Laconic Conditional Disclosure of Secrets and Applications

Nico Döttling* Sanjam Garg†, Vipul Goyal‡, Giulio Malavolta§
*CISPA Helmholtz Center for Information Security, Saarbrücken, Germany
†University of California Berkeley, Berkeley, USA
‡Carnegie Mellon University, Pittsbugh, USA
§Simons Institute for the Theory of Computing, Berkeley, USA

***Abstract*—In a Conditional Disclosure of Secrets (CDS) a verifier $\mathbf{V}$ wants to reveal a message $\mathbf{m}$ to a prover $\mathbf{P}$ conditioned on the fact that $\mathbf{x}$ is an accepting instance of some NP-language $\mathcal{L}$. An honest prover (holding the corresponding witness $\mathbf{w}$) always obtains the message $\mathbf{m}$ at the end of the interaction. On the other hand, if $\mathbf{x} \notin \mathcal{L}$ we require that no PPT $\mathbf{P}^*$ can learn the message $\mathbf{m}$. We introduce *laconic CDS*, a two round CDS protocol with *optimal* computational cost for the verifier $\mathbf{V}$ and *optimal* communication cost. More specifically, the verifier's computation and overall communication grows with $\mathsf{poly}(|\mathbf{x}|, \lambda, log(\mathbf{T}))$, where λ is the security parameter and $\mathbf{T}$ is the verification time for checking that $\mathbf{x} \in \mathcal{L}$ (given $\mathbf{w}$). We obtain constructions of laconic CDS under standard assumptions, such as CDH or LWE.**

Laconic CDS serves as a powerful tool for *maliciousifying* semi-honest protocols while preserving their computational and communication complexities. To substantiate this claim, we consider the setting of non-interactive secure computation: Alice wants to publish a *short* digest corresponding to a private *large* input $\mathbf{x}$ on her web page such that (possibly many) Bob, with a private input $\mathbf{y}$, can send a *short* message to Alice allowing her to learn $\mathcal{C}(\mathbf{x}, \mathbf{y})$ (where $\mathcal{C}$ is a public circuit). The protocol must be reusable in the sense that Bob can engage in arbitrarily many executions on the same digest. In this context we obtain the following new implications.

1) ***UC Secure Bob-optimized 2PC:*** **We obtain a UC secure protocol where Bob's computational cost and the communication cost of the protocol grows with $\mathsf{poly}(|\mathbf{x}|, |\mathbf{y}|, \lambda, \mathbf{d})$, where $\mathbf{d}$ is the depth of the computed circuit $\mathcal{C}$.**
2) ***Malicious Laconic Function Evaluation:*** **Next, we move on to the setting where Alice's input $\mathbf{x}$ is large. For this case, UC secure protocols must have communication cost growing with $|\mathbf{x}|$. Thus, with the goal of achieving better efficiency, we consider a weaker notion of malicious security. For this setting, we obtain a protocol for which Bob's computational cost and the communication cost of the protocol grows with $\mathsf{poly}(|\mathbf{y}|, \lambda, \mathbf{d})$, where $\mathbf{d}$ is the depth of the computed circuit $\mathcal{C}$.**

***Keywords*-Cryptography; Secure Computation.**

I. Introduction

Consider the following setting: Alice would like to publish a digest h corresponding to her private data x on her web page. Next, Bob (with a private input y) would like to send a *short* message to Alice, so that she learns $\mathcal{C}(x, y)$ for some public circuit $\mathcal{C}$. As is standard in secure computation, Alice and Bob want to achieve this result while keeping their respective inputs hidden from each other. Analogous to Bob's message, Charlie or any other party could reuse h provided by Alice on her web page. This *non-interactive* nature of Alice's message makes this setting highly desirable and has been extensively studied in cryptography [1]. This primitive is commonly referred to as *non-interactive secure computation*.

Bob-optimization. Unfortunately, traditional cryptographic techniques for realizing the above task require Bob's computational complexity to grow with $|\mathcal{C}|$ and $|x|$ (in addition to y), which is undesirable in several settings. In particular, since only Alice learns the output of the computation, Bob might find it unfair that he has to perform the computation of $|\mathcal{C}|$. In other words, Bob might be happy to see Alice learn the output of the computation, but may not be willing to take on doing the computation himself. For example, consider the case where Alice's input is her genomic sequence and Bob would like to help Alice to learn how likely she is to have cancer while keeping the logic that he is using to make predictions secret.

Making the problem more challenging, computational constraints on Bob imply additional constraints on the sizes of Alice's digest and Bob's message, which must now also be small. Constructing such Bob-optimized protocols has been the focus of several recent results such as laconic function evaluation [2] and laconic oblivious transfer [3]. Our setting is inspired by these works. In particular, prior work on laconic function evaluation [2] also considers the problem of non-interactive secure computation, which is computationally optimized for Bob. However, these results are limited to the semi-honest setting and need non-falsifiable assumptions [4], [5], [6], [7], [8] to upgrade security to the malicious setting. This brings us to the following question:

2575-8454/19/$31.00 ©2019 IEEE
DOI 10.1109/FOCS.2019.00046

Can we realize maliciously secure Bob-optimized non-interactive secure computation from standard assumptions?

Bob-optimization when Alice has a large input. Next, we consider the more demanding setting, where Alice's input x itself might be very *large*. As before, Bob might be happy to see Alice learning the output of the computation, but may find reading x himself prohibitively expensive. This could be especially prohibitive if Bob interacts with multiple users, all playing the role of Alice. For example, the FBI (playing as Alice) could hold a huge database of passenger names on the no-fly list. An an airline company (playing as Bob) might want to help the FBI check if any of the passengers on one of its flights is on the list. Here the airline company would prefer to perform computation independent of the FBI's database. This is a special case of the very important and well-studied problem of private set intersection (PSI) [9]. This brings us to the following question:

Can we realize maliciously secure Bob-optimized non-interactive secure computation when Alice has a large input from standard assumptions?

Reusablity of Alice's Digest. We allow multiple Bob's to be able to reuse Alice's digest for multiple computations. In this setting, a corrupt Bob could provide Alice with an ill-formed message and use Alice's output (that it may learn in some way) to cheat Alice or to learn something about Alice's input x. Thus, we aim for security against a malicious Bob that has access to Alice's outputs on its prior interactions.

Reverse Delegation. One can view our setting as a reverse delegation scheme. In a delegation scheme, a user outsources its computation to an untrusted server with the goal of learning the output (possibly while even keeping it private). In contrast, our applications allow a user (Bob) to delegate its computation to the untrusted server (Alice) while also letting it learn the output of the computation, and nothing beyond that.

A. Our Results

In this work we introduce the notion of laconic conditional disclosure of secrets (CDS). In a CDS protocol a verifier V wants to reveal a message m to a prover P conditioned on the fact that x is an accepting instance of some NP-languge $\mathcal{L}$. An honest prover (holding the corresponding witness w) always obtains the message m at the end of the interaction. On the other hand if $x \notin \mathcal{L}$ we require that no (possibly corrupted) P^* can learn the message m. CDS can be seen as the two round analog of witness encryption [10]. The new constraint that we impose is that the interaction has to be *laconic* in the sense that it should not depend on the size of the parties inputs, and in particular on the size of the witness w. Additionally, the verifier computation must be much smaller than recomputing the relation. More specifically, verifier's computation and overall communication grows with $\mathsf{poly}(|x|, \lambda, log(T))$, where λ is the security parameter and T is the verification time for checking that $x \in \mathcal{L}$ (given w).

Next, we obtain constructions of laconic CDS in the common reference string model under standard assumptions, such as CDH or (standard) LWE. Laconic CDS serves as a powerful tool for *maliciousifying* semi-honest protocols while preserving their computational and communication complexities. To demonstrate this power, we show application of laconic CDS to the above-mentioned settings:

1) *Bob-optimized 2PC:* We obtain a UC secure protocol where Bob's computational cost and the communication cost of the protocol grows with $\mathsf{poly}(|x|, |y|, \lambda, d)$, where d is the depth of the computed circuit C. This protocol achieves UC-security [11] and security is based on the LWE assumption.
2) *Malicious Laconic Function Evaluation:* Next, we consider the more demanding setting where Alice's input x is large. For this case, UC secure protocols must have communication cost growing with $|x|$. Thus, with the goal of achieving better efficiency, we consider a weaker notion of malicious security which we call *context security*. Context security allows us to rewind parts of the execution while at the same time providing some form of composability and subsumes the standard notion of indistinguishability-based security. For this setting, we obtain a protocol for which Bob's computational cost and the communication cost of the protocol grows with $\mathsf{poly}(|y|, \lambda, d)$, where d is the depth of the computed circuit C. The security of this construction is based on LWE as well. This result can be seen as the malicious variant of the recent work on laconic function evaluation [2].

The above applications demonstrate the power of laconic CDS. We see laconic CDS as a natural primitive and expect it to have other applications. As another example, laconic CDS allows a resource-constrained client to delegate the computation of a large circuit to an untrusted worker and to condition the payment of the worker on the fact that the provided output is the correct one. This does not achieve the notion of verifiable com-

putation in the traditional sense, since the client cannot *explicitly* check that the output is correct. However in certain scenarios, e.g., mining in cryptocurrencies, one is interested in rewarding under the condition that *some* computation has been performed. Then miners are free to give the wrong output and not claim the reward. However, there is no clear incentive for that.

B. Technical Overview

To understand and motivate our techniques it is instructive to discuss a few plausible approaches for constructing laconic CDS and highlight the barriers that they encounter. Below we start with such approaches.

Why Known Techniques Fail. If we were to omit the laconic constraint, then CDS admits a very simple two-round protocol based on two-round oblivious transfer (OT) and Yao's garbled circuits [12]: The prover encodes the binary representation of its witness as the choice bits in several parallel repetitions of the OT. The verifier garbles the circuit that hardcodes the statement x and the message m and returns m if and only if $\mathcal{R}(\cdot, x) = 1$. Then it sends the garbled circuit together with the second message of the OT on input each pair of input label. The prover can locally recover the labels corresponding to its witness and evaluate the circuit learning nothing beyond its output. It is clear that such a solution satisfies our security requirements but falls short in achieving laconic communication since the full circuit encoding of the NP-relation is exchanged between the two parties. At a first glance this seems to be inherent to garbled-circuit based solutions since secure two-party computation with low communication complexity usually requires more sophisticated tools, such as fully-homomorphic encryption [13].

Another plausible angle to attack the problem is to resort to techniques from the field of succinct arguments [4], [5]. Then the prover could simply augment a succinct proof with a public key and the verifier would then encrypt the message if the proof correctly verifies. Unfortunately known schemes require at least three rounds of interaction (four without assuming a setup) or rely on non-falsifiable assumptions [6], [7], [8]. Constructing succinct arguments from standard assumptions in less than three rounds seems to hit a roadblock: It has been shown [14] that non-interactive succinct arguments cannot be based on falsifiable assumptions, at least in a black-box sense.

Our Solution in a Nutshell. Our starting point is the classical garbled circuit-based solution. A closer look to the source of inefficiency reveals that there are two major challenges to overcome in order to achieve our goal: (1) We need to remove the linear dependency of the OT with respect to the size of the witness and (2) we cannot recompute the full blown NP-relation in the garbled circuit. Our first insight is to bypass the first obstacle using laconic OT [3]. A laconic OT allows one to hash a long database into a small digest, then the sender can compute the second message of an OT where the choice bit is set to be the value at an arbitrary location of the database. The important message here is that the communication complexity is independent of the size of the database (in our case the witness w). This primitive alone allows us to compress the size of the first message of the CDS. However, laconic OT alone does not buy us anything for the complexity of the second message. Overcoming the second obstacle requires us to borrow techniques from the domain of succinct arguments. The observation here is that checking the validity of an NP-instance does not necessarily require one to recompute the corresponding relation: Probabilistically checkable proofs (PCP) [15] encode a witness into a (longer) string such that the membership of the statement can be probabilistically verified by querying a constant amount of bits. Such verification algorithms are inherently erroneous, but the gap can be made negligibly small via standard amplification techniques.

This tool gives us the right leverage for a candidate laconic CDS construction: A prover can now hash the PCP encoding of its witness via a laconic OT and send the corresponding digest to the verifier. The latter then samples a few random locations and garbles a circuit that takes as input the bits of the PCP encoding at such locations and returns m if the PCP verifier accepts. Then it sends the garbled circuit in plain and the input labels via the laconic OT. This allows the prover to recover the labels corresponding to the bits of the PCP encoding, evaluate the circuit, and eventually recover the message. This introduces a negligible soundness gap, which does not make a difference in our settings as we rely anyway on computational assumptions. The stretch in the size of the witness is also not a problem since the communication complexity of a laconic OT is independent of the size of the database (in our case the PCP encoding).

While this construction seems to solve all problems at once, a closer look to the building blocks reveals the dependency on a *maliciously secure* laconic OT. Currently, the only known way to construct laconic OT resilient against active attackers is to compile a semi-honest one with succinct arguments. Due to its compressing nature, bypassing the usage of non-falsifiable

assumptions in laconic CDS might then appear to be out of reach of current techniques. However, we have now reduced the problem of constructing laconic CDS to that of building laconic OT: In this work we show how to construct maliciously secure laconic OT assuming the hardness of CDH or (standard) LWE with polynomial modulo-to-noise ratio. Most of the technical innovations of this work, and the remainder of this section, are devoted to instantiating malicious laconic OT from standard assumptions. Looking ahead, this will allow us to construct laconic CDS, which in turn will be used as the main technical component to build malicious LFE and UC-secure Bob-optimized 2PC. An outline of our results is given in Figure 1.

Towards Malicious Laconic OT. Before we delve into the actual instantiations of malicious laconic OT we further simplify the problem by lowering the efficiency and security requirements for a laconic OT.

Our starting point is a recent work by Döttling, Garg, Hajiabadi, Masny and Wichs [16] which constructs maliciously secure oblivious transfer from search assumptions. The main idea in [16] is to start with a very simple notion of security against malicious receivers and carefully bootstrap this notion to standard simulation based notions. The weakest notion considered in [16] is called *elementary OT*, and our basic notion of malicious laconic OT extends this notion to the setting of laconic OT. We refer to this primitive as *weak malicious* laconic OT (when it is clear from the context we drop the term malicious). Jumping ahead, we will then show, building on techniques developed in [16] how to generically upgrade a laconic OT that meets these conditions to a fully efficient and fully secure one. Concretely, we aim for the following guarantees.

1) Weak Functionality: The sender algorithm no longer takes as input two messages (m_0, m_1) to transfer but instead generates two fresh random strings (k_0, k_1) together with the sender output c. This can be seen as the analog of key-encapsulation mechanism, where (k_0, k_1) are then used as the session keys to transfer the desired messages.
2) Weak Security: We require that, given the sender message c, no efficient polynomial-time algorithm can output both k_0 and k_1 at the same time. This does not guarantee, for instance, that an attacker cannot output the first half of k_0 and the second half of k_1. This requirement is in fact identical to the notion of elementary sender security considered in [16].
3) Weak Efficiency: We consider a laconic OT where the only efficiency constraint is that the receiver message is 2-to-1 compressing. That is, we do not impose any bound on the size of the setup string and of the sender message, which are potentially as long as the database.

Pitfalls of Known Constructions. A natural question to ask is whether existing constructions of semi-honest laconic OT already satisfies any meaningful notion of security in presence of a corrupted receiver. To exemplify the issues that arise, we briefly recall a simplified version of the hash-encryption construction of [17]: The public parameters consist of a matrix of uniformly sampled group elements

$$crs = \begin{pmatrix} g_1^{(0)}, \ldots, g_m^{(0)} \\ g_1^{(1)}, \ldots, g_m^{(1)} \end{pmatrix} \leftarrow_{\$} \mathbb{G}^{2\times m}$$

and the hash of a database D is computed as

$$d = \prod_{i=1}^{m} g_i^{(D[i])}.$$

To generate two keys for a position L, one samples two random random integers (r, s), computes the ciphertext c as the concatenation of the matrices

$$\begin{pmatrix} \left(g_1^{(0)}\right)^r, & \ldots, & \left(g_L^{(0)}\right)^r, & \ldots, & \left(g_m^{(0)}\right)^r \\ \left(g_1^{(1)}\right)^r, & \ldots, & 1, & \ldots, & \left(g_m^{(1)}\right)^r \end{pmatrix}$$

and

$$\begin{pmatrix} \left(g_1^{(0)}\right)^s, & \ldots, & 1, & \ldots, & \left(g_m^{(0)}\right)^s \\ \left(g_1^{(1)}\right)^s, & \ldots, & \left(g_L^{(1)}\right)^s, & \ldots, & \left(g_m^{(1)}\right)^s \end{pmatrix},$$

then sets $k_0 = d^r$ and $k_1 = d^s$. Security stems from the fact that if $D[L] \neq b$, then the receiver has to find a different linear combination of elements that yields the same d in order to compute the corresponding key. This however crucially relies on the fact that the receiver always chooses at least one group element per column of the crs. Consider the following attack where d is set to be $d = g_1^{(0)}$. Then the receiver can recover the keys for both $b = \{0, 1\}$, when encrypting with respect to some $L \neq 1$. This is because the element $\left(g_1^{(0)}\right)^r$ is always given if $L \neq 1$, then k_0 and k_1 are simply the first element of each matrix. This breaks any meaningful notion of security and forces us to rethink even the most basic building block of the primitive.

Construction from CDH. Our approach diverges from prior works and takes a fresh look at the problem. In our scheme the public parameters consist of a vector of group elements sampled uniformly at random

$$crs = (g_1, \ldots, g_m) \leftarrow_{\$} \mathbb{G}^m$$

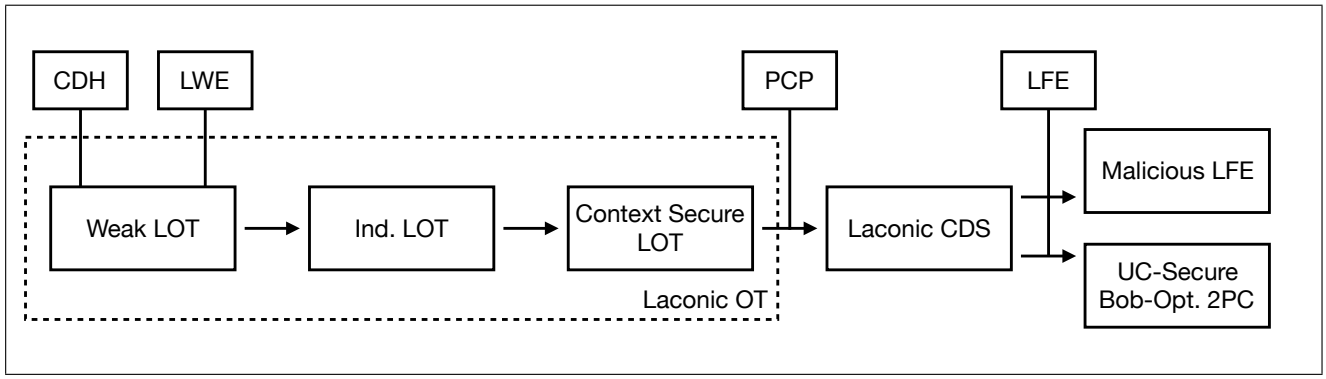

Figure 1: Overview of our Results

and the hash of a database D is defined as

$$d = \prod_{i=1}^{m} g_i^{D[i]}.$$

That is the i-th element of the vector is included in the product if $D[i] = 1$ and skipped if $D[i] = 0$. The subtle difference with respect to previous approaches lies in the fact that even setting $d = g_1$ gives us a perfectly valid hash, which corresponds to the pre-image $1\|0^{m-1}$. To generate two session keys (k_0, k_1) for a location L, the sender samples a random integer r and sets the ciphertext to

$$c = \left(g_1^r, \ldots, g_{L-1}^r, g_{L+1}^r, \ldots, g_m^r\right),$$

i.e., all powers are given except for g_L^r. We now want to define k_0 and k_1 in such a way that only $k_{D[i]}$ can be recovered using a pre-image D of d. To this end we set

$$k_0 = d^r \text{ and } k_1 = d^r / g_L^r.$$

Observe that if $D[L] = 0$, then c contains enough information to recompute

$$k_0 = d^r = \prod_{i=1}^{m} g_i^{D[i]\cdot r},$$

since g_L^r would be excluded from the product anyway. On the other hand if $D[L] = 1$, then one can only recompute

$$k_1 = d^r / g_L^r = \prod_{i=1, i\neq L}^{m} g_i^{D[i]\cdot r}.$$

To get an intuition why the scheme is secure for any (possibly maliciously generated) image d, observe that outputting both k_0 and k_1 also reveals $k_0/k_1 = g_L^r$, regardless of the value of d. However, g_L^r was not given as part of the ciphertext and therefore predicting it, given only c and crs, is as hard as solving a random instance of the CDH problem.

Construction from LWE. A CDH-based construction is not entirely satisfactory in terms of assumptions since most of the applications of malicious laconic OT (and consequently of laconic CDS) are only known under the hardness of lattice-related problems. If we were to plug in our machinery we would introduce an additional number-theoretic assumption (i.e., CDH), which is not ideal. An additional reason to look into lattice-based schemes is that current cryptanalytic techniques fail even in the quantum settings, whereas there are polynomial-time quantum algorithms to solve discrete logarithm-related problems. For these reasons, we turn our attention to constructing weak malicious laconic OT from the hardness of the LWE problem. The basic idea of our scheme is conceptually simple, but the noisy nature of the LWE problem introduces some complications. The public parameters consist of a single matrix

$$crs = \mathbf{A} \in \mathbb{Z}_q^{n\times m}$$

where m is the size of the database and n is a parameter that governs the hardness of the LWE problem. Hashing a database $D \in \{0,1\}^m$ (parsed as a column vector) is a simple multiplication

$$\mathbf{A}D = \mathbf{d} \in \mathbb{Z}_q^n$$

that is, we take a linear combination of the columns of $\mathbf{A}$ depending on the bits of D. For a given location L, the sender algorithm samples a column vector $\mathbf{s} \leftarrow_{\$} \mathbb{Z}_q^n$ and computes a noisy inner product with each column of $\mathbf{A}$, except for the L-th one. The sender message c consists of the set

$$\left\{\mathbf{s}^T\mathbf{a}_1 + e_1, \ldots, \mathbf{s}^T\mathbf{a}_m + e_m\right\} \setminus \left\{\mathbf{s}^T\mathbf{a}_L + e_L\right\}$$

where the noise vector is sampled from a discrete Gaussian and has small norm. As before we set the keys in such a way that it is easy to compute one, but

it is hard to compute both. This is done by setting

$$k_0 = \mathbf{s}^T\mathbf{d} \text{ and } k_1 = -\mathbf{s}^T\mathbf{d} + \mathbf{s}^T\mathbf{a}_L + e_L.$$

It is important to observe that $k_0 + k_1 = \mathbf{s}^T\mathbf{a}_L + e_L$, which is an LWE sample but it is not among those included in c. This suggests that outputting both values of k_0 and k_1 is equivalent to predicting an LWE sample and it is going to be our central leverage to show the security of the scheme. However, the more challenging part is how to ensure that the receiver is able to recover $k_{D[L]}$. In contrast to our previous scheme, the naive strategy fails to recover the key due to the presence of the noise. In fact

$$\sum_{i=1, i\neq L}^{m} c_i \cdot D[i] = k_{D[L]} + \sum_{i=1, i\neq L}^{m} e_i \cdot D[i] = k_{D[L]} + \tilde{e}.$$

To overcome this issue we exploit the fact that $\tilde{e}$ is relatively small when compared to $k_{D[L]}$. We then partition $\mathbb{Z}_q$ in c-many intervals and we set the actual session key to be k_b (as computed above) rounded at the closest multiple of q/c. The key property of this rounding function is that it is resilient to small perturbations, which means that

$$\left\lfloor k_{D[L]} + \tilde{e} \right\rceil_c = \left\lfloor k_{D[L]} \right\rceil_c$$

with high probability, for a small enough $\tilde{e}$. This strategy introduces a few issues, among others the fact that the output of the rounding function is not long enough to argue about unpredictability (i.e., it can be randomly guessed with high enough probability). Fortunately, this issue can be easily circumvented with standard amplification techniques.

Indistinguishability Laconic OT. As discussed before, our notion of security only guarantees that the adversary cannot output both k_0 and k_1 in their entirety but it does not guarantee that there exists a consistent bit b for which the adversary can always guess k_b and never $k_{b\oplus 1}$. For example, it could be possible that for a fixed value of d, the adversary is able to predict k_0 for half of the support of c and k_1 for the complement. To fix this issue, follow the blueprint of [16] and use techniques developed in the context of hardness amplification of puzzles [18]. The key idea of this step is to amplify the success probability of the adversary via parallel repetitions: By setting the new key to be the concatenation of many independent instances, with high probability at least one of elements will belong to the *wrong* partition of the support of c. Thus with high enough probability, the adversary is forced to simultaneously predict two values (k_0, k_1).

Once this obstacle is surpassed, turning our weak laconic OT into a fully functional one involves rather standard tools: First we turn the search problem into a decision one using Goldreich-Levin hardcore predicate [19]. At this point we can give as input to the sender algorithm any message pair (m_0, m_1) and implement the standard OT functionality by augmenting c with $k_0 \oplus m_0$ and $k_1 \oplus m_1$. It is easy to see that at least one of the two messages must be hidden to the eyes of the adversary as the key $k_{D[L]\oplus 1}$ acts as a one-time pad. At this point, our security definition looks as follows: We define two experiments (Exp_0 and Exp_1) where the adversary is allowed to choose a hash d, two messages (m_0, m_1), and a location L, then the challenger flips a coin b. If $b = 0$ then the ciphertext c is honestly computed via the sender algorithm on input (L, m_0, m_1), otherwise m_0 (m_1 for Exp_1) is replaced with a uniformly sampled message. The guarantee is that the adversary cannot succeed in both experiments with non-negligible probability.

This gives us a more intuitive security notion but it is still not sufficient for constructing laconic CDS. The reason is that we have no idea whether the adversary is going to be able to succeed in Exp_0 or Exp_1, for a given value of L. This information is needed by the simulator when using laconic OT in conjunction with garbled circuits: In a reduction against the security of the garbling scheme we need to know in which position we should plug the challenge label, to correctly simulate the view of the adversary. Even worse, determining whether the adversary is successful in Exp_0 or Exp_1, for all locations L, corresponds to extracting the whole database D! Recall that the hash d of a laconic OT is compressing, therefore its pre-image is not even information-theoretically determined, let alone efficiently computable.

Context-Secure Laconic OT. The final challenge towards constructing malicious laconic OT boils down to extracting the pre-image of any given (possibly maliciously computed) hash. As discussed before, the hash does not even determine the pre-image in an information theoretic sense, so trivial solutions are not applicable. An additional complication is that the protocol is non-interactive, i.e., the receiver outputs a single message and goes offline. Therefore rewinding the receiver also does not seem to help. Of course one could just *assume* the existence of an extractor, which would however place our solution within the category of non-falsifiable assumptions.

We address this problem by leveraging the distinguisher-dependent simulation technique recently

developed [20]. We introduce a new security notion called *context security* to provide a versatile toolkit which allows to deploy distinguisher-dependent simulation in much more complex settings. While on a technical level we use many of the same techniques as [16], context security will allow us to use the distinguisher-dependent simulation technique to its full effect.

We will use distinguisher-dependent simulation in the following way to extract the input of a malicious receiver. A careful look to the indistinguishability-based definition shows that, given black-box access to a malicious receiver, we can determine the bit $D[L]$, by approximating the success probability in Exp_0 and Exp_1: For a given location L, the experiment where the adversary is successful must correspond to the value of $D[L]$! Iterating over all locations, we can recover the complete database.

Our main insight is that this technique can be applied to a setting where the overall communication between receiver and sender is too small to information-theoretically specify the input of the receiver. This is in contrast to previous uses of this technique in [20], [16].

Somewhat counterintuitively, we are able to extract the full pre-image of a short hash in polynomial time without resorting to non-falsifiable assumptions. Our new notion of context *context security* is crafted to precisely characterize the security guarantees we can achieve via distinguisher dependend simulation.

The high level idea of context security is that protocols must remain secure against any context, where a context is defined as a (possibly corrupted) receiver and a distinguisher that takes as input the sender's message. Security is defined in terms of the existence of an extractor and a simulator: The extractor (with oracle access to the distinguisher) recovers the original input of the receiver and gives enough information to the simulator to simulate the sender message, where the simulator only interacts with the ideal functionality. Additionally, the runtime of the extractor is independent of that of the corrupted receiver (but depends on that of the distinguisher). This enables a meaningful notion of composability since the protocol is required to be remain secure for any context. Since the extractor depends on the distinguisher, context-security does not achieve the same strong guarantees of universal composability (UC) [11]. However, in contrast with UC, a protocol can be context-secure even when its transcript does not determine the inputs of the parties in an information theoretic sense. Thus we view context-security as a meaningful relaxation of UC and a versatile tool, instrumental to construct laconic CDS. Finally, we remark that context security implies the standard notion of indistinguishability-based secure computation.

Using ideas developed in [16] we show that any indistinguishability laconic OT is also context secure. The argument crucially relies on rewinding the distinguisher and measuring its success-probability.

Efficient Laconic OT. Equipped with a context-secure laconic OT, we can show a bootstrapping theorem in the same spirit as [3]: Given any 2-to-1 compressing context-secure laconic OT we can construct an equally efficient context-secure laconic OT where the size of the hash (and of the common reference string) is independent of the size of the database. The transformation is identical to that of [3] and works by hashing the database into a Merkle tree and using a chain of garbled circuit to access the leaf corresponding to the location L. However, the argument is fundamentally different since we cannot assume that the hash is honestly computed. The usefulness of the notion of context security is demonstrated by our bootstrapping theorem. We crucially rely on the recursive application of context extractors to extract an entire Merkle-tree. A critical aspect in this step which avoids an exponential blowup is that the runtime of the context-extractor is independent of the computation of the receiver.

C. Applications

We now discuss how laconic CDS can be used as a *maliciousifier* for two-round protocols: A two round protocol consists of a message m_1 from the receiver to the sender and a message m_2 from the sender to the receiver. At the end of the interaction, the receiver performs some computation to retrieve the output of the protocol. Using laconic CDS we can turn a semi-honest two-round protocol into a malicious one (for a corrupted receiver) by augmenting m_1 with the first message of a laconic CDS certifying that m_1 is well formed. Then the sender computes the second message of a laconic CDS where the secret is set to m_2. That is, the receiver will learn m_2 if and only if m_1 was well-formed. The new properties that our transformation enables are:

1) The communication complexity of the malicious protocol is identical to that of the semi-honest one.
2) The computational complexity of the sender is unchanged.

This comes particularly useful when the input and the computation of the receiver are particularly burdensome. As an example, a laconic CDS allows us to lift non-interactive secure computation for RAM programs [3] to the malicious settings in a very natural way. Another

interesting scenario is when the above transformation is applied to laconic function evaluation (LFE). In this case we obtain the first maliciously secure Bob-optimized non-interactive secure computation protocol where the communication complexity grows only with the depth of the circuit and with the size of Bob's input. In other words, the malicious protocol is (asymptotically) as efficient as the underlying semi-honest LFE. We can even lift the security to the UC-settings, however at the cost of the communication growing with the size of the receiver's input, which is unavoidable.

Delegating Computation. A less orthodox usage of a laconic CDS is in the context of non-interactive delegation of computation: A client wants to delegate the computation of a large circuit to an untrusted worker and wants some insurance that the output given by the worker is the correct one. The worker performs the computation and obtains the output z, then computes a laconic CDS that certifies that z is indeed the correct output. On input the first message of the laconic CDS and z, the client conditions the payment of the worker on the fact that z is computed correctly. This protocol is not verifiable in the traditional sense since the client cannot explicitly check that z is indeed the correct output. However we can envision scenarios where this is not a limitation. As an example, miners of cryptocurrencies often aggregate in pools, where each peer is rewarded basing on the amount of computation it performed. In this case it is not critical to verify that the output is correct but we are mostly interested in the fact that *some* large circuit has been computed. The miners could of course provide the wrong output and not claim the reward, but it is unclear what would be the incentive to do that.

II. Definitions

We denote by $\lambda \in \mathbb{N}$ the security parameter. We we say that a function $negl$ is negligible if it vanishes faster than any polynomial. Given a set S, we denote by $s \leftarrow\!\!\$\, S$ the uniform sampling of an element from S. We say that an algorithm is PPT if it can be implemented by a probabilistic Turing machine running in time polynomial in λ. We introduce two useful inequalities.

Theorem 1 (Hoeffding Inequality)**.** *Let $(X_1, \ldots, X_N)$ be independent and identically distributed random variables in $[0,1]$ with expectation $\mathsf{E}[\bar{X}]$. Then it holds that*

$$\Pr\left[\left|\frac{1}{N}\sum_{i=1}^{N} X_i - \mathsf{E}[\bar{X}]\right| > \delta\right] \leq 2e^{-2N\delta^2}.$$

Lemma 1 (Markov Inequality for Advantages [16])**.** *Let $A(Z)$ and $B(Z)$ be two random variables depending on a random variable Z and potentially additional random choices. Assume that*

$$\left|\Pr_Z[A(Z)=1] - \Pr_Z[B(Z)=1]\right| \geq \varepsilon \geq 0.$$

Then

$$\Pr_Z\left[|\Pr[A(Z)=1] - \Pr[B(Z)=1]| \geq \varepsilon/2\right] \geq \varepsilon/2.$$

A. Intractable Problems

In the following we introduce some hard problems in cryptography that are going to be useful for our work.

Computational Diffie-Hellman. We recall the search version of the classical computational Diffie-Hellman (CDH) problem [21]. Let $\mathcal{G}$ be a (prime-order) group generator that takes as input the security parameter and 1^λ and outputs $(\mathbb{G}, p, g)$, where $\mathbb{G}$ is the description of a multiplicative cyclic group, p is the order of the group, and g is a generator of the group.

Definition 1 (Computational Diffie-Hellman Assumption)**.** *We say that $\mathcal{G}$ satisfies the CDH assumption (or is CDH-hard) if for any PPT adversary $\mathcal{A}$ it holds that*

$$\Pr[\mathcal{A}(\mathbb{G}, p, g, g^{a_1}, g^{a_2}) = g^{a_1 a_2}] = negl(\lambda)$$

where $(\mathbb{G}, p, g) \leftarrow\!\!\$\, \mathcal{G}(1^\lambda)$ and $(a_1, a_2) \leftarrow\!\!\$\, \mathbb{Z}_p$.

Learning with Errors. The learning with errors (LWE) problem was introduced by Regev [22]. In this work we exclusively use the decisional version, since there exists a well known reduction to the search variant of the problem.

Definition 2 (Learning with Errors Assumption)**.** *The $\mathrm{LWE}_{n,\tilde{n},q,\chi}$ problem, for $(n, \tilde{n}, q) \in \mathbb{N}$ and for a distribution χ supported over $\mathbb{Z}$, is to distinguish between the distributions $(\mathbf{A}, \mathbf{sA}+\mathbf{e} \mod q)$ and $(\mathbf{A}, \mathbf{u})$, where $\mathbf{A} \leftarrow\!\!\$\, \mathbb{Z}_q^{n\times\tilde{n}}$, $\mathbf{s} \leftarrow\!\!\$\, \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow\!\!\$\, \chi^{\tilde{n}}$, and $\mathbf{u} \leftarrow\!\!\$\, \mathbb{Z}_q^{\tilde{n}}$. Then, the $\mathrm{LWE}_{n,\tilde{n},q,\chi}$ assumption is that the two distributions are computationally indistinguishable.*

The assumption is consider to hold for any $\tilde{n} = poly(n, log(q))$ and we denote this problem by $\mathrm{LWE}_{n,q,\chi}$. As shown in [22], [23], the $\mathrm{LWE}_{n,q,\chi}$ problem with χ being the discrete Gaussian distribution with parameter $\sigma = \alpha q \geq 2\sqrt{n}$ (i.e. the distribution over $\mathbb{Z}$ where the probability of x is proportional to $e^{-\pi(|x|/\sigma)^2}$), is at least as hard as approximating the shortest independent vector problem (SIVP) to within a factor of $\gamma = \tilde{O}(n/\alpha)$ in *worst case* dimension n lattices. This is proven using a quantum reduction. Classical reductions (to a slightly different problem) exist as well [24], [25] but with somewhat worse parameters. The best known (classical or quantum) algorithms

for these problems run in time $2^{\tilde{O}(n/log(\gamma))}$, and in particular they are conjectured to be intractable for $\gamma = poly(n)$.

A discrete gaussian with parameter αq is $B = \alpha q$ bounded, except with negligible probability. For parameter α the worst-to-average case reduction of [22] gives a worst-case approximation factor of $\tilde{O}(n/\alpha)$ for SIVP. Consequently, in terms of the bound B and the modulus q we get a worst-case approximation factor of $\tilde{O}(nq/B)$ for SIVP.

B. Laconic Oblivious Transfer

In the following we introduce laconic oblivious transfer (LOT), the main object of interest of our work [3]. For presentation purposes, we define two variants of LOT with different efficiency and security requirements. Looking ahead, we will show a generic transformation, thereby simplifying the task of designing a LOT scheme to its simplest flavour.

As discussed by Cho et al. [3], we can ignore any security requirement on the receiver side: Although some bits of the database could be leaked, any LOT scheme can be generically upgraded to achieve receiver security through a generic transformation, i.e., by running the LOT under a 2PC protocol. In contrast to our work, Cho et al. [3] also considers the property of updatability for a LOT, which allows one to modify a bit of the database and update the digest in a significantly more efficient way than computing it from scratch. Since it is not relevant for our applications, we omit this property.

Weak Laconic Oblivious Transfer. A weak LOT scheme is identical to the standard LOT [3] except that the sender algorithm does not take as input two messages but chooses itself two random keys (k_0, k_1). This can be seen as the analogous to key encapsulation for encryption schemes. The syntax is given below.

Definition 3 (Weak Laconic Oblivious Transfer)**.** *A weak LOT* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{KGen}, \mathsf{Receive})$ *is defined as the following tuple of algorithms.*

$\underline{\mathsf{Setup}(1^\lambda) \to crs :}$ *On input the security parameter* 1^λ*, the generation algorithm returns a common reference string* crs*.*

$\underline{\mathsf{Hash}(crs, D) \to (d, \tilde{D}) :}$ *On input the common reference string* crs *and a database* $D \in \{0,1\}^*$*, the hashing algorithm returns a digest* d *and a state* $\tilde{D}$*.*

$\underline{\mathsf{KGen}(crs, d, L) \to (c, k_0, k_1) :}$ *On input the common reference string* crs*, a digest* d*, and a location* $L \in \mathbb{N}$*, the key generation algorithm returns a ciphertext* c *and a pair of keys* (k_0, k_1)*.*

$\underline{\mathsf{Receive}^{\tilde{D}}(crs, c, L) \to m :}$ *On input the common reference string* crs*, a ciphertext* c*, and a location* $L \in \mathbb{N}$*, the receiver algorithm (with random access to* $\tilde{D}$*) returns a key* k*.*

The definition of correctness is given in the following.

Definition 4 (Correctness)**.** *A weak LOT* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{KGen}, \mathsf{Receive})$ *is correct if for all* $\lambda \in \mathbb{N}$*, for all databases* D *of size polynomial in* λ*, and for all memory locations* $L \in \{1, \ldots, |D|\}$ *it holds that*

$$\Pr\left[k_{D[L]} = \mathsf{Receive}^{\tilde{D}}(crs, c, L)\right] = 1$$

where

- $crs \leftarrow \mathsf{Setup}(1^\lambda)$,
- $(d, \tilde{D}) \leftarrow \mathsf{Hash}(crs, D)$,
- $(c, k_0, k_1) \leftarrow \mathsf{KGen}(crs, d, L)$,

and the probability is taken over the random coins of Setup *and* KGen*.*

We also consider a weaker notion of correctness where the above probability is bounded from below by $1 - 1/poly(\lambda)$, for some polynomial function $poly$. We refer to such notion as $(1/poly(\lambda))$-correctness. A weak LOT is required to satisfy only a rudimental notion of sender security that we define in the following. Loosely speaking, we require that the adversary is not able to predict both keys (k_0, k_1) in their entirety.

Definition 5 (Weak Sender Security)**.** *A weak LOT* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{KGen}, \mathsf{Receive})$ *is weakly sender secure if for all* $\lambda \in \mathbb{N}$ *and for all PPT adversaries* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ *there exists a negligible function* $negl$ *such that*

$$\Pr\left[\mathsf{ExpSendSec}_{(\mathsf{LOT},\mathcal{A})}(1^\lambda) = 1\right] = negl(\lambda)$$

where the experiment is defined as

$\underline{\mathsf{ExpSendSec}_{(\mathsf{LOT},\mathcal{A})}(1^\lambda) :}$

- *Compute* $crs \leftarrow \mathsf{Setup}(1^\lambda)$
- $(d, L, \mathsf{st}) \leftarrow \mathcal{A}_1(crs)$
- $(c, k_0, k_1) \leftarrow \mathsf{KGen}(crs, d, L)$
- $(k_0^*, k_1^*) \leftarrow \mathcal{A}_2(c, \mathsf{st})$
- *Output* $(k_0, k_1) = (k_0^*, k_1^*)$

and the probability is taken over the random coins of Setup*,* $\mathcal{A}_1$*, and* KGen*.*

Efficiency. For efficiency, it is only required that the size of the digest d is at most half of that of D, i.e, the Hash function is 2-to-1 compressing. In particular, there is no bound on the efficiency of the algorithms and on the size of the ciphertext, except for being polynomial in the security parameter. We refer to such notion as

semi-efficiency. Jumping ahead, we will show a generic transformation from any semi-efficient LOT to a full-fledged LOT without additional assumptions.

Indistinguishable Laconic Oblivious Transfer. This version of LOT is equivalent to that introduced in [3], except that we upgrade the definition of sender security to the malicious settings. Note that we consider without loss of generality schemes that transfer a single bit (bit-LOT), which can be turned into multi-bit schemes (string-LOT) by simply computing multiple ciphertexts over the same digest. The security is preserved by a standard hybrid argument.

Definition 6 (Laconic Oblivious Transfer)**.** *A LOT* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{Send}, \mathsf{Receive})$ *is defined as the following tuple of algorithms.*

$\underline{\mathsf{Setup}(1^\lambda) \to crs}$: *On input the security parameter* 1^λ*, the generation algorithm returns a common reference string* crs*.*

$\underline{\mathsf{Hash}(crs, D) \to (d, \tilde{D})}$: *On input the common reference string* crs *and a database* $D \in \{0,1\}^*$*, the hashing algorithm returns a digest* d *and a state* $\tilde{D}$*.*

$\underline{\mathsf{Send}(crs, d, L, m_0, m_1) \to c}$: *On input the common reference string* crs*, a digest* d*, a location* $L \in \mathbb{N}$*, and a pair of messages* $(m_0, m_1) \in \{0,1\}^2$*, the sender algorithm returns a ciphertext* c*.*

$\underline{\mathsf{Receive}^{\tilde{D}}(crs, c, L) \to m}$: *On input the common reference string* crs*, a ciphertext* c*, and a database location* $L \in \mathbb{N}$*, the receiver algorithm (with random access to* $\tilde{D}$*) returns a message* m*.*

The definition of correctness is given in the following.

Definition 7 (Correctness)**.** *A LOT* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{Send}, \mathsf{Receive})$ *is correct if for all* $\lambda \in \mathbb{N}$*, for all databases* D *of size polynomial in* λ*, for all memory locations* $L \in \{1, \ldots, |D|\}$*, and any pair of messages* $(m_0, m_1) \in \{0,1\}^2$ *it holds that*

$$\Pr\left[m_{D[L]} = \mathsf{Receive}^{\tilde{D}}(crs, c, L)\right] = 1$$

where

- $crs \leftarrow \mathsf{Setup}(1^\lambda)$*,*
- $(d, \tilde{D}) \leftarrow \mathsf{Hash}(crs, D)$*,*
- $c \leftarrow \mathsf{Send}(crs, d, L, m_0, m_1)$*,*

and the probability is taken over the random coins of Setup *and* Send*.*

The fact that the digest of a LOT is compressing, makes even defining sender security a non-trivial task. We put forward the following indistinguishability-based definition, which suffices for our purposes.

Definition 8 (Sender Indistinguishability)**.** *A LOT* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{Send}, \mathsf{Receive})$ *is sender secure if for all* $\lambda \in \mathbb{N}$*, for all PPT adversaries* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, and for all polynomial functions* $poly(\lambda)$ *there exists a negligible function* $negl$ *such that*

$$\Pr\left[\varepsilon_0 > \frac{1}{poly(\lambda)} \text{ and } \varepsilon_1 > \frac{1}{poly(\lambda)}\right] = negl(\lambda)$$

where, for $\beta \in \{0,1\}$*,* ε_β *is defined as*

$$\varepsilon_\beta = \left| \begin{array}{c} \Pr\left[\mathsf{ExpSendInd}^{(0,\beta)}_{(\mathsf{LOT},\mathcal{A})}(1^\lambda) = 1\right] - \\ \Pr\left[\mathsf{ExpSendInd}^{(1,\beta)}_{(\mathsf{LOT},\mathcal{A})}(1^\lambda) = 1\right] \end{array} \right|$$

where the experiment is defined as

$\underline{\mathsf{ExpSendInd}^{(\mathsf{b},\beta)}_{(\mathsf{LOT},\mathcal{A})}(1^\lambda)}$:

- *Compute* $crs \leftarrow \mathsf{Setup}(1^\lambda)$
- $(d, L, m_0, m_1, \mathsf{st}) \leftarrow \mathcal{A}_1(crs)$
- *If* $b = 0$: *Set* $(r_0, r_1) = (m_0, m_1)$
- *If* $b = 1$: *Set* $r_\beta \leftarrow_\$ \{0,1\}$ *and* $r_{1-\beta} = m_{1-\beta}$
- $c \leftarrow \mathsf{Send}(crs, d, L, r_0, r_1)$
- $b^* \leftarrow \mathcal{A}_2(c, \mathsf{st})$
- *Output* b^*

and the probabilities are taken over the random coins of Setup*,* Send*, and* $\mathcal{A}_1$*.*

Efficiency. Here it is required that the size of the digest d is a fixed polynomial in the security parameter. Furthermore, we also impose a bound on the running time of the Hash algorithm of $|D| \cdot poly(log(|D|), \lambda)$ and on the time complexity of Setup, Encode, and Receive of $poly(log(|D|), \lambda)$. These are the same efficiency requirements of Cho et al. [3].

C. Laconic Conditional Disclosure of Secrets

Conditional disclosure of secrets (CDS) [26] for a language $\mathcal{L}$ in NP with relation $\mathcal{R}$ is the two-round analog of witness encryption [10]: Given a statement x and a message m from the sender, the receiver is able to recover m if $x \in \mathcal{L}$, whereas m stays hidden if this is not the case. Furthermore, the witness w for x should be kept secret from the eyes of the sender.

Definition 9 (Conditional Disclosure of Secrets)**.** *A CDS scheme* $\mathsf{CDS} = (\mathsf{CDSSetup}, \mathsf{CDSReceive}, \mathsf{CDSSend}, \mathsf{CDSDecode})$ *for an NP-langauge* $\mathcal{L}$ *is defined as the following tuple of algorithms.*

$\underline{\mathsf{CDSSetup}(1^\lambda) \to crs}$: *On input the security parameter* 1^λ*, the generation algorithm returns a common reference string* crs*.*

$\underline{\mathsf{CDSReceive}(crs, x, w) \to (\mathsf{cds}_1, r)}$: *On input the common reference string* crs *and a statement-witness pair*

(x, w)*, the receiver algorithm returns a first message* cds_1 *and a state* r*.*

$\underline{\mathsf{CDSSend}(crs, x, m, \mathsf{cds}_1) \to \mathsf{cds}_2 :}$ *On input the common reference string* crs*, a statement* x*, a message* $m \in \{0,1\}^*$*, a first message* cds_1*, the sender algorithm returns a second message* cds_2*.*

$\underline{\mathsf{CDSDecode}(crs, \mathsf{cds}_2, r) \to m :}$ *On input the common reference string* crs*, a second message* cds_2*, and the receiver state* r*, the decoding algorithm returns a message* m*.*

Correctness requires that the decoding is successful if $x \in \mathcal{L}$.

Definition 10 (Correctness)**.** *A CDS scheme* $\mathsf{CDS} = (\mathsf{CDSSetup}, \mathsf{CDSSend}, \mathsf{CDSReceive}, \mathsf{CDSDecode})$ *is correct if for all* $\lambda \in \mathbb{N}$*, all* $(x, w) \in \mathcal{R}$*, and any message* $m \in \{0,1\}^*$ *it holds that*

$$\Pr\left[m = \mathsf{CDSDecode}(crs, \mathsf{cds}_2, r)\right] = 1$$

where

- $crs \leftarrow \mathsf{CDSSetup}(1^\lambda)$,
- $(\mathsf{cds}_1, r) \leftarrow \mathsf{CDSReceive}(crs, x, w)$,
- $\mathsf{cds}_2 \leftarrow \mathsf{CDSSend}(crs, x, m, \mathsf{cds}_1)$,

and the probability is taken over the random coins of CDSSetup*,* CDSReceive*, and* CDSSend*.*

The central requirement for a CDS scheme is that the message is hidden if the given statement is not in the corresponding NP-language.

Definition 11 (Message Indistinguishability)**.** *A CDS scheme* $\mathsf{CDS} = (\mathsf{CDSSetup}, \mathsf{CDSSend}, \mathsf{CDSReceive}, \mathsf{CDSDecode})$ *is message-indistinguishable if for all* $\lambda \in \mathbb{N}$*, for all admissible PPT adversaries* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, for all* $x \notin \mathcal{L}$ *there exists a negligible function* $negl$ *such that*

$$\left| \begin{array}{l} \Pr\left[\mathsf{ExpMesInd}^{(0)}_{(\mathsf{CDS},\mathcal{A})}(1^\lambda) = 1\right] - \\ \quad \Pr\left[\mathsf{ExpMesInd}^{(1)}_{(\mathsf{CDS},\mathcal{A})}(1^\lambda) = 1\right] \end{array} \right| = negl(\lambda)$$

where the experiment is defined as

$\underline{\mathsf{ExpMesInd}^{(\mathsf{b})}_{(\mathsf{CDS},\mathcal{A})}(1^\lambda) :}$

- *Compute* $crs \leftarrow \mathsf{CDSSetup}(1^\lambda)$
- $(\mathsf{cds}_1, m_0, m_1, \mathsf{st}) \leftarrow \mathcal{A}_1(crs)$
- $\mathsf{cds}_2 \leftarrow \mathsf{CDSSend}(crs, x, m_b, \mathsf{cds}_1)$
- $b^* \leftarrow \mathcal{A}_2(\mathsf{cds}_2, \mathsf{st})$
- *Output* b^*

where $\mathcal{A}$ *is admissible if* m_0 *and* m_1 *have equal length and the probabilities are taken over the random coins of* CDSSetup*,* CDSSend*, and* $\mathcal{A}_1$*.*

Finally, we require that no information is leaked to the sender by the message of the receiver.

Definition 12 (Receiver Simulation)**.** *A CDS scheme* $(\mathsf{CDSSetup}, \mathsf{CDSSend}, \mathsf{CDSReceive}, \mathsf{CDSDecode})$ *is receiver simulatable if there exists a PPT simulator* $\mathsf{CDSSim} = (\mathsf{CDSSim}_1, \mathsf{CDSSim}_2)$ *such that for all* $\lambda \in \mathbb{N}$*, all* $(x, w) \in \mathcal{R}$*, and all PPT adversaries* $\mathcal{A}$ *it holds that*

$$\left| \begin{array}{l} \Pr\left[\mathsf{ExpRecSim}^{(0)}_{(\mathsf{CDS},\mathcal{A})}(1^\lambda) = 1\right] - \\ \quad \Pr\left[\mathsf{ExpRecSim}^{(1)}_{(\mathsf{CDS},\mathcal{A})}(1^\lambda) = 1\right] \end{array} \right| = negl(\lambda)$$

where the experiment is defined as

$\underline{\mathsf{ExpRecSim}^{(\mathsf{b})}_{(\mathsf{CDS},\mathcal{A})}(1^\lambda) :}$

- *If* $b = 0$: *Compute* $crs \leftarrow \mathsf{CDSSetup}(1^\lambda)$
- $(\mathsf{cds}_1, \cdot) \leftarrow \mathsf{CDSReceive}(crs, x, w)$
- *If* $b = 1$: *Compute* $(crs, \mathsf{td}) \leftarrow \mathsf{CDSSim}_1(1^\lambda)$
- $\mathsf{cds}_1 \leftarrow \mathsf{CDSSim}_2(crs, \mathsf{td}, x)$
- $b^* \leftarrow \mathcal{A}(crs, \mathsf{cds}_1)$
- *Output* b^*

and the probabilities are taken over the random coins of CDSSetup*,* CDSReceive*,* CDSSim*, and* $\mathcal{A}$*.*

Efficiency. The standard requirement for a CDS scheme is to run in time polynomial in the security parameter and, possibly, in the size of the statement and of the witness. We say that a CDS is *laconic* if the communication complexity is a fixed polynomial $poly(\lambda)$ and in particular is independent of the size of w (up to a poly-logarithmic factor). This notion can be seen as the equivalent of succinct arguments [5], [4] with implicit verification.

III. Context Security

In this Section we provide our definition of context security, which is a UC-inspired security notion [11]. The notion is specifically geared towards non-interactive secure computation protocols for which the communication complexity is sublinear in the receiver's input.

UC security ensures that a protocol is secure in any environment or context. However, the way UC-security is formalized makes it necessary that a straight-line simulator is able to extract the inputs of malicious parties, which immediately implies that the communication complexity scales with the size of the inputs of the parties. An important aspect in (standard) UC-security is that the environment is chosen after the simulator. Specialized-Simulator UC [27] allows the simulator to depend on the environment, but not on its random coins. Context security aims for a compromise, while

salvaging the original motivation of UC, namely that protocols remain secure in any context.

Unlike other works which relax UC security, e.g. [28], [29], the purpose of context security is *not* to get rid of a trusted setup. Rather, our goal is to obtain meaningful composability guarantees in the malicious setting while allowing for round-optimal protocols with communication complexity that is sublinear in the inputs. This is neither possible with UC/straight-line extraction, nor in the standalone simulation framework: While rewinding does allow to extract large inputs from protocols with small communication, it does not help in the two message setting.

Our definition will only consider the single instance setting, but is crafted in a way that multi instance security can be dealt with via standard hybrid arguments. In this sense, we aim for conceptual simplicity rather than generality. We start by providing the definition of a context.

Definition 13 (Protocol Context)**.** *We say that a PPT machine $\mathcal{Z} = (\mathcal{Z}_1, \mathcal{Z}_2)$ is a context for two message protocol $\Pi = (\mathsf{Setup}, \mathsf{R}_1, \mathsf{S}, \mathsf{R}_2)$, if it has the following syntactic properties: The first stage $\mathcal{Z}_1$ takes as input a common reference string crs (generated by* Setup*) and random coins r_1 and outputs a receiver message* rec *and a state* st*. The second phase $\mathcal{Z}_2$ takes as input the state* st *and random coins r_2. The second phase is allowed to make queries y to a sender oracle $\mathcal{O}(y)$, which are answered by $\mathsf{S}(crs, \mathsf{rec}, y)$ (using fresh randomness from r_2). In the end the context outputs a bit b^*. Define the following experiment.*

$\underline{\mathcal{Z}(1^\lambda):}$

- *Choose random tapes* $(\mathsf{r}_1, \mathsf{r}_2)$
- *Compute* $crs \leftarrow \mathsf{Setup}(1^\lambda)$
- $(\mathsf{st}, \mathsf{rec}) \leftarrow \mathcal{Z}_1(crs, \mathsf{r}_1)$
- $b^* \leftarrow \mathcal{Z}_2^{\mathsf{S}(crs, \mathsf{rec}, \cdot)}(\mathsf{st}, \mathsf{r}_2)$.
- *Output* b^*

We now provide our definition of context security.

Definition 14 (Context Security)**.** *Let $\Pi = (\mathsf{Setup}, \mathsf{R}_1, \mathsf{S}, \mathsf{R}_2)$ be a two-message protocol realizing a two-party functionality $\mathcal{F}$. We say that Π is* context-secure *if the following holds for every Π-context $\mathcal{Z} = (\mathcal{Z}_1, \mathcal{Z}_2)$. We require that there exists a context extractor $\mathsf{Ext}_{\mathcal{Z}}$ and simulators* SimSetup *and* Sim *such that the following holds for every $\delta > 0$:*

1) SimSetup *runs in time polynomial in λ (but independent of the run-time of $\mathcal{Z}_1$ and $\mathcal{Z}_2$) and outputs a common reference string crs and a trapdoor* td*.*
2) $\mathsf{Ext}_{\mathcal{Z}}$ *takes as input $crs, \mathsf{st}, \mathsf{rec}$, random coins r^* and a parameter δ and outputs a value x^* and an auxiliary string* aux*. $\mathsf{Ext}_{\mathcal{Z}}$ has overhead $poly(\lambda, p(\lambda)) \cdot T_2$, where T_2 is the overhead of $\mathcal{Z}_2$ and the polynomial is independent of $\mathcal{Z}$, only depending on λ.*
3) Sim *takes as input* td, rec, aux *and a value z and outputs a sender-message* snd*. We require The overhead of* Sim *to be polynomial in the overhead of λ, but independent of $\mathcal{Z}_1$ and $\mathcal{Z}_2$.*
4) *Define the following experiment.*

   $\underline{\mathcal{EZ}(1^\lambda, \delta):}$

   - *Choose random tapes* $(\mathsf{r}_1, \mathsf{r}_2)$
   - *Compute* $(crs, \mathsf{td}) \leftarrow \mathsf{SimSetup}(1^\lambda)$.
   - $(\mathsf{st}, d) \leftarrow \mathcal{Z}_1(crs, \mathsf{r}_1)$
   - $(x^*, \mathsf{aux}) \leftarrow \mathsf{Ext}_{\mathcal{Z}}(crs, \mathsf{st}, \mathsf{rec}, \mathsf{r}^*, \delta)$
   - $b^* \leftarrow \mathcal{Z}_2^{\mathcal{O}'(\cdot)}(\mathsf{st}, \mathsf{r}_2)$*, where* $\mathcal{O}'(y)$ *computes and outputs* $\mathsf{Sim}(\mathsf{td}, \mathsf{rec}, \mathsf{aux}, \mathcal{F}(x^*, y))$.
   - *Output* b^*

   (Security) It holds for every inverse polynomial $\varepsilon = \varepsilon(\lambda)$ that

$$\left|\Pr\left[\mathcal{Z}(1^\lambda) = 1\right] - \Pr\left[\mathcal{EZ}(1^\lambda, \varepsilon) = 1\right]\right| < \varepsilon,$$

   except for finitely many λ.

Note specifically that we allow the ideal experiment $\mathcal{EZ}$ to depend on the distinguishing advantage ε, which looks unusual at first glance. However, it is precisely this dependence which allows us to prove context security for protocols with sublinear communication complexity. Specifically, this dependence will let us use the distinguisher-dependent simulation technique [20] to construct extractors $\mathsf{Ext}_{\mathcal{Z}}$.

The extractor $\mathsf{Ext}_{\mathcal{Z}}$ has a similar syntax as a knowledge extractor [30], however the main difference is we will be able to prove this notion under standard falsifiable assumptions. An important aspect which enables context security to be used in a meaningful way is that $\mathsf{Ext}_{\mathcal{Z}}$ does not get to see r_2.

The runtime requirement for $\mathsf{Ext}_{\mathcal{Z}}$ will make this notion compose benignly, unlike knowledge extractors [30]. More specifically, it will be natural and convenient to define experiments iteratively, and this runtime requirement will ensure that runtimes of the experiments do not explode. An important aspect will be that extractors are not run recursively. Extractors will only be used in the first phase of a modified context, but the extractor itself does not use the first phase. That is, in successive uses of the technique the runtime will only grow additively, but not multiplicatively.

We remark that context-security has to be used with care. If a context $(\mathcal{Z}_1, \mathcal{Z}_2)$ is such that $\mathcal{Z}_2$ uses a *long-term secret* generated by $\mathcal{Z}_1$, the extractor also needs to

use this secret when simulating $\mathcal{Z}_2$. In particular, this means that in a hybrid proof we cannot make any further modifications to $\mathcal{Z}_1$ after extracting the receiver input, as this could invalidate the extracted input. For example, consider a functionality where the sender provides a (long term) signing key for a signature and the receiver obtains a signature on his input. Moreover, assume that the receiver is also given the verification key of the signature (and can therefore check the validity of signatures). Context-security allows us to extract the inputs of the receiver, but in order to do so the extractor needs to know the signing key. This is problematic if we later want to argue about the unforgeability of the signature, where a reduction is not given the signing key but only access to a signing oracle. Yet, by the observation above the context-extractor cannot extract the signing queries without being given the signing key. Note that this issue does not exist in the setting of UC security.

Finally, a very convenient property of context security is that it immediately implies game-based security notions with an efficient experiment. Specifically, phrase a game as a context $\mathcal{Z}$, apply context security and reason that the advantage in $\mathcal{E}\mathcal{Z}$ is 0. Via the way we have defined context security this will immediately imply that also in the original experiment $\mathcal{Z}$ the advantage is negligible.

For instance, for non-interactive secure computation one can consider the following indistinguishability-based security definition. The malicious receiver first obtains a CRS, outputs a receiver message and two randomized functions F_0 and F_1 for which it holds that on any input x the distributions $F_0(x)$ and $F_1(x)$ are indistinguishable. The experiment flips a bit b and provides the adversary with a sender message which allows the receiver to evaluate F_b on his input. The adversary then has to guess the bit b.

Any context-secure NISC scheme immediately also satisfies this notion. We can phrase the experiment, including the adversary as a context $(\mathcal{Z}_1, \mathcal{Z}_2)$ where $\mathcal{Z}_1$ outputs the receiver message (together with a state), and $\mathcal{Z}_2$ chooses the random bit b and computes the sender message via access to a sender oracle. Now, context-security lets us argue that in the experiment $\mathcal{E}\mathcal{Z}$ the actual bit b used by the experiment is hidden from the view of the adversary, as the output of the oracle only depends on $F_b(x)$ (which is indistinguishable from $F_{1-b}(x)$) rather than on F_b.

Some Useful Functionalities. We define the functionalities corresponding to the primitives of interest of this work, defined in Section II. Note that all of the functionalities that we consider are single-output, in the sense that only one party learns an output at the end of the execution.

Laconic Oblivious Transfer: The ideal functionality $\mathcal{F}_{\mathsf{LOT}}(D, m_0, m_1, L)$ takes as input a database D, a pair of messages (m_0, m_1) and a location L. It returns $(m_{D[L]}, L)$.

Laconic Conditional Disclosure of Secrets: The ideal functionality $\mathcal{F}_{\mathsf{CDS}}(w, x, m)$ is parametrized by an NP-language $\mathcal{L}$ with relation $\mathcal{R}$ and takes as input a statement x, a witness w, and a message m. The functionality returns m if $\mathcal{R}(w, x) = 1$ and $\perp$ otherwise.

IV. Malicious Laconic Oblivious Transfer

In this section we present our constructions from different cryptographic hard problems. We first construct a 2-to-1 compressing weak LOT from the CDH or the LWE assumption, then we show how to generically bootstrap any 2-to-1 compressing weak LOT to a context-secure LOT with arbitrary compression.

A. From Computational Diffie-Hellman

We present our scheme for a weak LOT with malicious security assuming the hardness of the CDH problem over a prime-order multiplicative cyclic group $(\mathbb{G}, p, g)$. For conceptual simplicity, we assume that a group element can be represented using λ-many bits, but the scheme can be generalized to arbitrary representations in a natural way. The construction is given in Figure 2.

Analysis. We show that our scheme is correct.

Theorem 2 (Correctness)**.** *Let $(\mathbb{G}, p, g)$ be a prime-order multiplicative cyclic group. Then the construction $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{Send}, \mathsf{Receive})$ as defined in Figure 2 is correct.*

Proof of Theorem 2: The proof consists in the observation that

$$k_{D[L]} = \left(\frac{d}{g_L^{D[L]}}\right)^r = \left(\frac{\prod_{i=1}^{2\lambda} g_i^{D[i]}}{g_L^{D[L]}}\right)^r = \left(\prod_{i=1, i\neq L}^{2\lambda} g_i^{D[i]}\right)^r = \prod_{i=1, i\neq L}^{2\lambda} h_i^{D[i]}$$

which is exactly the output of the receiver. ■

Next we argue about weak sender security.

Theorem 3 (Weak Sender Security)**.** *Let $(\mathbb{G}, p, g)$ be a CDH-hard group. Then the construction $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{KGen}, \mathsf{Receive})$ as defined in Figure 2 is weakly sender secure.*

$\underline{\mathsf{Setup}(1^\lambda)}$:

- Sample $(\mathbb{G}, p, g) \leftarrow \mathcal{G}(1^\lambda)$.
- Sample a uniform vector $(g_1, \ldots, g_{2\lambda}) \leftarrow_\$ \mathbb{G}^{2\lambda}$
- Return $crs = (\mathbb{G}, p, g, g_1, \ldots, g_{2\lambda})$.

$\underline{\mathsf{Hash}(crs, D)}$:

- Parse crs as $(\mathbb{G}, p, g, g_1, \ldots, g_{2\lambda})$ and $D \in \{0,1\}^{2\lambda}$ as a bitstring.
- Compute $d = \prod_{i=1}^{2\lambda} g_i^{D[i]}$.
- Set $\tilde{D} = (D, d)$ and return $(d, \tilde{D})$.

$\underline{\mathsf{KGen}(crs, d, L)}$:

- Parse crs as $(\mathbb{G}, p, g, g_1, \ldots, g_{2\lambda})$ and $d \in \mathbb{G}$.
- Sample a uniform $r \leftarrow_\$ \mathbb{Z}_p$.
- Compute the vector $c = (g_1^r, \ldots, g_{L-1}^r, g_{L+1}^r, \ldots, g_{2\lambda}^r)$.
- Set $k_0 = d^r$ and $k_1 = (d/g_L)^r$.
- Return (c, k_0, k_1).

$\underline{\mathsf{Receive}^{\tilde{D}}(crs, c, L)}$:

- Parse $\tilde{D}$ as (D, d) and c as $(h_1, \ldots, h_{2\lambda})$.
- Return $\prod_{i=1, i\neq L}^{2\lambda} h_i^{D[i]}$.

Figure 2: CDH-Based Weak LOT.

Proof of Theorem 3: The theorem is shown with a reduction to the CDH assumption. Assume towards contradiction that there exists a PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that

$$\Pr\left[\mathsf{ExpSendSec}_{(\mathsf{LOT},\mathcal{A})}(1^\lambda) = 1\right] = \frac{1}{poly(\lambda)}$$

for some polynomial function $poly$. On input a CDH challenge $(\mathbb{G}, p, g, g^{a_1}, g^{a_2})$, the reduction samples a uniform vector $(x_1, \ldots, x_{2\lambda}) \leftarrow_\$ \mathbb{Z}_p^{2\lambda}$ and an index $\tilde{L} \leftarrow_\$ \{1, \ldots, 2\lambda\}$. Then it sets

$$crs = (\mathbb{G}, p, g, g^{x_1}, \ldots, g^{x_{L-1}}, g^{a_1}, g^{x_{L+1}}, \ldots, g^{x_{2\lambda}})$$

and runs $(d, L, \mathsf{st}) \leftarrow \mathcal{A}_1(crs)$. If $L \neq \tilde{L}$ the reduction aborts, else it sets

$$c = (g^{a_2 \cdot x_1}, \ldots, g^{a_2 \cdot x_{L-1}}, g^{a_2 \cdot x_{L+1}}, \ldots, g^{a_2 \cdot x_{2\lambda}}).$$

Then it runs $(k_0, k_1) \leftarrow \mathcal{A}_2(c, \mathsf{st})$ and returns k_0/k_1. The reduction runs in polynomial time and the inputs given to the adversary are identically distributed as those in the original game. Note that if $L = \tilde{L}$ and the adversary correctly guesses both k_0 and k_1, then we have that

$$\frac{k_0}{k_1} = \frac{d^{a_2}}{d^{a_2} \cdot g_L^{-a_2}} = g_L^{a_2} = g^{a_1 \cdot a_2}.$$

Since this happens with probability at least $\frac{1}{poly(\lambda) \cdot 2\lambda}$, it contradicts the CDH assumption. ∎

B. From Learning with Errors

In the following we present our weak LOT scheme assuming the hardness of the LWE problem. Our scheme is shown in Figure 3.

Rounding. Let $q = c \cdot p$ be an integer modulus. Define the rounding function $\lfloor \cdot \rceil_c : \mathbb{Z}_q \to \mathbb{Z}_c$ as

$$\lfloor x \rceil_c = \lfloor \bar{x} \cdot c/q \rceil \mod c$$

where $\bar{x} \in \mathbb{Z}$ is an arbitrary residue-class representative of $x \in \mathbb{Z}_q$. We are going to use the following lemma about the rounding function, implicitly proven in [31].

Lemma 2. *Let $q = c \cdot p$ be an integer modulus. Let $x \leftarrow_\$ \mathbb{Z}_q$ be distributed uniformly at random. Then it holds for all $v \in \{-B, \ldots, B\}$ that $\lfloor x + v \rceil_c = \lfloor x \rceil_c$, except with probability $(2B+1) \cdot c/q$ over the random choice of x.*

Proof of Lemma 2: Note that there are exactly c multiples of $p = q/c$ in $\mathbb{Z}_q$. Thus, the probability that a uniform $x \in \mathbb{Z}_q$ lands B-close from the nearest multiple of q/c is exactly $(2B+1) \cdot c/q$. ∎

Another property of the rounding function is that it partitions the domain $\mathbb{Z}_q$ in c-many well defined intervals of size exactly $q/c = p$. More concretely, the i-th interval is $[p \cdot (i-1), p \cdot i - 1]$. This fact is going to be useful for our analysis.

Analysis. Our scheme is parametrized by the following variables:

$\underline{\mathsf{Setup}(1^\lambda):}$

- For all $i \in \{1,\ldots,\lambda\}$: Sample $\mathbf{A}^{(i)} \leftarrow_\$ \mathbb{Z}_q^{n\times m}$.
- Return $crs = (\mathbf{A}^{(1)},\ldots,\mathbf{A}^{(\lambda)})$.

$\underline{\mathsf{Hash}(crs, D):}$

- Parse crs as $(\mathbf{A}^{(1)},\ldots,\mathbf{A}^{(\lambda)})$ and $D \in \{0,1\}^m$ as a column vector.
- For all $i \in \{1,\ldots,\lambda\}$: Compute $\mathbf{d}^{(i)} = \mathbf{A}^{(i)}D$.
- Set $d = (\mathbf{d}^{(1)},\ldots,\mathbf{d}^{(\lambda)})$, $\tilde{D} = (D, d)$, and return $(d, \tilde{D})$.

$\underline{\mathsf{KGen}(crs, d, L):}$

- Parse crs as $(\mathbf{A}^{(1)},\ldots,\mathbf{A}^{(\lambda)})$ and d as $(\mathbf{d}^{(1)},\ldots,\mathbf{d}^{(\lambda)})$.
- Sample $\mathbf{s} \leftarrow_\$ \mathbb{Z}_q^n$ and parse it as a column vector.
- For all $i \in \{1,\ldots,\lambda\}$:
  - Parse $(\mathbf{a}_1^{(i)},\ldots,\mathbf{a}_m^{(i)})$ as the columns of $\mathbf{A}^{(i)}$.
  - Sample $(e_1^{(i)},\ldots,e_m^{(i)}) \leftarrow_\$ \chi^m$.
  - For all $j \in \{1,\ldots,m\} \setminus \{L\}$: Compute $b_j^{(i)} = \mathbf{s}^T\mathbf{a}_j^{(i)} + e_j^{(i)}$.
  - Set $\mathbf{b}^{(i)}$ to be the row vector $(b_1^{(i)},\ldots,b_{m-1}^{(i)})$.
  - Sample $(t_0^{(i)}, t_1^{(i)}) \leftarrow_\$ \mathbb{Z}_q^2$.
  - Set $k_0^{(i)} = \left\lfloor \mathbf{s}^T\mathbf{d}^{(i)} + t_0^{(i)} \right\rfloor_c$ and $k_1^{(i)} = \left\lfloor -\mathbf{s}^T\mathbf{d}^{(i)} + \mathbf{s}^T\mathbf{a}_L^{(i)} + e_L^{(i)} + t_1^{(i)} \right\rfloor_c$.
- Set $c = ((\mathbf{b}^{(1)}, t_0^{(1)}, t_1^{(1)}),\ldots,(\mathbf{b}^{(\lambda)}, t_0^{(\lambda)}, t_1^{(\lambda)}))$.
- Set $k_0 = k_0^{(1)}\|\ldots\|k_0^{(\lambda)}$ and $k_1 = k_1^{(1)}\|\ldots\|k_1^{(\lambda)}$.
- Return (c, k_0, k_1).

$\underline{\mathsf{Receive}^{\tilde{D}}(crs, c, L):}$

- Parse c as $((\mathbf{b}^{(1)}, t_0^{(1)}, t_1^{(1)}),\ldots,(\mathbf{b}^{(\lambda)}, t_0^{(\lambda)}, t_1^{(\lambda)}))$ and $\bar{D} = D \setminus D[L]$ as a column vector.
- For all $i \in \{1,\ldots,\lambda\}$: Compute $k^{(i)} = \left\lfloor (-1)^{D[L]} \cdot \mathbf{b}^{(i)}\bar{D} + t_{D[L]}^{(i)} \right\rfloor_c$.
- Return $k = k^{(1)}\|\ldots\|k^{(\lambda)}$.

Figure 3: LWE-Based Weak LOT.

- m : The size of the database $D \in \{0,1\}^m$.
- n : The dimension of the LWE problem.
- q : The modulus of the LWE problem. For ease of the analysis we are going to assume that q is of the form $c \cdot p$, for some constant c and an arbitrary integer p
- B : The bound on the absolute value of the noise, i.e., $e \in \{-B,\ldots,B\}$.
- c : A constant that parametrizes the function $\lfloor\cdot\rfloor_c$, which is used in our construction.

In the analysis we set constraints on the parameters on demand. In the end of the section we are going to show that such a set of constraints forms a satisfiable system of relations and that the resulting $\mathrm{LWE}_{n,q,\chi}$ problem is still in the regime of parameters for which it is conjectured to be hard.

Note that setting $m \geq 2 \cdot \lceil log(q) \rceil \cdot n \cdot \lambda$ makes the first message 2-to-1 compressing. We now argue about the (weak) correctness of our scheme.

Theorem 4 (Correctness)**.** *The construction* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{KGen}, \mathsf{Receive})$ *as defined in Figure 3 is* $(1/\lambda^2)$*-correct.*

Proof of Theorem 4: For all $i \in \{1,\ldots,\lambda\}$, we expand the product

$$\begin{aligned}
\mathbf{b}^{(i)}\bar{D} &= \sum_{j\neq L} D[j]\cdot b_j^{(i)} \\
&= \sum_{j\neq L} D[j]\cdot(\mathbf{s}^T\mathbf{a}_j^{(i)} + e_j^{(i)}) \\
&= \sum_{j\neq L} D[j]\cdot\mathbf{s}^T\mathbf{a}_j^{(i)} + \sum_{j\neq L} D[j]\cdot e_j^{(i)} \\
&= \mathbf{s}^T\left(\sum_{j\neq L} D[j]\cdot\mathbf{a}_j^{(i)}\right) + \sum_{j\neq L} D[j]\cdot e_j^{(i)} \\
&= \mathbf{s}^T\mathbf{d}^{(i)} - D[L]\cdot\mathbf{s}^T\mathbf{a}_L^{(i)} + \sum_{j\neq L} D[j]\cdot e_j^{(i)}
\end{aligned}$$

by linearity. Applying the equality above for all $i \in$

$\{1,\dots,\lambda\}$ we obtain

$$
\begin{aligned}
k^{(i)} &= \left\lfloor (-1)^{D[L]} \cdot \mathbf{b}^{(i)} \bar{D} + t^{(i)}_{D[L]} \right\rfloor_c \\
&= \left\lfloor \begin{array}{l} (-1)^{D[L]} \cdot \left(\mathbf{s}^T \mathbf{d}^{(i)} - D[L] \cdot \mathbf{s}^T \mathbf{a}^{(i)}_L \right) + \\ (-1)^{D[L]} \cdot \sum_{j \neq L} D[j] \cdot e^{(i)}_j + t^{(i)}_{D[L]} + \\ D[L] \cdot e^{(i)}_L - D[L] \cdot e^{(i)}_L \end{array} \right\rfloor_c \\
&= \left\lfloor \begin{array}{l} (-1)^{D[L]} \cdot \mathbf{s}^T \mathbf{d}^{(i)} + D[L] \cdot \mathbf{s}^T \mathbf{a}^{(i)}_L + \\ (-1)^{D[L]} \cdot \sum_{j=1}^{m} D[j] \cdot e^{(i)}_j + t^{(i)}_{D[L]} + \\ D[L] \cdot e^{(i)}_L \end{array} \right\rfloor_c \\
&= \left\lfloor \begin{array}{l} (-1)^{D[L]} \cdot \mathbf{s}^T \mathbf{d}^{(i)} + \\ D[L] \left(\mathbf{s}^T \mathbf{a}^{(i)}_L + e^{(i)}_L \right) + \\ \underbrace{(-1)^{D[L]} \cdot \sum_{j=1}^{m} D[j] \cdot e^{(i)}_j}_{=\tilde{e}} + t^{(i)}_{D[L]} \end{array} \right\rfloor_c
\end{aligned}
$$

Note that the absolute value of $\tilde{e}$ is bounded by $m \cdot B$. Further note that $t^{(i)}_{D[L]}$ is uniform in $\mathbb{Z}_q$. Thus, by Lemma 2 we have that

$$
\begin{aligned}
k^{(i)} &= \left\lfloor \begin{array}{l} (-1)^{D[L]} \cdot \mathbf{s}^T \mathbf{d}^{(i)} + \\ D[L] \left(\mathbf{s}^T \mathbf{a}^{(i)}_L + e^{(i)}_L \right) + \tilde{e} + t^{(i)}_{D[L]} \end{array} \right\rfloor_c \\
&= \left\lfloor \begin{array}{l} (-1)^{D[L]} \cdot \mathbf{s}^T \mathbf{d}^{(i)} + \\ D[L] \left(\mathbf{s}^T \mathbf{a}^{(i)}_L + e^{(i)}_L \right) + t^{(i)}_{D[L]} \end{array} \right\rfloor_c \\
&= k^{(i)}_{D[L]}
\end{aligned}
$$

except with probability

$$\frac{(2mB+1) \cdot c}{q}.$$

Setting $log(q) \geq 3log(\lambda) + log(2mBc + c)$ then we have that the above equality holds for all $i \in \{1,\dots,\lambda\}$ except with probability at most $1/\lambda^3$. Taking a union bound over the λ-many parallel repetitions gives us the desired inequality. ∎

We now argue that our construction satisfies weak sender security.

Theorem 5 (Weak Sender Security)**.** *If the LWE$_{n,q,\chi}$ assumption holds, then the construction* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{KGen}, \mathsf{Receive})$ *as defined in Figure 3 is weakly sender secure.*

Proof of Theorem 5: We show the claim with a reduction against the LWE$_{n,q,\chi}$ problem. Assume towards contradiction that there exists a PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ such that

$$\Pr\left[\mathsf{ExpSendSec}_{(\mathsf{LOT},\mathcal{A})}(1^\lambda) = 1\right] = \frac{1}{poly(\lambda)}.$$

Then we construct a reduction ($\mathcal{R}$) that solves LWE$_{n,q,\chi}$ as follows. The reduction is given a challenge $(\mathbf{A}, \mathbf{u})$ where $\mathbf{A} \in \mathbb{Z}_q^{n \times m \cdot \lambda}$ and $\mathbf{u} \in \mathbb{Z}_q^{m \cdot \lambda}$ and parses $\mathbf{A}$ as the horizontal concatenation of λ-many $\mathbb{Z}_q^{n \times m}$ matrices $(\mathbf{A}^{(i)}, \dots, \mathbf{A}^{(\lambda)})$. The adversary $\mathcal{A}_1$ is invoked on input $crs = (\mathbf{A}^{(i)}, \dots, \mathbf{A}^{(\lambda)})$ and returns a tuple (d, L, st). For each $i \in \{1,\dots,\lambda\}$, the reduction sets the vector $\mathbf{b}^{(i)}$ to $(\mathbf{u}_{(i-1)m+1}, \dots, \mathbf{u}_{(i-1)m+m})$ except that it omits the element $\mathbf{u}_{(i-1)m+L}$. Then samples $(t^{(i)}_0, t^{(i)}_1) \leftarrow_\$ \mathbb{Z}_q^2$ uniformly at random. The reduction returns to the adversary $\mathcal{A}_2$ the ciphertext

$$c = ((\mathbf{b}^{(1)}, t^{(1)}_0, t^{(1)}_1), \dots, (\mathbf{b}^{(\lambda)}, t^{(\lambda)}_0, t^{(\lambda)}_1))$$

along with the state st. The adversary outputs a pair (k_0, k_1) that is parsed by the reduction as the concatenation of two λ-long vectors in $\mathbb{Z}_c$: $k_0 = k^{(1)}_0 \| \dots \| k^{(\lambda)}_0$ and $k_1 = k^{(1)}_1 \| \dots \| k^{(\lambda)}_1$. For all $i \in \{1,\dots,\lambda\}$, define $R^{(i)}_0$ as the interval of all elements $y_0 \in \mathbb{Z}_q$ such that $\lfloor y_0 \rfloor_c = k^{(i)}_0$ and define $R^{(i)}_1$ analogously. Let $R^{(i)}$ be the set of elements in $z \in \mathbb{Z}_q$ such that there exists a $y_0 \in R^{(i)}_0$ and a $y_1 \in R^{(i)}_1$ such that $y_0 + y_1 = z$. The reduction checks whether $\mathbf{u}_{(i-1)m+L} + t^{(i)}_0 + t^{(i)}_1 \in R^{(i)}$ and returns 1 if this is the case for all $i \in \{1,\dots,\lambda\}$. Else it returns 0.

Note that all steps are efficiently computable: In particular, for all $k^{(i)}_b \in \mathbb{Z}_c$ the corresponding interval $R^{(i)}_b$ consists of $[p \cdot k^{(i)}_b, p \cdot (k^{(i)}_b + 1) - 1]$ and consequently also $R^{(i)}$ is trivial to compute. We analyze the probability that the reduction outputs 1 for two separate cases.

Uniform Samples: In this case it is enough to observe that all elements $(\mathbf{u}_L, \dots, \mathbf{u}_{(\lambda-1)m+L})$ are uniform in $\mathbb{Z}_q$ and completely independent from the view of the adversary. Therefore, for all $i \in \{1,\dots,\lambda\}$, the probability that $\mathbf{u}_{(i-1)m+L} + t^{(i)}_0 + t^{(i)}_1 \in R^{(i)}$ is at most

$$\frac{\left|R^{(i)}\right|}{q} = \frac{\left|R^{(i)}_0\right| + \left|R^{(i)}_1\right|}{q} = \frac{2q}{c} \cdot \frac{1}{q} = \frac{2}{c}.$$

Setting $c = 4$ we obtain that the probability that the reduction outputs 1 is at most $1/2^\lambda$.

LWE Samples: Note that in this case the input given to the adversary are identically distributed as an honest run of the protocol where the keys are (implicitly) set to

$$k^{(i)}_0 = \left\lfloor \mathbf{s}^T \mathbf{d}^{(i)} + t^{(i)}_0 \right\rfloor_c$$

and

$$k^{(i)}_1 = \left\lfloor -\mathbf{s}^T \mathbf{d}^{(i)} + \mathbf{u}_{(i-1)m+L} + t^{(i)}_1 \right\rfloor_c$$

where $\mathbf{d}^{(i)}$ is part of the digest d given by the adversary. Let guess be the event that the adversary correctly guesses both k_0 and k_1. Then, by the law of total probability we have

$$\begin{aligned}\Pr[1 \leftarrow \mathcal{R}] = &\Pr[1 \leftarrow \mathcal{R}|\mathsf{guess}] \Pr[\mathsf{guess}] + \\ &\Pr[1 \leftarrow \mathcal{R}|\neg\mathsf{guess}] \Pr[\neg\mathsf{guess}] \\ \geq &\Pr[1 \leftarrow \mathcal{R}|\mathsf{guess}] \Pr[\mathsf{guess}] \\ = &\Pr[1 \leftarrow \mathcal{R}|\mathsf{guess}] \frac{1}{poly(\lambda)}\end{aligned}$$

by initial hypothesis. Conditioned on the fact that the adversary correctly guesses both k_0 and k_1, then for all $i \in \{1, \ldots, \lambda\}$ we have that

$$\mathbf{s}^T\mathbf{d}^{(i)} + t_0^{(i)} \in R_0^{(i)}$$

and

$$-\mathbf{s}^T\mathbf{d}^{(i)} + \mathbf{u}_{(i-1)m+L} + t_1^{(i)} \in R_1^{(i)}$$

by definition, and consequently

$$\mathbf{s}^T\mathbf{d}^{(i)} + t_0^{(i)} - \mathbf{s}^T\mathbf{d}^{(i)} + \mathbf{u}_{(i-1)m+L} + t_1^{(i)} = \mathbf{u}_{(i-1)m+L} + t_1^{(i)} + t_0^{(i)} \in R^{(i)}.$$

Thus the reduction outputs 1 with probability 1. It follows that $\Pr[1 \leftarrow \mathcal{R}] \geq 1/poly(\lambda)$.

The two bounds above show that the probability that the reduction outputs 1 differs by a non-negligible amount depending on whether $(\mathbf{A}, \mathbf{u})$ is an LWE sample or not. This contradicts the $\mathrm{LWE}_{n,q,\chi}$ assumption and concludes our proof. ■

Parameters. The above analysis fixes the following set of constraints:

- $m \geq 2 \cdot \lceil log(q) \rceil \cdot n \cdot \lambda$
- $log(q) \geq 3log(\lambda) + log(2mBc + c)$
- $c = 4$

Ignoring the constants, the constraint $O(log(q)) \geq O(log(\lambda) + log(m) + log(B))$ introduces a gap polynomial in λ in the modulo-to-noise ratio, since m is polynomially bounded. Note that the circular dependency of the first two constraints is always satisfied for a large enough q. The parameter n is free and can be set to the regime for which $\mathrm{LWE}_{n,q,\chi}$ is conjectured to be hard.

C. Upgrading Laconic Oblivious Transfer

We first upgrade the security and then the efficiency of a LOT, through generic transformations.

From Weak Sender Security to Sender Indistinguishability. We show how to generically upgrade any weak LOT into a LOT with sender indistinguishability. The compiler, shown in Figure 4, heavily relies on the tools introduced by Döttling et al. [16]: We first ensure that the adversary is not able to produce some digest d that allows him to predict k_0 for some values of c and k_1 for the others. This is done by amplifying the success probability of the adversary up to the point where it is no longer possible to be successful consistently on two different choices for the receiver's bit. Then we turn the search problem into a decision one, using the standard Goldreich-Levin hard-core predicate [19]. Here the function $\mathsf{GLEnc}(k; t)$ is defined as $\sum_{i=1}^{|k|} k_i \cdot t_i$, computed over $\mathbb{F}_2$.

The following theorem establishes our claim. The analysis is imported by the work of Döttling et al. [16] and is given in the full version.

Theorem 6 (Weak to Full Sender Security)**.** *Let* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{KGen}, \mathsf{Receive})$ *be a weak LOT. Then* $\overline{\mathsf{LOT}} = (\overline{\mathsf{Setup}}, \overline{\mathsf{Hash}}, \overline{\mathsf{KGen}}, \overline{\mathsf{Receive}})$ *as defined in Figure 4 is a standard LOT.*

From Indistinguishability Security to Context Security. We now show that any LOT with indistinguishabilty security also suffices context security. To establish this, we will use distinguisher-dependent simulation. Technically, the theorem uses the same ideas of Theorem 6.3 in [16] and some paragraphs are verbatim from [16].

Theorem 7 (Ind. to Context Security)**.** *Assume that* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{Send}, \mathsf{Receive})$ *satisfies indistinguishability security. Then it also satisfies context security.*

Proof of Theorem 7: We prove the theorem via several lemmas. In order to do so, we will first provide constructions of the relevant algorithms. Fix a PPT-context $\mathcal{Z} = (\mathcal{Z}_1, \mathcal{Z}_2)$ for LOT. We start by defining a hybrid sender algorithm as follows:

$\underline{\mathcal{S}_i(crs, d, (x_1^*, \ldots, x_i^*), L, m_0, m_1) :}$

- If $L > i$ output $\mathsf{Send}(crs, d, L, m_0, m_1)$
- Otherwise, set $m'_{x_L^*} = m_{x_L^*}$ and $m'_{1-x_L^*} = 0$, compute and output $\mathsf{Send}(crs, d, L, m'_0, m'_1)$

We introduce the following notation. For an efficiently sampleable random variable $T \in \{0, 1\}$ we use the shorthand "Compute an approximation $\tilde{\mu}$ of $\mathsf{E}[T]$ with error δ" to denote the following algorithm which computes a sample average:

- Set $N = \lceil \lambda/\delta^2 \rceil$
- For $j \in \{1, \ldots, N\}$ sample $t_j \leftarrow\!\!\$\, T$
- Output $\tilde{\mu} \leftarrow \frac{1}{N} \sum_{j=1}^{N} t_j$

We also define the following bit-extraction algorithm, which extracts an approximation of the i-th bit of the receiver's input.

$\overline{\mathsf{Setup}}(1^\lambda)$: Return $\mathsf{Setup}(1^\lambda)$.

$\overline{\mathsf{Hash}}(crs, D)$: Return $\mathsf{Hash}(crs, D)$.

$\overline{\mathsf{Send}}(crs, d, L, m_0, m_1)$:

- For all $i \in \{1, \ldots, \lambda\}$: Sample $(c^{(i)}, k_0^{(i)}, k_1^{(i)}) \leftarrow \mathsf{KGen}(crs, d, L)$.
- Set $\tilde{c} = (c^{(1)}, \ldots, c^{(\lambda)})$, $K_0 = (k_0^{(1)}, \ldots, k_0^{(\lambda)})$, and $K_1 = (k_1^{(1)}, \ldots, k_1^{(\lambda)})$.
- Sample two random strings $(t_0, t_1) \leftarrow\!\!\$\ \{0,1\}^{|K_0|+|K_1|}$.
- For all $b \in \{0,1\}$: Compute $C_b = \mathsf{GLEnc}(K_b; t_b) \oplus m_b$.
- Return $c = (\tilde{c}, C_0, C_1, t_0, t_1)$.

$\overline{\mathsf{Receive}}^{\tilde{D}}(crs, c, L)$:

- Parse c as $(\tilde{c}, C_0, C_1, t_0, t_1)$.
- For all $i \in \{1, \ldots, \lambda\}$: Compute $k^{(i)} \leftarrow \mathsf{Receive}^{\tilde{D}}(crs, c^{(i)}, L)$.
- Set $K = (k^{(1)}, \ldots, k^{(\lambda)})$.
- Return $\mathsf{GLEnc}(K; t_{D[L]}) \oplus \tilde{c}$.

Figure 4: From Weak to Full Sender Security.

$\mathsf{Extract}_i(crs, \mathsf{st}, d, \mathsf{r}_1, (x_1^*, \ldots, x_{j-1}^*), \varepsilon')$:

- Compute an approximation $\tilde{\mu}$ of $\mathsf{E}[\mathcal{Z}_2^{\mathcal{S}_{i-1}(crs,d,(x_1^*,\ldots,x_{i-1}^*),\cdot,\cdot,\cdot)}(crs, \mathsf{st})]$ with error $\varepsilon'/4$
- Compute an approximation $\tilde{\mu}_0$ of $\mathsf{E}[\mathcal{Z}_2^{\mathcal{S}_i(crs,d,(x_1^*,\ldots,x_{i-1}^*,0),\cdot,\cdot,\cdot)}(crs, \mathsf{st})]$ with error $\varepsilon'/4$
- Compute an approximation $\tilde{\mu}_1$ of $\mathsf{E}[\mathcal{Z}_2^{\mathcal{S}_i(crs,d,(x_1^*,\ldots,x_{i-1}^*,1),\cdot,\cdot,\cdot)}(crs, \mathsf{st})]$ with error $\varepsilon'/4$.
- Set $\tilde{\delta}_{i,0} \leftarrow |\tilde{\mu}_{i,0} - \tilde{\mu}_i|$
- Set $\tilde{\delta}_{i,1} \leftarrow |\tilde{\mu}_{i,1} - \tilde{\mu}_i|$
- If $\tilde{\delta}_{i,0} > \varepsilon'$ and $\tilde{\delta}_{i,1} > \varepsilon'$ abort and output $\perp$.
- else if $\tilde{\delta}_{i,1} > 2\varepsilon'$ output $x_i^* \leftarrow 0$
- Otherwise output $x_i^* \leftarrow 1$

We now define the context extractor $\mathsf{Ext}_{\mathcal{Z}}$ and the simulator Sim. The setup-simulator $\mathsf{SimSetup}$ is identical to the Setup algorithm, thus we do not specify it.

$\mathsf{Ext}_{\mathcal{Z}}(crs, \mathsf{st}, d, \mathsf{r}^*, \varepsilon)$:

- For $j \in \{1, \ldots, i\}$: Compute $x_j^* \leftarrow \mathsf{Extract}_j(crs, \mathsf{st}, d, (x_1^*, \ldots, x_{j-1}^*), \varepsilon')$
- Output $x^* \leftarrow (x_1^*, \ldots, x_n^*)$

$\mathsf{Sim}(crs, d, x^*, z = (L, m))$:

- Set $m'_{x_L^*} = m$ and $m'_{1-x_L^*} = 0$
- Compute and output $\mathsf{Send}(crs, d, L, m'_0, m'_1)$

Finally, we set choose $\varepsilon'(\varepsilon) = \varepsilon/n$. This choice of ε' will become meaningful in a moment. Assume towards contradiction that there exists an inverse polynomial ε such that it holds for infinitely many λ that

$$\left|\Pr[\mathcal{Z}(1^\lambda) = 1] - \Pr[\mathcal{EZ}(1^\lambda, \varepsilon') = 1]\right| > \varepsilon,$$

where $\varepsilon' = \varepsilon/n$. Consider the following sequence of hybrid experiments.

Hybrid $\mathcal{H}_0$: This experiment is real context $\mathcal{Z}(1^\lambda)$.

For $i \in \{1, \ldots, n\}$ define the following hybrids.

Hybrid $\mathcal{H}_i$:

- Choose random tapes $\mathsf{r}_1, \mathsf{r}_2$
- Compute $crs \leftarrow \mathsf{Setup}(1^\lambda)$
- Compute $(\mathsf{st}, d, \mathsf{aux}) \leftarrow \mathcal{Z}_1(crs, \mathsf{r}_1)$
- For $j \in \{1, \ldots, i\}$: Compute $x_j^* \leftarrow \mathsf{Extract}_j(crs, \mathsf{st}, d, (x_1^*, \ldots, x_{j-1}^*), \varepsilon')$
- Output $b^* \leftarrow \mathcal{Z}_2^{\mathcal{O}_i(\cdot,\cdot,\cdot)}(\mathsf{st}, \mathsf{r}_2)$, where the oracle $\mathcal{O}_i(L, m_0, m_1)$ returns $\mathcal{S}_i(crs, d, (x_1^*, \ldots, x_i^*), L, m_0, m_1)$

Notice that it holds that $\mathcal{H}_n$ is identically distributed to $\mathcal{EZ}(1^\lambda, \varepsilon')$. Consequently, it holds by the averaging principle that there must exist an index $i^* \in \{1, \ldots, b\}$ such that

$$|\Pr[\mathcal{H}_{i^*} = 1] - \Pr[\mathcal{H}_{i^*-1} = 1]| > \varepsilon/n.$$

We now construct an adversary against the indistinguishability security of LOT. The algorithm $\mathsf{Extract}_{i^*}(crs, \mathsf{st}, d, (x_1^*, \ldots, x_{i^*-1}^*), \varepsilon')$ computes approximations $\tilde{\delta}_{i^*,0}$ and $\tilde{\delta}_{i^*,1}$ of the advantages

$$\delta_{i^*,0} = \left| \Pr\left[\mathcal{Z}_2^{\mathcal{S}_i(crs,d,(x_1^*,\ldots,x_{i-1}^*,0),\cdot,\cdot,\cdot)}(crs, \mathsf{st}) = 1\right] - \Pr\left[\mathcal{Z}_2^{\mathcal{S}_{i^*-1}(crs,d,(x_1^*,\ldots,x_{i^*-1}^*),\cdot,\cdot,\cdot)}(crs, \mathsf{st}) = 1\right] \right|$$

$$\delta_{i^*,1} = \left| \Pr\left[\mathcal{Z}_2^{\mathcal{S}_i(crs,d,(x_1^*,\ldots,x_{i-1}^*,1),\cdot,\cdot,\cdot)}(crs, \mathsf{st}) = 1\right] - \Pr\left[\mathcal{Z}_2^{\mathcal{S}_{i^*-1}(crs,d,(x_1^*,\ldots,x_{i^*-1}^*),\cdot,\cdot,\cdot)}(crs, \mathsf{st}) = 1\right] \right|.$$

We first establish that the approximations $\tilde{\delta}_{i^*,0}$ and $\tilde{\delta}_{i^*,1}$ are close to $\delta_{i^*,0}$ and $\delta_{i^*,1}$ respectively, except with negligible probability over the coins used to compute the approximations. We establish this by a routine application of the Hoeffding bound.

Lemma 3. *Fix crs, st, d and $(x_1^*, \ldots, x_{i^*-1}^*)$. Then it holds that*

$$|\tilde{\delta}_{i^*,0} - \delta_{i^*,0}| \leq \varepsilon'$$
$$|\tilde{\delta}_{i^*,1} - \delta_{i^*,1}| \leq \varepsilon'$$

except with probability $2^{-\lambda}$ over the choice of $\mathsf{r}_{\mathsf{Extract},i^*}$.

Proof of Lemma 3: The random variable $\tilde{\mu}_{i^*}$ is the average of $N = \lceil \lambda/{\varepsilon'}^2 \rceil$ samples of

$$Y = \mathcal{Z}_2^{\mathcal{S}_{i^*-1}(crs,d,(x_1^*,\ldots,x_{i^*-1}^*),\cdot,\cdot,\cdot)}(crs, \mathsf{st}).$$

Analogously, for $b \in \{0,1\}$ the random variables $\tilde{\mu}_{i^*,b}$ is the average of $N = \lceil \lambda/{\varepsilon'}^2 \rceil$ samples of

$$Y_b = \mathcal{Z}_2^{\mathcal{S}_{i^*}(crs,d,(x_1^*,\ldots,x_{i^*-1}^*,b),\cdot,\cdot,\cdot)}(crs, \mathsf{st}).$$

We can thus write

$$\delta_{i^*,b} = |\mathsf{E}[Y_b] - \mathsf{E}[Y]|.$$

Consequently, it holds by the Hoeffding inequality (Theorem 1) that

$$\Pr[|\tilde{\mu}_{i^*} - \mathsf{E}[Y]| > \varepsilon'/2] \leq 2e^{-2N(\varepsilon'/2)^2} \leq 2e^{-\lambda}$$

and for $b \in \{0,1\}$

$$\Pr\left[|\tilde{\mu}_{i^*,b} - \mathsf{E}[Y_b]| > \varepsilon'/2\right] \leq 2e^{-2N(\varepsilon'/2)^2} \leq 2e^{-\lambda}.$$

Given that

$$|\tilde{\mu}_{i^*} - \mathsf{E}[Y]| \leq \varepsilon'/2$$
$$|\tilde{\mu}_{i^*,0} - \mathsf{E}[Y_0]| \leq \varepsilon'/2$$
$$|\tilde{\mu}_{i^*,1} - \mathsf{E}[Y_1]| \leq \varepsilon'/2$$

and using that

$$\tilde{\delta}_{i^*,0} = |\tilde{\mu}_{i^*,0} - \tilde{\mu}_{i^*}|$$
$$\tilde{\delta}_{i^*,1} = |\tilde{\mu}_{i^*,1} - \tilde{\mu}_{i^*}|$$

we get that

$$|\tilde{\delta}_{i^*,0} - \delta_{i^*,0}| \leq \varepsilon'$$

and

$$|\tilde{\delta}_{i^*,1} - \delta_{i^*,1}| \leq \varepsilon'.$$

Consequently, it holds by a union-bound that

$$\Pr\left[|\tilde{\delta}_{i^*,0} - \delta_{i^*,0}| > \varepsilon' \text{ or } |\tilde{\delta}_{i^*,1} - \delta_{i^*,1}| > \varepsilon'\right] \leq 6 \cdot e^{-\lambda} \leq 2^{-\lambda}$$

which concludes the proof. ∎

Lemma 4. *Assume that*

$$|\Pr[\mathcal{H}_{i^*} = 1] - \Pr[\mathcal{H}_{i^*-1} = 1]| \geq \varepsilon'$$

for an $i^* \in \{1, \ldots, n\}$. Then there exists a PPT adversary $\mathcal{B}$ which breaks the indistinguishability sender security of LOT.

Proof of Lemma 4: First note that $\mathcal{H}_{i^*-1}$ and $\mathcal{H}_{i^*}$ are identical until $x_{i^*}^* \leftarrow \mathsf{Extract}_j(crs, \mathsf{st}, d, (x_1^*, \ldots, x_{i^*-1}^*), \varepsilon')$ is computed in $\mathcal{H}_{i^*}$. We therefore first rephrase $\mathcal{H}_{i^*-1}$ and $\mathcal{H}_{i^*}$. Towards this goal, define an alternate first stage $\mathcal{Z}_1'$ by

$\underline{\mathcal{Z}_1'(crs, \mathsf{r}_1):}$

- Compute $(\mathsf{st}, d) \leftarrow \mathcal{Z}_1'(crs, \mathsf{r}_1)$
- For $j \in \{1, \ldots, i^* - 1\}$:Compute $x_j^* \leftarrow \mathsf{Extract}_j(crs, \mathsf{st}, d, (x_1^*, \ldots, x_{j-1}^*), \varepsilon')$
- Output $(\mathsf{st}, d, \mathsf{aux}, (x_1^*, \ldots, x_{i^*}^*))$

We can now rephrase $\mathcal{H}_{i^*-1}$ and $\mathcal{H}_{i^*}$ by cutting of the common prefix.

Hybrid $\mathcal{H}'_{i^*}(crs, \mathsf{st}, d, \mathsf{aux}, (x_1^*, \ldots, x_{i^*-1}^*))$:

- Compute $(\mathsf{st}, d, \mathsf{aux}, (x_1^*, \ldots, x_{i^*-1}^*)) \leftarrow \mathcal{Z}'_1(crs, \mathsf{r}_1)$
- Compute $b^* \leftarrow \mathcal{Z}'^{\mathcal{S}_{i^*-1}(crs,d,(x_1^*,\ldots,x_{i^*-1}^*),\cdot,\cdot,\cdot)}_2(\mathsf{st}, \mathsf{r}_2)$

Hybrid $\mathcal{H}'_{i^*}(crs, \mathsf{st}, d, \mathsf{aux}, (x_1^*, \ldots, x_{i^*-1}^*))$:

- Compute $(\mathsf{st}, d, \mathsf{aux}, (x_1^*, \ldots, x_{i^*-1}^*)) \leftarrow \mathcal{Z}_1'(crs, \mathsf{r}_1)$
- Compute $x_{i^*}^* \leftarrow \mathsf{Extract}_{i^*}(crs, \mathsf{st}, d, (x_1^*, \ldots, x_{i^*-1}^*), \varepsilon')$
- Compute $b^* \leftarrow \mathcal{Z}'^{\mathcal{S}_{i^*}(crs,d,(x_1^*,\ldots,x_{i^*}^*),\cdot,\cdot,\cdot)}_2(\mathsf{st}, \mathsf{r}_2)$

Further, fix crs, st, d, $(x_1^*, \ldots, x_{i^*-1}^*)$ and the random tape $\mathsf{r}_{i^*}^*$ used for $\mathsf{Extract}_{i^*}$ and collect them in a variable z. Note that the parameters in inp also determine $x_{i^*}^*$. We now define three events $\mathsf{gap}(\mathsf{inp})$, $\mathsf{approx}(\mathsf{inp})$ and $\mathsf{good}(\mathsf{inp})$ which only depend on inp.

1) $\mathsf{gap}(\mathsf{inp})$ holds, if and only if

$$|\Pr[\mathcal{H}'_{i^*}(\mathsf{inp}) = 1] - \Pr[\mathcal{H}'_{i^*-1}(\mathsf{inp}) = 1]| > 4\varepsilon',$$

where the random choices are over the random coins r_2 of $\mathcal{Z}_2$ and the random choices made by the oracles $\mathcal{S}_{i^*-1}$ and $\mathcal{S}_{i^*}$.

2) Let $\tilde{\delta}_{i^*,0}$ and $\tilde{\delta}_{i^*,1}$ be the approximated values computed during the execution of $\mathsf{Extract}_{i^*}(crs, \mathsf{st}, d, (x_1^*, \ldots, x_{i^*-1}^*), \varepsilon')$. $\mathsf{approx}(\mathsf{inp})$ holds, if and only if

$$|\tilde{\delta}_{i^*,0} - \delta_{i^*,0}| \leq \varepsilon'$$
$$|\tilde{\delta}_{i^*,1} - \delta_{i^*,1}| \leq \varepsilon'$$

3) $\mathsf{good}(\mathsf{inp})$ holds if and only if

$$\delta_{i^*,0} > \varepsilon'$$
$$\delta_{i^*,1} > \varepsilon'.$$

We first elaborate on the events in more detail. The event $\mathsf{gap}(\mathsf{inp})$ characterizes that for *the same* choice of inp, the hybrids $\mathcal{H}_{i^*}(\mathsf{inp})$ and $\mathcal{H}_{i^*-1}(\mathsf{inp})$ have distance at least $4\varepsilon'$. Notice that the extracted prefix $(\bar{x}_1, \ldots, \bar{x}_{i^*-1})$ is identical in both experiments $\mathcal{H}_{i^*}(\mathsf{inp})$ and $\mathcal{H}_{i^*-1}(\mathsf{inp})$. Consequently, $\mathsf{gap}(\mathsf{inp})$ immediately implies that none of the $x^*_1, \ldots, x^*_{i^*-1}$ was set to $\bot$, as this would imply that the two experiments are identically distributed.

The event $\mathsf{approx}(\mathsf{inp})$ ensures that the approximations $\tilde{\delta}_{i^*,0}$ and $\tilde{\delta}_{i^*,1}$ are sufficiently close to the true advantages. Finally, the event $\mathsf{good}(\mathsf{inp})$ ensures that inp is such that we will be able to mount a successful attack against indistinguishability sender security of LOT. Our first goal will be to show that the event $\mathsf{good}(\mathsf{inp})$ holds with reasonably high probability over the choice of inp. Once this is established, we will construct an adversary $\mathcal{B}$ against the indistinguishability sender security of LOT. Observe that by Lemma 3 it holds that

$$\Pr_{\mathsf{inp}}[\neg\mathsf{approx}(\mathsf{inp})] \leq 2^{-\lambda}. \tag{1}$$

Since

$$|\Pr_{\mathsf{inp}}[\mathcal{H}_{i^*}(\mathsf{inp}) = 1] - \Pr_{\mathsf{inp}}[\mathcal{H}_{i^*-1}(\mathsf{inp}) = 1]| =$$
$$|\Pr[\mathcal{H}_{i^*} = 1] - \Pr[\mathcal{H}_{i^*-1} = 1]| \geq 8 \cdot \varepsilon'$$

it holds by the Markov inequality for advantages (Lemma 1) that

$$\begin{aligned} &\Pr_{\mathsf{inp}}[\mathsf{gap}(\mathsf{inp})] \\ &= \Pr_{\mathsf{inp}}\left[\left| \begin{array}{l} \Pr[\mathcal{H}_{i^*}(\mathsf{inp}) = 1] - \\ \Pr[\mathcal{H}_{i^*-1}(\mathsf{inp}) = 1] \end{array} \right| > 4\varepsilon'\right] \\ &\geq 4\varepsilon'. \end{aligned} \tag{2}$$

We now show that if $\mathsf{gap}(\mathsf{inp})$ holds, then it must either hold $\mathsf{good}(\mathsf{inp})$ or not $\mathsf{approx}(\mathsf{inp})$. We will establish this by showing that $\neg\mathsf{good}(\mathsf{inp})$ and $\mathsf{approx}(\mathsf{inp})$ imply $\neg\mathsf{gap}(\mathsf{inp})$. Thus, fix $\mathsf{inp} = (crs, \mathsf{st}, d, (x^*_1, \ldots, x^*_{i^*-1}), \mathsf{r}^*_{i^*})$ with $\neg\mathsf{good}(\mathsf{inp})$ and $\mathsf{approx}(\mathsf{inp})$. From $\neg\mathsf{good}(\mathsf{inp})$ it follows that there is a $\beta \in \{0,1\}$ such that

$$\left| \begin{array}{l} \Pr[\mathcal{Z}_2^{\mathcal{S}_i(crs,d,(x^*_1,\ldots,x^*_{i-1},\beta),\cdot,\cdot,\cdot)}(crs, \mathsf{st}) = 1] - \\ \Pr[\mathcal{Z}_2^{\mathcal{S}_{i^*-1}(crs,d,(x^*_1,\ldots,x^*_{i^*-1}),\cdot,\cdot,\cdot)}(crs, \mathsf{st}) = 1] \end{array} \right| \leq \varepsilon'.$$

We are now going to show that under this condition $\mathsf{Extract}_{i^*}(crs, \mathsf{st}, d, (x^*_1, \ldots, x^*_{i^*-1}), \varepsilon')$ will be able to identify the correct $x^*_{i^*}$. Observe that since it holds that $\mathsf{approx}(\mathsf{inp})$, we get that

$$\tilde{\delta}_{i^*,\beta} \leq \delta_{i^*,\beta} + \varepsilon' \leq 2\varepsilon'.$$

Consequently, $\mathsf{Extract}_{i^*}(crs, \mathsf{st}, d, (x^*_1, \ldots, x^*_{i^*-1}), \varepsilon')$ will not output $\bot$. We distinguish two cases.

Case 1: In this case it holds that

$$\delta_{i^*,1-\beta} \leq 4\varepsilon'.$$

It follows immediately that

$$\delta_{i^*,x^*_{i^*}} \leq 4\varepsilon',$$

regardless which $x^*_{i^*} \in \{0,1\}$ is chosen.

Case 2: In this case it holds that

$$\delta_{i^*,1-\beta} > 4\varepsilon'.$$

Again since it holds that $\mathsf{approx}(\mathsf{inp})$, we get that

$$\tilde{\delta}_{i^*,1-\beta} \geq \delta_{i^*,1-\beta} - \varepsilon' \geq 3\varepsilon' > 2\varepsilon'.$$

Consequently, $\mathsf{Extract}_{i^*}(crs, \mathsf{st}, d, (x^*_1, \ldots, x^*_{i^*-1}), \varepsilon')$ will set $\bar{x}_{i^*} \leftarrow \beta$ and again we can conclude

$$\delta_{i^*,x^*_{i^*}} \leq 4\varepsilon',$$

Furthermore, observe that, since $\mathsf{Extract}_{i^*}(crs, \mathsf{st}, d, (x^*_1, \ldots, x^*_{i^*-1}), \varepsilon')$ will not output $\bot$, the output of $\mathcal{H}_{i^*}(\mathsf{inp})$ is distributed according to ${\mathcal{Z}'}_2^{\mathcal{S}_{i^*}(crs,d,(x^*_1,\ldots,x^*_{i^*}),\cdot,\cdot,\cdot)}(\mathsf{st}, \mathsf{r}_2)$. We also know that $\mathcal{H}_{i^*-1}(\mathsf{inp})$ is distributed according to ${\mathcal{Z}'}_2^{\mathcal{S}_{i^*-1}(crs,d,(x^*_1,\ldots,x^*_{i^*-1}),\cdot,\cdot,\cdot)}(\mathsf{st}, \mathsf{r}_2)$. This implies that

$$\begin{aligned} &|\Pr[\mathcal{H}_{i^*}(\mathsf{inp}) = 1] - \Pr[\mathcal{H}_{i^*-1}(\mathsf{inp}) = 1]| = \\ &\left| \begin{array}{l} {\mathcal{Z}'}_2^{\mathcal{S}_{i^*}(crs,d,(x^*_1,\ldots,x^*_{i^*}),\cdot,\cdot,\cdot)}(\mathsf{st}, \mathsf{r}_2) - \\ {\mathcal{Z}'}_2^{\mathcal{S}_{i^*-1}(crs,d,(x^*_1,\ldots,x^*_{i^*-1}),\cdot,\cdot,\cdot)}(\mathsf{st}, \mathsf{r}_2) \end{array} \right| \\ &\leq 4\varepsilon', \end{aligned}$$

which in turn implies that $\neg\mathsf{gap}(\mathsf{inp})$. Thus, we have established that

$$\mathsf{gap}(\mathsf{inp}) \Rightarrow \mathsf{good}(\mathsf{inp}) \text{ or } \neg\mathsf{approx}(\mathsf{inp}). \tag{3}$$

From (2), (3) and (1) we obtain that

$$\begin{aligned} 4\varepsilon' &\leq \Pr[\mathsf{gap}(\mathsf{inp})] \\ &\leq \Pr[\mathsf{good}(\mathsf{inp}) \text{ or } \neg\mathsf{approx}(\mathsf{inp})] \\ &\leq \Pr[\mathsf{good}(\mathsf{inp})] + \Pr[\neg\mathsf{approx}(\mathsf{inp})] \\ &\leq \Pr[(\mathsf{good}(\mathsf{inp})] + 2^{-\lambda}, \end{aligned}$$

where the third inequality follows by the union-bound. This implies that

$$\Pr_{\mathsf{inp}}[\mathsf{good}(\mathsf{inp})] \geq 4\varepsilon' - 2^{-\lambda} > \varepsilon'.$$

We are now ready to construct an adversary $\mathcal{B}$ against the indistinguishability sender privacy of LOT. The adversary $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ is given as follows. In abuse of notation, we assume that $\mathcal{B}$ is stateful, i.e. the second stage $\mathcal{B}_2$ remembers all variables of the first stage $\mathcal{B}_1$. $\mathcal{B}$ essentially simulates $\mathcal{H}_{i^*-1}$, with the modification that calls to $\mathsf{Send}(crs, d, i^*, \cdot, \cdot)$ of $\mathcal{S}_{i^*-1}$ will be forwarded to $\mathcal{B}$'s oracle.

$\underline{\mathcal{B}_1(crs; \mathsf{r}_1)}$:

- Compute $(\mathsf{st}, d) \leftarrow \mathcal{Z}_1'(crs, \mathsf{r}_1)$
- For $j \in \{1, \ldots, i^* - 1\}$: Compute $x_j^* \leftarrow \mathsf{Extract}_j(crs, \mathsf{st}, d, (x_1^*, \ldots, x_{j-1}^*), \varepsilon')$
- Output (d, i^*)

$\underline{\mathcal{B}_2^{\mathcal{O}(\cdot,\cdot)}(\mathsf{st}, \mathsf{r}_2)}$:

- Compute and output $b^* \leftarrow \mathcal{Z}_2^{\mathcal{S}_{i^*}(\cdot,\cdot,\cdot,\cdot)}$, where the oracle $\mathcal{S}_{i^*}(L, m_0, m_1)$ is implemented as follows:
  - If $L > i^*$ output $\mathsf{Send}(crs, d, L, m_0, m_1)$
  - If $L < i^*$, set $m'_{x_L^*} = m_{x_L^*}$ and $m'_{1-x_L^*} = 0$ output $\mathsf{Send}(crs, d, L, m'_0, m'_1)$
  - If $L = i^*$ query and output $\mathcal{O}(m_0, m_1)$

Fix crs and r_1. We distinguish 3 cases.

Case 1: The oracle $\mathcal{O}(m_0, m_1)$ computes the function $\mathsf{Send}(crs, d, i^*, m_0, m_1)$. It follows by inspection that in this case the output of $\mathcal{B}$ is distributed according to $\mathcal{Z}_2^{\mathcal{S}_{i^*-1}(crs,d,(x_1^*,\ldots,x_{i^*-1}^*),\cdot,\cdot,\cdot)}(crs, \mathsf{st})$.

Case 2: The oracle $\mathcal{O}(m_0, m_1)$ computes the function $\mathsf{Send}(crs, d, i^*, m_0, 0)$. It follows by inspection that in this case the output of $\mathcal{B}$ is distributed according to $\mathcal{Z}_2^{\mathcal{S}_{i^*}(crs,d,(x_1^*,\ldots,x_{i-1}^*,0),\cdot,\cdot,\cdot)}(crs, \mathsf{st})$.

Case 3: The oracle $\mathcal{O}(m_0, m_1)$ computes the function $\mathsf{Send}(crs, d, i^*, m_0, 1)$. It follows by inspection that in this case the output of $\mathcal{B}$ is distributed according to $\mathcal{Z}_2^{\mathcal{S}_{i^*}(crs,d,(x_1^*,\ldots,x_{i-1}^*,1),\cdot,\cdot,\cdot)}(crs, \mathsf{st})$.

We conclude that the advantage of the adversary (fixing the bit b) is equal to

$$\left| \begin{array}{l} \Pr[\mathcal{Z}_2^{\mathcal{S}_{i^*}(crs,d,(x_1^*,\ldots,x_{i-1}^*,b),\cdot,\cdot,\cdot)}(crs, \mathsf{st}) = 1] - \\ \Pr[\mathcal{Z}_2^{\mathcal{S}_{i^*-1}(crs,d,(x_1^*,\ldots,x_{i^*-1}^*),\cdot,\cdot,\cdot)}(crs, \mathsf{st}) = 1] \end{array} \right|.$$

This implies that for both $b \in \{0, 1\}$ the advantage of the adversary is

$$\Pr_{crs,\mathsf{r},\mathsf{r}^*}[\mathsf{good}(crs, \mathsf{r}_{\mathcal{A}}, \mathsf{r}_{\mathsf{Extract}})] > \varepsilon',$$

which contradicts the indistinguishability sender privacy of LOT. ∎

∎

Weak to Full Efficiency. The following construction shows that it suffices to consider LOT schemes with 2-to-1 compressing Hash and algorithms Send and $\mathsf{Receive}$ with polynomial running times in the size of the database. A similar statement already appeared in the work of Cho et al. [3], however their analysis crucially relies on the semi-honest settings and does not directly extend to malicious security. In the following we recall a simplified version of the transformation from [3]. We assume without loss of generality that the database has $2^d \cdot \lambda$ locations, which can be indexed with strings of the form $b_1 \| \ldots \| b_{d-1} \| t$ where the bits b_i define the root-to-leaf path and the string t define the position of the leaf. For ease of exposition we overload the notation for the sender algorithm and we write $\mathsf{Send}(crs, d, K)$ as $\mathsf{Send}(crs, d, 1, (K_0^1, K_1^1)) \| \ldots \| \mathsf{Send}(crs, d, 2\lambda, (K_0^{2\lambda}, K_1^{2\lambda}))$ where K is consists of 2λ pairs of λ-bit strings. The same shortcut is used for the receiver algorithm. The construction (given in Figure 7) consists of a chaining of garbled circuits that allows the receiver to traverse the tree downwards via the selected path. Since the sender only knows the root of the tree, the generation of subsequent Send algorithms is deferred by garbling the circuit $\mathcal{C}_{\mathsf{trav}}$ (Figure 6). Once the leaf is reached, the choice of the message can be constrained on the desired bit via the circuit $\mathcal{C}_{\mathsf{read}}$ (Figure 5).

Hardwired Values: (t, m_0, m_1).
Input: $(N_0^1 \| N_1^1, \ldots, N_0^{d-1} \| N_{d-1}^1, N^d)$.
- Return $m_{N^d[t]}$

Figure 5: Circuit $\mathcal{C}_{\mathsf{leaf}}[m_0, m_1]$.

Hardwired Values: (crs, b, K, r).
Input: (N_0, N_1).
- Compute $e \leftarrow \mathsf{Send}(crs, N_b, K; r)$.
- Return e.

Figure 6: Circuit $\mathcal{C}_{\mathsf{trav}}[crs, b, K, r]$.

The following theorem shows that considering a LOT with 2-to-1 compression factor suffices.

Theorem 8 (Weak to Full Efficiency)**.** *Let* $\mathsf{LOT} = (\mathsf{Setup}, \mathsf{Hash}, \mathsf{Send}, \mathsf{Receive})$ *be a LOT with* 2*-to-*1 *compression. Then* $\overline{\mathsf{LOT}} = (\overline{\mathsf{Setup}}, \overline{\mathsf{Hash}}, \overline{\mathsf{Send}}, \overline{\mathsf{Receive}})$ *as defined in Figure 7 is LOT with arbitrary compression.*

Proof of Theorem 8: We start with the high-level overview. Our strategy is to associate every node in the Merkle tree with two hybrids. In this sequence of hybrids we start in the root node traverse the tree in such a way that it is ensure that once we reach a node

$\overline{\mathsf{Setup}}(1^\lambda)$: Return $\mathsf{Setup}(1^\lambda)$.

$\overline{\mathsf{Hash}}(crs, D)$:

- Build a Merkle tree $\mathbf{D}$ of D using the function $\mathsf{Hash}(crs, \cdot)$.
- Return $(\rho, \mathbf{D})$, where ρ is the root of $\mathbf{D}$

$\overline{\mathsf{Send}}(crs, d, L, m_0, m_1)$:

- Parse L as $b_1 \| \ldots \| b_{d-1} \| t$
- For $j \in \{d, \ldots, 1\}$:
  - Sample $r_i \leftarrow\!\!\$\ \{0,1\}^\lambda$
  - If $j = 1$: Compute $e_0 \leftarrow \mathsf{Send}(crs, d, K^1)$
  - If $j = d$: Compute $(\tilde{\mathcal{C}}_d, K^d) \leftarrow \mathsf{Garble}(1^\lambda, \mathcal{C}_{\mathsf{leaf}}[t, m_0, m_1])$
  - Otherwise, compute $(\tilde{\mathcal{C}}_i, K^i) \leftarrow \mathsf{Garble}(1^\lambda, \mathcal{C}_{\mathsf{trav}}[crs, b_i, K^{i+1}, r_i])$
- Return $c = (e_0, \tilde{\mathcal{C}}_1, \ldots, \tilde{\mathcal{C}}_d)$.

$\overline{\mathsf{Receive}}^{\tilde{D}}(crs, c, L)$:

- Parse c as $(e_0, \tilde{\mathcal{C}}_1, \ldots, \tilde{\mathcal{C}}_d)$ and L as $b_1 \| \ldots \| b_{d-1} \| t$.
- Denote the end note of the path $b_1 b_2 \ldots b_i$ by $\mathbf{D}_{b_1 b_2 \ldots b_i}$.
- For all $i \in \{1, \ldots, d-1\}$, compute:
  - $M^i \leftarrow \mathsf{Receive}(crs, e_{i-1}, \mathbf{D}_{b_1 b_2 \ldots b_{i-1} 0} \| \mathbf{D}_{b_1 b_2 \ldots b_{i-1} 1})$.
  - $e_i \leftarrow \mathsf{Eval}(\mathcal{C}_i, M^i)$.
- Compute $M^d \leftarrow \mathsf{Receive}(crs, e_{d-1}, \mathbf{D}_{b_1 b_2 \ldots b_{d-1} 0} \| \mathbf{D}_{b_1 b_2 \ldots b_{d-1} 1})$.
- Return $m = \mathsf{Eval}(\tilde{\mathcal{C}}_{\mathsf{leaf}}, (M^i, \ldots, M^d))$.

Figure 7: From 2-to-1 to Arbitrary Compression.

ν, its parent has already been traversed, e.g. via breadth-first search or depth-first search from the root. This will ensure that once we reach hybrid $\mathcal{H}_\nu$, the LOT message corresponding to the parent of ν has been extracted. In each node ν, we will first switch the garbled circuit to simulation, then extract the digest d_ν at this node.

There are $2^d - 1$ nodes. Let the sequence of nodes, in the order in which they are interated be $\nu_1, \ldots, \nu_{2^d-1}$. For shorthand, let $\ell = 2^d - 1$. For ease of notation we omit random coins in the notation of $\mathcal{Z}_1, \mathcal{Z}_2$ and $\mathsf{Ext}_{\mathcal{Z}}$. We first provide the hybrids and derive the simulators and the extractor from the last hybrid. Fix a context $\mathcal{Z} = (\mathcal{Z}_1, \mathcal{Z}_2)$.

Hybrid $\mathcal{H}_0(\delta)$: This is the real experiment. Specifically,

- $crs \leftarrow \mathsf{Setup}(1^\lambda)$
- $(\mathsf{st}, d) \leftarrow \mathcal{Z}_1(crs)$
- Compute and output $b^* \leftarrow \mathcal{Z}_2^{\mathcal{O}^{(0)}(\cdot)}(\mathsf{st})$ where the oracle $\mathcal{O}^{(0)}(L, m_0, m_1)$ implements $\overline{\mathsf{Send}}(crs, d, L, m_0, m_1)$

Hybrid $\mathcal{H}_1(\delta)$: In this experiment we will extract d_{ν_1} as a LOT hash (i.e. 2-to-1). First we rephrase $\mathcal{H}_0$ as a context $(\mathcal{Z}_1^{(1)}, \mathcal{Z}_2^{(1)})$ for LOT. We will set $\mathcal{Z}_1^{(1)} = \mathcal{Z}_1$ but interpret $d = d_{\nu_1}$ as a hash value for LOT. Moreover, we interpret $\mathcal{Z}_2^{\mathcal{O}^{(0)}(\cdot)}$ as a single machine $\mathcal{Z}_2^{(1)}$, except for calls by $\mathcal{O}^{(0)}$ to $\mathsf{Send}(crs, d, \cdot)$, which will be implemented by oracle calls. Consequently, $\mathcal{Z}_2^{\mathcal{O}^{(0)}(\cdot)}$ and $\mathcal{Z}_2^{(1)^{\mathsf{Send}(crs,d,\cdot)}}$ compute identical functions. Context security of LOT yields an an extractor $\mathsf{Ext}_{\mathcal{Z}}^{(1)}$ and a simulator Sim, which we will use in the construction of $\mathcal{H}_1$. We can now define $\mathcal{H}_1$ as follows.

- $crs \leftarrow \mathsf{Setup}(1^\lambda)$
- $(\mathsf{st}, d) \leftarrow \mathcal{Z}_1(crs)$
- Let $\nu_{1,l}$ be the left child node of the root node ν_1 and $\nu_{1,r}$ be the right child node of the root node ν_1.
- $(d_{\nu_{1,l}} \| d_{\nu_{1,r}}, \mathsf{aux}_1) \leftarrow \mathsf{Ext}_{\mathcal{Z}}^{(1)}(crs, \mathsf{st}, d_{\nu_1}, \delta)$
- Compute and output $b^* \leftarrow \mathcal{Z}_2^{\mathcal{O}^{(1)}(\cdot)}(\mathsf{st})$ where $\mathcal{O}^{(1)}$ is identical to $\mathcal{O}^{(0)}$, except that we replace calls to $\mathsf{Send}(crs, d_{\nu_1}, K)$ for any K by calls to $\mathsf{Sim}(crs, d, \mathsf{aux}_{root}, K_{d_{\nu_{1,l}} \| d_{\nu_{1,r}}})$

Define the following hybrids $i \in \{1, \ldots, 2^d - 1\}$.

Hybrid $\mathcal{H}_{2i}(\delta)$: $\mathcal{O}^{(2i)}$ is identical to $\mathcal{O}^{(2i-1)}$, except for the following changes. For a given index L let $\nu'_1, \ldots, \nu'_d$ be the root-to-leaf path for L, where ν'_1 is the root node and ν'_l a leaf node. If $\nu_i = \nu'_d$ (i.e. ν_i is the leaf-node in the path for index L), we replace the instruction $(\tilde{\mathcal{C}}_d, K^d) \leftarrow \mathsf{Garble}(1^\lambda, \mathcal{C}_{\mathsf{leaf}}[t, m_0, m_1])$ by $(\tilde{\mathcal{C}}_d, K^d) \leftarrow \mathsf{GCSim}(1^\lambda, m_{D[t]})$.

Otherwise, at node $\nu_i = \nu'_j$ for some j we

replace the instruction $(\tilde{\mathcal{C}}_i, K^j) \leftarrow \mathsf{Garble}(1^\lambda, \mathcal{C}_{\mathsf{trav}}[crs, b_j, K^{j+1}, r_j])$ by $(\tilde{\mathcal{C}}_{\mathsf{trav}}, K^j_{\nu_{i,l}\|\nu_{i,r}}) \leftarrow \mathsf{GCSim}(1^\lambda, \mathsf{Send}(crs, N_{b_j}, K; r))$.

Hybrid $\mathcal{H}_{2i+1}(\delta)$: In this experiment we will extract d_{ν_i} as a LOT hash (i.e. 2-to-1). First we rephrase $\mathcal{H}_{2i}$ as a context for LOT. That is, we define a context $\mathcal{Z}_1^{(i)}$ which outputs a state st' comprising of a state st and all LOT hash preimages extracted so far. It also outputs the hash d_{ν_i} at node ν_i.

Moreover, we interpret $\mathcal{Z}_2^{\mathcal{O}^{(2i)}(\cdot)}$ as a single machine $\mathcal{Z}_2^{(i)}$, except for calls by $\mathcal{O}^{(2i)}$ to $\mathsf{Send}(crs, d_{\nu_i}, \cdot)$, which will be implemented by oracle calls. Consequently, $\mathcal{Z}_2^{\mathcal{O}^{(2i)}(\cdot)}$ and $\mathcal{Z}_2^{(i)^{\mathsf{Send}(crs, d_{\nu_i}, \cdot)}}$ compute identical functions. Context security of LOT yields an an extractor $\mathsf{Ext}_{\mathcal{Z}}^{(i)}$ and a simulator Sim, which we will use in the construction of $\mathcal{H}_{2i+1}$. We can now define $\mathcal{H}_{2i+1}$ as follows.

- $crs \leftarrow \mathsf{Setup}(1^\lambda)$
- $(\mathsf{st}, d) \leftarrow \mathcal{Z}_1(crs)$
- For $j \in \{0, \ldots, i\}$:
  - Let $\nu_{j,l}$ be the left child of node ν_j and $\nu_{j,r}$ be the right child of node ν_j.
  - $(d_{\nu_{j,l}}\|d_{\nu_{j,r}}, \mathsf{aux}_j) \leftarrow \mathsf{Ext}_{\mathcal{Z}}^{(j)}(crs, (\mathsf{st}, (d_{\nu_k}, d_{\nu_{k,l}}, d_{\nu_{k,r}}, \mathsf{aux}_k)_{k\le j}), d_{\nu_j}, \delta)$
- Compute and output $b^* \leftarrow \mathcal{Z}_2^{\mathcal{O}^{(2i+1)}(\cdot)}(\mathsf{st})$, where $\mathcal{O}^{(2i+1)}$ is identical to $\mathcal{O}^{(2i)}$, except that we replace calls to $\mathsf{Send}(crs, d_{\nu_i}, K)$ for any K by calls to $\mathsf{Sim}(crs, d, \mathsf{aux}_i, K_{d_{\nu_{i,l}}\|d_{\nu_{i,r}}})$

In the last hybrid $\mathcal{H}_{2\ell+1})(\delta)$, the oracle $\mathcal{O}^{2\ell+1}$ does not make any calls to Send, but only calls to Sim. We can therefore now define the extractor $\mathsf{Ext}_{\mathcal{Z}}$ and the simulator $\overline{\mathsf{Sim}}$.

$\mathsf{Ext}_{\mathcal{Z}}(crs, \mathsf{st}, d, \delta)$:

- For $j \in \{0, \ldots, \ell\}$:
  - Let $\nu_{j,l}$ be the left child of node ν_j and $\nu_{j,r}$ be the right child of node ν_j.
  - $(d_{\nu_{j,l}}\|d_{\nu_{j,r}}, \mathsf{aux}_j) \leftarrow \mathsf{Ext}_{\mathcal{Z}}^{(j)}(crs, (\mathsf{st}, (d_{\nu_k}, d_{\nu_{k,l}}, d_{\nu_{k,r}}, \mathsf{aux}_k)_{k\le j}), d_{\nu_j}, \delta)$
  - Set $D = (D_1, \ldots, D_{2^d})$
  - Set $\mathsf{aux} = (\mathsf{aux}_j)_j$
  - Output (D, aux)

$\overline{\mathsf{Sim}}(crs, d, \mathsf{aux}, (L, z))$:

- Let $\nu'_1, \ldots, \nu'_d$ be the root-to-leaf path corresponding to L
- Compute $(\tilde{\mathcal{C}}_d, K^d_{d_{\nu'_d}}) \leftarrow \mathsf{GCSim}(1^\lambda, m_{D[t]})$
- For $j \in \{d-1, \ldots, 1\}$:
  - Let $\nu_{j,l}$ be the left child of node ν'_j and $\nu_{j,r}$ be the right child of node ν'_j.
  - Compute $(\tilde{\mathcal{C}}_{\mathsf{trav}}, K^j) \leftarrow \mathsf{GCSim}(1^\lambda, \mathsf{Sim}(crs, d, \mathsf{aux}_k, K^{j+1}_{d_{\nu_{j,l}}\|d_{\nu_{j,r}}}))$
- Compute $e \leftarrow \mathsf{Sim}(crs, d, \mathsf{aux}_k, K^1_{d_{\nu_{0,l}}\|d_{\nu_{0,r}}})$

As there are $2\ell+1$ hybrid steps, we set $\delta = \varepsilon/(2\ell+1)$. Now assume towards contradiction that there exists an inverse polynomial ε such that it holds for infinitely many 1^λ that

$$|\Pr[\mathcal{Z}(crs) = 1] - \Pr[\mathcal{EZ}(crs, \varepsilon/(2\ell+1))]| > \varepsilon.$$

As $\mathcal{H}_0(\varepsilon/(2\ell+1))$ is identical to $\mathcal{Z}(crs)$ and $\mathcal{H}_{2\ell+1}(\varepsilon/(2\ell+1))$ is identical to $\mathcal{EZ}(crs, \varepsilon/(2\ell+1))$, by the averaging principle there exists an $k \in \{1, \ldots, 2\ell+1\}$ such that

$$|\Pr[\mathcal{H}_k(\varepsilon/(2\ell+1)) = 1] - \Pr[\mathcal{H}_{k-1}(\varepsilon/(2\ell+1))]| > \varepsilon/(2\ell+1).$$

We now show that this leads to a contradiction for every $k \in \{1, \ldots, 2\ell+1\}$.

Case $k = 1$: In this case we show that this leads to a contradiction against the context-security of LOT. For hybrid $\mathcal{H}_1$ we constructed an LOT context $\mathcal{Z}_1 = (\mathcal{Z}_1^{(1)}), \mathcal{Z}_2^{(1)}$ which produces an output that is identically distributed to $\mathcal{H}_0$. $\mathcal{H}_1$ is constructed in a way such that its output is identically distributed to that of $\mathcal{EZ}$. Consequently, it holds that

$$|\Pr[\mathcal{Z}_1^{(1)} = 1] - \Pr[\mathcal{EZ}^{(1)}(\varepsilon/(2\ell+1))) = 1]| = |\Pr[\mathcal{H}_1(\varepsilon/(2\ell+1)) = 1] - \Pr[\mathcal{H}_0(\varepsilon/(2\ell+1))]| > \varepsilon/(2\ell+1),$$

which contradicts the context-security of LOT.

Case $k = 2i$: For this case we get a contradiction via a routine application of simulation security of the garbling scheme. In more detail, in hybrid $\mathcal{H}_{2i-1}$ there is only one set of input labels at node ν_i which $\mathcal{Z}_2$ can obtain. Consequently, we can apply security of the garbling scheme in a hybrid fashion, i.e. once for each call to the oracle and the contradiction follows.

Case $k = 2i+1$: In this case we show that this leads to a contradiction against the context-security of LOT. We proceed as in the case $k = 1$. More specifically, to define hybrid $\mathcal{H}_{2i+1}$ we have rephrased $\mathcal{H}_{2i}$ as a context $\mathcal{Z}^{(i)} = (\mathcal{Z}_1^{(i)}, \mathcal{Z}_2^{(i)})$ for LOT. By the way we have constructed $\mathcal{H}_{2i+1}$ it holds that the output of $\mathcal{H}_{2i+1}$ is identically distributed to $\mathcal{EZ}^{(i)}$. Consequently, it holds

that

$$|\Pr[\mathcal{Z}_1^{(i)} = 1] - \Pr[\mathcal{E}\mathcal{Z}^{(i)}(\varepsilon/(2\ell+1))) = 1]| = |\Pr[\mathcal{H}_{2i}(\varepsilon/(2\ell+1)) = 1] - \Pr[\mathcal{H}_{2i-1}(\varepsilon/(2\ell+1))]| > \varepsilon/(2\ell+1),$$

which contradicts the context-security of LOT. This concludes the proof. ■

ACKNOWLEDGMENT

The second author is supported in part from DARPA/ARL SAFEWARE Award W911NF15C0210, AFOSR Award FA9550-15-1-0274, AFOSR Award FA9550-19-1-0200, AFOSR YIP Award, NSF CNS Award 1936826, DARPA and SPAWAR under contract N66001-15-C-4065, a Hellman Award and research grants by the Okawa Foundation, Visa Inc., and Center for Long-Term Cybersecurity (CLTC, UC Berkeley). The views expressed are those of the author and do not reflect the official policy or position of the funding agencies.

The third author is supported in part by a gift from Ripple, a gift from DoS Networks, a grant from Northrop Grumman, a Cylab seed funding award, and a JP Morgan Faculty Fellowship.

Part of the work was done when the fourth author was at Carnegie Mellon University.

REFERENCES

[1] Y. Ishai, E. Kushilevitz, R. Ostrovsky, M. Prabhakaran, and A. Sahai, "Efficient non-interactive secure computation," in *Advances in Cryptology – EUROCRYPT 2011*, ser. Lecture Notes in Computer Science, K. G. Paterson, Ed., vol. 6632. Tallinn, Estonia: Springer, Heidelberg, Germany, May 15–19, 2011, pp. 406–425.

[2] W. Quach, H. Wee, and D. Wichs, "Laconic function evaluation and applications," in *59th Annual Symposium on Foundations of Computer Science*, M. Thorup, Ed. Paris, France: IEEE Computer Society Press, Oct. 7–9, 2018, pp. 859–870.

[3] C. Cho, N. Döttling, S. Garg, D. Gupta, P. Miao, and A. Polychroniadou, "Laconic oblivious transfer and its applications," in *Advances in Cryptology – CRYPTO 2017, Part II*, ser. Lecture Notes in Computer Science, J. Katz and H. Shacham, Eds., vol. 10402. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 20–24, 2017, pp. 33–65.

[4] J. Kilian, "A note on efficient zero-knowledge proofs and arguments (extended abstract)," in *24th Annual ACM Symposium on Theory of Computing*. Victoria, BC, Canada: ACM Press, May 4–6, 1992, pp. 723–732.

[5] S. Micali, "CS proofs (extended abstracts)," in *35th Annual Symposium on Foundations of Computer Science*. Santa Fe, NM, USA: IEEE Computer Society Press, Nov. 20–22, 1994, pp. 436–453.

[6] J. Groth, "Short pairing-based non-interactive zero-knowledge arguments," in *Advances in Cryptology – ASIACRYPT 2010*, ser. Lecture Notes in Computer Science, M. Abe, Ed., vol. 6477. Singapore: Springer, Heidelberg, Germany, Dec. 5–9, 2010, pp. 321–340.

[7] N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer, "Recursive composition and bootstrapping for SNARKS and proof-carrying data," in *45th Annual ACM Symposium on Theory of Computing*, D. Boneh, T. Roughgarden, and J. Feigenbaum, Eds. Palo Alto, CA, USA: ACM Press, Jun. 1–4, 2013, pp. 111–120.

[8] R. Gennaro, C. Gentry, B. Parno, and M. Raykova, "Quadratic span programs and succinct NIZKs without PCPs," in *Advances in Cryptology – EUROCRYPT 2013*, ser. Lecture Notes in Computer Science, T. Johansson and P. Q. Nguyen, Eds., vol. 7881. Athens, Greece: Springer, Heidelberg, Germany, May 26–30, 2013, pp. 626–645.

[9] M. J. Freedman, K. Nissim, and B. Pinkas, "Efficient private matching and set intersection," in *Advances in Cryptology – EUROCRYPT 2004*, ser. Lecture Notes in Computer Science, C. Cachin and J. Camenisch, Eds., vol. 3027. Interlaken, Switzerland: Springer, Heidelberg, Germany, May 2–6, 2004, pp. 1–19.

[10] S. Garg, C. Gentry, A. Sahai, and B. Waters, "Witness encryption and its applications," in *45th Annual ACM Symposium on Theory of Computing*, D. Boneh, T. Roughgarden, and J. Feigenbaum, Eds. Palo Alto, CA, USA: ACM Press, Jun. 1–4, 2013, pp. 467–476.

[11] R. Canetti, "Universally composable security: A new paradigm for cryptographic protocols," in *42nd Annual Symposium on Foundations of Computer Science*. Las Vegas, NV, USA: IEEE Computer Society Press, Oct. 14–17, 2001, pp. 136–145.

[12] A. C.-C. Yao, "How to generate and exchange secrets (extended abstract)," in *27th Annual Symposium on Foundations of Computer Science*. Toronto, Ontario, Canada: IEEE Computer Society Press, Oct. 27–29, 1986, pp. 162–167.

[13] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *41st Annual ACM Symposium on Theory of Computing*, M. Mitzenmacher, Ed. Bethesda, MD, USA: ACM Press, May 31 – Jun. 2, 2009, pp. 169–178.

[14] C. Gentry and D. Wichs, "Separating succinct non-interactive arguments from all falsifiable assumptions," in *43rd Annual ACM Symposium on Theory of Computing*, L. Fortnow and S. P. Vadhan, Eds. San Jose, CA, USA: ACM Press, Jun. 6–8, 2011, pp. 99–108.

[15] S. Arora and S. Safra, "Probabilistic checking of proofs: A new characterization of NP," *Journal of the ACM (JACM)*, vol. 45, no. 1, pp. 70–122, 1998.

[16] N. Döttling, S. Garg, M. Hajiabadi, D. Masny, and D. Wichs, "Two-round oblivious transfer from CDH or LPN," in *(Under Submission)*, 2019.

[17] N. Döttling and S. Garg, "Identity-based encryption from the Diffie-Hellman assumption," in *Advances in Cryptology – CRYPTO 2017, Part I*, ser. Lecture Notes in Computer Science, J. Katz and H. Shacham, Eds., vol. 10401. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 20–24, 2017, pp. 537–569.

[18] R. Canetti, S. Halevi, and M. Steiner, "Hardness amplification of weakly verifiable puzzles," in *TCC 2005: 2nd Theory of Cryptography Conference*, ser. Lecture Notes in Computer Science, J. Kilian, Ed., vol. 3378. Cambridge, MA, USA: Springer, Heidelberg, Germany, Feb. 10–12, 2005, pp. 17–33.

[19] O. Goldreich and L. A. Levin, "A hard-core predicate for all one-way functions," in *21st Annual ACM Symposium on Theory of Computing*. Seattle, WA, USA: ACM Press, May 15–17, 1989, pp. 25–32.

[20] A. Jain, Y. T. Kalai, D. Khurana, and R. Rothblum, "Distinguisher-dependent simulation in two rounds and its applications," in *Advances in Cryptology – CRYPTO 2017, Part II*, ser. Lecture Notes in Computer Science, J. Katz and H. Shacham, Eds., vol. 10402. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 20–24, 2017, pp. 158–189.

[21] W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Transactions on Information Theory*, vol. 22, no. 6, pp. 644–654, 1976.

[22] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," in *37th Annual ACM Symposium on Theory of Computing*, H. N. Gabow and R. Fagin, Eds. Baltimore, MA, USA: ACM Press, May 22–24, 2005, pp. 84–93.

[23] C. Peikert, O. Regev, and N. Stephens-Davidowitz, "Pseudorandomness of ring-LWE for any ring and modulus," in *49th Annual ACM Symposium on Theory of Computing*, H. Hatami, P. McKenzie, and V. King, Eds. Montreal, QC, Canada: ACM Press, Jun. 19–23, 2017, pp. 461–473.

[24] C. Peikert, "Public-key cryptosystems from the worst-case shortest vector problem: extended abstract," in *41st Annual ACM Symposium on Theory of Computing*, M. Mitzenmacher, Ed. Bethesda, MD, USA: ACM Press, May 31 – Jun. 2, 2009, pp. 333–342.

[25] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé, "Classical hardness of learning with errors," in *45th Annual ACM Symposium on Theory of Computing*, D. Boneh, T. Roughgarden, and J. Feigenbaum, Eds. Palo Alto, CA, USA: ACM Press, Jun. 1–4, 2013, pp. 575–584.

[26] W. Aiello, Y. Ishai, and O. Reingold, "Priced oblivious transfer: How to sell digital goods," in *Advances in Cryptology – EUROCRYPT 2001*, ser. Lecture Notes in Computer Science, B. Pfitzmann, Ed., vol. 2045. Innsbruck, Austria: Springer, Heidelberg, Germany, May 6–10, 2001, pp. 119–135.

[27] Y. Lindell, "General composition and universal composability in secure multi-party computation," in *44th Annual Symposium on Foundations of Computer Science*. Cambridge, MA, USA: IEEE Computer Society Press, Oct. 11–14, 2003, pp. 394–403.

[28] R. Canetti, H. Lin, and R. Pass, "From unprovability to environmentally friendly protocols," in *54th Annual Symposium on Foundations of Computer Science*. Berkeley, CA, USA: IEEE Computer Society Press, Oct. 26–29, 2013, pp. 70–79.

[29] B. Broadnax, N. Döttling, G. Hartung, J. Müller-Quade, and M. Nagel, "Concurrently composable security with shielded super-polynomial simulators," in *Advances in Cryptology – EUROCRYPT 2017, Part I*, ser. Lecture Notes in Computer Science, J. Coron and J. B. Nielsen, Eds., vol. 10210. Paris, France: Springer, Heidelberg, Germany, Apr. 30 – May 4, 2017, pp. 351–381.

[30] M. Naor, "On cryptographic assumptions and challenges (invited talk)," in *Advances in Cryptology – CRYPTO 2003*, ser. Lecture Notes in Computer Science, D. Boneh, Ed., vol. 2729. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 17–21, 2003, pp. 96–109.

[31] A. Banerjee, C. Peikert, and A. Rosen, "Pseudorandom functions and lattices," in *Advances in Cryptology – EUROCRYPT 2012*, ser. Lecture Notes in Computer Science, D. Pointcheval and T. Johansson, Eds., vol. 7237. Cambridge, UK: Springer, Heidelberg, Germany, Apr. 15–19, 2012, pp. 719–737.