

(Efficient) Universally Composable Oblivious Transfer Using a Minimal Number of Stateless Tokens

Seung Geol Choi^{1,*}, Jonathan Katz^{2,**}, Dominique Schröder^{3,***},
Arkady Yerukhimovich^{4,†}, and Hong-Sheng Zhou^{5,‡}

¹ United States Naval Academy
choi@usna.edu

² University of Maryland
jkatz@cs.umd.edu

³ Saarland University
ds@ca.cs.uni-saarland.de

⁴ MIT Lincoln Laboratory
arkady@cs.umd.edu

⁵ Virginia Commonwealth University
hszhou@vcu.edu

Abstract. We continue the line of work initiated by Katz (Eurocrypt 2007) on using *tamper-proof hardware* for universally composable secure computation. As our main result, we show an efficient oblivious-transfer (OT) protocol in which two parties each create and exchange a single, stateless token and can then run an unbounded number of OTs. Our result yields what we believe is the most *practical* and *efficient* known approach for oblivious transfer based on tamper-proof tokens, and implies that the parties can perform (repeated) secure computation of arbitrary functions without exchanging additional tokens.

Motivated by this result, we investigate the minimal number of stateless tokens needed for universally composable OT/secure computation. We prove that our protocol is *optimal* in this regard for constructions making *black-box* use of the tokens (in a sense we define). We also show that nonblack-box techniques can be used to obtain a construction using only a single stateless token.

* Work done in part at the University of Maryland and Columbia University.

** Work supported in part by NSF awards #1111599 and #1223623.

*** Work done in part at the University of Maryland, and supported by the German Ministry for Education and Research (BMBF) through funding for the Center for IT-Security, Privacy, and Accountability (CISPA www.cispa-security.org) and also by an Intel Early Career Award.

† Work done in part at the University of Maryland.

‡ Work done at the University of Maryland, and supported by an NSF CI postdoctoral fellowship.

1 Introduction

The universal composability (UC) framework [6] provides a way of analyzing protocols while ensuring strong security guarantees. In particular, protocols proven secure in this framework remain secure when run concurrently with arbitrary other protocols in a larger networked environment. Unfortunately, most interesting cryptographic tasks are impossible to realize in the “plain” UC framework when an honest majority cannot be assumed and, in particular, in the setting of two-party secure computation [8,9,36]. This stark negative result has motivated researchers to explore various extensions/variants of the UC framework in which secure computation can be achieved [7], with notable examples being the assumption of a common reference string (CRS) [6,8,10] or a public-key infrastructure [6,3]. In the real world, implementing either of these approaches seems to require the existence of some *trusted entity* that parties agree to use (though see [11] for some ideas on using a naturally occurring high-entropy source in place of a CRS).

Katz [32] suggested using *tamper-proof hardware tokens* for UC computation. That is, Katz proposed a model where parties can construct hardware tokens to compute functions of their choice such that an adversary given a token \mathcal{T}_F for a function F can do no more than observe the input/output characteristics of this token. The motivation for this being that the existence of tamper-proof hardware can be viewed, in principle, as a *physical* assumption rather than an assumption of trust in some external entity. (In fact, in Katz’s model the parties may create the tamper-proof tokens themselves—rather than obtain them from a trusted provider [28]—and a malicious party can put any algorithm on a token it creates.) In addition, secure hardware may also potentially result in more efficient protocols; indeed, it has been suggested for improving efficiency in other settings (e.g., [17,13,14,5,27,31,34,22]). In addition to introducing the model, Katz showed that tamper-proof hardware tokens can be used for universally composable computation of arbitrary functions. His work motivated an extensive amount of follow-up work [12,37,23,15,24,25,19] that we discuss in detail later.

As our main result, we show here a new protocol for universally composable 1-out-of-2 string oblivious transfer (OT) based on tamper-proof hardware tokens, secure against a static, malicious adversary. Our work yields what we believe to be the most *practical* and *efficient* known protocol since it simultaneously achieves all the following (which are not achieved all at once by any other solution; see Table 1 for a detailed comparison):

- Our protocol is based on *stateless* tokens, which seem easier/cheaper to create in practice and are (automatically) resistant to resetting attacks.
- Our protocol requires the parties to exchange a *single* pair of tokens. This can be done in advance, before the parties’ inputs are known. Furthermore, the tokens can be used to implement an *unbounded* number of oblivious transfers, rather than requiring the parties to exchange a fresh pair of tokens for every oblivious transfer they want to compute. Thus, by relying on known completeness results [33,30], the parties can use the same tokens to perform

an unlimited number of secure computations (of possibly different functions, and on different inputs).

- Our protocol is efficient. It is *black-box*, and each OT needs mostly standard symmetric-key operations along with only a few (unique) digital signatures. Moreover, any desired number of OTs can be obtained (in parallel) in *constant* rounds.
- If the total number of OTs is a priori bounded, then a variant of our protocol can realize any bounded number of OTs in constant rounds based only on the existence of collision-resistant hash functions.

Inspired by our result, we investigate the *minimal* number of stateless tokens needed for universally composable OT/secure computation. We show that two tokens—one created by each party—are needed even to obtain a single universally composable OT as long as only “black-box techniques” are used. (We explain what we mean by “black-box techniques” in the relevant section of our paper.) Our protocol, above, is thus optimal in this regard. Our impossibility result is somewhat surprising, since a single *stateful* token suffices for OT [19]. Our results thus demonstrate an inherent difference between stateful and stateless tokens.

Since protocols based on nonblack-box techniques tend to be impractical, our work pins down the minimal number of stateless tokens needed as far as practical protocols are concerned. From a theoretical point of view, however, it is still interesting to completely resolve the question. In this vein, we show a protocol for carrying out an unbounded number of secure computations using only a *single* (stateless) token. Our construction uses a variant of the nonblack-box simulation technique introduced by Barak [2].

In summary, our work shows that efficient, universally composable oblivious transfer can be realized from two stateless tokens without any additional setup assumptions, and is unlikely using a single stateless token. On the other hand, using (inefficient) nonblack-box techniques, a single stateless token serves as a sufficient setup for general universally-composable two-party computation.

1.1 Related Work

Katz’s original protocol for secure computation using tamper-proof tokens [32] required stateful tokens and relied on number-theoretic assumptions (specifically, the DDH assumption). Subsequent work has mainly focused on improving one or both of these aspects of his work.

Several researchers have explored constructions using *stateless* tokens. Stateless tokens are presumably easier and/or cheaper to build, and are resistant to resetting attacks whereby an adversary cuts off the power supply and thus effectively “rewinds” the token. Chandran et al. [12] were the first to eliminate the requirement of stateful tokens. They construct UC commitments based on the existence of one-way functions, and oblivious transfer based on any enhanced trapdoor permutation (eTDP). They also introduce a variant security model in which an adversary need not know the code of the tokens he produces, thus

capturing scenarios where an adversary may pass along tokens whose code he doesn't know, e.g., via token replication. (We do not consider this model here.¹) From a practical perspective, however, their work has several drawbacks. Their OT protocol makes nonblack-box use of the underlying primitives, runs in $\Theta(\lambda)$ rounds (where λ is the security parameter), and uses the heavy machinery of concurrent non-malleable zero-knowledge. Improving upon their work, Goyal et al. [25] show a black-box construction of oblivious transfer; their protocol runs in constant rounds assuming a collision-resistant hash function, or $\Theta(\lambda/\log \lambda)$ rounds based on one-way functions. However, their protocol requires the parties to exchange $\Theta(\lambda)$ tokens for *every* oblivious transfer the parties wish to execute. Compared with these results, our protocol is much more efficient at the expense of the stronger assumption of existence of unique signature schemes.

A second direction has explored the possibility of eliminating computational assumptions altogether. This line of work was initiated by Moran and Segev [37], who showed how to realize statistically secure UC commitments using a single stateful token. (Note that commitment does not imply OT, or general secure computation, in the unconditional setting.) Their construction can be used for any bounded number of commitments, still using only one token, and the authors note that they can achieve an unbounded number of commitments (with computational security) based on one-way functions. Goyal et al. [25] show an unconditional construction of oblivious transfer (and hence general secure computation) using $\Theta(\lambda)$ stateful tokens. Recently, Döttling et al. [19] show how to construct unconditionally secure OT using only a single stateful token. Goyal et al. [24] showed that unconditional security from stateless tokens is impossible (unless the token model is extended to allow tokens to encapsulate each other). If such encapsulation is allowed, then they show how to realize statistically secure OT in constant rounds using $\Theta(\lambda)$ stateless tokens.

Kolesnikov [34] showed an efficient construction of oblivious transfer from stateless tokens. However, his work is not in the UC setting, and he achieves only covert security [1] rather than security against malicious parties. Dubovitskaya et al. [21] constructed an OT protocol from two stateful tokens. Their work assumes tokens are not reused (it is not clear how this is enforced), and is also not in the UC setting.

Our Work in Relation to Prior Work. For our main result, we carefully combine the techniques of [25] and [19] to achieve the most *practical* and *efficient* known protocol for universally composable OT based on tamper-proof hardware tokens. Our protocol uses two *stateless* tokens (one per party), and can be used for either a bounded number of OTs (assuming the existence of collision-resistant hash functions) or for an unbounded number of OTs (additionally assuming the existence of unique signatures or, equivalently, verifiable random

¹ Although we do not formally consider this model, it appears that our efficient, two-token protocol would remain secure in that model since the simulator in our security proof does not refer to the code of the tokens.

Table 1. Universally composable OT based on tamper-proof hardware tokens. The security parameter is denoted by λ and TE means token encapsulation.

	stateless tokens					stateful tokens		
	Here 1	Here 2	[12]	[25]	[24]	[32]	[25]	[19]
Tokens:	2	2	2	$\Theta(\lambda)$	$\Theta(\lambda)$	2	$\Theta(\lambda)$	1
Rounds:	$\Theta(1)$	$\Theta(1)$	$\Theta(\lambda)$	$\Theta(1)$	$\Theta(1)$	$\Theta(1)$	$\Theta(1)$	$\Theta(1)$
Asmpt.:	CRHF, VRF	CRHF	eTDP	CRHF	TE	DDH	none	none
# OTs:	unbounded	bounded	unbounded	1	1	unbounded	1	bounded

functions (VRFs)). Both instantiations run in constant rounds.² A detailed comparison of this protocol to relevant prior work is given in Table 1.

In addition to the above, we show two other results: there is no “black-box” construction of universally composable OT using fewer than two stateless tokens, but universally composable coin tossing (and hence OT) *can* be based on a single stateless token using nonblack-box techniques.

Concurrent Work. Independently, Döttling et al. [20] show a different nonblack-box construction of UC coin-tossing from a single stateless token, and argue (without proof) that nonblack-box techniques are needed. Here, we provide a rigorous version of their argument. Our efficient, black-box OT protocol using two stateless tokens—which we consider our primary contribution—has no analog in their work.

2 Preliminaries

Let λ be the security parameter. For a set S , we let $x \leftarrow S$ denote choosing x uniformly at random from S . We assume readers are familiar with pseudorandom functions, collision-resistant hash functions, strong extractors, commitment schemes, digital signature schemes, message-authentication codes (MACs), and witness-indistinguishable arguments of knowledge (WI-AoKs). Due to space restrictions, we omit the formal definitions; here, we mainly introduce notation.

Throughout the paper, a pseudorandom function is denoted by PRF, and it is assumed that the output length is sufficiently long so that it can be truncated to the appropriate length. We let $\text{MAC} = (\text{Sig}, \text{Vrfy})$ be a deterministic message-authentication code, where Vrfy is a canonical verification procedure that checks the validity of a tag τ by recomputing it. We slightly abuse the notation to let $(\text{Kg}, \text{Sig}, \text{Vrfy})$ also denote a digital signature scheme (the context should make it obvious whether the notation indicates a MAC or a signature). A digital signature scheme is called *unique* if for every possible verification key vk and

² The CRHF assumption is only used to make our protocols constant round by instantiating constant round statistically-hiding commitments. Thus, we can instead instantiate our protocols in $\Theta(\lambda/\log \lambda)$ rounds using only one-way functions.

every message m , there is a unique signature σ such that $\text{Vrfy}_{vk}(m, \sigma) = 1$. Dodis and Yampolskiy [18] give an efficient construction for unique signatures based on a certain number-theoretic assumption. Let $\text{Ext} : \{0, 1\}^{2\lambda} \times \{0, 1\}^d \rightarrow \{0, 1\}^\lambda$ denote a strong randomness extractor where the source has length 2λ and the seed has length d . If the min-entropy of the source is at least $2\lambda - O(\log \lambda)$, the output is statistically close to uniform.

We use SCom to denote a (possibly interactive) statistically-hiding and computationally-binding commitment scheme [16,26] and Com to denote a (possibly interactive) computationally-hiding and strongly-binding commitment scheme. We let $\text{com}_m \leftarrow \text{Com}(m; r_m, \check{r}_m)$ denote a commitment to a message m , where the sender (resp., receiver) uses uniform random coins r_m (resp., \check{r}_m) and the final transcript is com_m ; sometimes we omit \check{r}_m when it is clear from the context. (We also use similar notation for SCom .) In a strongly binding commitment scheme [29], with overwhelming probability over the receiver's coins \check{r}_m , for any commitment com there is at most one (m, r_m) such that $\text{com} = \text{Com}(m; r_m, \check{r}_m)$. Although this definition is stronger than usual, the Naor scheme [38] satisfies it. Without loss of generality, we assume a canonical decommit phase in which the sender sends m together with the randomness r_m (i.e., $\text{decom}_m = (m, r_m)$). Then the receiver runs the algorithm $\text{Open}(\text{com}_m, m, r_m)$ which checks if (m, r_m) is consistent with the transcript com_m . If so, Open outputs m ; otherwise it outputs \perp .

Linear Algebra. By \mathbb{F}_2 we denote the finite field with two elements. If $a \in \mathbb{F}_2^\lambda$ and $b \in \mathbb{F}_2^k$ are two column-vectors, then $(ab^T) = (a_i b_j)_{ij} \in \mathbb{F}_2^{\lambda \times k}$ is the outer product (or tensor product) of a and b , and $a^T b = \sum_{i=1}^\lambda a_i b_i \in \mathbb{F}_2$ the inner product.

Let $C \in \mathbb{F}_2^{\lambda \times 2\lambda}$. Then $\dim(\ker(C)) \geq \lambda$. Let $B = \{b_1, \dots, b_\lambda\} \subseteq \ker(C)$ be a linearly independent set. One can choose a set $B^* = \{b_{\lambda+1}, \dots, b_{2\lambda}\}$ such that $B \cup B^*$ is a basis of $\mathbb{F}_2^{2\lambda}$. Let $e_i \in \mathbb{F}_2^\lambda$ be the i th unit-vector. Then, there exists a matrix $G \in \mathbb{F}_2^{\lambda \times 2\lambda}$ such that $Gb_i = e_i$ for $i = 1, \dots, \lambda$ and $Gb_i = 0$ for $i = \lambda+1, \dots, 2\lambda$. This matrix is called the *complementary matrix* of C and we denote by $G \leftarrow \text{Comp}(C)$ its computation. It holds that $\text{rank}(G) = \lambda$ and $B^* \subseteq \ker(G)$. For such C and G , we can always solve the linear system $Cx = r$, $Gx = s$ by solving $Cx_{B^*} = r$ and $Gx_B = s$ independently with $x_B \in \text{span}(B) \subseteq \ker(C)$ and $x_{B^*} \in \text{span}(B^*) \subseteq \ker(G)$, and then setting $x := x_B + x_{B^*}$.

Token Functionality. We model tamper-proof hardware tokens as an ideal functionality in the UC framework, following Katz [32]; see Figure 1. Our ideal functionality models stateful tokens: although all our protocols use stateless tokens, an adversarially generated token may be stateful.

Oblivious-Transfer Functionality. The OT functionality is standard, but we wish here to model a multi-session variant where the sender and receiver repeatedly (in different sub-sessions) execute any agreed-upon number m of parallel OTs (in a given sub-session). We refer to this functionality as $\mathcal{F}_{\text{multi-OT}}$, and describe it in Figure 2. We note that the sub-sessions are executed sequentially. Additionally, as highlighted in [25], the sender is notified each time

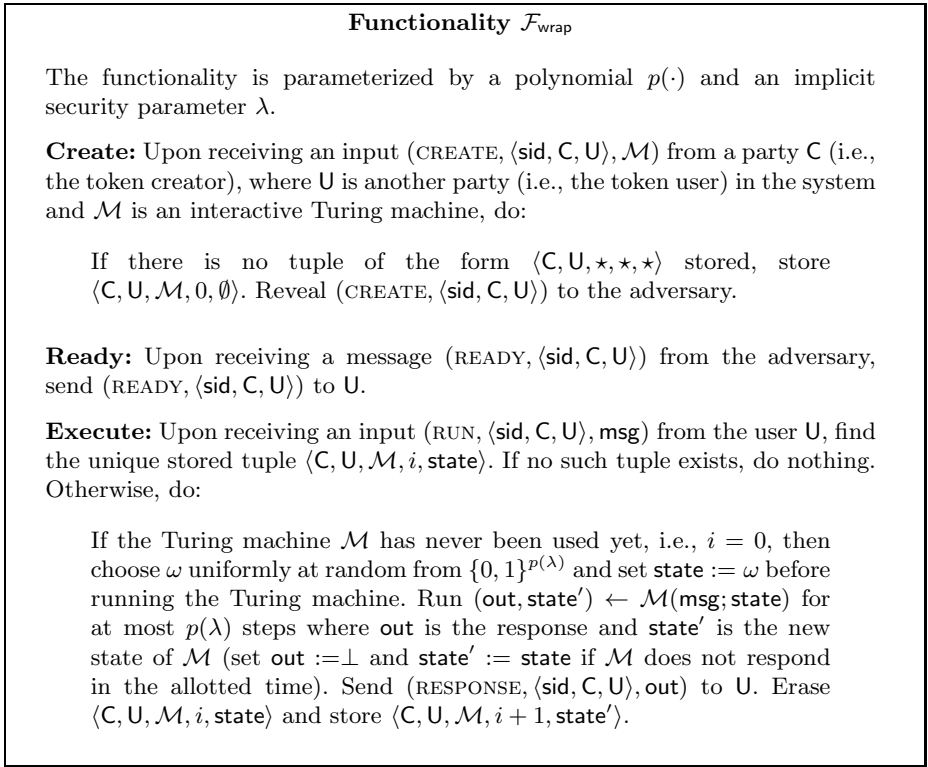


Fig. 1. The ideal $\mathcal{F}_{\text{wrap}}$ functionality for stateful tokens

the receiver obtains output. We define the bounded OT functionality similarly except that the sender and receiver only execute a single sub-session of m parallel OTs. This allows the sender and receiver to execute any bounded number of OTs.

3 Efficient Oblivious Transfer Using Two Stateless Tokens

In this section, we first give the details of our unbounded OT protocol. Then, in Section 3.3 we briefly sketch how this protocol can be modified to achieve a bounded OT protocol using only CRHFs. We provide some intuition and background before giving the details of our protocol. Our starting point is the unconditionally secure OT protocol from [19], which uses a single *stateful* token. We sketch a simplified version of their protocol for the case of a single OT carried out between the sender S with input $(x_0, x_1) \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$ and the receiver R with input $b \in \{0, 1\}$. We canonically identify the vector space \mathbb{F}_2^λ with the set $\{0, 1\}^\lambda$ of strings of length λ . The main steps of the protocol are as follows:

1. S creates a token \mathcal{T}_{S} holding random vector $a \in \mathbb{F}_2^{2\lambda}$ and matrix $B \in \mathbb{F}_2^{2\lambda \times 2\lambda}$ and gives it to R .

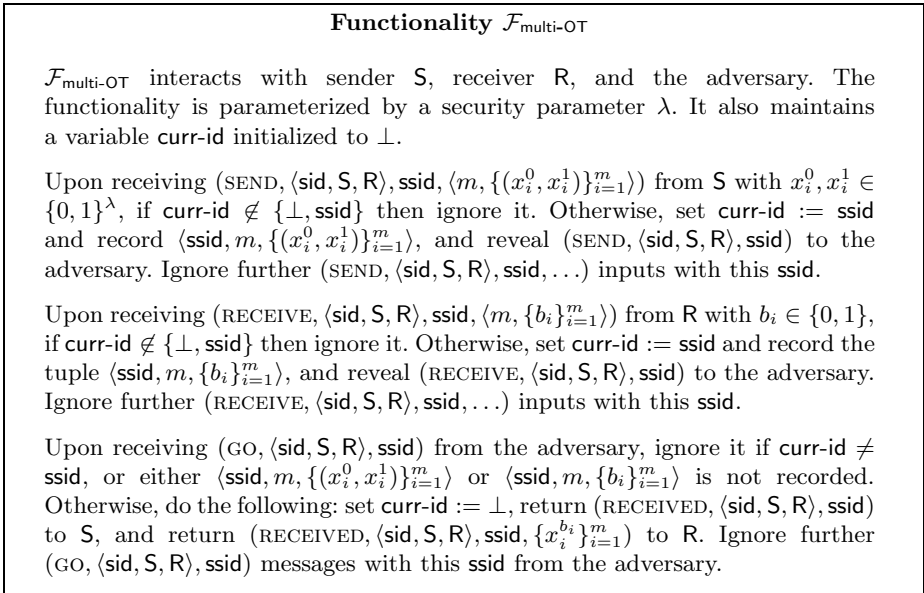


Fig. 2. The $\mathcal{F}_{\text{multi-OT}}$ functionality

2. R chooses a random matrix $C \in \mathbb{F}_2^{\lambda \times 2\lambda}$ and sends it to S. In turn, S computes $\tilde{a} := Ca$, $\tilde{B} := CB$, and a complementary matrix $G \in \mathbb{F}_2^{\lambda \times 2\lambda}$ to C and sends these to R.
3. R sends random $h \in \mathbb{F}_2^{2\lambda}$ to S. Then, S sends $\tilde{x}_0 := x_0 + GBh$ and $\tilde{x}_1 := x_1 + GBh + Ga$ to R.
4. R queries the token with a random $z \in \mathbb{F}_2^{2\lambda}$ such that $z^T h = b$. The token will in turn output $V := az^T + B$. Then R checks that $CV = \tilde{a}z^T + \tilde{B}$ and, if this is the case, outputs $x_b := \tilde{x}_b - GVh$. (Otherwise it detects that S was cheating.)

The basic idea of their protocol is as follows. The receiver R performs a secret sharing of its input b into shares z and h ; by using only h with the sender S and only z with the token, R maintains the privacy of b . In order to obtain the output x_b , the receiver R has to compute the mask (i.e., GBh or $GBh + Ga$). This is achieved by querying the token with z , since $GVh = G(az^T + B)h = b(Ga) + GBh$. To ensure that the token outputs the correct value V , the receiver checks that $CV = \tilde{a}z^T + \tilde{B}$. Since the token does not know C , incorrect behavior is detected with overwhelming probability. To achieve security against a malicious receiver, it is crucial that the receiver is only able to query the token once, and this is enforced by making the token *stateful*. In particular, the token “self destructs” after the first query by the receiver.

Using Stateless Tokens. We carefully combine the techniques of [19] with ideas from [25] so that we can use two *stateless* tokens instead of one *stateful* token. This entails several difficulties:

Multiple queries. A stateless token can be executed multiple times while the token remains oblivious about it (in contrast, stateful tokens “self destruct” after their use [23]), so we need to ensure that a malicious party queries the token only once. Motivated by similar techniques in [25], we handle this issue by modifying the token so that it only replies to *authenticated* inputs. That is, instead of querying z directly to the token, R queries $(\text{com}_z, z, r_z, \sigma_z)$ where com_z is a commitment to z , the value σ_z is a digital signature on com_z (received from S) with respect to a verification key vk_S , and z, r_z is the decommitment of com_z (see the discussion below about why this technique is useful).

Extracting the inputs. Using a stateless token additionally introduces the difficulty of extracting the sender’s inputs during the simulation. Extraction from a stateful token is possible by having the simulator rewind the token and query it multiple times. (The fact that the simulator can query the token multiple times is an advantage of the simulator over the real-world parties.) Once we move to a stateless token in the real world, and authenticate the queries as described above, even the simulator is no longer able to rewind and send multiple (authenticated) queries to the token.

To resolve this issue we introduce a second stateless token \mathcal{T}_R sent from the receiver to the sender which allows directly extracting the sender’s inputs.

At a high level, the above results in the following changes to the protocol:

1. The receiver sends a token \mathcal{T}_R to the sender. (The behavior of this token will become clear below.)
- 2(a). S generates a statistically hiding³ commitment $\text{com} \leftarrow \text{SCom}(a\|B)$ and sends com to R. Then, R chooses a matrix C and authenticates com by computing $\sigma := \text{Sig}_{sk_R}(\text{com}\|C)$ (where sk_R is a secret key also embedded in \mathcal{T}_R), and sends C and σ to S.
- 2(b). S queries \mathcal{T}_R on input $(C, \text{com}, \text{decom}, \sigma)$. The token checks that the signature σ is valid and that decom is a valid opening of com ; otherwise it aborts. \mathcal{T}_R returns $\tilde{a} := Ca$ and $\tilde{B} := CB$ together with a signature $\sigma' := \text{Sig}_{sk_R}(\tilde{a}\|\tilde{B})$. Then S sends $(\tilde{a}, \tilde{B}, \sigma')$ to R.

Intuitively, the value (a, B) remains hidden from a corrupted R as before. For a corrupted sender, the simulator will emulate $\mathcal{F}_{\text{wrap}}$ and observe the inputs to \mathcal{T}_R . To guarantee that the simulator extracts the sender’s inputs correctly, it is necessary that the sender queries the token exactly once with a value (a, B) . The binding property of the commitment scheme, together with the unforgeability of the signature scheme, guarantees that S can make at most one valid query to the token. The unforgeability of the signature scheme assures that S must query the token at least once to generate a valid next message of the protocol. A similar argument holds for the inputs $(\text{com}_z, z, r_z, \sigma_z)$ to the sender’s token \mathcal{T}_S .

³ We were not able to prove security using a computationally hiding commitment scheme. Indeed, it is an interesting open problem to achieve a constant-round OT protocol with two stateless tokens based only on one-way functions.

On input (key):
 output vk_S .

On input ($ssid, i, com_z, z, r_z, \sigma_z$):
 $v := \text{Vrfy}_{vk_S}(ssid \| i \| com_z, z, r_z, \sigma_z)$
 if $z = \text{Open}(com_z, z, r_z)$ and $v = 1$
 $a := \text{PRF}_{k_a}(ssid \| i)$
 $B := \text{PRF}_{k_B}(ssid \| i)$
 $V := az^T + B$
 $w := \text{Sig}_{sk_S}(ssid \| i)$
 output (V, w)
 else output (\perp, \perp)

Fig. 3. The Turing machine \mathcal{M}_S to be embedded in the sender-created token \mathcal{T}_S . It is initialized with (sk_S, vk_S, k_a, k_B) given by the creator.

On input (key):
 output vk_R .

On input ($ssid, i, com_{a\|B}, a, B, r_{a\|B}, \sigma_{a\|B}$):
 $v := \text{Vrfy}_{vk_R}(ssid \| i \| 0 \| com_{a\|B}, \sigma_{a\|B})$
 if $a\|B = \text{Open}(com_{a\|B}, a\|B, r_{a\|B})$
 and $v = 1$
 $\tilde{a} := Ca; \tilde{B} := CB;$
 $\sigma_{\tilde{a}\|\tilde{B}} := \text{Sig}_{sk_R}(ssid \| i \| 1 \| \tilde{a}\|\tilde{B})$
 output $(\tilde{a}, \tilde{B}, \sigma_{\tilde{a}\|\tilde{B}})$
 else output (\perp, \perp, \perp)

Fig. 4. The Turing machine \mathcal{M}_R to be embedded in the receiver-created token \mathcal{T}_R . It is initialized with (sk_R, vk_R, C) given by the creator.

Malicious tokens. There is one additional issue to be taken care of. The above protocol is secure assuming the tokens are generated as specified by the protocol. However, a malicious token may try to leak some information about the other party's queries to the token creator. For example, a malicious token \mathcal{T}_R^* may output $(\tilde{a}, \tilde{B}, \sigma')$ where the bits of σ' leak information about (a, B) . To protect against this, we require the underlying signature scheme to be *unique*; then, σ' carries no more information about (a, B) than (\tilde{a}, \tilde{B}) does. This ensures that the only information leakage that can occur is from a *token abort*. For \mathcal{T}_S , this type of leakage is already handled by the protocol of [19]. For \mathcal{T}_R , we handle the leakage using strong extractors; see the formal description of the protocol below.

The above suffices for a single OT from each pair of tokens. To achieve an unbounded number of OTs we replace all the secret keys and secret inputs with pseudorandom values output by a pseudorandom function.

3.1 The Protocol

The protocol π between a sender S and a receiver R consists of an initial token-exchange phase after which the parties can carry out an unlimited number of oblivious transfers. We describe π now; see also Figure 8.

Token-Exchange Phase. Each party generates a single token and sends it to the other party. The sender's token \mathcal{T}_S encapsulates the code described in Figure 3, where $k_a, k_B, \leftarrow \{0, 1\}^\lambda$ and $(sk_S, vk_S) \leftarrow \text{Kg}()$. The receiver's token \mathcal{T}_R encapsulates the code described in Figure 4, where $(sk_R, vk_R) \leftarrow \text{Kg}()$ and $C \leftarrow \mathbb{F}_2^{2\lambda \times 4\lambda}$. Formally, each party sends the relevant code to $\mathcal{F}_{\text{wrap}}$ using an appropriate (CREATE, \dots) message. Then, S runs $\mathcal{T}_R(key)$ to obtain vk_R ; likewise, R obtains vk_S by running $\mathcal{T}_S(key)$. Finally, R sends C to S, and in turn S computes the complementary matrix $G \leftarrow \text{Comp}(C)$ and sends it to R.

Oblivious Transfer Phase. Following the token-exchange phase, the parties can sequentially run an unbounded number of sub-sessions where in each sub-

session they carry out any desired number m of oblivious transfers. (Each sub-session uses only a constant number of rounds.) In each sub-session, S gets (SEND, $\langle \text{sid}, \text{S}, \text{R} \rangle$, $\text{ssid}, \langle m, \{(x_i^0, x_i^1)\}_{i=1}^m \rangle$) and R gets (RECEIVE, $\langle \text{sid}, \text{S}, \text{R} \rangle$, $\text{ssid}, \langle m, \{b_i\}_{i=1}^m \rangle$) from the environment, and they execute the following protocol:

S \rightarrow R: For $i \in [m]$, the sender computes $a_i := \text{PRF}_{k_a}(\text{ssid}||i)$ and $B_i := \text{PRF}_{k_B}(\text{ssid}||i)$. Here, we have $a_i \in \mathbb{F}_2^{4\lambda}$, $B_i \in \mathbb{F}_2^{4\lambda \times 4\lambda}$. Then, the sender commits to (a_i, B_i) by executing $\text{com}_{a_i||B_i} \leftarrow \text{SCom}(a_i||B_i; r_{a_i||B_i})$, and sends $\{\text{com}_{a_i||B_i}\}_{i=1}^m$ to R.

R \rightarrow S: R chooses $h_i \leftarrow \mathbb{F}_2^{4\lambda}$ and $z_i \leftarrow \mathbb{F}_2^{4\lambda}$ subject to the constraint that $b_i = z_i^T h_i$, and commits to z_i by executing $\text{com}_{z_i} \leftarrow \text{SCom}(z_i; r_{z_i})$. It next authenticates the commitment $\text{com}_{a_i||B_i}$ by computing $\sigma_{a_i||B_i} := \text{Sig}_{sk_R}(\text{ssid}||i||0||\text{com}_{a_i||B_i})$ for $i \in [m]$. Finally, it sends $\{(\sigma_{a_i||B_i}, \text{com}_{z_i})\}_{i=1}^m$ to S.

S \rightarrow R: The sender checks if $\text{Vrfy}_{vk_R}(\text{ssid}||i||0||\text{com}_{a_i||B_i}, \sigma_{a_i||B_i}) = 1$ for $i \in [m]$; if the check fails, then S aborts the protocol. For $i \in [m]$, the sender runs the token \mathcal{T}_R with $(\text{ssid}, i, \text{com}_{a_i||B_i}, a_i, B_i, r_{a_i||B_i}, \sigma_{a_i||B_i})$ and obtains in return $(\tilde{a}_i, \tilde{B}_i, \sigma_{\tilde{a}_i||\tilde{B}_i})$. Then S checks that $\tilde{a}_i = C a_i$, $\tilde{B}_i = C B_i$, $\text{Vrfy}_{vk_R}(\text{ssid}||i||1||\tilde{a}_i||\tilde{B}_i, \sigma_{\tilde{a}_i||\tilde{B}_i}) = 1$ for $i \in [m]$; if the check fails, S aborts the protocol. Otherwise, for $i \in [m]$ it authenticates the commitment com_{z_i} by computing $\sigma_{z_i} := \text{Sig}_{sk_S}(\text{ssid}||i||\text{com}_{z_i})$. Finally, it sends $\{(\tilde{a}, \tilde{B}, \sigma_{\tilde{a}_i||\tilde{B}_i}, \sigma_{z_i})\}_{i=1}^m$ to R.

R \rightarrow S: The receiver checks that $\text{Vrfy}_{vk_R}(\text{ssid}||i||1||\tilde{a}_i||\tilde{B}_i, \sigma_{\tilde{a}_i||\tilde{B}_i}) = 1$, and $\text{Vrfy}_{vk_S}(\text{ssid}||i||\text{com}_{z_i}, \sigma_{z_i}) = 1$ for $i \in [m]$. If this check fails, it aborts the protocol. Otherwise, R runs the token \mathcal{T}_S with $(\text{ssid}, i, \text{com}_{z_i}, z_i, r_{z_i}, \sigma_{z_i})$ and obtains in return (V_i, w_i) , for $i \in [m]$. The receiver checks that $\text{Vrfy}_{vk_S}(\text{ssid}||i, w_i) = 1$ and $C V_i = \tilde{a}_i z_i^T + \tilde{B}_i$ for $i \in [m]$. If the check fails, then it aborts the protocol. Otherwise, it sends $\{(h_i, w_i)\}_{i=1}^m$ to S.

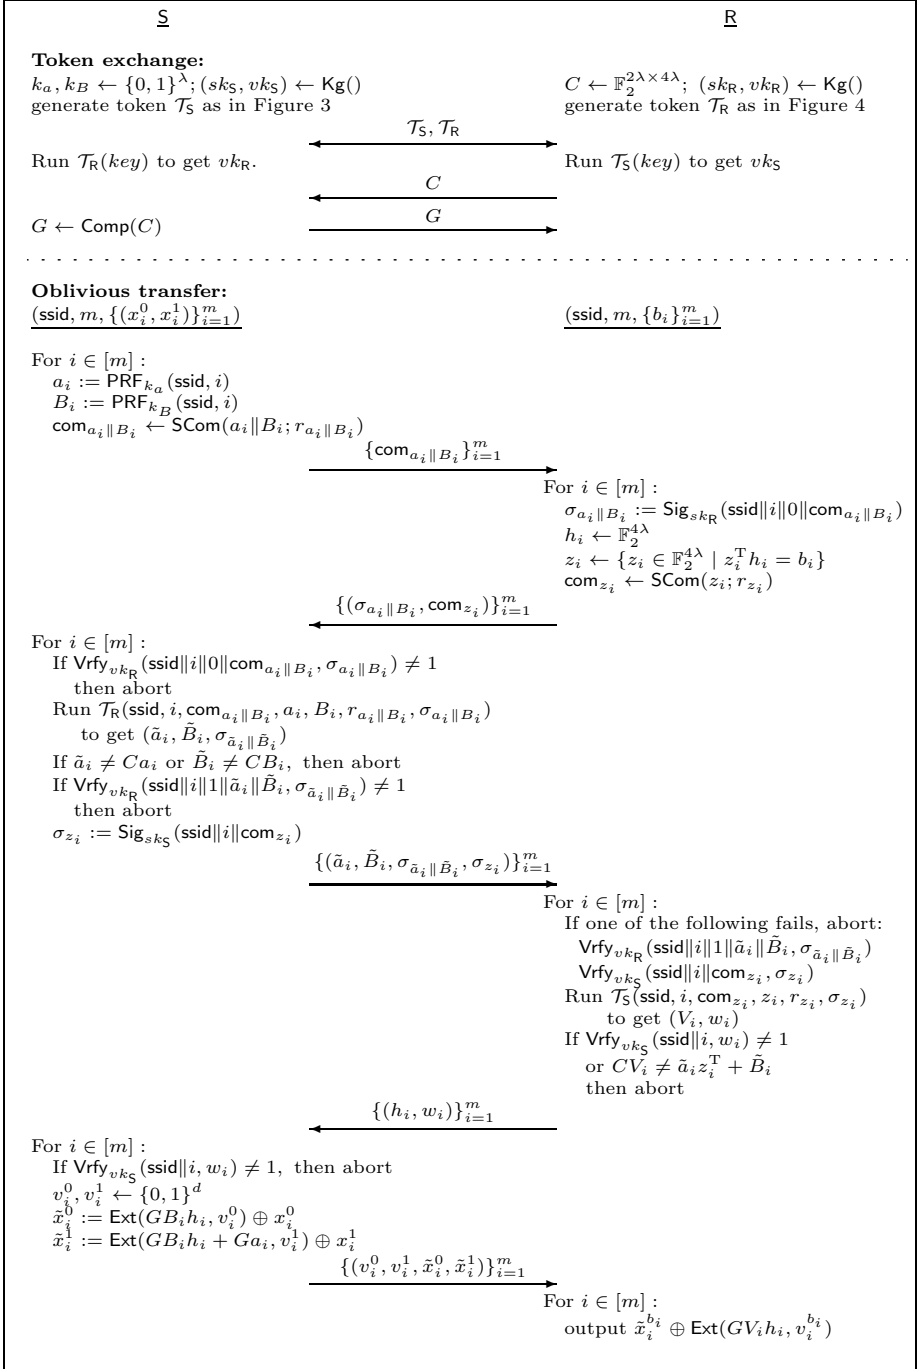
S \rightarrow R: The sender checks that $\text{Vrfy}_{vk_S}(\text{ssid}||i, w_i) = 1$ for $i \in [m]$; if not, it aborts the protocol. Otherwise, for $i \in [m]$ it chooses $v_i^0, v_i^1 \leftarrow \{0, 1\}^d$ and computes $\tilde{x}_i^0 := \text{Ext}(G B_i h_i, v_i^0) \oplus x_i^0$ and $\tilde{x}_i^1 := \text{Ext}(G B_i h_i + G a_i, v_i^1) \oplus x_i^1$. Here d is an appropriate seed length for the extractor. Finally it sends $\{(v_i^0, v_i^1, \tilde{x}_i^0, \tilde{x}_i^1)\}_{i=1}^m$ to the receiver.

R \rightarrow S: For $i \in [m]$, the receiver computes $x_i^{b_i} := \tilde{x}_i^{b_i} \oplus \text{Ext}(G V_i h_i, v_i^{b_i})$ and outputs it.

Theorem 1. *Assume PRF is a pseudorandom function, SCom is statistically hiding and computationally binding, and (Kg, Sig, Vrfy) is a unique signature scheme. Then protocol π securely realizes $\mathcal{F}_{\text{multi-OT}}$ in the $\mathcal{F}_{\text{wrap}}$ -hybrid model.*

3.2 Proof Idea

In this section, we briefly sketch the main ideas behind the proof of Theorem 1. A complete proof is deferred to the full version.

Fig. 5. An OT protocol π from two stateless tokens

To show security of the protocol, we need to construct a simulator Sim for any non-uniform PPT environment \mathcal{Z} such that $\text{EXEC}_{\pi, \mathcal{A}, \mathcal{Z}}^{\mathcal{F}_{\text{wrap}}} \approx \text{IDEAL}_{\mathcal{F}_{\text{multi-OT}}, \text{Sim}, \mathcal{Z}}$, where \mathcal{A} is the dummy adversary. Below, we briefly sketch the ideas used to construct such a simulator. Note that we assume w.l.o.g. that the adversary never asks the same query twice.

Sender Corruption. First, recall that the token \mathcal{T}_R takes as input $(\text{ssid}, i, \text{com}_{a_i \| B_i}, a_i, B_i, r_{a_i \| B_i}, \sigma_{a_i \| B_i})$ and returns $(\tilde{a}_i, \tilde{B}_i, \sigma_{\tilde{a}_i \| \tilde{B}_i})$. Our protocol forces the malicious sender S^* to query \mathcal{T}_R on exactly one input $(a_i \| B_i)$ for a fixed index i . This allows Sim to extract the input from the protocol transcript. To see this, note that the receiver authenticates the inputs to its token \mathcal{T}_R , which in turn authenticates its output. Thus, if the malicious sender does not query the token \mathcal{T}_R but sends a “valid” tuple $(\tilde{a}_i \| \tilde{B}_i, \sigma_{\tilde{a}_i \| \tilde{B}_i})$, then this immediately contradicts the unforgeability of the underlying signature scheme. Also, S^* cannot query \mathcal{T}_R with two different values $(a_i \| B_i), (a'_i \| B'_i)$ (for a fixed i) without contradicting the unforgeability of the signature $\sigma_{a_i \| B_i}$ or the binding of $\text{com}_{a_i \| B_i}$.

Next, consider the maliciously generated token \mathcal{T}_{S^*} . Its input consists of a tuple $(\text{ssid}, i, \text{com}_{z_i}, z_i, r_{z_i}, \sigma_{z_i})$, and its output is (V_i, w_i) . The input is chosen carefully in combination with the protocol execution such that the output of the token does not reveal any information about the choice bit $b_i = z_i^T h_i$ of the receiver. Observe that the signature σ_{z_i} depends only on vk_S and com_{z_i} in the information theoretic sense, since we use a unique signature scheme. Therefore, the token knows nothing about h_i . Also, the signature w_i depends only on vk_R and (ssid, i) in the information theoretic sense, and it does not contain any information about z_i .

Still, the token \mathcal{T}_{S^*} may send some limited amount of information about the *previous* executions to S^* by aborting the protocol. Observe that according to our definition of the functionality, OT sub-sessions are executed in a sequential manner, and aborting is allowed *only once*. Thus, the number ℓ of successful sub-sessions so far can encode some information of the input history of \mathcal{T}_{S^*} up to $O(\log \lambda)$ bits. Therefore, even with this leakage, the adversary S^* has only a negligible advantage in predicting b_i . To see this, note that $b_i = z_i^T h_i$ remains unfixed given $O(\log \lambda)$ bits of z_i .

Receiver Corruption. First, recall that the input to the token \mathcal{T}_S is $(\text{ssid}, i, \text{com}_{z_i}, z_i, r_{z_i}, \sigma_{z_i})$ and that its output is (V_i, w_i) . As in the malicious sender case, our protocol forces the malicious receiver R^* to query the token on exactly one input \hat{z}_i for each i . By observing this query, the simulator Sim can extract the input $\hat{b}_i = \hat{z}_i^T h_i$ of R^* . This property is achieved using the unforgeability of the signature scheme and the binding of the commitment scheme SCom . That is, the unforgeability guarantees that the token runs only on the input (including com_{z_i}) authenticated by the honest sender, and the binding of com_{z_i} disables the malicious receiver R^* to query the token on two distinct z_i s.

Next, consider the maliciously generated token \mathcal{T}_{R^*} . The input to the token and its output are carefully handled such that it does not reveal information about (x_i^0, x_i^1) of the sender. To see this, recall the token \mathcal{T}_{R^*} takes as input $(\text{ssid}, i, \text{com}_{a_i \| B_i}, a_i, B_i, r_{a_i \| B_i}, \sigma_{a_i \| B_i})$ and returns $(\tilde{a}_i, \tilde{B}_i, \sigma_{\tilde{a}_i \| \tilde{B}_i})$. In particular,

since we use a unique signature scheme, the signature $\sigma_{\tilde{a}_i \parallel \tilde{B}_i}$ information-theoretically depends only on $(vk_R, \tilde{a}_i, \tilde{B}_i)$, so any malicious attempt by R^* and \mathcal{T}_{R^*} to gain information about (a_i, B_i) beyond $(\text{com}_{a_i \parallel B_i}, \tilde{a}_i, \tilde{B}_i)$ is not possible.

Still, the token \mathcal{T}_{R^*} may reveal to the environment some limited amount of information about the previous executions to R^* by aborting the protocol (only allowed one-time), that is, the number ℓ of successful sub-sessions so far can encode some information of the input history of \mathcal{T}_{R^*} . This value ℓ can encode at most $O(\log \lambda)$ bits. Using appropriate extractors in computing the masks for \tilde{x}_i^0 and \tilde{x}_i^1 , we can handle this amount of leakage.

3.3 Bounded Oblivious Transfer from Collision Resistant Hash Functions

The bounded OT protocol between a sender S and a receiver R consists of an initial token-exchange phase after which the parties can carry out a bounded number of oblivious transfers in a single sub-session. We describe π now; see also Figure 8. Here, let $(\text{Sig}, \text{Vrfy})$ be a MAC scheme.

The main intuition for this protocol is that if we allow *only one sub-session* to be executed, then we only need to worry about covert channels that transmit information observed by each party/token so far up to the covert communication. We do not need to worry about a later sub-session sending information about a prior sub-session. This allows us to eliminate the need for unique signature and use MACs instead, thus giving a protocol based only on the existence of CRHFs.

We eliminate the unique signatures in two ways. First, some of the checks in the unbounded protocol can be removed. In particular, in the protocol, the receiver doesn't check the authentication on com_{z_i} before forwarding it to the sender-created token \mathcal{T}_S . This is because it cannot contain any useful information about z_i beyond com_{z_i} . For other checks, unique signatures are replaced with (roughly) the following technique: a party A commits to a MAC key, authenticates the message using the MAC key, and later sends the decommitment. Due to the unforgeability of MAC, the party A makes sure that B runs the token only once before receiving messages of B derived from the token execution result. Given the decommitment, the other party B can check that all the messages from A so far have been legitimate.

Token-Exchange Phase. Each party generates a single token and sends it to the other party. The sender's token \mathcal{T}_S encapsulates the code described in Figure 6, where $k_a, k_B, k_w, k_W, s_2 \leftarrow \{0, 1\}^\lambda$. The receiver's token \mathcal{T}_R encapsulates the code described in Figure 7, where $s \leftarrow \{0, 1\}^\lambda$ and $C \leftarrow \mathbb{F}_2^{2\lambda \times 4\lambda}$. Formally, each party sends the relevant code to $\mathcal{F}_{\text{wrap}}$ using an appropriate (CREATE, \dots) message. Finally, R sends C to S , and in turn S computes the complementary matrix $G \leftarrow \text{Comp}(C)$ and sends it to R .

Oblivious-Transfer Phase. Following the token-exchange phase, the parties enter the oblivious transfer phase, where m OT instances are executed in parallel. In this phase, S gets $(\text{SEND}, \langle \text{sid}, S, R \rangle, \langle m, \{(x_i^0, x_i^1)\}_{i=1}^m \rangle)$ and R gets

On input $(i, \text{com}_z, z, r_z, \tau_z)$:
 $v := \text{Vrfy}_{s_2}(i \| \text{com}_z, \tau_z)$
 if $z = \text{Open}(\text{com}_z, z, r_z)$ and $v = 1$
 $a := \text{PRF}_{k_a}(i)$; $B := \text{PRF}_{k_B}(i)$
 $w := \text{PRF}_{k_w}(i)$; $r_w := \text{PRF}_{k_W}(i)$
 $V := az^T + B$
 output (V, w, r_w)
 else output (\perp, \perp, \perp)

Fig. 6. The Turing machine \mathcal{M}_S to be embedded in the sender-created token \mathcal{T}_S . It is initialized with $(k_a, k_B, k_w, k_W, s_2)$ given by the creator.

On input $(i, \text{com}_{a\|B}, a, B, r_{a\|B}, \tau_{a\|B})$:
 $v := \text{Vrfy}_s(i \| 0 \| \text{com}_{a\|B}, \tau_{a\|B})$
 if $a\|B = \text{Open}(\text{com}_{a\|B}, a\|B, r_{a\|B})$
 and $v = 1$
 $\tilde{a} := Ca$; $\tilde{B} := CB$;
 $\tau_{\tilde{a}\|\tilde{B}} := \text{Sig}_s(i \| 1 \| \tilde{a}\|\tilde{B})$
 output $(\tilde{a}, \tilde{B}, \tau_{\tilde{a}\|\tilde{B}})$
 else output (\perp, \perp, \perp)

Fig. 7. The Turing machine \mathcal{M}_R to be embedded in the receiver-created token \mathcal{T}_R . It is initialized with (s, C) given by the creator.

(RECEIVE, $\langle \text{sid}, S, R \rangle, \langle m, \{b_i\}_{i=1}^m \rangle$) from the environment, and they execute the following protocol:

S \rightarrow R: For $i \in [m]$, the sender computes $a_i := \text{PRF}_{k_a}(i)$, $B_i := \text{PRF}_{k_B}(i)$, $w_i := \text{PRF}_{k_w}(i)$, and $r_{w_i} := \text{PRF}_{k_W}(i)$, and chooses $u_i \leftarrow \{0, 1\}^{8\lambda^2 + 3\lambda}$. Here, we have $a_i \in \mathbb{F}_2^{4\lambda}$, $B_i \in \mathbb{F}_2^{4\lambda \times 4\lambda}$, and $w_i \in \{0, 1\}^\lambda$. Then, the sender commits to (a_i, B_i) , w_i , and u_i by executing $\text{com}_{a_i\|B_i} \leftarrow \text{SCom}(a_i\|B_i; r_{a_i\|B_i})$, and $\text{com}_{w_i} \leftarrow \text{Com}(w_i; r_{w_i})$, and $\text{com}_{u_i} \leftarrow \text{Com}(u_i; r_{u_i})$, respectively. The sender sends $\{(\text{com}_{a_i\|B_i}, \text{com}_{w_i}, \text{com}_{u_i})\}_{i=1}^m$ to R.

R \rightarrow S: The receiver commits to s by executing $\text{com}_s \leftarrow \text{Com}(s; r_s)$. It also chooses $z_i \leftarrow \mathbb{F}_2^{4\lambda}$, and commits to z_i by executing $\text{com}_{z_i} \leftarrow \text{SCom}(z_i; r_{z_i})$. It next authenticates the commitment $\text{com}_{a_i\|B_i}$ by computing

$$\tau_{a_i\|B_i} := \text{Sig}_s(i \| 0 \| \text{com}_{a_i\|B_i}) \text{ for } i = 1, \dots, m.$$

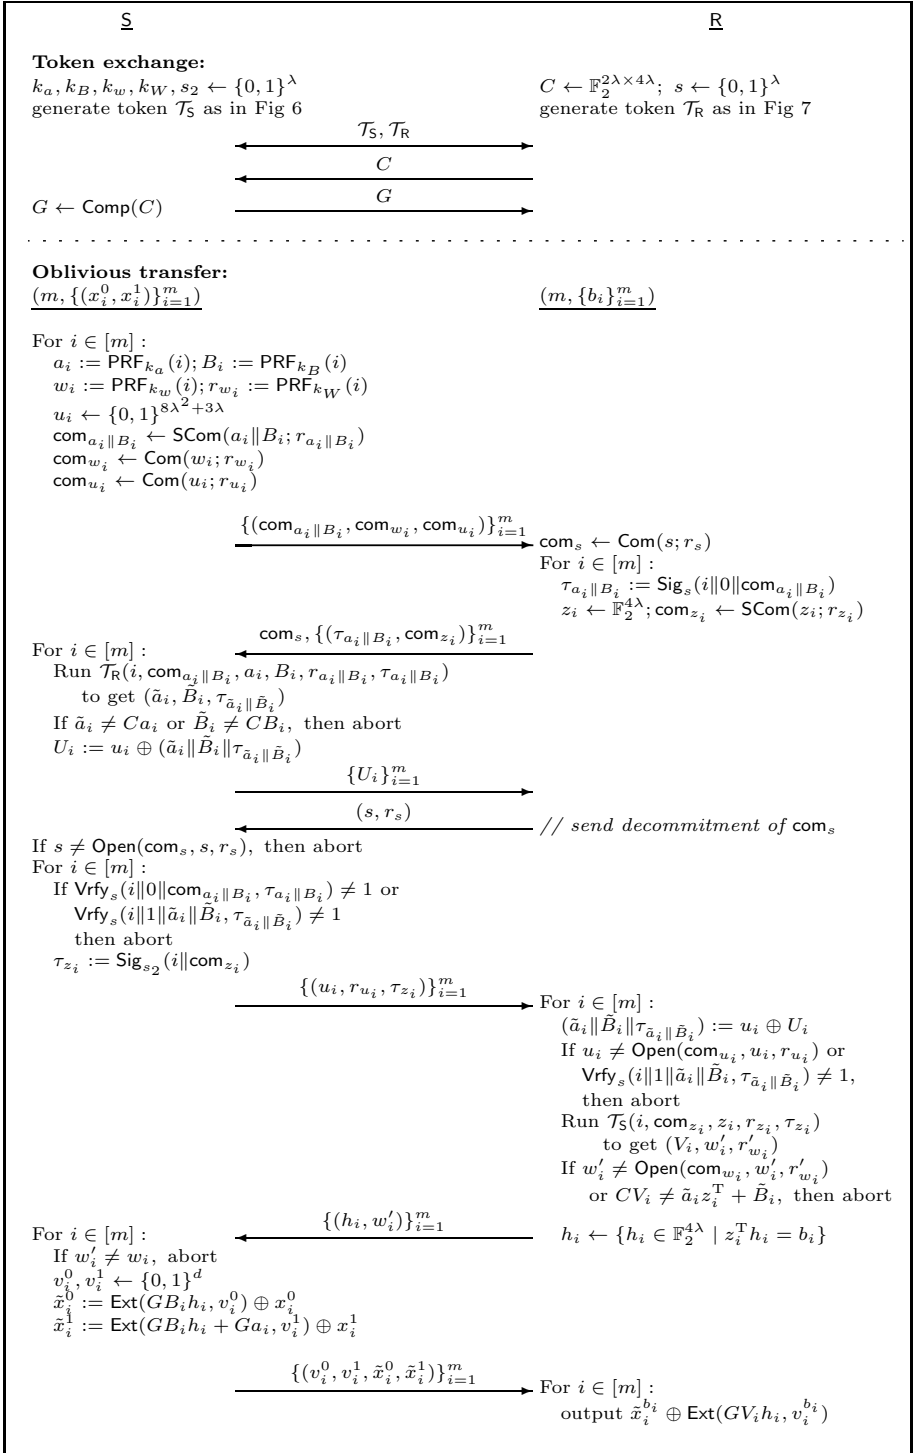
Finally, it sends $(\text{com}_s, \{(\tau_{a_i\|B_i}, \text{com}_{z_i})\}_{i=1}^m)$ to S.

S \rightarrow R: For $i \in [m]$, the sender runs the token \mathcal{T}_R with $(i, \text{com}_{a_i\|B_i}, a_i, B_i, r_{a_i\|B_i}, \tau_{a_i\|B_i})$ and obtains in return $(\tilde{a}_i, \tilde{B}_i, \tau_{\tilde{a}_i\|\tilde{B}_i})$. Then S checks that $\tilde{a}_i = Ca_i$ and $\tilde{B}_i = CB_i$ for $i \in [m]$. If the check fails, S aborts the protocol. Otherwise, S computes $U_i := u_i \oplus (\tilde{a}_i \| \tilde{B}_i \| \tau_{\tilde{a}_i\|\tilde{B}_i})$ for $i \in [m]$ and sends $\{U_i\}_{i=1}^m$ to R.

R \rightarrow S: R sends the decommitment (s, r_s) of com_s to S.

S \rightarrow R: The sender validates all values transmitted so far as follows: It checks that $s = \text{Open}(\text{com}_s, s, r_s)$ and that $\text{Vrfy}_s(i \| 0 \| \text{com}_{a_i\|B_i}, \tau_{a_i\|B_i}) = 1$, as well as $\text{Vrfy}_s(i \| 1 \| \tilde{a}_i\|\tilde{B}_i, \tau_{\tilde{a}_i\|\tilde{B}_i}) = 1$ for $i \in [m]$. If the check fails, then S aborts the protocol. Otherwise, for $i \in [m]$ it authenticates the commitment com_{z_i} by computing $\tau_{z_i} := \text{Sig}_{s_2}(i \| \text{com}_{z_i})$. Finally, it sends $(\{(u_i, r_{u_i}, \tau_{z_i})\}_{i=1}^m)$ to R.

R \rightarrow S: For $i \in [m]$, the receiver checks that $u_i = \text{Open}(\text{com}_{u_i}, u_i, r_{u_i})$; if not, it aborts. Then R sets $\tilde{a}_i\|\tilde{B}_i\|\tau_{\tilde{a}_i\|\tilde{B}_i} := U_i \oplus u_i$. It next checks that

Fig. 8. A bounded OT protocol π from two stateless tokens

$\text{Vrfy}_s(i \| 1 \| \tilde{a}_i \| \tilde{B}_i, \tau_{\tilde{a}_i \| \tilde{B}_i}) = 1$ for $i \in [m]$. If this check does not hold, it aborts the protocol. Otherwise, for $i \in [m]$, it runs the token \mathcal{T}_5 with $(i, \text{com}_{z_i}, z_i, r_{z_i}, \tau_{z_i})$ and obtains in return (V_i, w'_i, r'_{w_i}) . The receiver checks that $w'_i = \text{Open}(\text{com}_{w_i}, w'_i, r'_{w_i})$, that $CV_i = \tilde{a}_i z_i^T + \tilde{B}_i$, and that $w'_i \neq \perp$ for $i \in [m]$. If the check fails, then it aborts the protocol. Otherwise, it chooses $h_i \leftarrow \mathbb{F}_2^{4\lambda}$ subject to the constraint that $b_i = z_i^T h_i$ for $i \in [m]$, and sends $\{(h_i, w'_i)\}_{i=1}^m$ to S .

$\mathsf{S} \rightarrow \mathsf{R}$: The sender checks that $w_i = w'_i$ for $i \in [m]$; if not, it aborts the protocol. Otherwise, for $i \in [m]$ it chooses $v_i^0, v_i^1 \leftarrow \{0, 1\}^d$ and computes $\tilde{x}_i^0 := \text{Ext}(GB_i h_i, v_i^0) \oplus x_i^0$ and $\tilde{x}_i^1 := \text{Ext}(GB_i h_i + Ga_i, v_i^1) \oplus x_i^1$. Finally it sends $\{(v_i^0, v_i^1, \tilde{x}_i^0, \tilde{x}_i^1)\}_{i=1}^m$ to the receiver.

$\mathsf{R} \rightarrow \mathsf{S}$: For $i \in [m]$, the receiver computes $x_i^{b_i} := \tilde{x}_i^{b_i} \oplus \text{Ext}(GV_i h_i, v_i^{b_i})$ and outputs it.

4 Black-Box Impossibility of OT Using One Stateless Token

In the previous section we showed that an unbounded number of universally composable OTs (and hence universally composable secure computation) is possible by having the parties exchange two stateless tokens. We show here that a construction of (even a single) OT using *one* stateless token is impossible using black-box techniques alone. Here, by “black-box” we refer to the simulator’s access to the code \mathcal{M} encapsulated in the token. Note that the token model as defined by Katz [32] is inherently *nonblack-box* and, in particular, the simulator is given the code \mathcal{M} that the malicious party submits to $\mathcal{F}_{\text{wrap}}$. Nevertheless, an examination of the proof of security of our two-token protocol in Section 3— as well as the proofs of security in almost all prior work [12,37,23,15,24,25,19]— shows that the simulator only uses this code in a black-box fashion, namely, by running the code to observe its input/output behavior but without using any internal structure of the code itself. We formalize this in what follows.

Specifically, we consider simulators of the form $\text{Sim} = (\text{Sim}_{\text{code}}, \text{Sim}_{\text{bb}})$, where Sim_{code} gets the code \mathcal{M} that the adversary submits to $\mathcal{F}_{\text{wrap}}$, and then runs Sim_{bb} as a subroutine while giving it oracle access to \mathcal{M} . Inspired by [8,9] we show that, restricting to such simulators, constructions of OT from one stateless token are impossible. Intuitively, for any such protocol proven secure using black-box techniques, Sim_{bb} must be able to extract the input of a corrupted token-creating party by interacting with \mathcal{M} in a black-box fashion. The real-world adversary can then use Sim_{bb} to extract the input of the honest token-creating party by running Sim_{bb} and answering its oracle queries by *querying the token* itself; for *stateless* tokens, querying the token (in the real world) is equivalent to black-box access to \mathcal{M} (in the ideal world).

Theorem 2. *There is no protocol π that uses one stateless token to securely realize \mathcal{F}_{OT} in the $\mathcal{F}_{\text{wrap}}$ -hybrid model whose security is proven using a simulator $\text{Sim} = (\text{Sim}_{\text{code}}, \text{Sim}_{\text{bb}})$ as defined above.*

Proof (Sketch). For completeness, the (single-session) OT functionality \mathcal{F}_{OT} is given in Figure 9. Let π be a protocol between a sender S and a receiver R , in which a single token is sent from the sender to the receiver. (The other case is handled analogously.) Consider the following environment \mathcal{Z}' and dummy adversary \mathcal{A}' corrupting the sender: \mathcal{Z}' chooses random bits s_0, s_1 , and b , and instructs the sender to run π honestly on input (s_0, s_1) , including submitting some code to $\mathcal{F}_{\text{wrap}}$ to create a token. Once the honest receiver outputs s , then \mathcal{Z}' outputs 1 if $s = s_b$ and 0 otherwise.

Suppose that π securely realizes \mathcal{F}_{OT} , where security is proved via a simulator $\text{Sim} = (\text{Sim}_{\text{code}}, \text{Sim}_{\text{bb}})$ as previously described. In the course of the proof, Sim_{bb} plays the role of a receiver while interacting with \mathcal{A}' , and Sim_{code} provides Sim_{bb} with black-box access to whatever code \mathcal{A}' submits to $\mathcal{F}_{\text{wrap}}$. At some point during its execution, Sim_{bb} must send some inputs $(\tilde{s}_0, \tilde{s}_1)$ to the ideal functionality \mathcal{F}_{OT} . It is not hard to see that we must have $(\tilde{s}_0, \tilde{s}_1) = (s_0, s_1)$ with all but negligible probability.

We now consider a different environment \mathcal{Z} and an adversary \mathcal{A} corrupting the receiver. \mathcal{Z} chooses random bits s_0, s_1, b and provides (s_0, s_1) as input to the honest sender; it outputs 1 iff \mathcal{A} outputs (s_0, s_1) . Note that \mathcal{A} receives a token from the honest sender as specified by π . Adversary \mathcal{A} works as follows:

Run Sim_{bb} , relaying messages from the honest sender to this internal copy of Sim_{bb} . Whenever Sim_{bb} makes a query to Sim_{code} to run \mathcal{M} (the code created by the honest sender) on some input q , adversary \mathcal{A} runs the token with input q (formally, \mathcal{A} sends $(\text{RUN}, \langle \text{sid}, S, R \rangle, q)$ to $\mathcal{F}_{\text{wrap}}$ and gives the response to Sim_{bb}). At some point, Sim_{bb} sends $(\tilde{s}_0, \tilde{s}_1)$ to \mathcal{F}_{OT} , at which point \mathcal{A} outputs $(\tilde{s}_0, \tilde{s}_1)$ and halts.

The key point is that *because the token is stateless*, there is no difference between Sim_{code} running the code \mathcal{M} , and \mathcal{A} querying the token via $\mathcal{F}_{\text{wrap}}$. Thus, \mathcal{A} provides a perfect simulation for Sim_{bb} , and we conclude that $(\tilde{s}_0, \tilde{s}_1) = (s_0, s_1)$ with all but negligible probability in an execution of π in the $\mathcal{F}_{\text{wrap}}$ -hybrid world. But this occurs with probability at most $1/2$ in an ideal-world evaluation of \mathcal{F}_{OT} . Thus, \mathcal{Z} can distinguish between the real- and ideal-world executions, contradicting the claimed security of π . \square

5 Coin Tossing Using One Stateless Token

In the previous section we showed that universally composable OT cannot be realized from one stateless token if only “black-box techniques” are used. We complement that result by showing that UC secure computation from one stateless token *is* feasible via *non*black-box techniques. Here, we find it somewhat simpler to construct a protocol for universally composable coin tossing rather than OT. (The coin-tossing functionality $\mathcal{F}_{\text{coin}}$ is defined in the natural way; see Figure 10.) Note that coin tossing suffices for general secure computation under a variety of cryptographic assumptions [10,35].

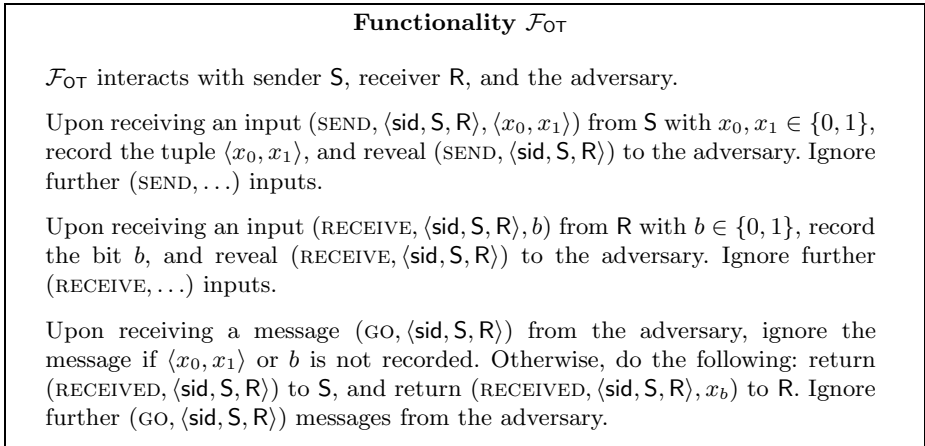


Fig. 9. The ideal \mathcal{F}_{OT} functionality for bit-OT

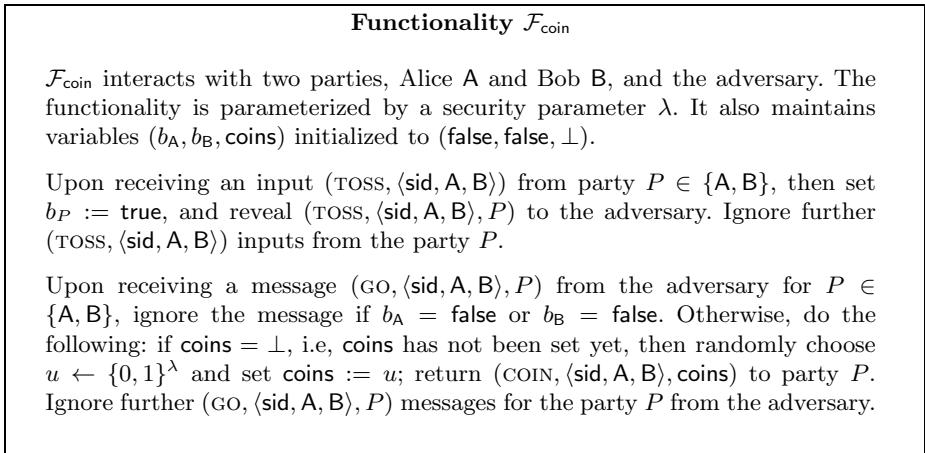


Fig. 10. The ideal $\mathcal{F}_{\text{coin}}$ functionality for coin tossing

At a high level, our protocol follows the general structure of Blum's coin-tossing protocol [4]. This protocol consists of three moves. In the first move (B1), Alice commits to a random x and sends com_x to Bob, who in return chooses a random value y and sends it to Alice (B2). In the last move (B3), Alice sends the decommitment to x and both parties output $x \oplus y$.

To obtain a UC coin-tossing protocol that follows this basic approach, we need to be able to *simulate* each party. In particular, if Alice is malicious then the simulator needs to extract the message contained in com_x (extractability), whereas when Bob is corrupted the simulator needs to open the commitment in an arbitrary way (equivocation). We achieve both goals by having Bob send a single stateless token \mathcal{T}_B to Alice. This token behaves in two different ways, depending on its input. The first task of the token is to generate a random value

e upon seeing com_x (we discuss the details later); the second task is to generate a notification t for Bob that Alice knows the decommitment x . The value e is used for equivocation, and the notification t gives the simulator the ability to extract x . We give further details next. In the high-level description here, we assume the token is created honestly; we deal with a potentially malicious token when we formally define the protocol, below.

Achieving Extractability. Similar to Section 3, the token only works on authenticated inputs. That is, whenever Alice wants to query the token on some inputs, she needs to first ask Bob to compute a MAC on those inputs. In order to achieve extractability, we let Alice query the token on input $(\text{com}_x, x, r_x, \tau_x)$, where $\text{com}_x \leftarrow \text{SCom}(x; r_x)$ comes from Alice and τ_x is a MAC tag on com_x . The output of the token is a random value t that can be seen as a notification to Bob that Alice knows the decommitment x . As in the previous OT protocol, the authentication of the inputs guarantees that Alice makes exactly one valid query to the token: more than one valid query would imply that Alice has violated security of the MAC or binding of com_x ; if Alice makes no query, then she cannot guess t . By forcing Alice to query the token exactly once, we can now easily construct a simulator that extracts the value x while emulating $\mathcal{F}_{\text{wrap}}$. Modifying Blum's coin-tossing protocol, we have the following step:

(B1) Alice commits to x by executing $\text{com}_x \leftarrow \text{SCom}(x; r_x)$. She sends the commitment com_x to Bob, who in turn computes a tag τ_x on com_x and sends τ_x to Alice. Alice runs the token with $(\text{com}_x, x, r_x, \tau_x)$ to obtain output t , and sends t to Bob. Bob checks if t is correct.

Achieving Equivocation. Toward this goal, we further modify the protocol and the token. Alice sends com_x and a dummy commitment com_M before getting the tag τ_x on $\text{com}_x \parallel \text{com}_M$ from Bob. The token also gets as an additional input this commitment com_M ; it outputs a random value e on input $(\text{com}_x, \text{com}_M, \tau_x)$ and a random value t on input $(\text{com}_x, \text{com}_M, x, r_x, \tau_x)$. In step (B3), instead of sending the decommitment (x, r_x) of com_x to Bob, Alice sends x together with a witness-indistinguishable (WI) proof that either x is a valid decommitment of com_x , or com_M contains code that outputs the actual output e of the token \mathcal{T}_{B} . (This is where we use the nonblack-box techniques of Barak [2].)

Due to the binding property of com_M , and because the notification e is unpredictable, Alice cannot commit to such code in com_M . The simulator, however, takes advantage of the fact that it obtains the code of the token generated by Bob while emulating $\mathcal{F}_{\text{wrap}}$. Then, as in [2], the simulator's ability to predict the output of the token beforehand can be used to achieve equivocation. We remark that in contrast to Barak's work, we do not need to use universal arguments; this is because $\mathcal{F}_{\text{wrap}}$ is parameterized with a fixed polynomial bounding the running time of the token (whereas Barak had to handle any polynomial running time).

More formally, we change the protocol as follows:

Token. The token input is either $(\text{com}_x, \text{com}_M, \tau_x)$ or $(\text{com}_x, x, r_x, \text{com}_M, \tau_x)$ and in both cases it checks the validity of τ_x before responding. In the first

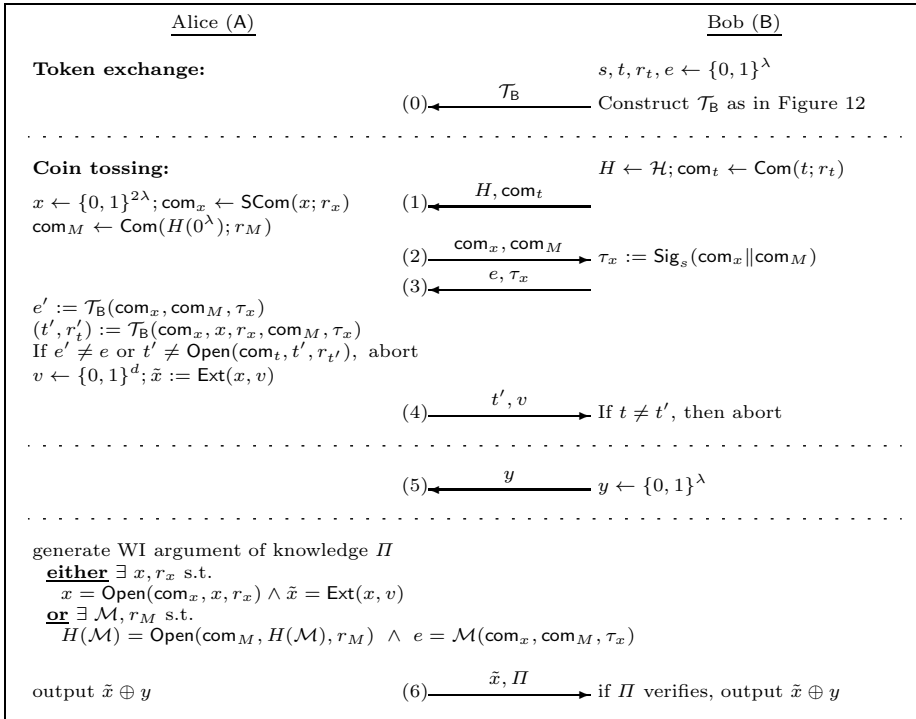


Fig. 11. A coin-tossing protocol ψ from a single stateless token

case, the token outputs a random value e . In the second case, it returns a random value t .

Protocol. Once Alice obtains the token, she runs the following protocol with Bob:

(B1). Alice commits to x and sends com_x together with a dummy commitment com_M to Bob. In turn, Bob authenticates $\text{com}_x || \text{com}_M$ using a MAC and sends (τ_x, e) to Alice. Alice invokes the token twice, first with $(\text{com}_x, \text{com}_M, \tau_x)$ and then with $(\text{com}_x, x, r_x, \text{com}_M, \tau_x)$, obtaining the values e' and t' , respectively. She checks if $e = e'$ and, if so, sends t' to Bob. Finally, Bob checks if $t' = t$.

(B2). Bob sends Alice a random value y .

(B3). Alice sends x together with a WI proof that either (i) (x, r_x) is the decommitment of com_x , or (ii) com_M is a commitment to a Turing machine \mathcal{M} such that $\mathcal{M}(\text{com}_x, \text{com}_M, \tau_x) = e$.

5.1 Formal Description of the Protocol

The protocol ψ between Alice A and Bob B consists of an initial token-exchange phase, followed by a coin-tossing phase for generating a random λ -bit string. We now describe ψ formally; see also Figure 11.

On input $(\text{com}_x, \text{com}_M, \tau_x)$ do: if $\text{Vrfy}_s(\text{com}_x \parallel \text{com}_M, \tau_x) = 1$ output e else output \perp	On input $(\text{com}_x, x, r_x, \text{com}_M, \tau_x)$ do: if $\text{Vrfy}_s(\text{com}_x \parallel \text{com}_M, \tau_x) = 1$ and $x = \text{Open}(\text{com}_x, x, r_x)$ output (t, r_t) else output (\perp, \perp)
---	--

Fig. 12. The Turing machine \mathcal{M} embedded in the sender-created token \mathcal{T}_B . Here, $e, s, t, r_t \in \{0, 1\}^\lambda$ are chosen uniformly at random and embedded in the token.

Token-Exchange Phase. Bob generates a single token \mathcal{T}_B and sends it to Alice. Bob's token \mathcal{T}_B encapsulates the code \mathcal{M} described in Figure 12, where s, e, t, r_t are each chosen uniformly from $\{0, 1\}^\lambda$.

Coin-Tossing Phase. In this phase, Alice and Bob proceed as follows.

- B \rightarrow A: Bob chooses a collision-resistant hash function $H \leftarrow \mathcal{H}$. He also commits to the value t (that he used when creating the token) by executing $\text{com}_t \leftarrow \text{Com}(t; r_t)$. He sends H, com_t to Alice.
- A \rightarrow B: Alice chooses a value $x \leftarrow \{0, 1\}^{2\lambda}$ and commits to x and $H(0^\lambda)$ by executing $\text{com}_x \leftarrow \text{SCom}(x; r_x)$ and $\text{com}_M \leftarrow \text{Com}(H(0^\lambda); r_M)$. She sends com_x and com_M to Bob.
- B \rightarrow A: Bob generates a MAC tag $\tau_x := \text{Sig}_s(\text{com}_x \parallel \text{com}_M)$, and sends (e, τ_x) to Alice. Recall that Bob has already chosen e and embedded the value in the token \mathcal{T}_B in the token exchange phase.
- A \rightarrow B: Alice runs the token \mathcal{T}_B with $(\text{com}_x, \text{com}_M, \tau_x)$ and obtains e' in response. Then she runs the token with $(\text{com}_x, x, r_x, \text{com}_M, \tau_x)$ and obtains (t', r'_t) . Alice checks if $e' = e$ and $t' = \text{Open}(\text{com}_t, t', r'_t)$. If not, she aborts the protocol. Otherwise, she chooses $v \leftarrow \{0, 1\}^d$ and sends (t', v) to Bob, where d is an appropriate seed length for the extractor.
- B \rightarrow A: Bob checks that $t = t'$, and aborts if not. Otherwise he chooses $y \leftarrow \{0, 1\}^\lambda$ and sends it to Alice.
- A \rightarrow B: Alice sends $\tilde{x} := \text{Ext}(x, v)$ and gives a WI argument of knowledge that $(H, \text{com}_x, \text{com}_M, e, \tau_x, v, \tilde{x})$ belongs to the \mathcal{NP} language L defined by the following relation R_L :

$R_L((H, \text{com}_x, \text{com}_M, e, \tau_x, v, \tilde{x}), (\alpha, \beta)) = 1$ if either one of the following holds:

- (i) $\alpha = \text{Open}(\text{com}_x, \alpha, \beta)$ and $\tilde{x} = \text{Ext}(\alpha, v)$.
- (ii) $H(\alpha) = \text{Open}(\text{com}_M, H(\alpha), \beta)$. In addition, treating α as the description of a Turing machine, the execution of $\alpha(\text{com}_x, \text{com}_M, \tau_x)$ outputs e in time at most $p(\lambda)$, where p is the polynomially bounded running time defined by $\mathcal{F}_{\text{wrap}}$.

If the proof succeeds, both parties output $\tilde{x} \oplus y$.

Theorem 3. Assume Com is computationally hiding and strongly binding, SCom is statistically hiding and computationally binding, MAC is a deterministic, unforgeable message-authentication code, \mathcal{H} is a family of collision-resistant hash functions, and the proof system is a witness-indistinguishable argument of knowledge. Then ψ securely realizes $\mathcal{F}_{\text{coin}}$ in the $\mathcal{F}_{\text{wrap}}$ -hybrid model.

5.2 Proof Idea

In this section, we briefly sketch the main ideas behind the proof of Theorem 3. A complete proof is deferred to the full version.

To show the security of the protocol, we need to construct a simulator Sim for any PPT environment \mathcal{Z} such that $\text{EXEC}_{\mathcal{A}, \pi, \mathcal{Z}}^{\mathcal{F}_{\text{wrap}}} \approx \text{IDEAL}_{\mathcal{F}_{\text{coin}}, \text{Sim}, \mathcal{Z}}$, where \mathcal{A} is the dummy adversary. Below, we briefly provide ideas for simulation.

Corrupting Alice. Note that Alice cannot commit to a TM \mathcal{M} in com_M that will output e , since e is chosen at random independently. Therefore, from the collision-resistance of H , Alice has to show that $\tilde{x} = \text{Ext}(x, v)$ in the WI proof. Since v appears in the communication transcript, by forcing malicious Alice A^* to query \mathcal{T}_B with x exactly once, the simulator Sim can extract the value \tilde{x} from binding of com_x . Then, upon receiving the random string coins from $\mathcal{F}_{\text{coin}}$, the simulator can send $y = \tilde{x} \oplus \text{coins}$ to A^* . To see how the protocol forces exactly one query to the token, note that malicious Alice is not able to generate a valid t without querying \mathcal{T}_B , since t is random and com_t is hiding. Also, note that A^* cannot query \mathcal{T}_B with different x s without contradicting the unforgeability of the tag τ_x and/or the binding of com_x . This allows the simulator to extract the value of x .

Corrupting Bob. The simulator Sim needs to equivocate the value \tilde{x} so that it may hold that $\tilde{x} = \text{coins} \oplus y$, where coins is the random string from $\mathcal{F}_{\text{coin}}$. This is achieved as follows. While emulating $\mathcal{F}_{\text{wrap}}$, the simulator obtains the token code \mathcal{M} generated by Bob, then it generates $\text{com}_M \leftarrow \text{Com}(H(\mathcal{M}); R_M)$. Given the hiding property of com_M , the simulated transcript is indistinguishable from that in the $\mathcal{F}_{\text{wrap}}$ -hybrid world. Now, with the witness (\mathcal{M}, R_M) in the WI proof, the simulator can send any value \tilde{x} .

References

1. Aumann, Y., Lindell, Y.: Security against covert adversaries: Efficient protocols for realistic adversaries. In: TCC 2007. LNCS, vol. 4392, pp. 137–156. Springer, Heidelberg (2007)
2. Barak, B.: How to go beyond the black-box simulation barrier. In: 42nd Annual Symposium on Foundations of Computer Science (FOCS), pp. 106–115. IEEE (2001)
3. Barak, B., Canetti, R., Nielsen, J.B., Pass, R.: Universally composable protocols with relaxed set-up assumptions. In: 45th Annual Symposium on Foundations of Computer Science (FOCS), pp. 186–195. IEEE (2004)
4. Blum, M.: Coin flipping by telephone. In: Proc. IEEE COMPCOM, pp. 133–137 (1982)
5. Brands, S.: Untraceable off-line cash in wallets with observers (extended abstract). In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 302–318. Springer, Heidelberg (1994)
6. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: 42nd Annual Symposium on Foundations of Computer Science (FOCS), pp. 136–145. IEEE (2001) (December 2005), <http://eprint.iacr.org/2000/067>

7. Canetti, R.: Obtaining universally composable security: Towards the bare bones of trust (invited talk). In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 88–112. Springer, Heidelberg (2007)
8. Canetti, R., Fischlin, M.: Universally composable commitments. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 19–40. Springer, Heidelberg (2001)
9. Canetti, R., Kushilevitz, E., Lindell, Y.: On the limitations of universally composable two-party computation without set-up assumptions. *Journal of Cryptology* 19(2), 135–167 (2006)
10. Canetti, R., Lindell, Y., Ostrovsky, R., Sahai, A.: Universally composable two-party and multi-party secure computation. In: 34th Annual ACM Symposium on Theory of Computing (STOC), pp. 494–503. ACM Press (2002)
11. Canetti, R., Pass, R., Shelat, A.: Cryptography from sunspots: How to use an imperfect reference string. In: 48th Annual Symposium on Foundations of Computer Science (FOCS), pp. 249–259. IEEE (2007)
12. Chandran, N., Goyal, V., Sahai, A.: New constructions for UC secure computation using tamper-proof hardware. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 545–562. Springer, Heidelberg (2008)
13. Chaum, D., Pedersen, T.P.: Wallet databases with observers. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 89–105. Springer, Heidelberg (1993)
14. Cramer, R.J.F., Pedersen, T.P.: Improved privacy in wallets with observers (extended abstract). In: Helleseth, T. (ed.) EUROCRYPT1993. LNCS, vol. 765, pp. 329–343. Springer, Heidelberg (1994)
15. Damgård, I., Nielsen, J.B., Wichs, D.: Universally composable multiparty computation with partially isolated parties. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 315–331. Springer, Heidelberg (2009)
16. Damgård, I., Pedersen, T.P., Pfitzmann, B.: On the existence of statistically hiding bit commitment schemes and fail-stop signatures. *Journal of Cryptology* 10(3), 163–194 (1997)
17. Desmedt, Y., Quisquater, J.-J.: Public-key systems based on the difficulty of tampering (is there a difference between DES and RSA?). In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 111–117. Springer, Heidelberg (1987)
18. Dodis, Y., Yampolskiy, A.: A verifiable random function with short proofs and keys. In: Vaudenay, S. (ed.) PKC 2005. LNCS, vol. 3386, pp. 416–431. Springer, Heidelberg (2005)
19. Döttling, N., Kraschewski, D., Müller-Quade, J.: Unconditional and composable security using a single stateful tamper-proof hardware token. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 164–181. Springer, Heidelberg (2011)
20. Döttling, N., Mie, T., Müller-Quade, J., Nilges, T.: Basing obfuscation on simple tamper-proof hardware assumptions. *Cryptology ePrint Archive*, Report 2011/675 (2011), <http://eprint.iacr.org/>
21. Dubovitskaya, M., Scafuro, A., Visconti, I.: On efficient non-interactive oblivious transfer with tamper-proof hardware. *Cryptology ePrint Archive*, Report 2010/509 (2010)
22. Fischlin, M., Pinkas, B., Sadeghi, A.-R., Schneider, T., Visconti, I.: Secure set intersection with untrusted hardware tokens. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 1–16. Springer, Heidelberg (2011)
23. Goldwasser, S., Kalai, Y.T., Rothblum, G.N.: One-time programs. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 39–56. Springer, Heidelberg (2008)
24. Goyal, V., Ishai, Y., Mahmoody, M., Sahai, A.: Interactive locking, zero-knowledge PCPs, and unconditional cryptography. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 173–190. Springer, Heidelberg (2010)

25. Goyal, V., Ishai, Y., Sahai, A., Venkatesan, R., Wadia, A.: Founding cryptography on tamper-proof hardware tokens. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 308–326. Springer, Heidelberg (2010)
26. Haitner, I., Reingold, O.: Statistically-hiding commitment from any one-way function. In: 39th Annual ACM Symposium on Theory of Computing (STOC), pp. 1–10. ACM Press (2007)
27. Hazay, C., Lindell, Y.: Constructions of truly practical secure protocols using standard smartcards. In: ACM CCS 2008: 15th ACM Conf. on Computer and Communications Security, pp. 491–500. ACM Press (2008)
28. Hofheinz, D., Müller-Quade, J., Unruh, D.: Universally composable zero-knowledge arguments and commitments from signature cards. In: Proc. 5th Central European Conference on Cryptology, MoraviaCrypt (2005)
29. Horvitz, O., Katz, J.: Bounds on the efficiency of “black-box” commitment schemes. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 128–139. Springer, Heidelberg (2005)
30. Ishai, Y., Prabhakaran, M., Sahai, A.: Founding cryptography on oblivious transfer - efficiently. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 572–591. Springer, Heidelberg (2008)
31. Järvinen, K., Kolesnikov, V., Sadeghi, A.-R., Schneider, T.: Embedded SFE: Offloading server and network using hardware tokens. In: Sion, R. (ed.) FC 2010. LNCS, vol. 6052, pp. 207–221. Springer, Heidelberg (2010)
32. Katz, J.: Universally composable multi-party computation using tamper-proof hardware. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 115–128. Springer, Heidelberg (2007)
33. Kilian, J.: Founding cryptography on oblivious transfer. In: STOC, pp. 20–31 (1988)
34. Kolesnikov, V.: Truly efficient string oblivious transfer using resettable tamper-proof tokens. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 327–342. Springer, Heidelberg (2010)
35. Lin, H., Pass, R., Venkatasubramanian, M.: A unified framework for concurrent security: Universal composable from stand-alone non-malleability. In: 41st Annual ACM Symposium on Theory of Computing (STOC), pp. 179–188. ACM Press (2009)
36. Lindell, Y.: General composition and universal composable in secure multiparty computation. *Journal of Cryptology* 22(3), 395–428 (2009)
37. Moran, T., Segev, G.: David and Goliath commitments: UC computation for asymmetric parties using tamper-proof hardware. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 527–544. Springer, Heidelberg (2008)
38. Naor, M.: Bit commitment using pseudorandomness. *Journal of Cryptology* 4(2), 151–158 (1991)