudorandom function; if the circuit generator is malicious (and in particular if the PUF is malicious), it cannot violate privacy of the other party. (We remark that Brzuska et al. [4] also observed that PUFs can be used to implement a pseudorandom function, though in a different context and assuming honest PUFs.)

Since we showed in the previous section that \mathcal{F}_{OT} can be realized in the $\mathcal{F}_{\text{PUF}}^0$ -hybrid model, we thus have (we continue to omit explicit mention of an authenticated channel):

Theorem 11. *Let f be any functionality. There is a (constant-round) protocol that securely computes f for semi-honest adversaries in the $\mathcal{F}_{\text{PUF}}^0$ -hybrid model.*

Universally composable two-party computation. In the next step we apply the IPS compiler [12]. This is a black-box compiler that relies on protocols of the following types:

1. An “outer” multi-party computation protocol Π with security against a constant fraction of malicious parties.
2. An “inner” two-party protocol ρ , in the \mathcal{F}_{OT} -hybrid model, secure against semi-honest parties.

The result of the IPS compiler is a two-party protocol, in the \mathcal{F}_{OT} -hybrid model, that is universally composable for malicious adversaries.

In our setting, we instantiate the “outer” protocol with the BGW protocol [3], which is unconditionally secure in the presence of a malicious minority. We instantiate the “inner” protocol with the protocol from Theorem 11. Using Theorems 10 and 11, along with the UC composition theorem, we thus obtain the following result:

Theorem 12. *Let f be any functionality. There is a protocol that securely computes f for malicious adversaries in the $\mathcal{F}_{\text{PUF}}^0$ -hybrid model.*

Acknowledgments

Work of Nils Fleischhacker and Dominique Schröder was done in part while at Saarland University and while visiting the University of Maryland. Their work was supported by the German Federal Ministry of Education and Research (BMBF) through funding for the Center for IT-Security, Privacy, and Accountability (CISPA; see www.cispa-security.org). The visit of Nils Fleischhacker was supported by the Saarbrücken Graduate School of Computer Science funded by the German National Excellence Initiative, and the visit of Dominique Schröder was supported by NSF award #1223623. Work of Dominique Schröder was also supported by an Intel Early Career Faculty Honor Program Award. Work of Jonathan Katz was supported in part by NSF award #1223623. Work of Anna Lysyanskaya was supported by NSF awards #0964379 and #1012060.

References

- [1] Frederik Armknecht, Roel Maes, Ahmad-Reza Sadeghi, François-Xavier Standaert, and Christian Wachsmann. A formalization of the security features of physical functions. In *IEEE Symposium on Security and Privacy*, pages 397–412. IEEE, 2011.
- [2] Boaz Barak and Mohammad Mahmoody-Ghidary. Merkle puzzles are optimal—an $O(n^2)$ -query attack on any key exchange from a random oracle. *J. Cryptology*, 30(3):699–734, 2017.
- [3] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for noncryptographic fault-tolerant distributed computations. In *20th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1–10. ACM Press, 1988.
- [4] Christina Brzuska, Marc Fischlin, Heike Schröder, and Stefan Katzenbeisser. Physically cloneable functions in the universal composition framework. In *Advances in Cryptology—Crypto 2011*, volume 6841 of *LNCS*, pages 51–70. Springer, 2011.
- [5] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 136–145. IEEE, 2001. Full version available at <http://eprint.iacr.org/2000/067/>.

- [6] Ran Canetti and Marc Fischlin. Universally composable commitments. In *Advances in Cryptology—Crypto 2001*, volume 2139 of *LNCS*, pages 19–40. Springer, 2001.
- [7] Ran Canetti, Eyal Kushilevitz, and Yehuda Lindell. On the limitations of universally composable two-party computation without set-up assumptions. *J. Cryptology*, 19(2):135–167, 2006.
- [8] Ivan Damgård and Alessandra Scafuro. Unconditionally secure and universally composable commitments from physical assumptions. In *Advances in Cryptology—Asiacrypt 2013, Part II*, volume 8270 of *LNCS*, pages 100–119. Springer, 2013.
- [9] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM Journal on Computing*, 38(1):97–139, 2008.
- [10] Vipul Goyal, Yuval Ishai, Mohammad Mahmoody, and Amit Sahai. Interactive locking, zero-knowledge PCPs, and unconditional cryptography. In *Advances in Cryptology—Crypto 2010*, volume 6223 of *LNCS*, pages 173–190. Springer, 2010.
- [11] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 44–61. ACM Press, May 1989.
- [12] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer—efficiently. In *Advances in Cryptology—Crypto 2008*, volume 5157 of *LNCS*, pages 572–591. Springer, 2008.
- [13] Jonathan Katz. Universally composable multi-party computation using tamper-proof hardware. In *Advances in Cryptology—Eurocrypt 2007*, volume 4515 of *LNCS*, pages 115–128. Springer, 2007.
- [14] Stefan Katzenbeisser, Ünal Koçabas, Vladimir Rozic, Ahmad-Reza Sadeghi, Ingrid Verbauwhede, and Christian Wachsmann. PUFs: Myth, fact or busted? A security evaluation of physically unclonable functions (PUFs) cast in silicon. In *Cryptographic Hardware and Embedded Systems—CHES 2012*, volume 7428 of *LNCS*, pages 283–301. Springer, 2012.
- [15] Yehuda Lindell and Benny Pinkas. A proof of security of Yao’s protocol for two-party computation. *J. Cryptology*, 22(2):161–188, 2009.
- [16] Rafail Ostrovsky, Alessandra Scafuro, Ivan Visconti, and Akshay Wadia. Universally composable secure computation with (malicious) physically uncloneable functions. In *Advances in Cryptology—Eurocrypt 2013*, volume 7881 of *LNCS*, pages 702–718. Springer, 2013.
- [17] Ravikanth S. Pappu. *Physical One-Way Functions*. PhD thesis, Massachusetts Institute of Technology, 2001.
- [18] Ravikanth S. Pappu, Ben Recht, Jason Taylor, and Neil Gershenfeld. Physical one-way functions. *Science*, 297:2026–2030, 2002.
- [19] Ulrich Rührmair. Oblivious transfer based on physical uncloneable functions. In *Trust and Trustworthy Computing*, volume 6101 of *LNCS*, pages 430–440. Springer, 2010.

- [20] Ulrich Rührmair, Stefan Katzenbeisser, and Heike Busch. Strong PUFs: Models, constructions, and security proofs. In *Towards Hardware-Intrinsic Security*, pages 79–96. Springer, 2010.
- [21] Ulrich Rührmair and Marten van Dijk. PUFs in security protocols: Attack models and security evaluations. In *IEEE Symposium on Security and Privacy*, pages 286–300. IEEE, 2013.
- [22] Marten van Dijk and Ulrich Rührmair. Physical unclonable functions in cryptographic protocols: Security proofs and impossibility results. Cryptology ePrint Archive, Report 2012/228, 2012.