

# Efficient Non-Interactive Oblivious Transfer with Tamper-Proof Hardware

MARIA DUBOVITSKAYA\*    ALESSANDRA SCAFURO†    IVAN VISCONTI‡

August 31st, 2010

## Abstract

Oblivious transfer (OT, for short) [Rab81] is a fundamental primitive in the foundations of Cryptography. While in the standard model OT constructions rely on public-key cryptography, only very recently Kolesnikov in [Kol10] showed a truly efficient string OT protocol by using tamper-proof hardware tokens. On the positive side, his construction only needs few evaluations of a block cipher and requires stateless (therefore resettable) tokens. On the negative side, security against malicious sender is only achieved in a *covert* sense (a malicious sender can actually obtain the private input of the receiver but the receiver can detect this malicious behavior with probability  $1/2$ ), the protocol is interactive and does not enjoy forward security (by breaking a token one violates the security of all previously played OTs).

In this work, we propose new techniques to achieve truly efficient string OT using tamper-proof hardware tokens. While from one side our tokens need to be stateful, our protocol enjoys several appealing features: 1) it is secure against malicious receivers and unconditionally secure against malicious senders, 2) it is non-interactive, 3) it is forward secure, 4) it enjoys adaptive input, therefore tokens can be sent before parties know their private inputs, and can (virtually) be reused any polynomial number of times. This gracefully fits a large number of client-server settings (digital TV, e-banking) and thus many practical applications.

We finally notice that as shown by Goyal et al. in [GIS<sup>+</sup>10], non-interactive OT with stateless tokens is impossible. Therefore our use of a stateful token is not avoidable.

**Keywords:** OT, Tamper-Proof Hardware Tokens.

## 1 Introduction

Oblivious transfer (OT) [Rab81, EGL85] is among the most investigated primitives in modern Cryptography. Its importance is due to several reasons. OT is conceptually useful on his own, and as subprotocol, for realizing secure multi-party computation [IPS08, Kil88] and in general for cryptographic protocol design. Unfortunately in the standard model OT constructions rely on public-key encryption, which limits its practical applicability when OT is used as subprotocol and several instances of it have to be executed by the larger protocol. This efficiency problem is even more critical when one would like to use lightweight cryptographic hardware as smart cards or RFID chips. Furthermore achieving OT secure under concurrent composition in the standard model has been proved to be impossible [CKL03].

---

\*IBM Research, Zurich, SWITZERLAND. E-mail: [mdu@zurich.ibm.com](mailto:mdu@zurich.ibm.com).

†Dip. di Informatica ed Appl., University of Salerno, ITALY. E-mail: [scafuro@dia.unisa.it](mailto:scafuro@dia.unisa.it).

‡Dip. di Informatica ed Appl., University of Salerno, ITALY. E-mail: [visconti@dia.unisa.it](mailto:visconti@dia.unisa.it).

Among different assumptions proposed in the past to overcome the impossibility results in the standard model, the use of tamper proof hardware tokens has obtained an increasing interest for its capability of designing elegant and extremely powerful protocols. First of all, tamper-resistant hardware tokens implementing some well-known cryptographic tasks are widely available (e.g., smart cards) and have been used in practice in the last decades. Second, it has been shown that protocols benefit from the use of such tokens either in practice obtaining better computational and communication complexities, and/or from a theoretic point of view, obtaining feasibility results otherwise impossible to achieve.

Hardware tokens are supposed to carry on some publicly known tasks such that any player can initialize them with their own secrets. Then tokens are distributed to the other players (honest or malicious) that can use them in a “black box” way, i.e., by sending an input to the token and getting back only the output without obtaining any information about its internal state. Obviously nothing can prevent a malicious player from constructing adversarially designed tokens on its own, and thus it is challenging to design advanced cryptographic protocols that benefit from the use of honest tokens, and simultaneously maintain all the desired security guarantees in presence of adversarially designed tokens.

## 1.1 Evaluating Protocols with Hardware Tokens

There are several physical properties that could be enjoyed by tokens used in a cryptographic protocol. Subtle differences can heavily simplify or complicate the design of a cryptographic protocol. The literature could quickly become a jungle if the precise assumptions on the used hardware and the power of the adversary are not clearly stated. Before discussing related work, here we take a chance to summarize and discuss requirements, and features of tamper-proof hardware tokens proposed in literature, and their impact on the efficiency and security of the protocols. This will simplify our discussion on previous work and will clearly place our contribution in the state-of-the-art.

**Token requirements.** We list below specific features/assumptions that differentiate heterogeneous tokens.

**Resettability:** a resettable token can be reset to a previous state, i.e., the adversary is able to manipulate the token and reset its internal variables (e.g., a counter) to its initial state. This is a positive property for a protocol, as it decreases the hardware assumption (i.e., the protocol does not need the additional hardware assumption that a token can not be reset).

**State assumption:** it is sometimes the case that next computations of a token are based on previous ones, in particular depending on the content of the answered queries. Indeed, when answering queries, a token could for instance erase some data, increment a counter and so on. Such tokens have to reliably keep their state even without the power supply, and are referred to as *stateful* in contrast to *stateless* tokens, that instead do not require any permanent updatable memory. Obviously a stateless token require a less demanding hardware assumption than a stateful token.

**Knowledge of the code:** sometimes security proofs rely on the ability of the simulator to rewind the tokens received from potentially malicious parties. This means that the adversary that produced the token is supposed to “know” its code. Such an assumption prevents to model real-life adversaries who may simply pass on hardware tokens obtained from one party to another party, or may construct a *super*-token by encapsulating received tokens, without actually knowing its content. This assumption is therefore problematic to

justify when one wants to use such tokens to prove results under concurrent compositions. Indeed in these scenarios, it is natural to consider a man-in-the-middle adversary that passes on tokens, or that builds super tokens on top of some other tokens.

**Re-usability:** given the overhead of exchanging tokens and in some cases, their manufacturing cost, an extremely important property for the design of practical protocols is that tokens should be reusable any polynomial number of times. A consequence of this property is that protocols should be secure under adaptive inputs, which means that when the token is sent, the inputs for the computations that will be later performed are not known yet and will be dynamically decided according to the outputs of the already performed computations.

*Efficiency.* Here we stress some more standard measures.

**Number of rounds** represents the number of communication rounds played by parties. This is a classical measure in cryptographic protocols, and non-interactiveness makes a protocol suitable for a much larger number of applications.

**Number of queries** represents the number of times a token is required to answer a query. For practical reasons tokens are usually lightweight devices with a slow communication interface, therefore minimizing the number of queries is worthy.

**Amount of computations** represents the amount of computations required to both parties and tokens. This is also a classical measure with a specific focus here because of the power constraints of available tamper-proof tokens. In this context even the use of symmetric key encryptions in place of public key encryption turns out to be a big improvements in efficiency.

**Quantity and directions:** the introduction of physical tokens in a distributed computation becomes hard to justify when the number of tokens required to carry on the computation is large, and tokens have to be generated by many parties. Indeed, given the problematic aspects of producing and exchanging tokens one would like to see only a small number of tokens circulating in the system, possibly in one direction only, therefore matching several client-server scenarios (e.g., often users after a subscription to a service receive a smart card).

*Security issues.* We now point out some security issues when dealing with hardware tokens.

**Protocol composition:** an important measure for the security of a protocol is whether it maintains the security guarantees also when played in composition with itself, or with other protocols. Therefore one has to argue whether a protocol enjoys one-time, sequential, self-concurrent, general concurrent or universal composition.

**Adversary's power:** an adversary can be semi-honest (i.e., a protocol guarantees input/output privacy as long as the adversary follows honestly the protocol and then tries to analyze its view), covert (i.e., a protocol can leak some input/output privacy of honest players, as long as there is a reasonable chance to catch the malicious party) and malicious (i.e., a protocol must protect honest players' input/output privacy from any adversarial behavior). Moreover, security can be based on computational assumptions, therefore against efficient adversaries only, or unconditional, therefore against unbounded adversaries.

**Cryptographic assumptions and models:** the use of general assumptions should be preferred since a protocol can then be instantiated under various candidate assumptions. Moreover, the use of random oracles (which limits the adversary to a black-box use of a collision-resistant hash function) and in general of controversial conjectures should be avoided.

**Forward security:** in case the adversary at some point breaks the tamper resistance of the token and obtains its current state, previous computations should remain secure.

## 1.2 Related Work

The use of tamper-proof hardware tokens for cryptographic purposes was investigated by Goldreich and Ostrovsky focusing on oblivious RAM [GO96]. Moran and Naor [MN05] considered a relaxation of tamper-proof hardware called “tamper-evident seals”, and demonstrated the possibility of implementing cryptographic primitives based on this relaxed notion. Then Katz in [Kat07] put forth a more general formalization of a tamper-proof token model. He provides a scheme achieving Universally Composable (UC-secure) two-party computation under the DDH assumption. His construction uses tokens assumed to be equipped with a built-in source of randomness (this can be more concretely implemented by using stateful tokens and pseudo-random functions) that cannot be reset and requires that tokens be sent by both parties. In [Kat07], it is assumed that once a party creates a hardware token and sends it off, it can not send any messages to the token (but can receive messages from it). Also it is assumed that all parties (including the malicious ones) know the code run by the hardware token that they distributed, which technically means that tokens can be rewound.

Chandran, Goyal and Sahai [CGS08] extended the results of Katz and suggested three main improvements: first, they considered resettable tokens, second, their construction is based on general assumption (enhanced trapdoor permutations) instead of DDH. And finally, their security proof does not rely on the simulator’s ability to rewind hardware tokens, thus they make no assumptions on how malicious parties create the hardware token which they distribute. As for communication between token and its creator they require the opposite assumption: token can not send any messages to its creator, but can potentially receive messages from it.

Moran and Segev [MS08] realized the UC commitment functionality in the same model of [Kat07] but providing several improvements. In their constructions tokens are passed in one direction only, therefore fitting all real-life scenarios which are inherently asymmetric (e.g, voting systems). Furthermore they showed a protocol realizing the commitment functionality without relying on computational assumptions and proved that the existence of one-way functions suffices to realize the multiple commitment functionality.

We notice that both the results of [CGS08, MS08] focus on constructing UC commitments. Then the compiler of [CLOS02] must be used to obtain security for general functionalities. Therefore it turns out that both constructions do not achieve efficiently the oblivious transfer functionality.

Goyal, Ishai, Sahai and Wadia in [GIS<sup>+</sup>10] considered the general question of basing cryptography on tamper-proof tokens, under minimal computational assumptions. They considered both stateful and stateless tokens and, in particular, they showed that by exchanging simple stateful hardware tokens, any functionality can be realized with unconditional security against malicious parties, and that stateless tokens and one-way functions are sufficient for UC-secure computation.

Very recently, Goyal, Ishai, Mahmood and Sahai in [GIMS10] focusing on feasibility results addressed the challenging issue of achieving unconditional security by using only stateless token. On top of the model proposed in [CGS08] where parties do not need to know the token’s code, they prove that if token encapsulation is possible (i.e., a token can be used to encapsulate other tokens) then unconditional UC-secure OT is possible. Furthermore, sticking on feasibility results they prove that if encapsulation is not possible then there exist no protocols using stateless tokens and achieving statistically secure OT.

The question of using simple tamper proof hardware tokens for practical secure computation

has been addressed recently by Kolesnikov [Kol10]. On top of some new techniques and a careful use of strong PRPGs, Kolesnikov presented an efficient protocol for string OT, relying on resettable (and actually, stateless) tamper-proof tokens. This protocol is secure against covert sender and malicious receiver under sequential composition. All parties, including the tokens only have to run few evaluations of a block cipher. Security against covert sender means that the input/output privacy of the receiver can be compromised, but there is a deterrence factor that can expose the cheating behavior of the sender, and this can be sufficient in some applications. If the token is semi-honest (e.g., if it is provided by a trusted entity, but adversarially initialized), then his protocol is secure against malicious adversaries under concurrent general composition. However this last model assumes that an adversary can not create adversarially designed tokens, which is unrealistic in practice.

Concerning the covert setting, note that the deterrence factor in practice would be more effective in presence of evidence of cheating, that is, once a covert sender get caught there should be some proof of the cheating that is publicly verifiable (or at least verifiable by some specific entity). We notice that the protocol of Kolesnikov does not provide evidence of cheating. Indeed in case the sender prepare a stateful token cheating only once, even if the receiver discovers the attempt of cheating it has no chance to prove it.

### 1.3 Our Contribution

We propose a truly efficient non-interactive string oblivious transfer (OT) protocol based on stateful tamper-proof tokens and the existence of one-way functions. From a practical point of view, the only computations required consist of few evaluations of forward-secure PRGs. We obtain the optimal round complexity by means of a careful use of *two stateful* hardware tokens. These tokens are created by one party (the sender) and then distributed to the other party (the receiver), like in standard client-server setting (digital TV, e-banking). Furthermore our protocol guarantees unconditional security against malicious senders.

Our result is obtained by means of a new technique: the receiver will play separately with both tokens essentially emulating an OT protocol between the tokens that are not able to communicate directly with each other. Jumping ahead, we will show that one can move the interaction from the communication network that connects remote parties to the location of one party only that by means of his computing device can run a protocol between the two tokens. We stress that tokens are sent from the sender at the very beginning, when parties do not have to know their inputs yet. We believe that this new technique can lead to further improved constructions in the design of protocols with tamper proof hardware tokens.

We achieve a more general definition of OT (adaptive-input OT [Lin04]) that allows parties to choose their inputs adaptively on the outputs obtained from the previous protocol executions.

In our protocol the communication between parties is done only in one direction (from sender to receiver), so the protocol is *non-interactive*. After the receiver gets tokens from the sender it queries each token once and obtains the selected string. Neither a token nor the sender know about the receiver's choice.

We prove our protocol secure against malicious adversaries under sequential composition. Furthermore, our protocol is forward secure, i.e., if the adversary at some point breaks the token and obtains its current state, all future computations will be insecure, but the previous outputs of the token still remain completely hidden from the adversary. Our protocol also provides token re-usability, which means that once the token is released to the receiver it can be used to perform polynomial many protocol executions. We assume that one of the tokens can be rewound, i.e., the sender needs to know the content of the token. This in turn does not guarantee security against man-in-the-middle attacks. Concerning communication between the token and its creator, we

allow the communication between the creator and the token.

Properties	This Paper	Katz [Kat07]/Moran et al. [MS08]	Chandran et al. [CGS08]	Goyal et al. [GIS <sup>+</sup> 10]	Kolesnikov [Kol10]
Security	malicious, seq.	UC	UC	UC	covert, seq.
Assumptions	OWF	DDH/OWF	ETP	Unconditional/OWF	OWF
State assumption	stateful	stateful	stateless	stateful/stateless	stateless
Rounds	non-interactive	$O(1)$	$O(\log n^{1+\epsilon})$	non-interactive/ $O(n)$	$O(1)$
No. of Tokens	2	2/1	2	$O(n)/2$	1
Token rewind	yes	yes	no	yes/no	no
Re-usability	yes	yes	yes	no/yes	yes
Efficiency	yes	no	no	no/no	yes

**ETP** = enhanced trapdoor permutations      **seq.** = sequential composition

Figure 1: OT Protocols with Tamper-Proof Hardware (comparison).

In Fig. 1 we compare our and related work based on the properties we described above. In terms of efficiency, our protocol can be compared to Kolesnikov’s protocol. Our proposal is computationally slightly more efficient, and moreover non-interactive. Although we use stateful tokens, we prove that our protocol is forward secure against malicious static adversaries. The results of [Kol10] either have weaker security guarantees (covert sender), or rely on a semi-honest token and do not enjoy forward security. Finally, we stress that in light of the results of Goyal et al. [GIS<sup>+</sup>10], non-interactiveness can not be obtained with a stateless token.

In our security proof the simulator as in [Kat07, MS08, GIS<sup>+</sup>10] rewinds the token, which means that the sender needs to know the code executed by the token that it sent. This implies that an adversary never forwards to a party a token received from another party, and moreover it never constructs a new token on top of a received one. Given the above evident limitations in a truly concurrent setting, we therefore prove the protocol secure under sequential composition (as in [Kol10]). Instead in [Kat07, MS08, GIS<sup>+</sup>10] security under general concurrent composition has been proved still under the above limitation about forwarding or extending some received tokens. The construction given in [CGS08] instead does not suffer of this limitation and gives a proof of UC security without further weaknesses in the adversary model. However, as shown in Fig. 1, this is obtained by losing other important properties.

## 2 Definitions and Tools

**Notation.** We denote by  $n$  the security parameter and by PPT the property of an algorithm of running in probabilistic polynomial-time. A function  $\text{negl}(\cdot)$  is negligible in  $n$  (or just negligible) if for every polynomial  $p(\cdot)$  there exists a value  $N$  such that for all  $n > N$  it holds that  $\text{negl}(n) < 1/p(n)$ .

Let  $X = \{X(n, a)\}_{n \in N, a \in \{0,1\}^*}$  and  $Y = \{Y(n, a)\}_{n \in N, a \in \{0,1\}^*}$  be distribution ensembles. We say that  $X$  and  $Y$  are computationally indistinguishable, (i.e.,  $X \stackrel{c}{\equiv} Y$ ), if for every non-uniform PPT distinguisher  $D$  there exists a negligible function  $\text{negl}(\cdot)$ , such that for every  $a \in \{0,1\}^*$ ,  $\left| \Pr[D(X(n, a)) = 1] - \Pr[D(Y(n, a)) = 1] \right| < \text{negl}(n)$ .

We denote by  $y \stackrel{\$}{\leftarrow} \mathbf{B}(x)$  the value assigned to variable  $y$  as the output of the PPT algorithm  $\mathbf{B}$  on input  $x$ , while with  $y \leftarrow \mathbf{B}(x)$  we indicate the output of the deterministic polynomial-time algorithm  $\mathbf{B}$ . For a finite set  $\mathbf{A}$ ,  $x \stackrel{\$}{\leftarrow} \mathbf{A}$  denotes the assignment of a uniformly chosen element of

A to variable  $x$ .

**Pseudo-random generators.** A pseudo-random generator (PRG)  $\mathcal{G}: \{0, 1\}^n \rightarrow \{0, 1\}^{n+k}$  is a deterministic polynomial-time algorithm that receives a short truly random seed and stretches it into a longer string (expansion) that is indistinguishable from a truly random string (pseudo-randomness).

Here we consider a more powerful adversary that at some point is able to break into the computing device and thus to obtain its current state. While in this case all future computations will obviously be completely insecure, here we require that the security (pseudo-randomness of the generated bits) of previous computations still holds. This is a stronger security notion that is referred to as forward-security and requires that the previous output of the PRG remains secure even after the adversary gets the current state. Bellare and Yee in [BY03] provide a comprehensive treatment of forward-security in the context of shared-key based cryptographic primitives investigating on definitions and constructions of forward secure PRGs. We follow their definition of forward secure pseudo-random generator that we recall below.

**Forward secure pseudo-random generators.** A quadruple of efficient algorithms  $\text{GEN} = (\text{GEN.key}, \text{GEN.next}, k, t)$  is a stateful generator, when  $k$  and  $t$  are two positive integers indicating respectively the stretch factor and the maximum number of blocks the generator can produce,  $\text{GEN.key}$  is a probabilistic key generation algorithm that outputs the initial state (the seed) and  $\text{GEN.next}$  is the (deterministic) next step algorithm that on input the current state returns a pair consisting of a  $k$ -bit string output block and the next state. A sequence  $Out_1, Out_2, \dots, Out_t$  of  $k$ -bit output blocks is obtained by picking a seed  $St_0 \xleftarrow{\$} \text{GEN.key}$  and then iterating  $(Out_i, St_i) \leftarrow \text{GEN.next}(St_{i-1})$  for  $i = 1, \dots, t$ .  $St_{i-1}$  can be seen as the “key” or “seed” at time  $i$ . Forward security will require that this key is erased as soon as the next one has been generated, so that someone breaking into the machine gets only the current key.

**Forward security.** The forward security property is modeled by the following security experiment. The adversary  $\mathcal{A}$  runs in two stages: in the “find” stage it receives output blocks, one at a time, until it (adaptively) decides to break in and obtains the current state. In the “guess” stage, it must decide whether the output blocks it had been fed were outputs of the generator or were independent random bits. Let  $\mathcal{A}(\text{find}, Out, h)$  denote  $\mathcal{A}$  in the find stage, taking an output block  $Out$  and current history  $h$  and returning a pair  $(d, h)$  where  $h$  is an updated history and  $d \in \{\text{find}, \text{guess}\}$ . This stage continues until  $d = \text{guess}$  or all  $t$  output blocks have been generated (in the latter case the adversary is given the final state in the guess stage). More formally, the adversary will play in one of the following experiments:

<p>Experiment <math>\text{Exp}_{\text{GEN}}^{\text{fprg-1}}(\mathcal{A})</math></p> <p><math>St_0 \xleftarrow{\\$} \text{GEN.key}</math>  <math>i \leftarrow 0; h \leftarrow \epsilon</math>  <b>Repeat</b>  <math>i \leftarrow i + 1</math>  <math>(Out_i, St_i) \leftarrow \text{GEN.next}(St_{i-1})</math></p> <p><math>(d, h) \xleftarrow{\\$} A(\text{find}, Out_i, h)</math>  <b>Until</b> <math>(d = \text{guess})</math> or <math>(i = t)</math>  <math>g \xleftarrow{\\$} A(\text{guess}, St_i, h)</math>  <b>Return</b> <math>g</math></p>	<p>Experiment <math>\text{Exp}_{\text{GEN}}^{\text{fprg-0}}(\mathcal{A})</math></p> <p><math>St_0 \xleftarrow{\\$} \text{GEN.key}</math>  <math>i \leftarrow 0; h \leftarrow \epsilon</math>  <b>Repeat</b>  <math>i \leftarrow i + 1</math>  <math>(Out_i, St_i) \leftarrow \text{GEN.next}(St_{i-1})</math>  <math>Out_i \xleftarrow{\\$} \{0, 1\}^k</math>  <math>(d, h) \xleftarrow{\\$} A(\text{find}, Out_i, h)</math>  <b>Until</b> <math>(d = \text{guess})</math> or <math>(i = t)</math>  <math>g \xleftarrow{\\$} A(\text{guess}, St_i, h)</math>  <b>Return</b> <math>g</math></p>
--	--

**Definition 1** (Forward Secure Pseudo-Random Generator). *Let GEN be a stateful generator. GEN is a forward secure pseudo-random generator if for all non-uniform PPT adversary  $\mathcal{A}$  there exists a negligible function  $\text{negl}(\cdot)$  such that  $\text{Adv}_{\text{GEN}}^{\text{fprg}}(\mathcal{A}) = \left| \Pr[\text{Exp}_{\text{GEN}}^{\text{fprg-0}}(\mathcal{A}) = 1] - \Pr[\text{Exp}_{\text{GEN}}^{\text{fprg-1}}(\mathcal{A}) = 1] \right| < \text{negl}(n)$ .*

**Symmetric-key encryption.** We will be using the One-Time-Pad (OTP) encryption scheme, where the encryption and decryption functions correspond to an XOR computation. We will indeed need one encryption only per key, therefore OTP will suffice. Moreover, the keys used in our scheme will be pseudorandom, and we will show in the proof that this is sufficient for our purposes (obviously we will lose the information-theoretic security of OTP).

## 2.1 1-out-of-2 Oblivious Transfer

1-out-of-2 string Oblivious Transfer (OT) is a two-party protocol between a sender  $S$  and a receiver  $R$ . The sender  $S$  has two secret strings  $s_0, s_1$ , and the receiver  $R$  has a selection bit  $i$  from  $\{0, 1\}$ . Upon completion,  $R$  learns  $s_i$ , but nothing about  $s_{(1-i)}$ , and  $S$  learns nothing about  $i$ . Following the notation of [Lin04] the Oblivious Transfer functionality is a function:  $\mathcal{F}_{OT} : ((s_0, s_1), i) \mapsto (\varepsilon, s_i)$  where the sender gets no output. In this work we consider the case in which parties run sequentially a polynomial number of OT protocol executions and we require that at each execution players are allowed to select their inputs adaptively on outputs obtained in all previous executions. We refer to this variant of OT as  $\mathcal{F}_{adpt}^{OT}$  and we follow the definition of functionality with adaptively chosen inputs adopted by Lindell in [Lin04].

**Sequential self-composition and stateful parties.** The notation used in the above  $\mathcal{F}_{adpt}^{OT}$  allows through session ids the execution of multiple OT instances. When multiple instances are possible, a natural question is whether the adversary can interleave the execution of different instances or not. Here we will consider an adversary that mounts a sequential attack, therefore multiple instances of an OT protocol will not be interleaved, and thus they will be played sequentially. The adversary will use his control of the communication network. Among his features, it will be able to ask the honest sender to play OT sessions concerning different senders, as also the same stateful senders. The only restriction concerns the fact that before a new OT starts, all the messages of the previous execution must have been exchanged, or otherwise the previous execution will never be completed.

**Static corruption.** Concerning the corruption capabilities of the adversary, we will consider a static adversary only, which means that the adversary selects the party to be impersonated when the experiment starts. This means that in all executions, the adversary will always play as sender or will always play as receiver.

**Adaptively chosen inputs.** When multiple executions of a protocol are possible, each party needs an input for each execution. It is often the case that a party chooses the input of each execution adaptively basing its decision on the outputs of previous executions. Following the definition of Lindell in [Lin04] this property is formalized by requiring that inputs of each execution be provided by a PPT input-selecting Turing Machine  $M$ . More formally, each honest party  $P_i$  is augmented by a PPT input-selecting Turing Machine  $M_i$  and the input played in the session  $\text{sid}$  is the result of the computation of  $M_i(\text{init}_{P_i}, \text{sid}, \alpha_{\text{sid},1}, \dots, \alpha_{\text{sid},j})$  where  $\text{init}_{P_i}$  is the initial state of  $P_i$ ,  $\text{sid}$  is the session id of the current session, and  $\alpha_{\text{sid},1}, \dots, \alpha_{\text{sid},j}$  are



outputs obtained from the sessions concluded so far. Note that `init` contains the initial inputs along with some possible auxiliary information.

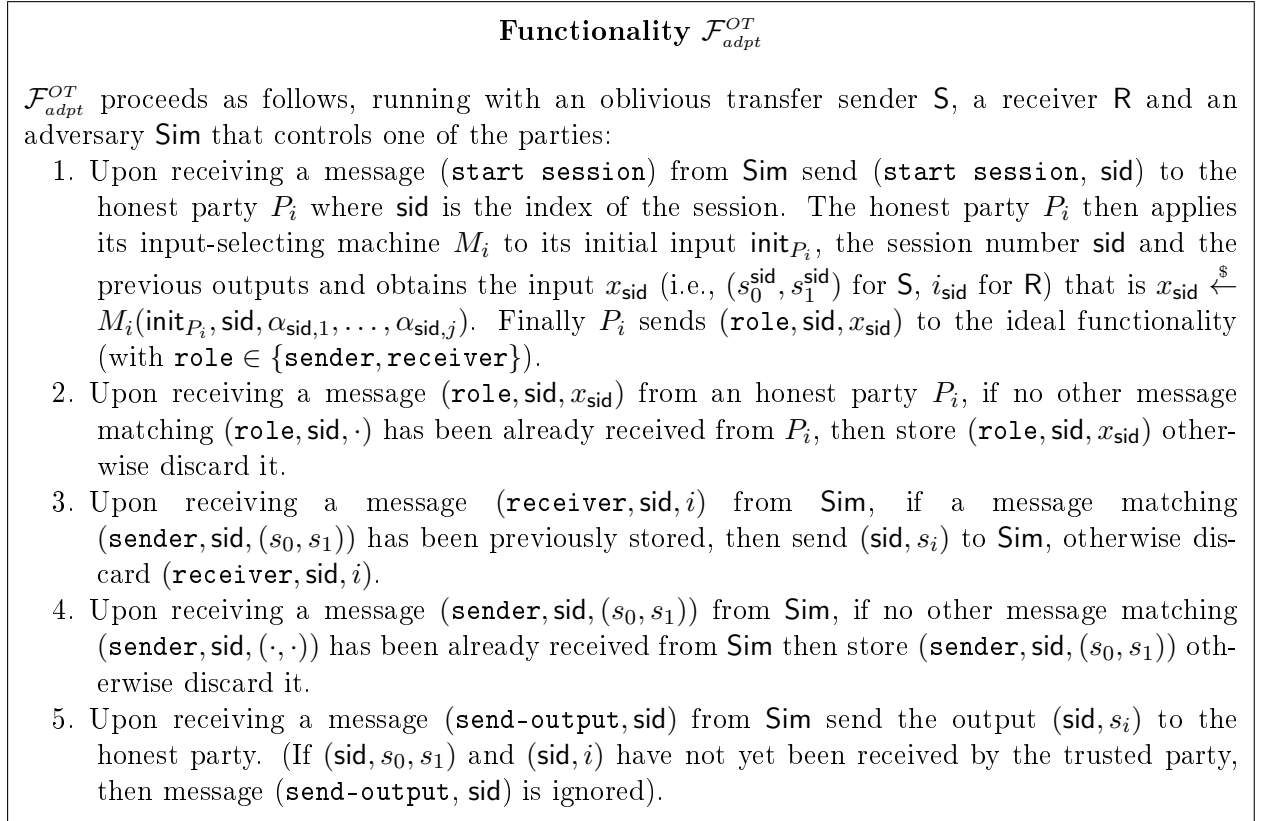


Figure 2: The functionality  $\mathcal{F}_{adpt}^{OT}$  for OT with adaptively chosen inputs.

**Ideal model execution.** The ideal functionality is shown in Fig. 2. We let **Sim** be the non-uniform PPT ideal world adversary. Then, the ideal execution of  $\mathcal{F}_{adpt}^{OT}$  (with security parameter  $n$ , input-selecting machines  $\overline{M} = (M_1, M_2)$ , initial inputs (`initS`, `initR`) and auxiliary input  $z$  to **Sim**), denoted  $\mathbf{Ideal}_{\mathcal{F}_{adpt}^{OT}, \text{Sim}, \overline{M}}(n, \text{init}_S, \text{init}_R, z)$  is the random variable defined as the output pair of the honest party and **Sim** from the above ideal execution.

**Real model execution.** In the real model execution the  $\mathcal{F}_{adpt}^{OT}$  functionality is computed by a protocol  $\pi$  in which **S** and **R** are defined as two sets of instructions  $\pi_S$  and  $\pi_R$  respectively, that are computable in polynomial time. We let  $\mathcal{A}$  be a non-uniform PPT adversary that controls either **S** or **R** and the scheduling of all messages throughout the (sequential) executions of the sessions. The (sequential) execution of  $\pi$  (with security parameter  $n$ , input-selecting machines  $\overline{M} = (M_1, M_2)$ , initial inputs (`initS`, `initR`), and auxiliary input  $z$  to  $\mathcal{A}$ ), denoted by  $\mathbf{Real}_{\pi, \mathcal{A}, \overline{M}}(n, \text{init}_S, \text{init}_R, z)$ , is the random variable defined as the output pair of the honest party and  $\mathcal{A}$  resulting from the following process.

The session `sid` is initiated by the adversary by sending a *start-session* message to the honest party. The honest party then applies its input-selecting machine on its initial input, the session number `sid` and its previously received outputs, and obtains the input for this new session. Upon the conclusion of an execution of  $\pi$ , the honest party writes its output from that execution on its output tape.

In order to obtain token re-usability we will assume that once a party sends tokens to the other party, it also keeps some state information regarding the tokens. Therefore multiple executions of a protocol that recycle the same tokens will require some state information shared among the executions. Without such a requirement token re-usability would not be possible.

**Security Definition.** The security is defined as the emulation of the real execution in the ideal model. A protocol is secure if for every real-model adversary  $\mathcal{A}$  and pair of input-selecting machines  $(M_1, M_2)$ , there exists an ideal model adversary  $\text{Sim}$  such that for all initial inputs  $\text{inits}, \text{init}_R$ , the outcome of an ideal execution with  $\text{Sim}$  is computationally indistinguishable from the outcome of a real protocol execution with  $\mathcal{A}$ . Notice that  $\text{Sim}$  knows the strategy used by the honest parties to choose their inputs. However,  $\text{Sim}$  does not know the initial input of the honest party, nor the random tape used by its input-selecting machine (any “secrets” used by the honest parties are included in the initial input, not the input-selecting machine).

**Definition 2** (Security under adaptive-input sequential composition (adapted from [Lin04])). *Let  $\mathcal{F}_{\text{adpt}}^{\text{OT}}$  and  $\pi$  be as above. Protocol  $\pi$  is said to securely compute  $\mathcal{F}_{\text{adpt}}^{\text{OT}}$  under sequential self composition if for every real-model non-uniform PPT adversary  $\mathcal{A}$  controlling  $P_i$  ( $i \in \{1, 2\}$ ) and every pair of PPT input-selecting machines  $\bar{M} = (M_1, M_2)$ , there exists an ideal-model non-uniform PPT adversary  $\text{Sim}$  controlling  $P_i$ , such that:*

$$\{\text{Ideal}_{\mathcal{F}_{\text{adpt}}^{\text{OT}}, \text{Sim}, \bar{M}}(n, \text{inits}, \text{init}_R, z)\}_{\{n \in N; \text{inits}, \text{init}_R, z \in \{0, 1\}^*\}} \stackrel{c}{\equiv} \{\text{Real}_{\pi, \mathcal{A}, \bar{M}}(n, \text{inits}, \text{init}_R, z)\}_{\{n \in N; \text{inits}, \text{init}_R, z \in \{0, 1\}^*\}}.$$

### 3 Efficient Non-Interactive 1-Out-Of-2 OT

In this section we show a protocol for efficient non-interactive OT with stateful tamper proof tokens. We first start with an high-level description of our result, and then we give the formal protocol and proof.

**High-level description of techniques and protocol.** In the protocol proposed by Kolesnikov in [Kol10] there is a first round where the sender creates and sends a token to the receiver. Then, the receiver interacts with the token and uses the obtained outputs to play two more rounds with sender. The protocol is therefore interactive which is a limitation in several applications (e.g., when the sender broadcasts messages through a satellite and the receiver can not transmit messages back to the sender), and when a protocol is used as subprotocol (e.g., for protocol composition issues and round complexity optimality).

Our main technique to overcome this limitation, therefore obtaining a non-interactive protocol, consists in asking sender  $S$  to send two tokens to the receiver  $R$ . The two tokens will be used by  $R$  to play locally (i.e., without exchanging messages with the sender) an execution of OT. Obviously, we want to achieve token re-usability and adaptive-input OT, therefore it is not possible for  $S$  to encode the usual inputs  $(s_0, s_1)$  in the token<sup>1</sup>. Instead, we will ask  $S$  to send a one-time pad encryption of  $s_0$  and  $s_1$  computed with two independent keys  $k_0, k_1$  given in output by two forward-secure PRGs. One of the two tokens, that we denote  $T_s$ , internally includes both keys. Clearly we do not want that  $T_s$  gives  $k_i$  to  $R$  as answer to a query  $i$ . Indeed, this would allow  $T_s$  to learn the bits of  $R$  and thus to adaptively change its behavior. For instance  $T_s$  could be programmed in a way such that it answers to any sequence of queries with the exception of a specific pattern 101001001, so that in the real model execution  $R$  will get answers to all queries up to the last query 1 of the sequence 101001001. Instead the simulator will never be able to replicate this attack in the ideal model execution for two main reasons: 1) discovering

<sup>1</sup>This implies that one can not use the one-time memory (OTM, for short) tokens used in [GIS<sup>+</sup>10].

such a pattern by internally running the adversary and querying and rewinding the token would require to explore all possible patterns, and this clearly require exponential time in the size of the pattern; 2) the simulator in the ideal model execution is not aware of the input of  $\mathbf{R}$  and thus can not send adaptive aborts to the ideal world functionality.

Therefore the above simple solution does not work since it does not allow one to get sequential composition. The mere use of a stateful token therefore does not seem to make trivial the problem of designing efficient non-interactive OT.

To overcome this limitation, we implement an OT execution between  $\mathsf{T}_s$  that will play the role of a sender having the two keys mentioned above as input, and  $\mathbf{R}$  that is interested in getting  $k_i$  without revealing  $i$  to  $\mathsf{T}_s$ . Here one can think that there is a circularity since we wanted to achieve efficient OT, but now we are back to the problem of achieving efficient OT. This is only partially true. Indeed, by using this technique, we have at least solved the problem with the round complexity. We have only one message sent by  $\mathbf{S}$  in the communication network, and now we need an efficient OT that is played locally and thus does not need to be non-interactive.

We implement the efficient (local) OT between  $\mathsf{T}_s$  and  $\mathbf{R}$ , using again the help of hardware tokens. This is why we require in our solution a second token that we denote as  $\mathsf{T}_k$ . The role of  $\mathsf{T}_k$  is to help for the design of a 2-message efficient OT protocol, to be played locally.

The local execution of OT goes as follow:  $\mathsf{T}_s$  playing as a sender will keep the two keys  $(k_0, k_1)$  encrypted by means of two additional keys  $\widehat{k}_0, \widehat{k}_1$ .  $\mathbf{R}$  will send a random bit  $b$  to  $\mathsf{T}_s$  that will answer with  $(\widehat{k}_0 \oplus k_0, \widehat{k}_1 \oplus k_1)$  if  $b = 0$  and with  $(\widehat{k}_0 \oplus k_1, \widehat{k}_1 \oplus k_0)$  if  $b = 1$ . Then,  $\mathbf{R}$  sends  $b \oplus i$  to  $\mathsf{T}_k$  that will answer sending  $\widehat{k}_{b \oplus i}$ . It is easy to see that both  $\mathsf{T}_s$  and  $\mathsf{T}_k$  only see random bits, but however  $\mathbf{R}$  obtains the correct key that allows it to obtain  $k_i$  first, and then  $s_i$ .

We give the intuition behind the security proof. When playing as malicious sender the view of the adversarial sender is empty since the protocol is non-interactive. Moreover, the view of each malicious token sent by the adversarial sender only consists of random bits (i.e., the queries of the token). We stress that tokens do not communicate with each other. Therefore the simulator can perfectly simulate the adversarial view and the protocol provides unconditional security against a malicious sender. When playing as malicious receiver, the adversarial view consists of the tuple containing encryptions of the two strings (obtained by the sender), encryptions of the two keys (obtained by  $\mathsf{T}_s$ ) and one decryption key (obtained by  $\mathsf{T}_k$ ) that will determine the final decrypted string. Due to the use of the one-time-pad encryption, any tuple is such that it is always possible to have four values of the view that are completely random and to set the remaining value so that one would compute the output string  $s_i$ . This property, along with the ability of observing all queries made by an adversarial receiver to the tokens (this coordination is not possible between malicious tokens and the malicious sender because they can not communicate with each other), allows the simulator to answer with randomly chosen strings and equivocate the encryptions by setting the remaining value of the view properly.

### 3.1 Non-Interactive OT

The formal protocol is depicted in Fig. 3.

**Theorem 1.** *If  $\text{GEN} = (\text{GEN.key}, \text{GEN.next}, n, t)$  is a forward-secure PRG then protocol 3.1 securely and efficiently implements  $\mathcal{F}_{\text{adpt}}^{\text{OT}}$  functionality. Furthermore the protocol is forward secure.*

*Proof.* Correctness is straightforward and relies on the properties of the  $\oplus$  operator.  $\mathbf{R}$  with input  $i$  aims at decrypting  $s_i$  from  $e_i$ . Thus it needs to get the key  $k_i$  by interrogating tokens  $\mathsf{T}_k$  and  $\mathsf{T}_s$ . Depending on the random bit  $b$  chosen for the query to  $\mathsf{T}_s$ , due to token's correctness, the desired key  $k_i$  is hidden in  $\widehat{e}'_{i \oplus b}$  through an OTP encryption with  $\widehat{k}_{i \oplus b}$ . Hence, by querying

the token  $\mathsf{T}_k$  with input  $i \oplus b$ ,  $\mathsf{R}$  retrieves the correct key  $\widehat{k_{i \oplus b}}$ . Finally, by the properties of the  $\oplus$  operator, from  $\widehat{e_{i \oplus b}}$  using  $\widehat{k_{i \oplus b}}$ ,  $\mathsf{R}$  gets  $k'_i$  and it uses  $k'_i$  to obtain  $s_i$  from  $e_i$ .

We now show the ideal world adversary  $\mathsf{Sim}$ , by considering the following cases.

**Case 1: sender is corrupted.** Let  $\mathcal{A}$ ,  $\mathcal{A}_{\mathsf{T}_s}$  and  $\mathcal{A}_{\mathsf{T}_k}$  be PPT adversaries controlling  $\mathsf{S}$ ,  $\mathsf{T}_s$  and  $\mathsf{T}_k$  respectively. We show a PPT ideal world adversary  $\mathsf{Sim}$  that simulates the honest receiver  $\mathsf{R}$  in the real world experiment to the adversary  $\mathcal{A}$  in order to extract the pair  $(s_0, s_1)$  and be able to carry out the same attack in the ideal world.  $\mathsf{Sim}$  runs  $\mathcal{A}$  internally and works as follows. It obtains tokens  $\mathcal{A}_{\mathsf{T}_s}$  and  $\mathcal{A}_{\mathsf{T}_k}$ . Upon receiving the *start-session* message from  $\mathcal{A}$ ,  $\mathsf{Sim}$  sends the message (**start session**) to  $\mathcal{F}_{adpt}^{OT}$ . This leads the ideal world receiver to get the message (**start session**, **sid**) from  $\mathcal{F}_{adpt}^{OT}$  and to activate the input-selecting machine  $M_2$  in order to obtain its next input. Upon receiving the pair  $(e_0, e_1)$  from  $\mathcal{A}$ ,  $\mathsf{Sim}$  interrogates token  $\mathcal{A}_{\mathsf{T}_s}$  with a randomly chosen bit  $b$  and obtains the pair  $(\widehat{k_0 \oplus k_b}, \widehat{k_1 \oplus k_{1-b}})$ . Let  $st$  be the state of the simulation at this point.  $\mathsf{Sim}$  then rewinds  $\mathcal{A}_{\mathsf{T}_s}$  and queries it with bit  $1 - b$  obtaining  $(\widehat{k_0 \oplus k_{1-b}}, \widehat{k_1 \oplus k_b})$ . The simulator continues the execution from  $st$  (i.e., the rewind is played as a look-ahead and no query computed during the rewind will appear in the future view of  $\mathcal{A}$ ). Note that it is crucial that the simulator chooses a random bit for the token's queries. Indeed if the simulator just queries the token with bit 0 and then bit 1 we have that the main thread of the simulation always includes 0 as query for  $\mathsf{T}_s$ , which is clearly different from what the honest receiver sends in the real world. Finally  $\mathsf{Sim}$  continues from state  $st$  (which is the main thread of the simulation), queries token  $\mathcal{A}_{\mathsf{T}_k}$  with a randomly chosen bit  $c$  obtaining  $\widehat{k_c}$ . At this point  $\mathsf{Sim}$  gets both decryption keys and it can decrypt strings  $s_0, s_1$  to play in the ideal world. Hence,  $\mathsf{Sim}$  sends the message (**sender**, **sid**,  $(s_0, s_1)$ ) followed by (**send-output**, **sid**) message to  $\mathcal{F}_{adpt}^{OT}$ . Note that the simulator can fail the decryption of both strings due to aborts in the following cases. In case  $\mathsf{T}_k$  does not abort and  $\mathsf{T}_s$  aborts then the simulator will not be able to extract at least one string to play in the ideal functionality. Thus it will play one of the following three pairs  $(s_0, 0^n)$ ,  $(0^n, s_1)$  or  $(0^n, 0^n)$ , depending on the query answered by  $\mathsf{T}_s$ . In case  $\mathsf{T}_k$  aborts then the simulator will play  $(0^n, 0^n)$  in the ideal world.

The intuition behind the security proof lies on the fact that all queries in the view of a token correspond to uniformly chosen random bit. This makes the output generated in the ideal execution distributed identically to the real execution. Thus, the security against a malicious sender is unconditional.

**Case 2: receiver is corrupted.** In this case we show a PPT ideal world adversary  $\mathsf{Sim}$  that simulates  $\mathsf{S}$  in the real world in order to extract the bit from the malicious receiver  $\mathcal{A}$  and carry out the same attack in the ideal world.  $\mathsf{Sim}$  runs  $\mathcal{A}$  internally and works as follows. Upon receiving the *start-session* message from  $\mathcal{A}$ ,  $\mathsf{Sim}$  sends the message (**start session**) to  $\mathcal{F}_{adpt}^{OT}$ . Then the ideal world sender will receive the message (**start session**, **sid**) from  $\mathcal{F}_{adpt}^{OT}$  and it will activate its input-selecting machine  $M_1$  to get the input to send to the functionality. At this point  $\mathsf{Sim}$  randomly chooses the pair  $(e_0, e_1)$  and sends it to  $\mathcal{A}$ . Then  $\mathsf{Sim}$  must answer to the queries made by the adversary to tokens  $\mathsf{T}_s$  and  $\mathsf{T}_k$ . We distinguish two cases. Case 2.1: the adversary queries  $\mathsf{T}_s$  first. In this case, as the adversary interrogates  $\mathsf{T}_s$  with input  $b$ ,  $\mathsf{Sim}$  outputs a randomly chosen pair of strings  $\widehat{e_0}, \widehat{e_1}$ . Then, when  $\mathcal{A}$  queries the token  $\mathsf{T}_k$  with input  $c$ ,  $\mathsf{Sim}$  computes  $i \leftarrow c \oplus b$  and plays (**receiver**, **sid**,  $i$ ) in the ideal functionality in order to get the string  $s_i$ . Now the simulator computes  $\widehat{k_c} \leftarrow \widehat{e_c} \oplus e_i \oplus s_i$  and returns it to  $\mathcal{A}$ . Case 2.2: the adversary queries  $\mathsf{T}_k$  first. In this case, as the adversary interrogates  $\mathsf{T}_k$  with input  $c$ ,  $\mathsf{Sim}$  outputs a randomly chosen string  $k_c$ . Then, when  $\mathcal{A}$  queries the token  $\mathsf{T}_s$  with input  $b$ ,  $\mathsf{Sim}$  computes  $i \leftarrow c \oplus b$  and plays (**receiver**, **sid**,  $i$ ) in the ideal functionality in order to get the string  $s_i$ . Then the simulator sets  $\widehat{e_c} \leftarrow e_i \oplus s_i \oplus k_c$  and  $\widehat{e_{(1-c)}}$  as a randomly chosen string,

and outputs  $(\hat{e}_0, \hat{e}_1)$ . Finally the simulator will send the message (`send-output`, `sid`) in the ideal world. Note that, the case where  $\mathcal{A}$  queries tokens many times for the same execution (i.e., without having received new pairs from the sender) is not harmful since `Sim` is only required to choose random keys to answer the queries consistently for both tokens and records the keys according to the order they were queried to make consistent the future messages to be sent. The case in which  $\mathcal{A}$  first queries one token, then gets the message from `Sim` and then queries the last token is managed following the previous approach (only the order with which `Sim` has to react to  $\mathcal{A}$ 's behavior changes). The last case regards `Sim` that gets the string from the ideal execution before actually sending the pair  $(e_0, e_1)$  to  $\mathcal{A}$ . However, this can be easily managed as `Sim` can again compute  $(e_0, e_1)$  properly, so that  $\mathcal{A}$  gets  $s_{b \oplus c}$ . Finally, when  $\mathcal{A}$  performs token corruption, `Sim` chooses random values  $K = (k_0, k_1, \text{St}_0, \text{St}_1)$ ,  $\hat{K} = (\hat{k}_0, \hat{\text{St}}_0, \hat{k}_1, \hat{\text{St}}_1)$  and returns the value according to the token that was corrupted. We finally stress that after corruption of a token, no other OT has to be simulated w.r.t. the sender that sent that token.

**The simulator runs in polynomial time.** When playing as receiver `Sim` is required to perform the rewind of a token. Since the malicious sender's attack (that involves  $\mathcal{A}, \mathcal{A}_{\tau_s}, \mathcal{A}_{\tau_k}$ ) runs in polynomial time, then interacting with a token requires polynomial time and thus the simulator interacting with the token twice will be still polynomial. Notice that a rewind does not involve to repeat work done for other sessions, and thus there is no blow up in the running time of the simulator. Now consider the case of a malicious sender that starts sequentially many executions with different receivers. Since the protocol is non-interactive (a malicious sender has no message in his view) and the extraction process of the simulator involves only rewinding of tokens (and not of the sender) each execution can be resolved independently of the others, and thus this part can be simulated in polynomial time. When playing as a sender the simulator is required to extract the bit of the receiver. The trick is that in this case the simulator has complete control on the tokens, and can see both queries before answering to the second query (this can not be done by a malicious sender). Using the properties of the  $\oplus$  operator, it can disclose at its wish the previously sent message by answering properly to the queries. Thus each execution requires only straight-line simulation, which is obviously done in polynomial time.

**Achieving adaptive-input property.** The adaptive-input property allows players to choose the next input to be provided in the protocol/functionality adaptively on the outputs obtained previously. Proving security in this setting requires a simulator able to reproduce the adversarial view without leading players of the simulated experiment to initiate more protocol executions (and thus more queries to the input-selecting machine) than the real world experiment. This could happen when the simulator's extraction strategy relies on the rewind of the malicious player. In such a case a malicious player could coordinate the executions and choose its inputs so that a simulator needs to open more ideal-world execution in order gets the inputs needed to simulate the adversarial view. This would make a deviation between real world and ideal world.

However, our simulator is not affected by such a problem, since it never rewinds the sender and the receiver, and the simulation is based on rewinding tokens in one case and in coordinating the answers to queries in another cases. Rewinding a token does not allow the adversary to open new sessions, and thus the input selecting machines do not need extra executions in the ideal world.

**Indistinguishability of the views.** We informally sketch the proof of security showing that the view of the adversary interacting with `Sim` is indistinguishable from the view of  $\mathcal{A}$  interacting with the honest players. In our setting the adversary can decide to play as a malicious sender in

all sessions or as a malicious receiver in all sessions. We will therefore show that in both cases the simulation goes through.

**$\mathcal{A}$  is a malicious sender.** In this case the view of the malicious sender is made only by queries to the tokens. The honest receiver queries the token with randomly chosen bits. The simulator deviates from the honest receiver only for the rewind of the token. However the query done after the rewind is random too for the token. Since each rewind leaves the token  $T_s$  as it had seen a randomly chosen bit, and the simulation continues from that state, the view of  $T_s$  interacting with  $\text{Sim}$  is distributed identically to the view of  $T_s$  when interacting with the honest receiver. Therefore, the views are perfectly indistinguishable, and thus our protocol achieves unconditional security against malicious senders.

**$\mathcal{A}$  is a malicious receiver.** In this case the view of the malicious receiver in one protocol execution consists of the pair of messages sent by the sender and the outputs of the tokens. Furthermore, dealing with a malicious receiver we have to prove that views remain indistinguishable even after the adversary gets the token's state performing the corruption. We notice that the simulator deviates from the honest sender in two ways: 1) it does not use **GEN** to perform the encryptions and the answers of tokens; 2) it sends the encryptions and it answers the token's queries with randomly chosen values except for one (depending on the order of the queries made to the tokens by the  $\mathcal{A}$ ) that is properly set allowing the adversary to decrypt its selected string. For lack of space, a detailed sequence of hybrid arguments, showing that the view of the adversary in the real world does not change during the ideal world, is shown in Appendix A. We now informally argue the two main steps that characterize the hybrid arguments. In the first step, starting from the real world execution we replace the use of **GEN** to obtain the keys, by selecting randomly chosen keys. Moreover, once the adversary performs the corruption of the token the simulator returns randomly chosen values. It follows from the forward-security of **GEN** that the view of  $\mathcal{A}$  interacting with the sender can not be distinguished in polynomial time from the view of  $\mathcal{A}$  interacting in this experiment. In the second step we show that the strategy used by **Sim**, answering queries and setting the last values properly is perfectly indistinguishable from the previous experiment. Indeed, this is due to the perfect security of OTP, since in both cases (i.e., the above experiment and during the simulation) all keys are used once and have uniform distribution. Summing up, we conclude that the view generated by the malicious receiver  $\mathcal{A}$  playing with the honest sender **S** is computational indistinguishable from the view generated by  $\mathcal{A}$  interacting with the simulator **Sim**.  $\square$

## 4 Acknowledgments

The work described in this paper has been supported in part by the European Commission through the ICT program under contract 216676 ECRYPT II and 215270 FRONTS, and in part by the MIUR Project PRIN "PEPPER: Privacy E Protezione di dati PERSONALI" (prot. 2008SY2PH4).

The authors wish to thank the participants of the MAYA-ECRYPT II meeting that took place in IBM Research Zurich in May 2010, for several interesting discussions on cryptographic protocols with tamper proof hardware tokens.

## References

- [BY03] Mihir Bellare and Bennet S. Yee. Forward-security in private-key cryptography. In Marc Joye, editor, *Topics in Cryptology – CT-RSA 2003*, volume 2612 of *Lecture Notes in Computer Science*, pages 1–18, San Francisco, CA, USA, April 13–17, 2003. Springer, Berlin, Germany.
- [CGS08] Nishanth Chandran, Vipul Goyal, and Amit Sahai. New constructions for UC secure computation using tamper-proof hardware. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 545–562, Istanbul, Turkey, April 13–17, 2008. Springer, Berlin, Germany.
- [CKL03] Ran Canetti, Eyal Kushilevitz, and Yehuda Lindell. On the limitations of universally composable two-party computation without set-up assumptions. In Eli Biham, editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 68–86, Warsaw, Poland, May 4–8, 2003. Springer, Berlin, Germany.
- [CLOS02] Ron Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *34th Annual ACM Symposium on Theory of Computing*, Lecture Notes in Computer Science, pages 494–503, Montréal, Québec, Canada, May 19–21, 2002. ACM Press.
- [EGL85] Shimon Even, Oded Goldreich, and Abraham Lempel. A randomized protocol for signing contracts. *Commun. ACM*, 28(6):637–647, 1985.
- [GIMS10] Vipul Goyal, Yuval Ishai, Mohammad Mahmoody, and Amit Sahai. Interactive locking, zero-knowledge pcps, and unconditional cryptography. In *Advances in Cryptology – CRYPTO 2010*, Lecture Notes in Computer Science, pages 173–190, Santa Barbara, CA, USA, August 2010. Springer, Berlin, Germany.
- [GIS<sup>+</sup>10] Vipul Goyal, Yuval Ishai, Amit Sahai, Ramarathnam Venkatesan, and Akshay Wadia. Founding cryptography on tamper-proof hardware tokens. In Daniele Micciancio, editor, *TCC 2010: 7th Theory of Cryptography Conference*, volume 5978 of *Lecture Notes in Computer Science*, pages 308–326, Zurich, Switzerland, February 9–11, 2010. Springer, Berlin, Germany.
- [GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious rams. *Journal of the ACM*, 43:431–473, 1996.
- [IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 572–591, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Berlin, Germany.
- [Kat07] Jonathan Katz. Universally composable multi-party computation using tamper-proof hardware. In Moni Naor, editor, *Advances in Cryptology – EUROCRYPT 2007*, volume 4515 of *Lecture Notes in Computer Science*, pages 115–128, Barcelona, Spain, May 20–24, 2007. Springer, Berlin, Germany.
- [Kil88] Joe Kilian. Founding cryptography on oblivious transfer. In *20th Annual ACM Symposium on Theory of Computing*, pages 20–31, Chicago, Illinois, USA, May 2–4, 1988. ACM Press.

- [Kol10] Vladimir Kolesnikov. Truly efficient string oblivious transfer using resettable tamper-proof tokens. In Daniele Micciancio, editor, *TCC 2010: 7th Theory of Cryptography Conference*, volume 5978 of *Lecture Notes in Computer Science*, pages 327–342, Zurich, Switzerland, February 9–11, 2010. Springer, Berlin, Germany.
- [Lin04] Yehuda Lindell. Lower bounds for concurrent self composition. In Moni Naor, editor, *TCC 2004: 1st Theory of Cryptography Conference*, volume 2951 of *Lecture Notes in Computer Science*, pages 203–222, Cambridge, MA, USA, February 19–21, 2004. Springer, Berlin, Germany.
- [MN05] Tal Moran and Moni Naor. Basing cryptographic protocols on tamper-evident seals. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *ICALP 2005: 32nd International Colloquium on Automata, Languages and Programming*, volume 3580 of *Lecture Notes in Computer Science*, pages 285–297, Lisbon, Portugal, July 11–15, 2005. Springer, Berlin, Germany.
- [MS08] Tal Moran and Gil Segev. David and Goliath commitments: UC computation for asymmetric parties using tamper-proof hardware. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 527–544, Istanbul, Turkey, April 13–17, 2008. Springer, Berlin, Germany.
- [Rab81] Michael O. Rabin. How to exchange secrets by oblivious transfer. Technical Report TR-81, Harvard Aiken Computation Laboratory, 1981.

## A Details on the Indistinguishability of the Views

In this section we provide a detailed proof of the indistinguishability of the views of the adversarial receiver  $\mathcal{A}$  when playing against an honest sender and when playing during the simulation (Section 3.1 includes a sketch of this proof).

Let  $m = p(n)$ , where  $p$  is a polynomial, consider an adversary performing a total of  $m$  protocol executions with possibly different senders:  $S_1, S_2, \dots, S_p$  in any order. Finally consider that  $\mathcal{A}$  ends its attack by making token corruptions. Since the security experiment with each different sender is anyway concluded after the corruption of the corresponding token then wlog we assume that  $\mathcal{A}$  ends corrupting all tokens. The proof continues by showing the indistinguishability of the following hybrids games.

$H_0$ : in this hybrid **Sim** with initial input  $\text{inits}$  will internally simulates a real world execution between  $\mathcal{A}$  and  $S$  and outputs whatever  $\mathcal{A}$  outputs. This is identical to  $\mathbf{Real}_{\pi, \mathcal{A}, \overline{M}}(n, \text{inits}, \text{init}_R, z)$ .

$H_{1,1}$ : in this hybrid **Sim** works exactly as in  $H_0$  except that, when playing as the sender  $S_1$ , it replaces the use of  $\text{gen}_0$  by giving in output randomly chosen keys when required (i.e., for computing  $(e_0, e_1)$  and for answering queries to the token  $T_s$ ). Obviously the same random value will be used to replace the  $j$ -th output of  $\text{gen}_0$  both when computing the  $j$ -th message and when answering to the  $j$ -th query of  $T_s$ .

Suppose that there exists a PPT distinguisher  $D$  distinguishing the views generated in hybrid  $H_0$  and  $H_{1,1}$  then we can construct a PPT distinguisher  $B$  for the forward-security of  $\text{gen}_0$ . We let  $q$  denote the number of protocol executions between  $\mathcal{A}$  and sender  $S_1$  ( $\mathcal{A}$  performs corruption at the end of the  $q$ -th execution obtaining the state that would have been used in the  $q + 1$  execution). The algorithm  $B$  runs the external forward-security experiment and has to guess whether it is in  $\mathbf{Exp}_{\text{GEN}}^{\text{fprg-0}}$  or  $\mathbf{Exp}_{\text{GEN}}^{\text{fprg-1}}$ . At the  $i$ -th iteration of this experiment,  $B$  obtains  $\text{Out}_i$ .  $B$  internally runs the experiment played by



honest senders and  $\mathcal{A}$  except when simulating the  $i$ -th (for  $i = 1, \dots, q$ ) protocol execution between  $\mathcal{A}$  and  $\mathsf{S}_1$ . In this case it replaces the key generated by  $\mathbf{gen}_0$  with  $Out_i$  (keeping the usual consistency, it will use  $Out_i$  also for the  $i$ -th answer of  $\mathsf{T}_s$ ) and returns  $\mathbf{find}$  to the external experiment obtaining the next output block  $Out_{i+1}$ . Then, after receiving the block  $Out_{q+1}$  in the last interaction,  $B$  returns  $\mathbf{guess}$  to the external experiment and obtain the state  $St_{q+1}$ . Finally, when adversary  $\mathcal{A}$  performs the corruption of token  $\mathsf{T}_s$ ,  $B$  answers to  $\mathcal{A}$  by sending the pair  $(St_{q+1}, Out_{q+1})$  obtained from the external experiment, along with the state of  $\mathbf{gen}_1$ . At the end of the simulation  $B$  provides the view of  $\mathcal{A}$  to  $D$  and outputs whatever  $D$  outputs. Now if the sequence of  $Out_i$  was computed by the pseudo-random generator  $\mathbf{GEN}$ , then the view generated by  $B$  is distributed identically to experiment  $H_0$ , otherwise if it was randomly chosen, the view is distributed identically to experiment  $H_{1,1}$ . By the forward security of  $\mathbf{GEN}$ ,  $H_0$  and  $H_{1,1}$  are computational indistinguishable.

$H_{1,2}, \dots, H_{1,p}$ : in the hybrid  $H_{1,j}$  (for  $j = 2, \dots, p$ ),  $\mathbf{Sim}$  works exactly as in the hybrid  $H_{1,j-1}$  except that, when playing as the honest sender  $\mathsf{S}_j$ , it replaces the use of  $\mathbf{gen}_0$  by selecting randomly chosen keys when required (keeping the usual consistency, it will use  $Out_i$  also for the  $i$ -th answer of  $\mathsf{T}_s$ ). By the same arguments as above it follows that hybrids  $H_{1,j-1}$  and  $H_{1,j}$  are computationally indistinguishable.

$H_{2,1}$ : in this hybrid,  $\mathbf{Sim}$  works exactly as in  $H_{1,p}$  except that, when playing as the sender  $\mathsf{S}_1$ , it replaces the use of  $\mathbf{gen}_1$  by outputting randomly chosen keys when required (i.e., for computing  $(e_0, e_1)$  and for answering queries to the token  $\mathsf{T}_s$ ). Obviously the same random value will be used to replace the  $j$ -th output of  $\mathbf{gen}_1$  both when computing the  $j$ -th message and when answering to the  $j$ -th query of  $\mathsf{T}_s$ . Precisely as we showed the indistinguishability of  $H_0$  and  $H_{1,1}$ , it follows that  $H_{1,p}$  and  $H_{2,1}$  are computationally indistinguishable.

$H_{2,2}, \dots, H_{2,p}$ : in the hybrid  $H_{2,j}$  (for  $j = 2, \dots, p$ ),  $\mathbf{Sim}$  works exactly as in the hybrid  $H_{2,j-1}$  except that, when playing as the honest sender  $\mathsf{S}_j$ , it replaces the use of  $\mathbf{gen}_1$  by selecting randomly chosen keys when required (keeping the usual consistency, it will use  $Out_i$  also for the  $i$ -th answer of  $\mathsf{T}_s$ ). Precisely as we showed the indistinguishability of  $H_{1,j-1}$  and  $H_{1,j}$ , it follows that  $H_{2,j-1}$  and  $H_{2,j}$  are computationally indistinguishable.

$H_{3,1}$ : in this hybrid,  $\mathbf{Sim}$  works exactly as in  $H_{2,p}$  except that, when playing as the sender  $\mathsf{S}_1$ , it replaces the use of  $\widehat{\mathbf{gen}}_0$  by selecting randomly chosen keys when required (i.e., for answering queries made to tokens  $\mathsf{T}_s$  and  $\mathsf{T}_k$ ). Precisely as we showed the indistinguishability of  $H_0$  and  $H_{1,1}$ , it follows that  $H_{2,p}$  and  $H_{3,1}$  are computationally indistinguishable.

$H_{3,2}, \dots, H_{3,p}$ : in the hybrid  $H_{3,j}$  (for  $j = 2, \dots, p$ ),  $\mathbf{Sim}$  works exactly as in the hybrid  $H_{3,j-1}$  except that, when playing as the honest sender  $\mathsf{S}_j$ , it replaces the use of  $\widehat{\mathbf{gen}}_0$  by selecting randomly chosen keys when required (keeping the usual consistency, it will use  $Out_i$  for the  $i$ -th answers of  $\mathsf{T}_s$  and  $\mathsf{T}_k$ ). Precisely as we showed the indistinguishability of  $H_{1,j-1}$  and  $H_{1,j}$ , it follows that  $H_{3,j-1}$  and  $H_{3,j}$  are computationally indistinguishable.

$H_{4,1}$ : in this hybrid,  $\mathbf{Sim}$  works exactly as in  $H_{3,p}$  except that, when playing as the sender  $\mathsf{S}_1$ , it replaces the use of  $\widehat{\mathbf{gen}}_1$  by selecting randomly chosen keys when required (i.e., for answering queries made to  $\mathsf{T}_s$  and  $\mathsf{T}_k$ ). Precisely as we showed the indistinguishability of  $H_0$  and  $H_{1,1}$ , it follows that  $H_{3,p}$  and  $H_{4,1}$  are computationally indistinguishable.

$H_{4,2}, \dots, H_{4,p}$ : in the hybrid  $H_{4,j}$  (for  $j = 2, \dots, p$ ),  $\mathbf{Sim}$  works exactly as in the hybrid  $H_{4,j-1}$  except that, when playing as the honest sender  $\mathsf{S}_j$ , it replaces the use of  $\widehat{\mathbf{gen}}_1$  by selecting randomly chosen keys when required. Precisely as we showed the indistinguishability of  $H_{1,j-1}$  and  $H_{1,j}$ , it follows that  $H_{4,j-1}$  and  $H_{4,j}$  are computationally indistinguishable.

$H_{5,1}$ : in this hybrid, **Sim** works exactly as in hybrid  $H_{4,p}$  except that in the first session **Sim** plays  $(e_0, e_1)$  and answers to the tokens queries sending  $\widehat{k}_c$  and  $(\widehat{e}_0, \widehat{e}_1)$  according to the description of the simulator (i.e., all values are random except the last one that is properly computed so that the adversary obtains  $s_{b \oplus c}$  where  $b$  and  $c$  are the bits of the two queries). We stress that in  $H_{4,p}$  all keys used for such a session are random, and **Sim** sends to the adversary the following messages:

1.  $e_0 = s_0 \oplus k_0$ , where  $k_0$  is a random key, as first part of the non-interactive message;
2.  $e_1 = s_1 \oplus k_1$ , where  $k_1$  is a random key, as second part of the non-interactive message;
3.  $\widehat{e}_0 = \widehat{k}_0 \oplus k_b$ , where  $\widehat{k}_0$  is a random key, as first part of the answer of  $T_s$  to query  $b$ ;
4.  $\widehat{e}_1 = \widehat{k}_1 \oplus k_{1-b}$ , where  $\widehat{k}_1$  is a random key, as second part of the answer of  $T_s$  to query  $b$ ;
5.  $\widehat{k}_c$  as answer of  $T_k$  to query  $c$ .

Notice that the adversary does not obtain neither  $k_{1-(b \oplus c)}$  nor  $\widehat{k}_{(1-c)}$ . From the above equations this means that in the view of the adversary  $e_{1-(b \oplus c)}$  and  $\widehat{e}_{1-c}$  are random strings. The remaining values obtained by the adversary are  $e_{b \oplus c}$ ,  $\widehat{e}_c$  and  $k_c$  where the last two strings are random and the first string corresponds to  $s_{b \oplus c} \oplus \widehat{e}_c \oplus k_c$ .

In  $H_{5,1}$ , according to the description of the simulator, depending on  $\mathcal{A}$ 's behavior, we should distinguish the following cases.

- Case 1:** first **Sim** sends  $(e_0, e_1)$ , then it plays as  $T_s$  and finally it plays as  $T_k$ .
- Case 2:** first **Sim** sends  $(e_0, e_1)$ , then it plays as  $T_k$  and finally it plays as  $T_s$ .
- Case 3:** first **Sim** plays as  $T_k$ , then it sends  $(e_0, e_1)$  and finally it plays as  $T_s$ .
- Case 4:** first **Sim** plays as  $T_k$ , then it plays as  $T_s$  and finally it sends  $(e_0, e_1)$ .
- Case 5:** first **Sim** plays as  $T_s$ , then it sends  $(e_0, e_1)$  and finally it plays as  $T_k$ .
- Case 6:** first **Sim** plays as  $T_s$ , then it plays as  $T_k$  and finally it sends  $(e_0, e_1)$ .

We now show that the output of  $H_{5,1}$  in **Case 1** is perfectly indistinguishable from the one of  $H_{4,p}$ , the other cases follow the same approach and trivially fit the equations.

So we will assume that first **Sim** sends  $(e_0, e_1)$ , then it plays as  $T_s$  and finally it plays as  $T_k$ . Notice that  $(e_0, e_1)$  will be played as two random strings, as well as  $\widehat{e}_0$  and  $\widehat{e}_1$  as answer to a query  $b$  of token  $T_s$ . Then the answer given by  $T_k$  to a query  $c$  will be  $s_{b \oplus c} \oplus e_{b \oplus c} \oplus \widehat{e}_c$ .

Now observe that again, as discussed previously for  $H_{4,p}$ , the adversary does not obtain neither  $k_{1-(b \oplus c)}$  nor  $\widehat{k}_{(1-c)}$ . Therefore in the view of the adversary both  $e_{1-(b \oplus c)}$  and  $\widehat{e}_{1-c}$  are random strings. Also here, it happens that the remaining values obtained by the adversary are  $e_{b \oplus c}$ ,  $\widehat{e}_c$  and  $k_c$  where the last two strings are random, indeed  $k_c$  has been computed as the  $\oplus$  operation over three strings, of which two were random. Moreover the first string corresponds to  $s_{b \oplus c} \oplus \widehat{e}_c \oplus k_c$  which is precisely what happens in  $H_{4,p}$ .

Therefore the views are identically distributed.

$H_{5,2}, \dots, H_{5,m}$ : in the hybrid  $H_{5,j}$ , (for  $j = 2, \dots, m$ ), **Sim** works exactly as in the hybrid  $H_{5,j-1}$  except that in the  $j$ -th session **Sim** computes the pair  $(e_0, e_1)$  and the answers of  $T_s$  and  $T_k$  as described above. By the same argument as before hybrids  $H_{5,j}$   $H_{5,j+1}$  are identically distributed.

Since  $H_{5,m}$  corresponds to the simulated game, the claim holds.

**Non-Interactive Initialization Phase:**

S prepares and sends tokens to R as follows:

1. compute  $K = \{(k_0, \text{St}_0) \leftarrow \text{gen}_0.\text{next}(\text{GEN.key}(1^n)); (k_1, \text{St}_1) \leftarrow \text{gen}_1.\text{next}(\text{GEN.key}(1^n))\}$ ;
2. compute  $\widehat{K} = \{(\widehat{k}_0, \widehat{\text{St}}_0) \leftarrow \widehat{\text{gen}}_0.\text{next}(\text{GEN.key}(1^n)); (\widehat{k}_1, \widehat{\text{St}}_1) \leftarrow \widehat{\text{gen}}_1.\text{next}(\text{GEN.key}(1^n))\}$ ;
3. create token  $\mathsf{T}_s$  seeded with the pair  $(\widehat{K}, K)$  and token  $\mathsf{T}_k$  seeded with  $\widehat{K}$  such that:
  - $\mathsf{T}_k$  on input a bit  $c$ :
    1. output  $\widehat{k}_c$ ;
    2. update keys:  $(\widehat{k}_0, \widehat{\text{St}}_0) \leftarrow \widehat{\text{gen}}_0.\text{next}(\widehat{\text{St}}_0)$ ;  $(\widehat{k}_1, \widehat{\text{St}}_1) \leftarrow \widehat{\text{gen}}_1.\text{next}(\widehat{\text{St}}_1)$ .
  - $\mathsf{T}_s$  on input a bit  $b$ :
    1. output  $\widehat{e}_0 \leftarrow \widehat{k}_0 \oplus k_b$ ;  $\widehat{e}_1 \leftarrow \widehat{k}_1 \oplus k_{(1-b)}$ ;
    2. update keys:
$$(\widehat{k}_0, \widehat{\text{St}}_0) \leftarrow \widehat{\text{gen}}_0.\text{next}(\widehat{\text{St}}_0); (\widehat{k}_1, \widehat{\text{St}}_1) \leftarrow \widehat{\text{gen}}_1.\text{next}(\widehat{\text{St}}_1);$$

$$(k_0, \text{St}_0) \leftarrow \text{gen}_0.\text{next}(\text{St}_0); (k_1, \text{St}_1) \leftarrow \text{gen}_1.\text{next}(\text{St}_1);$$
4. send tokens  $\mathsf{T}_s$  and  $\mathsf{T}_k$  to R.

**Non-Interactive Oblivious Transfer:**

S on input  $(s_0, s_1)$  works as follows:

1. compute  $e_0, e_1$  as:  $e_0 \leftarrow k_0 \oplus s_0$ ;  $e_1 \leftarrow k_1 \oplus s_1$ ;
2. update keys:  $(k_0, \text{St}_0) \leftarrow \text{gen}_0.\text{next}(\text{St}_0)$ ;  $(k_1, \text{St}_1) \leftarrow \text{gen}_1.\text{next}(\text{St}_1)$ ;
3. send  $\{e_0, e_1\}$  to R.

R on input  $i$ , upon receiving  $\{e_0, e_1\}$  works as follows:

1. choose a random bit  $b$  and computes  $c \leftarrow b \oplus i$ ;
2. query  $\mathsf{T}_k$  with input  $c$  and obtains  $\widehat{k}'_c$ , in case of abort output  $0^n$ ;
3. query  $\mathsf{T}_s$  with input  $b$  and obtains  $\widehat{e}'_0, \widehat{e}'_1$ , in case of abort output  $0^n$ ;
4. compute  $k'_i \leftarrow \widehat{k}'_c \oplus \widehat{e}'_c$ ;
5. compute  $s'_i \leftarrow k'_i \oplus e_i$  and return  $s'_i$ .

Figure 3: Efficient Non-Interactive String OT with 2 Tokens.