

Secure Multiparty Computation from SGX^{*}

Raad Bahmani¹, Manuel Barbosa², Ferdinand Brasser¹, Bernardo Portela², Ahmad-Reza Sadeghi¹, Guillaume Scerri³, and Bogdan Warinschi⁴

¹ Technische Universität Darmstadt

² HASLab – INESC TEC & DCC-FCUP

³ Université de Versailles St-Quentin – INRIA

⁴ University of Bristol

Abstract. Isolated Execution Environments (IEE) offered by novel commodity hardware such as Intel’s SGX deployed in Skylake processors permit executing software in a protected environment that shields it from a malicious operating system; it also permits a remote user to obtain strong interactive attestation guarantees on both the code running in an IEE and its input/output behaviour. In this paper we show how IEEs provide a new path to constructing general secure multiparty computation (MPC) protocols. Our protocol is intuitive and elegant: it uses code within an IEE to play the role of a trusted third party (TTP), and the attestation guarantees of SGX to bootstrap secure communications between participants and the TTP. In our protocol the load of communications and computations on participants only depends on the size of each party’s inputs and outputs and is thus small and independent from the intricacy of the functionality to be computed. The remaining computational load— essentially that of computing the functionality — is moved to an untrusted party running an IEE-enabled machine, an appealing feature for Cloud-based scenarios. However, as often the case even with the simplest cryptographic protocols, we found that there is a large gap between this intuitively appealing solution and a protocol with rigorous security guarantees. We bridge this gap through a comprehensive set of results that include: i. a detailed construction of a protocol for secure computation for arbitrary functionalities; ii. formal security definitions for the security of the overall protocol and that of its components; and iii. a modular security analysis of our protocol that relies on a novel notion of labeled attested computation. We implemented and extensively evaluated our solution on SGX-enabled hardware, providing detailed measurements of our protocol as well as comparisons with software-only MPC solutions. Furthermore, we show the cost induced by using constant-time, i.e., timing side channel resilient, code in our implementation.

1 Introduction

Secure multiparty computation (MPC) allows a set of parties to collaboratively execute a distributed computation using a cryptographic protocol, with the same security and privacy guarantees that would result from relying on a Trusted Third Party (TTP) to compute the same functionality in an ideally secure setting. Functionalities can be very simple, for example allowing one party to commit to a secret value and later on reveal it; or they can be highly complex, for example an electronic voting system. Functionalities may also be reactive (i.e., keep a state and respond to external stimuli) or they can be stateless, as in the case of the one-shot evaluation of a function (this scenario is usually called secure function evaluation). In any case, participants expect that their inputs remain as secret as in the computation of the functionality computed by the TTP.

In a seminal paper, Katz [28] shows how to bootstrap (universally composable) MPC in a setting where users have access to tamper-proof tokens on which they can load arbitrary code. The required guarantee is that anyone in possession of the token learns nothing beyond the input/output behaviour of the embedded code. For this setting, Katz shows that one can implement universally composable commitments [13], and hence arbitrary multiparty computation [14]. This result is theoretically important as it makes no additional setup assumptions, yet from a practical perspective the solution relies on unavailable and unfeasible to deploy hardware.

In this paper we study general MPC from novel trusted hardware that is currently shipped on commodity PCs: Intel’s Software Guard Extensions [27]. The main security capability that such hardware offers are Isolated Execution Environments (IEE) – a powerful tool for boosting trust in remote systems under the total or partial control of malicious parties (hijacked boot, corrupt OS, running malicious software, or simply a dishonest service provider). Specifically, code loaded in an IEE is executed in isolation from other software present in the system,⁵

^{*} This work was supported by the European Union’s 7th Framework Program (FP7/2007-2013) under grant agreement n. 609611 (PRACTICE).

⁵ We discuss the issue of side-channels that may disrupt the isolation barrier at the end of the paper.

and built-in cryptographic attestation mechanisms guarantee the integrity of the code and its I/O behaviour to a remote user.

The functionality outlined above suggests a simple, natural, and effective design for *general multiparty computation*: load the functionality to be computed into an IEE (which plays the role of a trusted third party) and have users provide inputs and receive outputs via secure channels to the IEE. Attestation ensures that the function and its inputs have not been tampered with and that the users receive, untampered, the outputs of the computation. The resulting protocol is extremely efficient when compared to existing solutions that cannot rely on such hardware assumptions. Indeed, the load of communications and computations on protocol participants is very small and independent of the intricacy of the functionality that is being computed; it depends only on the size of each party’s inputs and outputs. The remaining computational load — essentially that of computing the functionality expressed as a transition function in a standard programming language — is moved to an untrusted party running an IEE-enabled machine. This makes the protocol very appealing for Cloud scenarios.

The appealing intuition obscures the significant gap between this simple idea and a cryptographic protocol with rigorous security guarantees. There are multiple obstacles to overcome, including: i. the lack of private channels between the users and the remote machine; ii. the need to authenticate/agree on a computation in a setting where communication between parties is inherently asynchronous and only mediated by the IEE; iii. the need to ensure that the “right” parties are engaged in the computation; iv. dealing with the interaction between different parts of the code that coexist within the same IEE, sharing the same memory space, each potentially corresponding to different users; and v. ensuring that the code running inside an IEE does not leak sensitive information to untrusted code running outside.

In this paper we bridge this gap through a comprehensive set of results that include: i. a detailed construction of a protocol for MPC computation for arbitrary functionality; ii. formal security definitions for the security of the overall protocol and that of its components; iii. a modular security analysis of our protocol that relies on a novel notion of labelled attested computation; and iv. an open-source implementation of our protocol and a detailed experimental analysis in SGX-enabled hardware.

1.1 Our results

GENERIC SECURE MULTIPARTY COMPUTATION. The main result of the paper is a highly efficient protocol for the secure multiparty computation of an arbitrary reactive functionality F . The protocol follows the simple design outlined earlier. We load into an IEE code that first executes (in parallel) key-exchanges with the intended users; the code then executes whatever (reactive) functionality F one would like to implement. In this latter stage, the input/output communication is over secure channels initialized with the keys exchanged earlier. The communication and computational load for a user comprises a standard key exchange, a constant number of public key signature verifications, and the secure transfer of inputs/outputs to/from the functionality using authenticated encryption. This makes the client-side of the protocol suitable for deployment in (modest) mobile devices. Importantly, the computational overhead on the *server side* is also very small. Finally, our protocol is non-interactive in the sense that each user can perform an initial reusable set-up, and then provide its inputs and receive outputs independently of other protocol participants, which means that it provides a solution for “secure computation on the web” [24] with standard MPC security.

We provide a rigorous analysis of security for our protocol in a simulation-based security model where the adversary controls an IEE-enabled remote machine. We model this machine as an incorruptible party, which can be seen as a new hardware assumption. Our execution model allows the adversary to control all the communications between protocol participants and this remote machine, but it does not allow the adversary to corrupt the machine itself. The security assumptions on this type of hardware differ from those considered in [28]: the honest parties never access trusted hardware directly. Access to hardware is always through the attacker, who can manipulate the code submitted by an honest participant. The downside of our approach, in addition to the obvious requirement that one trusts the hardware manufacturer, is that we still require a PKI to authenticate public parameters (namely those of the SGX machine) and we focus on the static corruption model. Our choices also have implications in the composability of our results and, indeed, recent independent work [36] sheds some light on this issue. We expand on this in the related work section.

LABELLED ATTESTED COMPUTATION. Our protocol relies on ideal functionalities viewed as programs written as transition functions in a programming language compatible with the IEE-enabled machine. We instrument these programs to run inside an IEE and add bootstrapping code that permits protocol participants to establish independent secure channels with the functionality so that they can send inputs and receive outputs from it. The crux of the protocol is to ensure the attestation guarantees which convince the parties that they are involved in the “right” run

of the protocol (i.e. with the right parties all interacting with the same IEE). The required technical component is inspired by the recent work of Barbosa et al.[2] who show how to use IEEs to securely outsource computation. Part of their work develops the concept of Attested Computation (AC) as an abstraction of the basic integrity assurances specific to IEEs. As far as our protocol is concerned, their definition has two important limitations. First, the local user needs to know all of the input/output that occurred at the IEE interface and, secondly, attestation only works for a (entirely known) prefix of the computation. Both of these restrictions make the AC notion of [2] inapplicable to a setting where multiple parties are involved in the communication with the IEE and where, for efficiency, it is critical that each party is able to asynchronously provide input and retrieve outputs from a functionality running within.

To overcome these problems, we propose a novel notion of *labelled attested computation* (LAC). In LAC, attested outputs provide integrity guarantees on well-defined sequences of computational steps, even when these are interleaved with other (potentially) unknown computations. Furthermore, the integrity of these outputs is publicly verifiable. We achieve this level of fine-grained control by appending labels to inputs, and binding together the inputs and outputs that correspond to a particular label. To understand the power of this generalized version of attested computation, consider the problem of n parties establishing n independent secure channels with an IEE. Then, the IEE should be running code of the form $(K_1 | K_2 | \dots | K_n); Q$ where each K_i is a key exchange protocol that corresponds to a different user, and Q is some code that will use the derived keys to communicate with each party. Using LAC, one assigns a different label to the key exchange of each user (say the identity of the user that should be running that key exchange); LAC then guarantees, for each user, that the steps of its key exchange (and only those) have been executed properly. Attestation only requires the user’s own communication with the IEE and is independent of the other exchanges. Furthermore, since the attestation guarantees provided by LAC bind an attested trace to the code running within the IEE, a local user can be sure that the other key exchanges are being executed with specific parties: the code running inside the IEE will contain public cryptographic material that enables the IEE to authenticate all the local parties and hence authenticating the code implicitly authenticates all the parties participating in the protocol.

We provide syntax and a formal security model for LAC and show how this primitive can be used to deploy arbitrary (labelled) programs to remote IEEs with flexible attestation guarantees. Our provably secure LAC protocol relies on hardware equipped with SGX-like IEEs. Our results allow for compositional proofs and, for that, we formalize the interplay between LAC and program composition. In particular, we formalize what it means for a key exchange program to be composed with other arbitrary code and outsourced with desirable attestation guarantees. Indeed, we show that by enforcing syntactic restrictions on the composition pattern and a *minimal leakage* security property on the LAC scheme,⁶ one can use the specific flavour of key exchange for attested computation introduced in [2] to bootstrap an arbitrary number of independent secure channels between local users and an IEE; we call this a *key exchange utility theorem*. This is an extension of the utility theorem in [2], where we leverage the additional power of LAC to tackle a (labelled) parallel/sequential program composition pattern.

Our solution combines our SGX-specific authenticated key exchange, the LAC scheme, and hardwiring participant’s keys in the isolated code. Analysing this solution requires subtle reasoning about attested traces of specific sub-components of the isolated program, provided independently to different parties; this type of reasoning is beyond the results in [3], which only considers the attestation of a prefix of the full program trace. LAC is the key stepping stone for dealing with fine-grained attestation guarantees of arbitrarily composed programs, which not only solves this particular problem, but also paves the way for the modular analysis of other protocols relying on IEEs.

We also expect LAC to find applications beyond the specific one that we consider in this paper. For example, an application which should be a direct application of LACs is the implementation of a secure bulletin board. Consider an IEE which receives inputs (which can optionally be authenticated) and appends them to an internally maintained bulletin board. Upon request of an arbitrary party, the IEE returns the current content of the IEE. The desired security is obtained by attesting (only) this communication step.

IMPLEMENTATION AND EXPERIMENTAL VALIDATION. To obtain an accurate estimate of the performance of our protocol in the real world, we have constructed two implementations—`sgx-mpc-mbed` and `sgx-mpc-nacl`— which are open-source and will be made publicly available. The core difference between the two implementations is the underlying cryptography: `sgx-mpc-mbed` relies on the `mbed TLS`⁷ (formerly PolarSSL) library, whereas `sgx-mpc-nacl`

⁶ Minimal leakage ensures that the outsourced instrumented program P^* reveals no information about its internal state beyond what the normal input/output behavior of the original program P would reveal.

⁷ <https://tls.mbed.org/>

relies on the NaCl⁸ cryptographic library [8]. They differ both on the underlying cryptographic technology and on the level of protection against timing side-channel attacks:

- `sgx-mpc-mbed` inherits the standard trust model underlying TLS implementations such as mbed TLS, OpenSSL or AWS’s S2N library, which places the adversary on the network.⁹ This means that the end-point machine is assumed to be trusted, so defending against timing attacks launched from within the machine in which an implementation is running is usually seen as an overkill. The standard approach is to eliminate *significant* variations in execution time that might be observable from outside the machine. Furthermore, `sgx-mpc-mbed` uses standard RSA technology for the key exchange stage, and an AES128-CTR and HMAC-SHA256 Encrypt-Then-Mac construction for authenticated encryption.
- `sgx-mpc-nacl` inherits the more ambitious trust model underlying the NaCl library in what concerns timing attacks. The entire implementation follows a stricter side-channel countermeasure coding policy called *constant-time*—this excludes *all* control-flow or memory access patterns that depends on secret data. This coding policy is generally accepted as the best software-based countermeasure against timing attacks that one can adopt, eliminating attack vectors that may arise, not only from direct measurements of execution time, but also from indirect ones relying for example on cache and page fault correlations. On the downside, it typically implies a performance penalty that may be minimized by aggressive optimizations at the assembly level. `sgx-mpc-nacl` uses elliptic-curve based technology for both key exchange (Diffie-Hellman) and digital signatures, and a combination of the Salsa20 and Poly1305 encryption and authentication schemes [8] for authenticated encryption.

We conclude the paper with an experimental evaluation of our protocol based on these implementations. In particular, we include a detailed comparison of our solution to state-of-the-art multiparty computation protocols using a series of benchmarks commonly used in the literature. The experimental results confirm the theoretical performance advantages that we have highlighted above in comparison to non hardware-based solutions. We also include a detailed discussion of the issue of side-channels in SGX-like systems and an assessment of the impact of adopting stricter coding policies aiming to thwart timing attacks.

STRUCTURE OF THE PAPER. The document structure is as follows. Section 2 presents preliminary definitions. Section 3 formalizes notions of IEEs, Programs and Machines. Sections 4 and 5 describe our primitive for IEE-enabled machines LAC and propose a construction. Section 6 and 7 detail a model for secure computation using LAC and propose a protocol. Section 8 is dedicated to implementation and evaluation analysis of the protocol. Section 9 concludes the document and suggests future research directions.

1.2 Related work

The closest works to ours are that of Katz [28] who considers the use of specific trusted hardware to bootstrap multiparty computation and that by Barbosa et al. [2] who are specifically concerned with employing SGX-like capabilities for securely outsourcing arbitrary computation. We next review other related literature.

ATTESTATION. In recent independent work Pass, Shi and Tramer [36] formalize attestation guarantees offered by trusted hardware in the Universal Composability setting, and consider the feasibility of achieving UC-secure MPC starting from a hardware assumption formalized in very much the same style as we model SGX. Interestingly, they show that in the setting that they consider (Universal Composability with a Global Setup (GUC) [12]) multiparty computation is impossible to achieve without additional assumptions, unless *all* parties have access to trusted hardware. They also provide UC secure protocols that solve the problem by both providing access to SGX to all parties or introducing additional set-up assumptions. The resulting protocols are more intricate and less efficient than the one we propose here. The work in this paper can therefore be seen as providing a practice-oriented description and security proof of the most natural MPC protocol relying on SGX, where we forsake composability for efficiency, while still preserving strong privacy guarantees for the inputs to the computation. Furthermore, as was mentioned in [36], composability may be achieved by introducing an independent set-up of the trusted hardware for each protocol that is executed. This may be problematic for one-shot functionalities, but here we deal with arbitrarily complex reactive functionalities that may execute over a long period of time.

A relevant line of research leverages trusted hardware to bootstrap entire platforms for secure software execution (e.g. Flicker [33], Trusted Virtual Domains [15], Haven [4]). These are large systems that are currently outside the scope of provable-security techniques. Smaller protocols which solve specific problems (secure disk encryption [34], one-time password authentication [26] outsourced Map-Reduce computations [37], Secure Virtual Disk Images [21],

⁸ <https://nacl.cr.yp.to>

⁹ A good description of this model is available on the web site of the recent CacheBleed attack (<https://ssrg.nicta.com.au/projects/TS/cachebleed/>).

two-party computation [23], secure embedded devices [35,30]) are more susceptible to rigorous analysis. Although some protocols (e.g., those of Hoekstra et al. [26]) come only with intuition regarding their security, others—most notably those by Schuster et. al [37] which uses SGX platforms to outsource map-reduce computation—come with a proof of security. The use of attestation in those protocols is akin to our use of attestation in our general MPC protocol. Provable security of realistic protocols that use trusted hardware-based protocols has considered protocols based on the Trusted Platform Module (TPM) [11,38,10,20,19]. The weaker capabilities offered by the TPM makes them more suitable for static attestation than for a dynamic setting like the one we consider in this paper.

MULTIPARTY COMPUTATION. The overarching goal of our work is shared with the rich literature on software-only multiparty secure computation. We mention here the works that are close to ours in the sense that they aim to bring secure multiparty computation to practice.

Fairplay is a system originally developed to support two-party computation [32] and then extended to FairplayMP to support multiparty computation [6]: Fairplay implements a two party computation protocol in the manner suggested by Yao; FairplayMP is based on the Beaver-Micali-Rogaway protocol. Sharemind [9] is a secure service platform for data collection and analysis, employing a 3-party additive secret sharing scheme and provably secure protocols in the honest-but-curious security model with no more than one passively corrupted party. TASTY (Tool for Automating Secure Two-party computations) is a tool suite addressing secure two-party computation in the semi-honest model [25] whose main feature allows to compile and evaluate functions not only using garbled circuits, but also homomorphic encryption schemes, at the same time. SPDZ [17] is a protocol for general multi party computations considering active adversaries and tolerating the corruption of $n - 1$ out of the n parties, leveraging a pre-processing stage for exchanging randomness between participants, towards reducing communication requirements associated with the on-line stage.

The main advantage of our solution with respect to the previous systems is its efficiency both for the parties that provide inputs and collect outputs from the computation, and those that perform the computation. In all of the above solutions, the computation is distributed, and the communication load for parties performing the computation grows (with varying degrees of scalability) with the complexity of the computed functionality (often expressed as a circuit). In our solution, a single party (the owner of the IEE-enabled machine) performs the computation which is run essentially as fast as the program that computes it in the clear, so the overhead is reduced to establishing secure channels with all other participants. For these parties, the overhead is a single key exchange, and then all the inputs and outputs are transferred using standard authenticated encryption. On the downside, although our protocol is secure in the presence of active adversaries, we only consider static corruptions and rely on a strong trust assumption in idealizing the IEE-enabled machine.

The document structure is as follows. Section 2 presents preliminary definitions. Section 3 formalizes notions of IEEs, Programs and Machines. Sections 4 and 5 describe our primitive for IEE-enabled machines LAC and propose a construction. Section 6 and 7 detail a model for secure computation using LAC and propose a protocol. Section 8 is dedicated to implementation and evaluation analysis of the protocol. Section 9 concludes the document and suggests future research directions.

2 Preliminaries

2.1 Message Authentication Codes

SYNTAX. A message authentication code scheme Π is a triple of PPT algorithms $(\text{Gen}, \text{Auth}, \text{Ver})$. On input 1^λ , where λ is the security parameter, the randomized key generation algorithm returns a fresh key. On input key and message m , the deterministic MAC algorithm Auth returns a tag t . On input key, m and t , the deterministic verification algorithm Ver returns T or F indicating whether t is a valid MAC for m relative to key. We require that, for all $\lambda \in \mathbb{N}$, all key $\in [\text{Gen}(1^\lambda)]$ and all m , it is the case that $\text{Ver}(\text{key}, m, (\text{Auth}(\text{key}, m))) = \text{T}$.

SECURITY. We use the standard notion of existential unforgeability for MACs [5]. We say that Π is existentially unforgeable if $\text{Adv}_{\mathcal{A}, \Pi}^{\text{Auth}}(\lambda)$ is negligible for every ppt adversary \mathcal{A} , where advantage is defined as the probability that the game in Figure 1 (top) returns T .

2.2 Digital Signature Schemes

SYNTAX. A signature scheme Σ is a triple of PPT algorithms $(\text{Gen}, \text{Sign}, \text{Vrfy})$. On input 1^λ , where λ is the security parameter, the randomized key generation algorithm returns a fresh key pair (pk, sk) . On input secret key sk and message m , the possibly randomized signing algorithm Sign returns a signature σ . On input public key pk , m and σ ,

Game $\text{Auth}^{\mathcal{A}}(1^\lambda)$: List $\leftarrow \square$ key $\leftarrow_s \text{Gen}(1^\lambda)$ $(m, t) \leftarrow_s \mathcal{A}^{\text{Auth}}(1^\lambda)$ Return $\text{Ver}(\text{key}, m, t) = \text{T} \wedge m \notin \text{List}$	Oracle $\text{Auth}(m)$: List $\leftarrow (m : \text{List})$ $t \leftarrow \text{Mac}(\text{key}, m)$ Return t
Game $\text{UF}^{\Sigma, \mathcal{A}}(1^\lambda)$: List $\leftarrow \square$ $(pk, sk) \leftarrow_s \text{Gen}(1^\lambda)$ $(m, \sigma) \leftarrow_s \mathcal{A}^{\text{Sign}}(1^\lambda, pk)$ Return $\text{Vrfy}(pk, m, \sigma) = \text{T} \wedge m \notin \text{List}$	Oracle $\text{Sign}(m)$: List $\leftarrow (m : \text{List})$ $\sigma \leftarrow \text{Sign}(sk, m)$ Return σ

Fig. 1. Games defining the security of a MAC scheme (top) and a signature scheme (bottom).

Game $\text{IND}^{\mathcal{A}, \mathcal{A}}(1^\lambda)$: List $\leftarrow \square$ key $\leftarrow_s \text{Gen}(1^\lambda)$ $(m_0, m_1) \leftarrow_s \mathcal{A}_1^{\text{Encrypt}, \text{Decrypt}}(1^\lambda)$ $b \leftarrow_s \{0, 1\}$ $m' \leftarrow_s \text{Enc}(\text{key}, m_b)$ $b' \leftarrow_s \mathcal{A}_2^{\text{Encrypt}, \text{Decrypt}}(m_0, m_1, m')$ If $m' \in \text{List}$: $b' \leftarrow_s \{0, 1\}$ Return $b = b'$	Oracle $\text{Encrypt}(m)$: Return $\text{Enc}(\text{key}, m)$ Oracle $\text{Decrypt}(m')$: List $\leftarrow (m' : \text{List})$ $m \leftarrow \text{Dec}(\text{key}, m')$ Return m
Game $\text{UF}^{\mathcal{A}, \mathcal{A}}(1^\lambda)$: List $\leftarrow \square$ key $\leftarrow_s \text{Gen}(1^\lambda)$ $m' \leftarrow_s \mathcal{A}^{\text{Encrypt}}(1^\lambda)$ If $\text{Dec}(\text{key}, m') \neq \perp \wedge m' \notin \text{List}$: Return T Return F	Oracle $\text{Encrypt}(m)$: $m' \leftarrow \text{Sign}(sk, m)$ List $\leftarrow (m' : \text{List})$ Return m'

Fig. 2. Games defining ciphertext indistinguishability (top) and existential unforgeability (bottom) of an authenticated encryption scheme.

the deterministic verification algorithm Vrfy returns T or F indicating whether σ is a valid signature for m relative to pk . We require that, for all $\lambda \in \mathbb{N}$, all $(pk, sk) \in [\text{Gen}(1^\lambda)]$ and all m , it is the case that $\text{Vrfy}(pk, m, (\text{Sign}(sk, m))) = \text{T}$. **SECURITY.** We use the standard notion of existential unforgeability for signature schemes [22]. We say that Σ is existentially unforgeable if $\text{Adv}_{\mathcal{A}, \Sigma}^{\text{UF}}(\lambda)$ is negligible for every ppt adversary \mathcal{A} , where advantage is defined as the probability that the game in Figure 1 (bottom) returns T .

2.3 Authenticated Encryption Schemes

SYNTAX. An authenticated encryption scheme \mathcal{A} is a triple of PPT algorithms $(\text{Gen}, \text{Enc}, \text{Dec})$. On input 1^λ , where λ is the security parameter, the randomized key generation algorithm returns a fresh key. On input key key and message m , the randomized encryption algorithm Enc returns a ciphertext m' . On input key key and ciphertext m' , the deterministic decryption algorithm Dec returns the decrypted message m , or \perp if the ciphertext is found to be invalid. We require that, for all $\lambda \in \mathbb{N}$, all $key \in [\text{Gen}(1^\lambda)]$ and all m , it is the case that $m = \text{Dec}(key, \text{Enc}(key, m))$.

SECURITY. We use the standard notions of indistinguishability and existential unforgeability for authenticated encryption schemes [29]. We say that \mathcal{A} provides ciphertext indistinguishability if $\text{Adv}_{\mathcal{A}, \mathcal{A}}^{\text{IND}}(\lambda)$ is negligible for every ppt adversary \mathcal{A} , where advantage is defined as the probability that the game in Figure 2 (top) returns T over the random guess. We say that \mathcal{A} is existentially unforgeable if $\text{Adv}_{\mathcal{A}, \mathcal{A}}^{\text{UF}}(\lambda)$ is negligible for every ppt adversary \mathcal{A} , where advantage is defined as the probability that the game in Figure 2 (bottom) returns T .

2.4 Key Exchange

We recall here the notion of Key Exchange for Attested Computation (AttKE) from [2]. This notion is tailored to establish a secure key with a remote program running inside an IEE; the remote program includes a first stage where the key is derived using a key exchange subprogram called Rem_{KE} and an arbitrary second stage program that uses the derived key. Active security in the key exchange protocol is achieved by running Rem_{KE} using attestation mechanisms provided by the IEE. We will see in this paper that the exact same notion of AttKE can be used

together with Labelled Attested Computation in a much wider program composition context and, particularly, to enable multiple parties to establish independent secure channels to the same IEE.

SYNTAX. An AttKE is defined as a pair of algorithms ($\text{Setup}, \text{Loc}_{\text{KE}}$). On input the security parameter 1^λ and the local party identity id , Setup generates the part of the key exchange intended for remote execution Rem_{KE} , and the initial state st_L of the local part of the key exchange Loc_{KE} . Setup is intended to generate a fresh instance of the AttKE protocol between the party with identity id and a remote IEE. The dynamically generated Rem_{KE} will be run remotely in an IEE, under the protection of a LAC scheme, following some composition pattern. On input m and local state st_L , Loc_{KE} returns the next message intended for the remote part of the key exchange and an updated state.

We require the local and remote parts of an AttKE to keep a set of variables in their states. The execution state of the AttKE is kept as $\delta \in \{\text{derived}, \text{accept}, \text{reject}, \perp\}$. The derived key is stored as key . It is supposed to be \perp if $\delta \notin \{\text{derived}, \text{accept}\}$. The identity of the owner of the instance is represented as oid . This will be initialised on the fly for the remote part. The identity of the partner of the session is stored in variable pid . This will typically be set at generation for Rem_{KE} and constructed during execution for Loc_{KE} . Finally the session identifier sid , will typically be constructed on the fly. An AttKE is correct if, after an honest run of a local instance and the corresponding remote instance, both accept, derive the same key and agree on sid , pid and oid .

SECURITY. The adversary model for AttKE security is tailored so that active security is provided when adding attestation to the remote part. This adversary is a middle ground between a passive and an active adversary. The security model considers an adversary which has access to oracles whose behaviour depend on a bit b and a list of pairs of real and fake keys, one for each instance. In addition to the oracles initialising new local and remote instances (multiple instances of the same remote instance can be created), the following oracles are provided:

- **Reveal:** when queried on a local or remote instance, it returns the corresponding derived key.
- **Test:** when queried on a local or remote instance, it returns \perp if $\delta \neq \text{accept}$; otherwise, if $b = 0$ it returns key and it returns $\text{fake}(\text{key})$ if $b = 1$.
- **Send** allows delivering messages between instances, making sure that messages from remote instances to local instances are reliably delivered. The messages delivered to remote instances, however, are arbitrary. This oracle updates the instances according to the message delivered and returns the response, together with the corresponding pid , sid and δ .

The model keeps track of local instances Loc_{KE}^l created by the local identity id and remote instances $\text{Rem}_{\text{KE}}^{i,j}$ (there can be many copies of Rem_{KE} for each locally initialized session. A local and a remote instance are partnered if $\delta^{i,j}, \delta^l \in \{\text{derived}, \text{accept}\}$ and they agree on sid . We further restrict the adversary, by disallowing **Test** queries if **Reveal** was queried for this instance or a partnered instance.

We say that a protocol ensures valid partners if, for every partnered Loc_{KE}^l and $\text{Rem}_{\text{KE}}^{i,j}$, $\text{pid}^l = \text{oid}^{i,j}$, $\text{oid}^l = \text{pid}^{i,j}$ and the derived key is the same for both instances. We say that a protocol ensures confirmed partners if when an instance of the key exchange accepts, it has at least one partner. We say that a protocol ensure unique partners if each instance is partnered with at most one other instance. A protocol ensures two sided authentication if it ensures these three properties with overwhelming probability, in the presence of an adversary with access to the aforementioned oracles.

Definition 1 (AttKE security). *An AttKE protocol is secure if the protocol ensures two sided authentication, and for any ppt adversary \mathcal{A} , as described above, and for any local identity id , the probability of \mathcal{A} guessing b is overwhelmingly close to $1/2$.*

3 IEEs, Programs, and Machines

The models that we develop in this paper rely on the cryptographic model for IEEs introduced in [2]. Here we recall the key features of that model – for more details we refer the reader to [2]. An IEE is viewed as an idealised machine running some fixed program P and which exposes an interface through which one can pass inputs and receive outputs to/from P . The I/O behaviour of a process running in an IEE is determined by the program it is running, the semantics of the language in which the program is written, and the inputs it receives. The interface essentially guarantees that the only information that is revealed about a program running within an IEE is contained in its input-output behaviour and models the strict isolation between processes running in different IEEs (and any other program running on the machine).

PROGRAMS. The programs are assumed to be written in some programming language \mathcal{L} enriched with IEE system calls. These calls give access to different cryptographic functionalities offered by *security module* interface. The

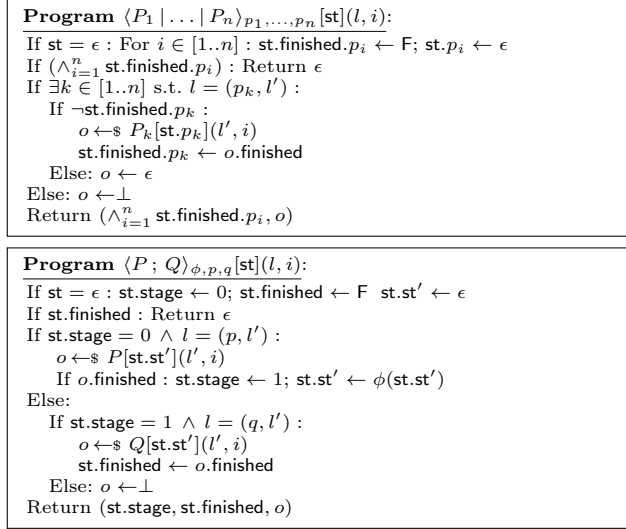


Fig. 3. Parallel (top) and sequential (bottom) program composition.

cryptography offered may differ between IEEs. The language \mathcal{L} is assumed to be deterministic modulo the operation of system calls; in particular we assume a system call `rand` which gives access to fresh random coins sampled uniformly at random. As mentioned above, it is important for our results that system calls cannot be used by a program to store additional implicit state that would escape our control. To this end, we impose that the results of system calls within an IEE can depend only on: i. an initially shared state that is defined when a program is loaded (e.g., the cryptographic parameters of the machine, and the code of the program); ii. the input explicitly passed on that particular call; and iii. fresh random coins. As a consequence of this, we may assume that system calls placed by different parts of a program are identically distributed, assuming that the same input is provided. This is particularly important when we consider program composition below.

We use the same model for programs as [2] (extended to the settings where inputs/outputs are labeled): transition functions which take a current state st and a label-input pair (l, i) , and produce a new output o and an updated state. We write $o \leftarrow P[\text{st}](l, i)$ for each such action and refer to it as an *activation*. Throughout the paper we restrict our attention to programs (even if they are adversarially created) for which the transition function is guaranteed to run in polynomial-time.¹⁰ Unless otherwise stated, st is assumed to be initially empty. We impose that every output produced by a program includes a Boolean flag `finished` that indicates whether the transition function will accept further input. We will denote by $o.\text{finished}$ the value of this flag in some output o . The transition function may return arbitrary outputs until it produces an output where `finished` = `T`, at which point it can return no further output or change its state. Some programs may not use labels internally and, in that case, we simply pass it the empty string at the label input.

We extend our notation to account for probabilistic programs that invoke the `rand` system call. We write $o \leftarrow P[\text{st}; r](l, i)$ for the activation of P which when invoked on labeled input (l, i) (with internal state st and random coins r) produced output o . We write a sequence of activations as $(o_1, \dots, o_n) \leftarrow P[\text{st}; r](l_1, i_1, \dots, l_n, i_n)$ and denote by $\text{Trace}_{P[\text{st}; r]}(l_1, i_1, \dots, l_n, i_n)$ the corresponding input/output trace $(l_1, i_1, o_1, \dots, l_n, i_n, o_n)$. When dealing with a trace T , we will use $\text{filter}[L](T)$ to denote the projection of the trace that retains only I/O pairs that correspond to labels in L . We abuse notation and use $\text{filter}[l]$ when L is a singleton.

PROGRAM COMPOSITION. We extend the basic notion of program composition in [2] to consider the two general label-based forms of program composition shown in Fig. 3 that can be applied recursively and interchangeably to create arbitrarily complex programs in a modular way.

By parallel composition of programs P_1, \dots, P_n , denoted $\langle P_1 \mid \dots \mid P_n \rangle_{p_1, \dots, p_n}$, we mean the transition function that takes inputs with extended labels of the form (p_i, l) ¹¹—here p_i are bitstrings used to identify the target program, where we assume $p_i \neq p_j$ for i, j distinct—and dispatches incoming label-input pairs to the appropriate program. In parallel composition we exclude the possibility of state sharing between programs, and define termination to occur when all composed programs have terminated. By sequential composition of two programs P and Q via

¹⁰ In particular we assume that adversarially generated programs cannot *blow up* the execution time of an experiment beyond poly-time in the security parameter.

¹¹ We assume some form of non-ambiguous encoding of composed labels and output strings, but in our presentation we simply present these encoded values as tuples.

projection function ϕ , denoted $\langle P; Q \rangle_{\phi, p, q}$, we mean the transition function that has two execution stages, which are signaled in its output via an additional **stage** flag. As above, we will denote by $o.\text{stage}$ the value of this flag in some output o . For consistency, we again assume labels of the form (p, l) and (q, l) where $p \neq q$ are used to identify the target program. In the first stage, every label-input pair will be checked for consistency (i.e., that it indicates P as the target program) and dispatched to program P . This will proceed until P 's last output indicates it has finished (inclusively). The next activation will trigger the start of the second stage, at which point the composed program initialises the state of Q using $\phi(\text{st}_P)$ before activating it for the first time. Additionally we require that a constant indicating the current stage is appended to any output of a composition. We do not admit other state sharing between P and Q in addition to that fixed by ϕ .

MACHINES. As in [2] we the basis of our model for IEEs is a *machine* \mathcal{M} which we view as an abstract computational device that captures the resources offered by a real world computer or group of computers. These machines contain hardware security functionalities which are initialised by a specific manufacturer before being deployed, possibly for different end-users. For example, a machine may represent a single computer produced by a manufacturer, configured with a secret signing key for a public key signature scheme, and whose public key is authenticated via some public key infrastructure, possibly managed by the manufacturer itself. Similarly, a machine may represent a group of computers, each configured with secret signing keys associated with a group signature scheme; again, the public parameters for the group would then be authenticated by some appropriate infrastructure.¹² Inspired by the functionality offered by SGX, in this paper we consider the case where standard public key signatures are used; our models extend to more complex group management schemes.

We model machines via a simple external interface, which we see as both the functionality that higher-level cryptographic schemes can rely on when using the machine, and the adversarial interface that will be the basis of our attack models. This interface can be thought of as an abstraction of Intel's SGX [27]. The interface is as follows:

- $\text{Init}(1^\lambda)$ is the global initialisation procedure which, on input the security parameter, outputs the global parameters prms . This algorithm represents the machine's hardware initialisation procedure, which is out of the user's and the adversary's control. Intuitively, it initialises the internal security module, the internal state of the remote machine and returns any public cryptographic parameters that the security module releases. The global parameters of machines are assumed to be authenticated using external mechanisms, such as a PKI.
- $\text{Load}(P)$ is the IEE initialisation procedure. On input a program/transition function P , the machine produces a fresh handle hdl , creates a new IEE with handle hdl , loads P into the new IEE and returns hdl . The machine interface does not provide direct access to either the internal state of an IEE nor to its randomness input. This means that the only information that is leaked about internal state and randomness input is that revealed (indirectly) via the outputs of the program.
- $\text{Run}(\text{hdl}, l, i)$ is the process activation procedure. On input a handle hdl and a label-input pair (l, i) , it will activate process running in isolated execution environment of handle hdl with (l, i) as the next input. When the program/transition function produces the next output o , this is returned to the caller.

We define the I/O trace $\text{Trace}_{\mathcal{M}}(\text{hdl})$ of a process hdl running in some machine \mathcal{M} as the tuple $(l_1, i_1, o_1, \dots, l_n, i_n, o_n)$ that includes the entire sequence of n inputs/outputs resulting from all invocations of Run on hdl ; $\text{Program}_{\mathcal{M}}(\text{hdl})$ is the code (program) running inside the process with handle hdl ; $\text{Coins}_{\mathcal{M}}(\text{hdl})$ represents the coins given to the program by the rand system call; and $\text{State}_{\mathcal{M}}(\text{hdl})$ is the internal state of the program. Finally, we will denote by $\mathcal{A}^{\mathcal{M}}$ the interaction of some algorithm with a machine \mathcal{M} , i.e., having access to the Load and Run oracles defined above.

4 Labelled Attested Computation

We now formalize a cryptographic primitive that generalizes the notion of Attested Computation proposed in [2], called Labelled Attested Computation. The main difference to the original proposal is that, rather than fixing a particular form of program composition for attestation, Labelled Attested Computation is agnostic of the program's internal structure; on the other hand, it permits controlling data flows and attestation guarantees via the label information included in program inputs.

SYNTAX. A *Labelled Attested Computation* (LAC) scheme is defined by the following algorithms:

- $\text{Compile}(\text{prms}, P, L^*)$ is the deterministic program compilation algorithm. On input global parameters for some machine \mathcal{M} , program P and an attested label set L^* , it will output program P^* . This algorithm is run locally.

¹² If the possibility of removing elements from the group is not needed, then even sharing the same signing key for a public key encryption scheme between multiple computers could be a possibility.

P^* is the code to be run as an isolated process in the remote machine, whereas L^* defines which labelled inputs should be subject to attestation guarantees.

- $\text{Attest}(\text{prms}, \text{hdl}, l, i)$ is the stateless attestation algorithm. On input global parameters for \mathcal{M} , a process handle hdl and label-input pair (l, i) , it will use the interface of \mathcal{M} to obtain attested output o^* . This algorithm is run remotely, but in an unprotected environment: it is responsible for interacting with the isolated process running P^* , providing it with inputs and recovering the attested outputs that should be returned to the local machine.
- $\text{Verify}(\text{prms}, l, i, o^*, \text{st})$ is the public (stateful) output verification algorithm. On input global parameters for \mathcal{M} , a label l , an input i , an attested output o^* and some state st it will produce an output value o and an updated state, or the failure symbol \perp . This failure symbol is encoded so as to be distinguishable from a valid output of a program, resulting from a successful verification. This algorithm is run locally on claimed outputs from the Attest algorithm. The initial value of the verification state is set to be (prms, P, L^*) , the same inputs provided to Compile .

CORRECTNESS. Intuitively, a LAC scheme is correct if, for any given program P and attested label set L^* , assuming an honest execution of all components in the scheme, both locally and remotely, the local user is able to accurately reconstruct a partial view of the I/O sequence that took place in the remote environment, for an arbitrary set of labels L (which may or may not be related to L^*). For non-attested labels, i.e., labels in $L \setminus L^*$, we restrict the I/O behaviour of the compiled program inside the IEE imposing that it is identical to that of the original program. For this reason, the set of labels L should be seen as a parameter that can be used by higher level protocols relying on LAC to specify the *partial* local view that may interest a particular party interacting with a remote machine. Different parties may be interested in different partial views, including both attested and non-attested labels, and the protocol should be correct for all of them. More technically, suppose the compiled program is run under handle hdl^* in remote machine \mathcal{M} , with random coins $\text{Coins}_{\mathcal{M}}(\text{hdl}^*)$ and on labelled input sequence $(l_1, i_1, \dots, l_n, i_n)$. Suppose also that, running the original program on the same random coins and inputs yields

$$\text{Trace}_{R[\text{st}; \text{Coins}_{\mathcal{M}}(\text{hdl}^*)]}(l_1, i_1, \dots, l_n, i_n) = (l_1, i_1, o_n, \dots, l_n, i_n, o_n)$$

Then, for any set of labels L , if a local user recovers outputs (o'_1, \dots, o'_m) corresponding to labelled inputs (l_{k_1}, i_{k_1}) to (l_{k_m}, i_{k_m}) , where $l_{k_j} \in L$, it must be the case that $(o'_1, \dots, o'_m) = (o_{k_1}, \dots, o_{k_m})$. Outputs for attested labels are passed through Attest and Verify , whereas inputs and outputs for non-attested labels are processed independently of these algorithms. The following definition formalizes the notion of a local user *correctly remotely executing program* P using labelled attested computation.

Definition 2 (Correctness). *A labelled attested computation scheme LAC is correct if, for all λ and all adversaries \mathcal{A} , the experiment in Fig. 4 (top) always returns \top .*

The adversary in this correctness experiment definition is choosing inputs, hoping to find a sequence that causes the attestation protocol to behave inconsistently with respect to the semantics of P (when these are made deterministic by hardwiring the same random coins used remotely). We use this approach to defining correctness because it makes explicit what is an honest execution of an attested computation scheme, when compared to the security experiment introduced next.

STRUCTURAL PRESERVATION. To simplify analysis of our constructions and proofs, we extend the correctness requirements on labelled attested computation schemes to preserve the structure of the input program when dealing with sequential composition, and to modify only the part of the code that will be attested. Formally, we impose that, for all global parameters, given any program $R = \langle P; Q \rangle_{\phi, p, q}$, and an attested label set L^* that contains only labels of the form (p, l) , then there exists a (unique) compiled program P^* , such that, $\langle P^*; Q \rangle_{\phi, p, q} = \text{Compile}(\text{prms}, R, L)$. Note that this implies that the state of compiled program P^* somehow encodes the state of P in a way that is transparent for ϕ , and we will loosely rely on this when referring to the execution state of P^* and extracting values from it. Note also that, for composed programs compiled in this way, the unattested I/O behaviour of the second program will be identical to that of the original program.

SECURITY. Security of labelled attested computation imposes that an adversary with control of the remote machine cannot convince the local user that some arbitrary remote (partial) execution of a program P has occurred, when it has not. It says nothing about the parts of the execution trace that are hidden from the client or are not in the attested label set L^* . Formally, we allow the adversary to freely interact with the remote machine, whilst providing a sequence of (potentially forged) attested outputs for a specific label $l \in L^*$. The adversary wins if the local user reconstructs an execution trace without aborting (i.e., all attested outputs must be accepted by the verification algorithm) and one of two conditions occur: i. there does not exist a remote process hdl^* running a compiled version

<p>Game $\text{Corr}_{\text{LAC}, \mathcal{A}}(1^\lambda)$:</p> <pre> prms \leftarrow $\mathcal{M}.\text{Init}(1^\lambda)$ $(P, L^*, L, n, \text{st}_{\mathcal{A}}) \leftarrow$ $\mathcal{A}_1(\text{prms})$ $P^* \leftarrow \text{Compile}(\text{prms}, P, L^*)$ $\text{hdl}^* \leftarrow \mathcal{M}.\text{Load}(P^*)$ $\text{st}_V \leftarrow (\text{prms}, P, L^*)$ For $k \in [1..n]$: $(l_k, i_k, \text{st}_{\mathcal{A}}) \leftarrow$ $\mathcal{A}_2(o_1^*, \dots, o_{k-1}^*, \text{st}_{\mathcal{A}})$ If $l_k \in L \cap L^*$: $o_k^* \leftarrow$ $\mathcal{A}.\text{Attest}^{\mathcal{M}}(\text{prms}, \text{hdl}^*, l_k, i_k)$ $(o_k, \text{st}_V) \leftarrow \text{Verify}(\text{prms}, l_k, i_k, o_k^*, \text{st}_V)$ If $o_k = \perp$ Then Return F Else If $l_k \in L \setminus L^*$: $o_k \leftarrow$ $\mathcal{M}.\text{Run}(\text{hdl}^*, l_k, i_k)$ Else Return F $T' \leftarrow \text{filter}[L](\text{Trace}_{P[\text{st}, \text{Coins}_{\mathcal{M}}](\text{hdl}^*)}(l_1, i_1, \dots, l_n, i_n))$ $T \leftarrow \text{filter}[L](l_1, i_1, o_1, \dots, l_n, i_n, o_n)$ // get only labels in L Return $T = T'$ </pre>
<p>Game $\text{Att}_{\text{LAC}, \mathcal{A}}(1^\lambda)$:</p> <pre> prms \leftarrow $\mathcal{M}.\text{Init}(1^\lambda)$ $(P, L^*, l, n, \text{st}_{\mathcal{A}}) \leftarrow$ $\mathcal{A}_1(\text{prms})$ $P^* \leftarrow \text{Compile}(\text{prms}, P, L^*)$ $\text{st}_V \leftarrow (\text{prms}, P, L^*)$ For $k \in [1..n]$: $(i_k, o_k^*, \text{st}_{\mathcal{A}}) \leftarrow$ $\mathcal{A}_2^{\mathcal{M}}(\text{st}_{\mathcal{A}})$ $(o_k, \text{st}_V) \leftarrow \text{Verify}(\text{prms}, l, i_k, o_k^*, \text{st}_V)$ If $o_k = \perp$ Return F $T \leftarrow (l, i_1, o_1, \dots, l, i_n, o_n)$ For hdl^* s.t. $\text{Program}_{\mathcal{M}}(\text{hdl}^*) = P^*$ $(l'_1, i'_1, o'_1, \dots, l'_m, i'_m, o'_m) \leftarrow \text{Trace}_{\mathcal{M}_{\mathcal{B}_y}}(\text{hdl}^*)$ $T' \leftarrow \text{filter}[l](\text{Trace}_{P[\text{st}, \text{Coins}_{\mathcal{M}}](\text{hdl}^*)}(l'_1, i'_1, \dots, l'_m, i'_m))$ If $T \subseteq T'$ Return F Return \top </pre>

Fig. 4. Games defining the correctness (top) and security (bottom) of LAC.

of P where a consistent set of inputs was provided *for label* l ; or ii. the outputs recovered by the local user for those inputs are not consistent with the semantics of P if it were run locally.

Technically, these conditions are checked in the definition by retrieving the full sequence of *label-input pairs* and random coins passed to all compiled copies of P running in the remote machine and running P on the same inputs to obtain the expected outputs. One then checks that for at least one of these executions, when the traces are restricted to special label l , that the expected trace matches the locally recovered trace via `Verify`. Since the adversary is free to interact with the remote machine as it pleases, we cannot hope to prevent it from providing arbitrary inputs to the remote program at arbitrary points in time, while refusing to deliver the resulting (possibly attested) outputs to the local user. This justifies the winning condition referring to a prefix of the execution in the remote machine, rather than imposing trace equality. Indeed, the definition's essence is to impose that, if the adversary delivers attested outputs for a particular label in the attested label set, then the subtrace of verified outputs for that label will be an exact prefix of the projection of the remote trace for that label.

We note that a higher-level protocol relying on LAC will be able to fully control the semantics of labels, as these depend on the semantics of the compiled program. In particular, adopting the specific forms of parallel and sequential composition presented in Section 3, it is possible to use labels to get the attested execution of a sub-program that is fully isolated from other programs that it is composed with. This provides a much higher degree of flexibility than what was available in the original notion of Attested Computation.

Definition 3 (Security). *A labelled attested computation scheme is secure if, for all ppt adversaries \mathcal{A} , the probability that experiment in Fig. 4 (bottom) returns \top is negligible.*

We note that the adversary loses the game as long as there exists at least one remote process that matches the locally reconstructed trace. This should be interpreted as the guarantee that IEE resources are indeed being allocated in a specific remote machine to run at least one instance of the remote program (note that if the program is deterministic, many instances could exist with exactly the same I/O behaviour, which is *not* seen as a legitimate attack). Furthermore, our definition imposes that the compiled program uses essentially the same randomness as the source program (except of course for randomness that the security module internally uses to provide its cryptographic functionality), as otherwise it will be easy for the adversary to make the (idealized) local trace diverge from the remote. This is a consequence of our modelling approach, but in no way does it limit the applicability of the primitive we are proposing: it just makes it explicit that the transformation that is performed on the code

Game $\text{Leak-Real}_{\text{LAC},\mathcal{A}}(1^\lambda)$: $\text{PrgList} \leftarrow []$ $\text{prms} \leftarrow \mathcal{M}.\text{Init}(1^\lambda)$ $b \leftarrow \mathcal{A}^\mathcal{O}(\text{prms})$ Return b	Oracle $\text{Compile}(P, L)$: $P^* \leftarrow \text{Compile}(\text{prms}, P, L)$ $\text{PrgList} \leftarrow P^* : \text{PrgList}$ Return P^*
Oracle $\text{Run}(\text{hdl}, l, i)$: Return $\mathcal{M}.\text{Run}(\text{hdl}, l, i)$	Oracle $\text{Load}(P)$: Return $\mathcal{M}.\text{Load}(P)$

Game $\text{Leak-Ideal}_{\text{LAC},\mathcal{A},\mathcal{S}}(1^\lambda)$: $\text{PrgList} \leftarrow []$; $\text{List} \leftarrow []$ $\text{hdl} \leftarrow 0$ $(\text{prms}, \text{st}_\mathcal{S}) \leftarrow \mathcal{S}_1(1^\lambda)$ $b \leftarrow \mathcal{A}^\mathcal{O}(\text{prms})$ Return b	Oracle $\text{Compile}(P, L)$: $P^* \leftarrow \text{Compile}(\text{prms}, P, L)$ $\text{PrgList} \leftarrow (P^*, L, P) : \text{PrgList}$ Return P^*
Oracle $\text{Run}(\text{hdl}, l, i)$: $(P^*, \text{st}) \leftarrow \text{List}[\text{hdl}]$ If $(P^*, L, P) \in \text{PrgList}$: $o \leftarrow P[\text{st}](l, i)$ $(o^*, \text{st}_\mathcal{S}) \leftarrow \mathcal{S}_2(\text{hdl}, P, L, l, i, o, \text{st}_\mathcal{S})$ Else: $(o^*, \text{st}_\mathcal{S}) \leftarrow \mathcal{S}_3(\text{hdl}, P^*, l, i, \text{st}, \text{st}_\mathcal{S})$ $\text{List}[\text{hdl}] \leftarrow (P^*, \text{st})$ Return o^*	Oracle $\text{Load}(P^*)$: $\text{hdl} \leftarrow \text{hdl} + 1$ $\text{List}[\text{hdl}] \leftarrow (P^*, \epsilon)$ Return hdl

Fig. 5. Games defining minimum leakage of LAC.

for attestation will typically consist of an instrumentation of the code by applying cryptographic processing to the inputs and outputs it receives.

MINIMAL LEAKAGE. The above discussion shows that a LAC scheme guarantees that the I/O behaviour of the program in the remote machine includes at least the information required to reconstruct an hypothetical local execution of the source program. Next, we require that a compiled program does not reveal any information beyond what the original program would reveal. The following definition imposes that nothing from the internal state of the source programs (in addition to what is public, i.e., the code and I/O sequence) is leaked in the trace of the compiled program.

Definition 4 (Minimal leakage). *A labelled attested computation scheme LAC ensures security with minimal leakage if it is secure according to Definition 3 and there exists a ppt simulator \mathcal{S} that, for every adversary \mathcal{A} , the following distributions are identical:*

$$\{\text{Leak-Real}_{\text{LAC},\mathcal{A}}(1^\lambda)\} \approx \{\text{Leak-Ideal}_{\text{LAC},\mathcal{A},\mathcal{S}}(1^\lambda)\}$$

where games $\text{Leak-Real}_{\text{LAC},\mathcal{A}}$ and $\text{Leak-Ideal}_{\text{LAC},\mathcal{A},\mathcal{S}}$ are shown in Fig. 5.

Notice that we allow the simulator to replace the global parameters of the machine with some value prms for which it can keep some trapdoor information. Intuitively this means that one can construct a perfect simulation of the remote trace by simply appending cryptographic material to the local trace. This property is important when claiming that the security of a cryptographic primitive is preserved when it is run within an attested computation scheme (one can simply reduce the advantage of an adversary attacking the attested trace, to the security of the original scheme using the minimum leakage simulator).

5 LAC from SGX-like systems

Our labelled attested computation protocol relies on the capabilities offered by the security module of Secure Guard Extensions (SGX) architecture proposed by Intel [1] (i.e. MACs for authenticated communication between IEEs, and digital signatures for inter-platform attestation of executions). Our security module formalization is the same as the one adopted in [2].

SECURITY MODULE. The security module relies on a signature scheme $\Sigma = (\text{Gen}, \text{Sign}, \text{Vrfy})$ and a MAC scheme $\Pi = (\text{Gen}, \text{Mac}, \text{Ver})$, and it operates as follows:

- When the host machine is initialised, the security module generates a key pair (pk, sk) using $\Sigma.\text{Gen}$ and a symmetric key key using $\Pi.\text{Gen}$. It also creates a special process running code S^* (see below for a description of S^*) in an IEE with handle 0. The security module then securely stores the key material for future use, and outputs the public key. In this case we will have that the output of $\mathcal{M}.\text{Init}$ will be $\text{prms} = \text{pk}$.

- The operation of IEE with handle 0 will be different from all other IEEs in the machine. Program S^* will permanently reside in this IEE, and it will be the only one with direct access to both sk and key .
- The code of S^* is dedicated to transforming messages authenticated with key into messages signed with sk . On each activation, it expects an input (m, t) . It obtains key from the security module and verifies the tag using $\Pi.Ver(key, t, m)$. If the previous operation was successful, it obtains sk from the security module, signs the message using $\sigma \leftarrow \Sigma.Sign(sk, m)$ and writes σ to the output. Otherwise, it writes \perp in the output.
- The security module exposes a single system call $mac(m)$ to code running in all other IEEs. On such a request from a process running program P , the security module returns a MAC tag t computed using key over both the code of P and the input message m .

LABELLED ATTESTED COMPUTATION SCHEME. We now define a LAC scheme that relies on a remote machine supporting a security module with the above functionality. Basic replay protection using a sequence number does not suffice to bind a remote process to a subtrace, since the adversary could then run multiple copies of the same process and *mix and match* outputs from various traces. This is similar to the reasoning in [2]. However, in this paper we are interested in validating traces for specific attested labels, independently from each other, rather than the full remote trace.

Our LAC scheme works as follows:

- $Compile(prms, P, L)$ will generate a new program P^* and output it. Program P^* is instrumented as follows:
 - in addition to the internal state st of P , it maintains a list ios_l of all the I/O pairs it has previously received and computed for each label $l \in L$.
 - On input (l, i) , P^* computes $o \leftarrow P[st_P](l, i)$ and verifies if $l \in L$. If this is not the case, then P^* simply outputs non-attested output o .
 - Otherwise, it updates the list ios by appending (l, i, o) , computes the subset of ios for label l : $ios_l \leftarrow filter[l](ios)$ and requests from the security module a MAC of for that list. Due to the operation of the security module, this will correspond to a tag t on the tuple (P^*, ios_l) .
 - It finally outputs (o, t, P^*, ios_l) . We note that we include (P^*, ios_l) explicitly in the outputs of P^* for clarity of presentation only. This value would be kept in an insecure environment by a stateful *Attest* program.
- $Attest(prms, hdl, l, i)$ invokes $\mathcal{M}.Run(hdl, (l, i))$ using the handle and input value it has received. When the process produces an output o , *Attest* parses it into (o', t, P^*, ios_l) . It may happen that parsing fails, e.g., if the label is not to be attested, in which case *Attest* simply produces o as its own output. Otherwise, it uses $\mathcal{M}.Run(0, (P^*, ios_l, t))$ to convert the tag into a signature σ on the same message. If this conversion fails, then *Attest* produces the original output o as its own output. Otherwise, it outputs (o', σ) .
- $Verify(prms, l, i, o^*, st)$ is the stateful verification algorithm. The original (public) value of the state st includes uncompiled program P and the list of attested labels L (this naturally extends to including compiled program P^* since *Compile* is deterministic); it also includes a (initially empty) list of previously attested input-output pairs ios . *Verify* returns o^* if $l \notin L$. Otherwise, it first parses o^* into (o, σ) , appends (l, i, o) to ios and verifies the digital signature σ using $prms$ and $(P^*, filter[l](ios))$. If parsing or verification fails, *Verify* outputs \perp . If not, then *Verify* terminates outputting o .

CORRECTNESS. It is easy to see that our LAC scheme is correct, provided that the underlying signature and message authentication schemes are correct, and that it preserves the structure of compiled programs. To see that this is the case, note that during the execution of P^* for $l_k \in L$, unless a MAC or signature verification fails, the I/O sequence provided by *Verify* will match the one reconstructed in T' (the inputs are the same, and the associated randomness tapes are fixed by $Coins_{\mathcal{M}}(hdl^*)$), and therefore $T = T'$. Since these algorithms are only used for attested labels, we only need to consider this possibility for labels $l \in L^* \cap L$. Now, observe that if the message authentication code scheme is correct, then the MAC verification will never fail, and if the message signature scheme is correct, then the signature verification will never fail. This is the case because the combined operations of P^* , *Attest*, the signing IEE running S^* and the security module lead to tags and signatures on pairs (P^*, ios_l) that exactly match the inputs provided to the verification algorithms in $\Pi.Ver$ and $\Sigma.Verify$. This gives us that the received trace and the reconstructed trace will be the same for all labels in L .

SECURITY.

Theorem 1. *The LAC scheme presented above provides secure attestation if the underlying MAC scheme Π and signature scheme Σ are existentially unforgeable. Furthermore, it unconditionally ensures minimum leakage.*

The proof of the following theorem can be found in Appendix A. The proof intuition is a generalization of the argument for the attested computation scheme in [2]. All attested outputs are bound to a partial execution trace that contains the entire I/O sequence associated with the corresponding attested label, so all messages accepted

by `Verify` must exist as a prefix for a remote trace of some instance of P^* . The adversary can only cause an inconsistency in $T \sqsubseteq T'$ if the signature verification performed by `Verify` accepts a message of label $l \in L^*$ that was never authenticated by an IEE running P^* . However, in this case the adversary is either breaking the MAC scheme (and dishonestly executing `Attest`), or breaking the signature (directly forging attested outputs).

6 Secure computation from LAC

FUNCTIONALITIES. We want to securely execute a functionality \mathcal{F} defined by a four-tuple $(n, F, \text{Lin}, \text{Lout})$, where F is a deterministic stateful transition function that takes inputs of the form (id, i) . Here, id is a party identifier, which we assume to be an integer in the range $[1..n]$, and n is the total number of participating parties. On each transition, F produces an output that is intended for party id , as well as an updated state. We associate to F two leakage functions $\text{Lin}(k, i, \text{st})$ and $\text{Lout}(k, o, \text{st})$ which define the public leakage that can be revealed by a protocol about a given input i or output o for party k , respectively; for the sake of generality, both functions may depend on the internal state st of the functionality, although this is not the case in the examples we consider in this paper. Arbitrary reactive functionalities formalized in the Universal Composability framework can be easily recast as transition function such as this. The upside of our approach is that one obtains a precise code-based definition of what the functionality should do (this is central to our work since these descriptions give rise to concrete programs); the downside is that the code-based definitions may be less clear to a human reader, as one cannot ignore the tedious *book-keeping* parts of the functionality.

EXECUTION MODEL. We assume the existence of a machine \mathcal{M} allowing for the usage of isolated execution environments, such as those defined in Section 3. In secure computation terms, this machine should *not* be seen as an ideal functionality that enables some hybrid model of computation, but rather an additional party that comes with a specific setup assumption, a fixed internal operation, and which cannot be corrupted.¹³ More in detail, and following [2], this machine is first initialized via the `Init` algorithm, which defines public parameters that one assumes can be independently authenticated by all parties. The machine \mathcal{M} is assumed to be adversarially controlled, but it does include isolated execution environments in which programs can be loaded (via the `Load` mechanism) and then interactively fed with new inputs (via the `Run` mechanism) to obtain attested outputs. All the code that is run in \mathcal{M} but outside these execution environments is considered to be adversarially controlled. However, the adversary controls the interaction with \mathcal{M} and the goal is to guarantee that a set of parties can use the IEE capabilities of \mathcal{M} securely (bar the possibility that \mathcal{M} refuses to allow the protocol to proceed, which would amount to a DoS attack).

SYNTAX. A protocol π for functionality \mathcal{F} is a seven-tuple of algorithms as follows:

- **Setup** – This is the party local set-up algorithm. Given the security parameter, the public parameters `prms` for machine \mathcal{M} and the party’s identifier `id`, it returns the party’s initial state `st` (including its secret key material) and its public information `pub`.
- **Compile** – This is the (deterministic) code generation algorithm. Given the description of a functionality F , and the public parameters `(prms, Pub)` for both the remote machine and the entire set of public parameters for the participating parties, it generates the instrumented program that will run inside an IEE.
- **Remote** – This is the untrusted code that will be run in \mathcal{M} and which ensures the correctness of the protocol by controlling its scheduling and input collection order. It has oracle access to \mathcal{M} , and is run on public parameters `prms`, the handle to the IEE in which the compiled program is running and input message `m`; it returns the output message `m'`, the identity `id` of the party for which `m'` is intended, and a flag `inreq` that indicates whether party `id` is expected to provide an input at this step of the protocol. Its initial state describes the order in which inputs of different parties should be provided to the functionality.
- **Init** – This is the party local protocol initialization algorithm. Given the party’s state `st` produced by `Setup` and the public information of all participants `Pub` it outputs an updated state `st`. We note that a party can choose to engage in a protocol by checking if the public parameters of all parties are correct and assigned to roles in the protocol that match the corresponding identities.
- **AddInput** – This is the party local input providing algorithm. Given the party’s current state `st` and an input `in`, it outputs an updated state `st`.
- **Process** – This is the party local message processing algorithm. Given its internal state `st`, and an input message `m`, it runs the next protocol stage, updates the internal state and returns output message `m'`. Protocol termination will be locally signalled with an output message `m' = ⊥`.

¹³ Relating this to the Universal Composability framework, this special party \mathcal{M} does not take inputs or outputs, and is accessible only via its communications tape, which is assumed to be controlled by the adversary.

```

Game  $\text{Corr}_{\mathcal{F}, \pi, r, m, \mathcal{A}, \mathcal{M}}(1^\lambda)$ :
// Trusted setup of machine and parties
 $(n, F, \text{Lin}, \text{Lout}) \leftarrow \mathcal{F}$ 
 $\text{prms} \leftarrow \mathcal{M}.\text{Init}(1^\lambda)$ 
For  $\text{id} \in [1..n]$ :
   $(\text{st}_{\text{id}}, \text{pub}_{\text{id}}) \leftarrow \mathcal{S}.\text{Setup}(\text{prms}, \text{id})$ 
 $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_n)$ 
For  $\text{id} \in [1..n]$ :
   $\text{st}_{\text{id}} \leftarrow \mathcal{S}.\text{Init}(\text{st}_{\text{id}}, \text{Pub})$ 
   $\text{out}_{\text{id}_i} \leftarrow \epsilon$ 

// Adversarially scheduled ideal execution
 $\text{st}_{\mathcal{F}} \leftarrow \epsilon$ ;  $\text{st}_{\mathcal{A}} \leftarrow \epsilon$ 
For  $i \in [1..m]$ :
   $(\text{id}_i, \text{in}_i, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}(\text{prms}, \text{Pub}, \text{st}_{\mathcal{A}})$ 
   $\text{out}_{\text{id}_i} \leftarrow \text{out}_{\text{id}_i} \parallel F[\text{st}_{\mathcal{F}}](\text{id}_i, \text{in}_i)$ 

// Protocol execution
 $F^* \leftarrow \text{Compile}(\text{prms}, \mathcal{F}, \text{Pub})$ 
 $\text{hdl} \leftarrow \mathcal{M}.\text{Load}(F^*)$ 
 $t \leftarrow \top$ ;  $m \leftarrow \epsilon$ ;  $j \leftarrow 0$ 
 $\text{st}_R \leftarrow (\text{id}_1, \dots, \text{id}_m)$  // input schedule
For  $i \in [1..r]$ :
  If  $t$ : // Remote step
     $(\text{id}, \text{inreq}, \text{m}', \text{st}_R) \leftarrow \mathcal{R}.\text{Remote}^{\mathcal{M}}(\text{prms}, \text{hdl}, \text{m}, \text{st}_R)$ 
  Else: // Local step
    If  $\text{inreq} = \top$ :
      If  $\text{id} \neq \text{id}_j$  Return  $F$ 
       $\text{st}_{\text{id}} \leftarrow \mathcal{S}.\text{AddInput}(\text{in}_j, \text{st}_{\text{id}})$ 
       $j \leftarrow j + 1$ 
     $(\text{st}_{\text{id}}, \text{m}) \leftarrow \mathcal{S}.\text{Process}(\text{st}_{\text{id}}, \text{m}')$ 
   $t \leftarrow \neg t$ 
For  $\text{id} \in [1..n]$ :
   $\text{out}'_{\text{id}} \leftarrow \text{Output}(\text{st}_{\text{id}})$ 
Return  $(\text{out}_1, \dots, \text{out}_n) = (\text{out}'_1, \dots, \text{out}'_n)$ 

```

Fig. 6. Game defining protocol correctness.

– **Output** – This is the party local output retrieval algorithm. Given internal state st , it returns the current output o .

CORRECTNESS. The following definition formalizes the notion of n users *correctly running a function evaluation protocol* π for F .

Definition 5. *We say π is correct for functionality \mathcal{F} on m inputs in r rounds if, for all λ , and all adversaries \mathcal{A} , the experiment in Figure 6 always returns \top .*

DISCUSSION. Our correctness definition considers an honest execution environment, but includes a correctness adversary that is in charge of finding problematic inputs for the protocol and potentially erroneous execution schedules. It is parametrized by a number of inputs m and a number of rounds r .

The first stage of the experiment executes the **Setup** and **Init** algorithms that initialize both the remote machine and the parties' local states, and collects the public parameters for all of these participants (which we assume to be authenticated through the paper, e.g., using a PKI). In the second part of the experiment, the adversary chooses a sequence of m inputs for the functionality, interleaving different parties in an arbitrary way. The sequence $(\text{id}_1, \dots, \text{id}_m)$ implicitly defines a schedule for the execution of our protocol, which should ensure that the inputs of each party are provided to the functionality in precisely this order.

The last stage of the correctness experiment emulates the protocol execution, alternating between local steps and remote steps. The **Remote** algorithm commands the scheduling of message exchanges; this algorithm is always invoked first and its output indicates the next party to be activated, the message this party will receive, and whether or not the party is expected to provide an input. The protocol is run for r rounds, at which point its outputs are retrieved via **Output**. The adversary wins the game if it can force the game to produce a set of outputs that wouldn't be obtained by simply running the functionality F with the given inputs in the provided order.

REMARK. The correctness experiment shows the crucial scheduling role of the **Remote** algorithm, which is run in an untrusted environment in the remote machine. Here we deviate from the standard approach in the UC setting, where the simulation-based definition of security is taken as providing sufficient detail to evaluate correctness of the protocol. Indeed, as other simulation-based definitions, our security experiment below will impose some input/output consistency conditions on the protocol. Intuitively, these must hold for any adversarially chosen **Remote**

<p>Game $\text{Real}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}}(1^\lambda)$:</p> <p>$(n, F, \text{Lin}, \text{Lout}) \leftarrow \mathcal{F}$ $\text{prms} \leftarrow \mathcal{M}.\text{Init}(1^\lambda)$ $(\text{st}_{\mathcal{A}}, k) \leftarrow \mathcal{A}(\text{prms})$ For $\text{id} \in [1..k]$: $(\text{st}_{\text{id}}, \text{pub}_{\text{id}}) \leftarrow \text{Setup}(\text{prms}, \text{id})$ $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_k)$ For $\text{id} \in [k + 1..n]$: $(\text{st}_{\mathcal{A}}, \text{pub}_{\text{id}}) \leftarrow \mathcal{A}(\text{st}_{\mathcal{A}}, \text{id}, \text{Pub})$ $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_n)$ For $\text{id} \in [1..k]$: $\text{st}_{\text{id}} \leftarrow \mathcal{I}.\text{Init}(\text{st}_{\text{id}}, \text{Pub})$ $b \leftarrow \mathcal{A}^{\mathcal{O}}(\text{st}_{\mathcal{A}})$</p> <p>Oracle $\text{Send}(\text{id}, m)$:</p> <p>If $\text{id} \notin [1..k]$ Return \perp $(\text{st}_{\text{id}}, m') \leftarrow \mathcal{S}.\text{Process}(\text{st}_{\text{id}}, m)$ Return m'</p>	<p>Oracle $\text{SetInput}(\text{in}, \text{id})$:</p> <p>If $\text{id} \notin [1..k]$ Return \perp $\text{st}_{\text{id}} \leftarrow \mathcal{S}.\text{AddInput}(\text{in}, \text{st}_{\text{id}})$</p> <p>Oracle $\text{Load}(P)$:</p> <p>Return $\mathcal{M}.\text{Load}(P)$</p> <p>Oracle $\text{Run}(\text{hdl}, l, x)$:</p> <p>Return $\mathcal{M}.\text{Run}(\text{hdl}, l, x)$</p> <p>Oracle $\text{GetOutput}(\text{id})$:</p> <p>If $\text{id} \notin [1..k]$ Return \perp Return $\text{Output}(\text{st}_{\text{id}})$</p>
<p>Game $\text{Ideal}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{S}}(1^\lambda)$:</p> <p>$(n, F, \text{Lin}, \text{Lout}) \leftarrow \mathcal{F}$ $\text{st}_{\mathcal{F}} \leftarrow \epsilon$ $(\text{st}, \text{prms}) \leftarrow \mathcal{S}(1^\lambda)$ $(\text{st}_{\mathcal{A}}, k) \leftarrow \mathcal{A}(\text{prms})$ For $\text{id} \in [1..k]$: $(\text{st}, \text{pub}_{\text{id}}) \leftarrow \mathcal{S}(\text{st}, \text{id})$ $\text{ListIn}_{\text{id}} \leftarrow []$ $\text{ListOut}_{\text{id}} \leftarrow []$ $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_k)$ For $\text{id} \in [k + 1..n]$: $(\text{st}_{\mathcal{A}}, \text{pub}_{\text{id}}) \leftarrow \mathcal{A}(\text{st}_{\mathcal{A}}, \text{id}, \text{Pub})$ $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_n)$ For $\text{id} \in [1..k]$: $\text{st} \leftarrow \mathcal{S}(\text{st}, \text{id}, \text{Pub})$ $b \leftarrow \mathcal{A}^{\mathcal{O}}(\text{st}_{\mathcal{A}})$</p> <p>Oracle $\text{Fun}(\text{id}, \text{in})$:</p> <p>If $\text{id} \in [1..k]$: $(\text{in}_1, \dots, \text{in}_k) \leftarrow \text{ListIn}_{\text{id}}$ $\text{ListIn}_{\text{id}} \leftarrow (\text{in}_1, \dots, \text{in}_{k-1})$ $\text{out} \leftarrow F[\text{st}_{\mathcal{F}}](\text{id}, \text{in}_k)$ $\text{ListOut}_{\text{id}} \leftarrow \text{out} : \text{ListIn}_{\text{id}}$ Return $\text{Lout}(\text{out}, \text{id}, \text{st}_{\mathcal{F}})$ Else $\text{out} \leftarrow F[\text{st}_{\mathcal{F}}](\text{id}, \text{in})$ Return out</p>	<p>Oracle $\text{SetInput}(\text{in}, \text{id})$:</p> <p>If $\text{id} \notin [1..k]$ Return \perp $\ell \leftarrow \text{Lin}(\text{in}, \text{id}, \text{st}_{\mathcal{F}})$ $\text{st} \leftarrow \mathcal{S}(\text{st}, \ell, \text{id})$ $\text{ListIn}_{\text{id}} \leftarrow \text{in} : \text{ListIn}_{\text{id}}$</p> <p>Oracle $\text{Send}(\text{id}, m)$:</p> <p>$(\text{st}, \text{out}) \leftarrow \mathcal{S}^{\text{Fun}}(\text{st}, \text{id}, m)$ Return out</p> <p>Oracle $\text{Load}(P)$:</p> <p>$(\text{st}, \text{out}) \leftarrow \mathcal{S}(\text{st}, P)$ Return out</p> <p>Oracle $\text{Run}(\text{hdl}, l, x)$:</p> <p>$(\text{st}, \text{out}) \leftarrow \mathcal{S}^{\text{Fun}}(\text{st}, \text{hdl}, l, x)$ Return out</p> <p>Oracle $\text{GetOutput}(\text{id})$:</p> <p>If $\text{id} \notin [1..k]$ Return \perp $i \leftarrow \mathcal{S}(\text{st}, \text{id})$ $(\text{out}_1, \dots, \text{out}_k) \leftarrow \text{ListOut}_{\text{id}}$ Return $\text{out}_1 \parallel \dots \parallel \text{out}_i$</p>

Fig. 7. Real and Ideal security games.

scheduling algorithm, as the adversary has full control of the remote machine and scheduling can be arbitrarily controlled by the attacker. However, we believe that there is added value in including a separate correctness definition, where the scheduling tasks of the non-security critical parts of the protocol can be specified as a first class feature of the protocol syntax. This also clarifies the envisioned execution model and makes it explicit that untrusted code running in an adversarially run machine is only relevant for correctness purposes.

SECURITY. Our security definition is based on the experiments shown in Figure 7, where \mathcal{O} represents access to all oracles except Fun , i.e., the adversary has oracle access to Run , Setup , SetInput , GetOutput and Send in both games.

Definition 6. We say π is secure for \mathcal{F} if, for any ppt adversary \mathcal{A} , there exists a ppt simulator \mathcal{S} such that the following definition of advantage is a negligible function in the security parameter.

$$|\Pr[\text{Real}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}}(1^\lambda) \Rightarrow b = 1] - \Pr[\text{Ideal}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{S}}(1^\lambda) \Rightarrow b = 1]|$$

DISCUSSION. As is customary in secure computation models, we take the ideal world versus real world approach. In the real world we have a remote machine \mathcal{M} following the description in Section 3, which is under the control of the adversary. We consider also n parties, the first k of which are honest, and the rest corrupt. The experiment begins with the trusted setup of the machine, the attacker selectively choosing the number of corrupt parties, and the trusted initialization of the initial states of all honest parties. The adversary gets the public parameters of all honest parties (including the machine) and chooses those of the corrupt parties. All public parameters are assumed to be authenticated and available to all (e.g., via a PKI). We note that the adversary has all the information it needs to honestly execute the protocol and, in particular, it may properly run Compile and Remote on its own or choose to arbitrarily deviate from it. The honest parties will, of course, run Compile correctly and will therefore

be able to take advantage of the attestation guarantees provided by the remote machine. Intuitively, this is the underlying setup assumption that enables secure computation to be done efficiently.

The adversary then takes control of the experiment execution and it has access to a series of oracles that allow it to fully control the remote machine (`Load` and `Run`) and interact with the honest parties: setting their inputs via `SetInput`, checking the outputs via `GetOutput` and delivering messages via `Send`. In this sense, our real world experiment is similar to the Universal Composability framework: our adversary represents the UC environment that controls party inputs and outputs combined with a *dummy adversary* that conveys communications between parties and the environment. When the adversary terminates, it outputs a bit b with its guess of whether it is executing in the real or ideal worlds.

As in the real world experiment, the ideal experiment begins with the global set-up procedure. We give the simulator full control of the remote machine, which is always assumed to be honest, and this includes the ability to generate its global parameters. We also give the simulator the ability to control the honest party parameter generation. The remaining parties are initialized in the standard way, and the adversary is run on an equivalent set of parameters as in the real world.

The central component in the ideal experiment is the functionality oracle `Fun`. This oracle represents an idealized component, very much like the ideal functionality in the UC setting. The adversary can provide inputs and read outputs from the functionality via the `SetInput` and `GetOutput` oracles; this ability applies only to the inputs of honest parties. The scheduling of the computation of the functionality F is controlled by the simulator, as is the setting of the inputs corresponding to corrupt parties. As in the UC framework, this can be interpreted as an idealization of the operation of the functionality where some of its participants are dishonest: from this perspective, the simulator should be seen as an adversary *attacking* the ideal functionality and, as such, it should control the dishonest inputs.

Succinctly, our model is inspired in the UC framework, and can be derived from it when natural restrictions are imposed: PKI, static corruptions, and a distinguished non-corruptible party modeling an SGX-enabled machine.¹⁴ A security proof for a protocol in our model can be interpreted as translation of any attack against the protocol in the real world, as an attack against the ideal functionality in the ideal world. The simulator performs this translation by presenting an execution environment to the adversary that is consistent with what it is expecting in the real world. It does this by simulating the operations of the `Load`, `Run` and `Send` oracles, which represent the operation of honest parties in the protocol. While the adversary is able to provide the inputs and read the outputs for honest parties directly from the functionality, the simulator is only able obtain partial leakage about this values via the `Lin` and `Lout` functions, which in this paper are assumed to just reveal the lengths of their inputs. Conversely, it can obtain the functionality outputs for corrupt parties via the `Fun` oracle and, furthermore, it is also able to control the rate and order in which all inputs are provided to the functionality. Were this not the case, the adversary would be able to distinguish the two worlds by manipulating scheduling in such a way that the simulator could not possibly match.

Similarly to the real world, the adversary will finish the interaction by outputting a bit b containing its guess of which world it is executing in.

7 A New MPC Protocol from SGX

We describe here a secure computation protocol based on LAC that works for any functionality. This protocol starts by running bootstrap code in an isolated execution environment in the remote machine, which exchanges keys with each of the participants in the protocol. These key exchange programs are composed in parallel, as seen in Section 3. Once this bootstrap stage is concluded, the code of the functionality, which is composed sequentially with the bootstrap stage, starts executing and it uses the secure channels established with each party to ensure that the collection of inputs and delivery of outputs is secure. In order to define the protocol, we first present a utility theorem very similar to that given in [2] for the use of key exchange in the context of attestation. This theorem shows that, under the specific program composition pattern that we require for our MPC protocol, which guarantees `AttKE` isolation from other programs, each party obtains a secret key that is indistinguishable from a random string and can therefore be used to construct a secure channel that connects it to code emulating the functionality within an IEE.

`ATTKE` AND ITS UTILITY. We define utility almost as in the AC case from [2]. The main difference is the composition context we allow for. In [2] the composition context is restricted to a key exchange composed sequentially with

¹⁴ This particular choice in our model has implications for the composability properties of our results, as discussed in the related work section.

another program. Here we allow for the key exchange to be composed in parallel with other arbitrary programs and then sequentially with another program. The proof follows the same lines. The utility security experiment intuitively states that the adversary cannot distinguish between a derived key and a random key, whenever the key exchange has been performed between an honest party and a remote machine running Rem_{KE} within an IEE, and Rem_{KE} is composed with other arbitrary programs as described above. The reason this parallel/sequential composition pattern does not harm security is that the parts of the state belonging to different parallel-composed programs are disjoint, and sequential composition only reveals controlled information to other programs. Indeed, as in [2], we restrict sequential composition in the utility theorem to pass only specific parts of the state of the key exchange program to the following phase: mapping function ϕ_{key} passes on the derived key, the session and party identifiers, and the state (derived or accept) of the key exchange. Contrary to [2] this is not enough to define our mapping function, as other programs composed in parallel with the remote key exchange need to pass states to the next phase as well. To that extent, if ϕ_1, \dots, ϕ_n are mapping functions, we define $\phi_1^{l_1} | \dots | \phi_n^{l_n}$ as $\phi^* := \text{st}.l_i \mapsto \phi_i(\text{st}.l_i)$. If the state comes from the program $\langle P_1 | \dots | P_n \rangle_{(l_1, \dots, l_n)}$, this mapping function maps the state belonging to each P_i using ϕ_i . In our composition context, we take the ϕ_i corresponding to the key exchange to be ϕ_{key} . This ensures that only the key is transmitted to the following stage of the protocol, and not information supposed to be local to the key exchange protocol and not intended for further use.

In the experiment in Figure 8 the adversary has to distinguish between an ideal machine and a real world machine where an AttKE is run in parallel with other programs in the first phase of a LAC-compiled protocol. The machine \mathcal{M} represents the remote machine expected by the LAC protocol and the machine \mathcal{M}' is a modification of machine \mathcal{M} in which the key derived by a key-exchange session is magically replaced by a fresh key. In order to maintain consistency between the tested keys and the keys used in \mathcal{M}' , oracle $\mathcal{M}'.Run takes two additional parameters: a list fake of pairs of keys and a flag tweak. If the flag is activated, the following modifications in the behaviour of \mathcal{M} occur in \mathcal{M}' :$

- It expects the sub-program being activated due to input label l to be a key exchange Rem_{KE} instance. After running its transition function, \mathcal{M}' checks if it has reached the `derived` or `accept` state. If so, it retrieves the derived key and if there is no association (key, $_$) in fake it generates a fresh key key^* and appends (key, key^*) to fake.
- Furthermore, if the key exchange process has entered `accept` state, it performs $\text{st}.l.\text{key} \leftarrow \text{fake}(\text{key})$, i.e., it replaces the derived key with a fake random one. Note that this will cause the fake key to be passed to the next stage of the sequentially composed program.

The oracles provided to the adversary provide access to the remote machine. Additionally the adversary can create new sessions of the key exchange using the `NewSession` oracle, where the remote key exchange is composed in parallel (with label l^*) with programs P_1, \dots, P_n , followed with Q and compiled for LAC. Note that, given the structure of our parallel composition, the position in which a program is listed in the composition expression is irrelevant. The adversary makes the local part of the key exchange progress by using the `Send` oracle, provided that the message passes the LAC verification step for the relevant label. Finally the adversary can challenge a session by executing the `Test` oracle, which return either the real key of a fake key according to b (provided that the key exchange has reached a `derived` or `accept` state).

Theorem 2 (Local AttKE utility). *If the AttKE is correct and secure, and the LAC protocol is correct, secure and ensures minimal leakage, then for all ppt adversaries in the labelled utility experiment: the probability that the adversary violates the AttKE two-sided entity authentication is negligible; and the key secrecy advantage $2 \cdot \Pr[\text{guess}] - 1$ is negligible.*

BOXING USING AUTHENTICATED ENCRYPTION. After the bootstrapping stage of our protocol, we will be running the ideal functionality within an isolated execution environment, and using secure channels to communicate with each participating party. The availability of these secure channels will follow from the utility theorem we presented above. We now formalized the concept of *boxing a functionality*, which defines a program that executes the functionality code, but receives inputs and delivers outputs using secure channels. This is done in the form of our `Box` construction presented in Figure 9. This construction takes a functionality \mathcal{F} for n parties and a secure authenticated encryption encryption scheme Λ . We construct a labelled program $\text{Box}(\mathcal{F}, \Lambda)$ whose initial state is assumed to contain n symmetric keys compatible with scheme Λ , denoted sk_1 to sk_n (one for each participating party) and the empty initial state for the functionality $\text{st}_{\mathcal{F}}$. To avoid replays of encrypted messages, we keep one sequence number seq_{id} per communicating party id. On input (i^*, l) , the program interprets the label as party identity $\text{id} = l$, decrypts input i^* with sk_{id} to obtain input $i = (\text{in}, \text{seq})$, and checks the corresponding sequence number. It then passes the decrypted input to the functionality, and subsequently encrypts the output of the functionality back to the same user, together with the updated sequence number. The updated state includes the updated sequence number

```

Game AttAttKE, $\mathcal{A}$ ( $1^\lambda$ , id):
prms0  $\leftarrow$   $\mathcal{M}$ .Init( $1^\lambda$ )
prms1  $\leftarrow$   $\mathcal{M}'$ .Init( $1^\lambda$ )
PrgList  $\leftarrow$  []
fake  $\leftarrow$  []
i  $\leftarrow$  0
b  $\leftarrow$   $\{0, 1\}$ 
b'  $\leftarrow$   $\mathcal{A}^{\mathcal{O}}$ (prmsb)
Return b = b'

Oracle Load( $R^*$ ):
hdl0  $\leftarrow$   $\mathcal{M}$ .Load( $R^*$ )
hdl1  $\leftarrow$   $\mathcal{M}'$ .Load( $R^*$ )
Return hdlb

Oracle Run(hdl, l, x):
o0  $\leftarrow$   $\mathcal{M}$ .Run(hdl, l, x)
tweak  $\leftarrow$  F
If (Program $_{\mathcal{M}'}$ (hdl), l)  $\in$  PrgList then tweak  $\leftarrow$  T
(o1, fake)  $\leftarrow$   $\mathcal{M}'$ .Run(hdl, l, x, tweak, fake)
Return ob

Oracle Test(i):
If stKEi. $\delta \neq$  accept: Return  $\perp$ 
If b = 0: Return stKEi.key
Else: Return fake(stKEi.key)

Oracle NewSession( $P_1, l_1, \phi_1, \dots, P_n, l_n, \phi_n, l^*, Q, L^*$ ):
If  $\exists j, k$  such that  $j \neq k \wedge l_j = l_k$ : Return  $\perp$ 
If  $(p, (l^*, \epsilon)) \notin L^*$ : Return  $\perp$ 
i  $\leftarrow$  i + 1
li*  $\leftarrow$   $(p, (l^*, \epsilon))$ 
(stKEi, RemKEi)  $\leftarrow$  Setup( $1^\lambda$ , id)
inlasti  $\leftarrow$   $\epsilon$ 
RemComp := (RemKEi | P1 | ... | Pn)(l^*, l_1, \dots, l_n)
 $\phi^*$  :=  $\phi_{\text{key}}^l | \phi_{l_1}^{l_1} | \dots | \phi_{l_n}^{l_n}$ 
Ri := (RemComp ; Q) $\phi, p, q$ 
Ri*  $\leftarrow$  LAC.Compile(prmsb, Ri, L*)
stVi  $\leftarrow$  (Ri*, L*)
PrgList  $\leftarrow$  (Ri*, li*) : PrgList
Return Ri*

Oracle Send(o*, i):
o  $\leftarrow$  LAC.Verify[stVi](prms, li*, inlasti, o*)
If o =  $\perp$ : Return  $\perp$ 
m*  $\leftarrow$  LockKE[stKEi](o)
inlasti  $\leftarrow$  m*
If stKEi. $\delta \in$  {derived, accept}  $\wedge$  stKEi.key  $\notin$  fake:
key*  $\leftarrow$   $\{0, 1\}^\lambda$ 
fake  $\leftarrow$  (key, key*) : fake
Return m*

```

Fig. 8. Utility of adversarially composed AttKE.

```

Program Box( $\mathcal{F}, \Lambda$ )[st](i*, l):
(n, F, Lin, Lout)  $\leftarrow$   $\mathcal{F}$ 
id  $\leftarrow$  l
If id  $\notin$  [1..n] : Return  $\perp$ 
If st.seqid =  $\epsilon$  :
st.seqid  $\leftarrow$  0
i  $\leftarrow$   $\Lambda$ .Dec(st.keyid, m)
If m = (in, st.seqid) :
o  $\leftarrow$  F[st.stF](id, in)
st.seqid  $\leftarrow$  st.seqid + 1
c  $\leftarrow$   $\Lambda$ .Enc(st.keyid, (seq, o))
st.seqid  $\leftarrow$  st.seqid + 1
Return c
Else: Return  $\perp$ 

```

Fig. 9. Boxing using Authenticated Encryption

and the updated state of the functionality. The transition function finally externalizes this updated state and the encrypted output.

THE PROTOCOL. Building on top of a LAC scheme, an AttKE scheme and our Box construction we define in Figure 10 a general secure multiparty computation protocol that works for any (possibly reactive) functionality F. The core of the protocol is the execution of an AttKE for each participant in parallel, followed by the execution

of the functionality F on the remote machine, under a secure channel with each participant as specified in the `Box` construct. More precisely:

- `Setup` derives the code for a remote key exchange program Rem_{KE} using the `AttKE` setup procedure. This code (which intuitively includes cryptographic public key material) is set to be the public information for this party. The algorithm also stores various parameters in the local state for future usage.
- `Compile` uses the LAC compilation algorithm on a program that results from the parallel composition of all the remote key exchange programs for all parties, which is then sequentially composed with the boxed functionality. Sequential composition uses the special ϕ_{key}^* function that maps the keys derived by *all* the key exchange Rem_{KE} instances into the initial state of the `Box` construction. The set of attested labels is restricted to those of form $(p, (\text{id}_i, \epsilon))$, corresponding to the `AttKE` subprograms.
- `Init` locally recomputes the program that is intended for remote execution, as this is needed for attestation verification. The set of labels that define the locally recovered trace is set to those of the form $\{(p, (\text{id}, \epsilon)), (q, \text{id})\}$, which correspond to those exactly matching the parts of the remote trace that are relevant for this party, namely its own key exchange and its own input/output relation with the functionality. Various parts of the local state that are used by `Process` are also initialized.
- `Process` is split into two stages. In the first stage it uses LAC with attested labels of the form $(p, (\text{id}, \epsilon))$ to execute `AttKE` protocol and establish a secure channel with the remote program. In the second stage, it uses non-attested labels of the form (q, id) , and it provides inputs to the remote functionality (on request) and recovers the corresponding outputs when they are delivered. The input sending process is initiated by passing an empty message into the algorithm, which triggers the encryption of the next input using the derived secret key from the first stage. A non-empty message input to this stage will trigger decryption and recovery of an output.
- `Output` reads the output in the state of the participant and returns it.
- `AddInput` adds an input to the end of the list of inputs that have to be transmitted by the participant.

The (untrusted) scheduling algorithm is shown in Figure 11. It is in charge of dispatching messages to/from the remote machine `IEE` using the `Attest` algorithm provided by the LAC, and animating the protocol to generate a correct execution for an arbitrary sequence of input-party interactions provided externally as an input schedule which is stored in its initial state. During the bootstrap stage, the `Remote` procedure interacts with one party at a time,¹⁵ moving from one party to the next when the previous party has moved to its second stage. When all parties have completed the key exchange, the `Remote` procedure detects this in the output of the `IEE` (consistently with the properties of our sequential composition), and moves to the functionality execution stage.

In this second stage, the algorithm simply follows the provided input schedule. Moving to the next input is triggered by feeding the algorithm with an empty input provided by the previous party (this is syntactic book-keeping to match our correctness requirement, and it signals the fact that the previous output was correctly delivered to the previous party). The consequence of such an action is that `Remote` signals that a new input should be requested from the next party in the schedule. When an actual input is received, this is passed into the `IEE` using an unattested label of the form (q, id) . The output is sent back to the same party.

For proving security, we restrict the functionalities we consider to a particular leakage function: size of inputs/outputs. We say that a functionality $(n, F, \text{Lin}, \text{Lout})$ leaks size if it is such that `Lin` and `Lout` return the length of the inputs/outputs (i.e. $\text{Lin}(k, x, \text{st}) = \text{Lout}(k, x, \text{st}) = |x|$ for every k, x, st).

Theorem 3. *If LAC is a correct and secure LAC scheme, AttKE is a secure AttKE scheme and Λ a secure authenticated encryption scheme, then the protocol in Figure 10 and Figure 11 is correct and secure for any functionality that leaks size.*

PROOF SKETCH. We build the required simulator \mathcal{S} as follows. For dishonest parties, the simulator executes the protocol normally while for the honest parties instead of encrypting the inputs/outputs the simulator encrypts dummy messages of the correct length (obtained through the leakage function) under freshly generated keys.

More precisely, simulator \mathcal{S} creates the public parameters of the machine honestly, and it stores all the information required to accurately execute all the system calls performed by all programs to the security module. It then generates the public parameters of the honest agents following the protocol, and generates a fresh key key_{id} for each honest party. Simulator \mathcal{S} also maintains a list of undelivered input and output messages for each `id`, and local state st_{id} for every honest party. Let F^* be the result of `Compile`(`prms`, F , `Pub`). When asked to simulate oracle queries, \mathcal{S} behaves as follows:

¹⁵ Other options were of course possible for implementing `Remote`, and the core of our protocol is actually compatible with a totally asynchronous scheduling. Dealing with such issues is out of the scope of this paper.

```

algorithm Setup(prms, id):
  stid.id ← id; stid.prms ← prms
  (stL, RemKE) ← SetupKE(1λ, id)
  stid.stL ← stL; stid.pub ← RemKE
  Return (stid, stid.pub)

algorithm Compile(prms,  $\mathcal{F}$ , Pub):
  (RemKE1, ..., RemKEn) ← Pub
  P ← ⟨ (RemKE1, ..., RemKEn)1, ..., n; Box⟨ $\mathcal{F}$ ,  $\Lambda$ ⟩φkey, p, q
  L* ← {(p, (1, ε)), ..., {(p, (n, ε))} // Labels for RemKEi are empty
  P* ← LAC.Compile(prms, P, L*)
  Return P*

algorithm Init(stid, Pub):
  stid.InList ← []; stid.stage ← 0;
  stid.seqin ← 0; stid.seqout ← 1; stid.inlast ← ε
  If Pub[stid.id] ≠ stid.pub : Return ⊥
  (RemKE1, ..., RemKEn) ← Pub
  P ← ⟨ (RemKE1, ..., RemKEn)1, ..., n; Box⟨ $\mathcal{F}$ ,  $\Lambda$ ⟩φkey, p, q
  L ← {(p, (stid.id, ε)), (q, stid.id)}
  stid.stV ← (P, L)
  Return stid

algorithm Process(stid, m):
  // Bootstrap (attested labels)
  if stid.stage = 0 :
    (i, stid.stV) ← LAC.Verify(stid.prms, (p, (stid.id, ε)), inlast, m, stV)
    If i = ⊥: Return ⊥
    (o, stid.stL) ← s LocKE(stid.stL, i)
    stid.inlast ← o
    If (stid.stL.stKE.δ) = accept : Then stage ← 1
    m' ← (stid.stage, stid.id, o)
    Return (stid, m')

  // Execution (non-attested labels)
  if stid.stage = 1 :
    If m = ε : // Input requested (empty message signal)
      in ← stid.InList[0]
      (in1, ..., ink) ← stid.ListInid
      stid.ListInid ← (in1, ..., ink-1)
      o ← s  $\Lambda$ .Enc(stid.stL.key, (stid.seqin, in))
      stid.inlast ← o
      stid.seqin ← stid.seqin + 2
      m' ← (stid.stage, stid.id, o)
      Return (stid, m')
    Else: // Process received output
      m' ←  $\Lambda$ .Dec(stid.stL.key, m)
      If m' = (stid.seqout, out') :
        stid.seqout ← stid.seqout + 2
        stid.out ← out'
        m' ← (stid.stage, stid.id, ε)
        Return (stid, m')
      Else: Return ⊥

algorithm AddInput(in, stid):
  stid.InList ← stid.InList + [in]
  Return stid

algorithm Output(stid):
  Return stid.out

```

Fig. 10. General SMPC protocol.

- When queried $\text{Send}(\text{id}, m)$ (with id honest), if $\text{st}_{\text{id}}.\text{stage} = 0$ this party is still in the key exchange phase, and \mathcal{S} executes **Process** honestly. Otherwise, if $m = \epsilon$, \mathcal{S} looks-up the length of the next unsent input, encrypts a dummy message of the same length using key_{id} ; it then remembers the result as an undelivered input message and returns it. Note that, since in the second stage we are dealing with non-attested labels, the adversary is expecting this message to be an authenticated encryption (under the derived key) of the correct input for that party. Finally, if $m \neq \epsilon$, \mathcal{S} checks its integrity and whether m corresponds to the next undelivered output message, if so it increments the count of delivered outputs for id , and fails otherwise.
- When queried $\text{Run}(\text{hdl}, l, x)$, if the queried program is not P^* , \mathcal{S} simply executed the program as \mathcal{M} would. If the queried program is P^* and it is still in the bootstrap phase, \mathcal{S} computes the key exchanges honestly as \mathcal{M} would (recording the keys derived by corrupt agents). If the program is not in the bootstrap phase anymore, and the message originates from a corrupt party, \mathcal{S} decrypts the message, performs the relevant checks, queries

```

algorithm RemoteM(prms, hdl, m, stR):
// Initial message
If m = ε :
  m ← (0, 1, ε) // Force bootstrap start
  stR.ldList ← stR // Input schedule
  stR.stage ← 0

// Bootstrap (attested labels)
If stR.stage = 0 :
  (stageid, id, i) ← m
  o ← LAC.AttestM(prms, hdl, (p, (id, ε)), i)
  If o.stage = 1 : // IEE just finished bootstrap
    stR.stage = 1; inreq ← T; m' ← ε
    (id1, . . . , idk) ← stR.ldList; id = id1
    stR.ldList ← (id2, . . . , idk)
    Return (id, inreq, m', stR)
  Else: // Just continue bootstrap
    If stageid = 1 : // This id finished bootstrap
      id ← id + 1
      o ← LAC.AttestM(prms, hdl, (p, (id, ε)), ε)
      inreq ← F; m' ← o
      Return (id, inreq, m', stR)
    Else:
      inreq ← F; m' ← o
      Return (id, inreq, m', stR)

// Execution (non-attested labels)
If stR.stage = 1 :
  (stageid, id, i) ← m
  If i = ε : // Move to next input (empty incoming message)
    (id1, . . . , idk) ← stR.ldList
    inreq ← T; m' ← ε
    id = id1; stR.ldList ← (id2, . . . , idk)
    Return (id, inreq, m', stR)
  Else: // Process input and send output
    o ← M.Run(hdl, (q, id), i)
    inreq ← F; m' ← o
    If ldList = [] : Then stR.stage ← 2 // No additional inputs
    Return (id, inreq, m', stR)

```

Fig. 11. SMPC protocol untrusted scheduler.

Fun for this id and input and returns a proper encryption of the correct output using the correct key for that corrupt party. If the message originates from an honest party, \mathcal{S} checks its integrity and that this was the next undelivered input message for this id (if not it aborts the simulation) and updates the count of delivered inputs. It then queries Fun for this id to get the size of the output and returns the encryption of a dummy encrypted message of the correct length under the fake key key_{id} , recording the result as an undelivered output message. Again, note that, since in the second stage we are dealing with non-attested labels, the adversary is expecting this message to be an authenticated encryption (under the derived key) of the correct output for that party.

We now sketch a proof of indistinguishability between the real world and the ideal world when instantiated with the simulator described above. A detailed proof can be found in Appendix C. The proof is performed in 3 hops, the first of which corresponds to a hybrid argument over the honest parties in the protocol. In this hybrid argument one gradually replaces the secret key derived by each honest party by a random one. In each step, the AttKE utility theorem can be used to show that this change cannot be noticed by the adversary. In the second hop, we replace the encrypted inputs/outputs for honest parties by encrypted dummy payloads of the correct length. This hop is correct by the indistinguishability of authenticated encryption ciphertexts. After this last game hop, the resulting game is *identical until bad* to the ideal world, where the bad event corresponds to the simulator aborting due to an inconsistent message being accepted as the next undelivered input or output. Due to the use of sequence numbers, this bad event can be reduced to the authenticity of the encryption scheme and the Theorem follows.

8 Implementation and Evaluation

We experimented with two implementations of our protocol—`sgx-mpc-mbed` and `sgx-mpc-nacl`— which differ in the underlying cryptography: `sgx-mpc-mbed` relies on the mbed TLS (formerly PolarSSL) library and `sgx-mpc-nacl` which relies on the NaCl. Furthermore, `sgx-mpc-mbed` uses standard RSA technology for the key exchange stage, and an AES128-CTR and HMAC-SHA256 Encrypt-Then-Mac construction for authenticated encryption, whereas `sgx-mpc-nacl` uses elliptic-curves both for key exchange (Diffie-Hellman) and digital signatures, and a combination of the Salsa20 and Poly1305 encryption and authentication schemes [8] for authenticated encryption.

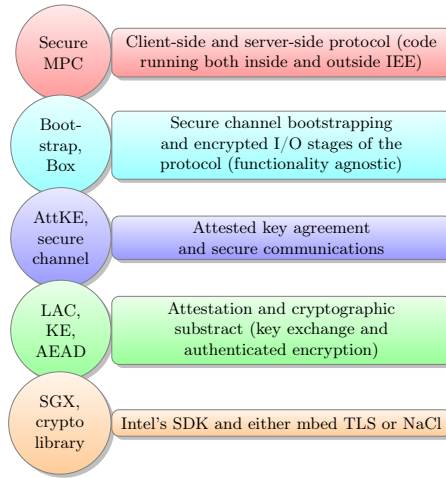


Fig. 12. Bird’s eye view of our implementations.

Both implementations rely on Intel’s Software Development Kit (SDK) for dealing with the SGX low-level operations. These include loading a piece of code into an IEE (our Load abstraction), calling a top-level function within the IEE (our Run abstraction), and constructing an attested message (first getting a MAC’ed message within the IEE, and then using the quoting enclave to convert it into a digital signature). Furthermore, both implementations build on top of this cryptographic underpinning and share the structure in Figure 12. They employ the LAC scheme proposed in this paper, and include wrappers that match our abstractions of digital signatures and authenticated encryption. These are then used to construct the secure bootstrapping protocol (AttKE) that enables each party to establish an independent secret key, and a secure channel that uses this key to communicate with the **Box** construction running inside the enclave. Finally, both implementations of the **Box** are agnostic of the functionality that should be computed by the protocol, and can be linked to arbitrary functionality implementations, provided that these comply with a simple labelled I/O interface. The top-level interface to our protocol includes the code that should be run inside the IEE, the code that runs outside the IEE in the remote machine to perform the book-keeping operations and the client-side code that permits bootstrapping a secure channel and then send inputs/receive outputs from the functionality.

Below, we present the experimental evaluation of both implementations. We first describe our methodology and some micro-benchmarks, which are SGX-specific and not really intrinsic to our protocol. Afterwards, we discuss the issue of side channel attacks, and assess the cost of adopting more stringent countermeasures against timing attacks. Finally, we compare our protocol with a non hardware-based state-of-the-art secure two-party computation framework.

8.1 Evaluation methodology and micro-benchmarks

We performed the evaluation on a SGX-enabled platform, equipped with an Intel Core i7-6700 processor (3.4 MHz) and 8 GB DDR4 RAM, running 64 bit Ubuntu 14.04. Performing fine grained performance measurements for SGX is a challenging task due to the lack of methods to directly measure the runtime of subcomponents *inside* an enclave, leading to noisy results. To eliminate the noise, or at least minimize its effect, we repeated each measurement 100 times and report the average value.

MICRO-BENCHMARKS. Our first results are protocol agnostic, and show the SGX-specific inherent overhead of creating an enclave, invoking an enclave, generating MAC-based reports and signature generation by the quoting enclave (QE). Table 1 lists the runtime of those components with fixed run times. The invocation of an enclave cannot be measured on its own. We measured the time to enter an enclave and immediately returning, specified as *SGX context switch* in Table 1. However, this measurement comprises both the time required for entering the enclave, as well as the time required for leaving the enclave.

The creation time of an SGX enclave is dependent on its size. The initial enclave state must be copied into the enclave memory area; the execution time of this operation is, as would be expected, linear in the size of the enclave. Additionally, the initial state of the enclave is “measured” (i.e., hashed with SHA-256) during enclave creation.

Table 1. SGX Micro-Benchmarks

Component	Time (ms)
SGX context switch	0.0054
MAC-based report (ERPORT)	0.2887
QE signatur	23.930

Again, this operation takes an amount of time that is linear in the size of the enclave.¹⁶ We measured the creation time of an enclave to be $202.55\mu s + S \cdot 0.72\mu s$ with enclave size S in KB.

8.2 Side channels and software resilient against timing attacks

Recent works [39,16] have pointed out that IEE-enabled systems such as Intel’s SGX do not offer more protection against side-channel attacks than traditional microprocessors. This is a relevant concern, since the IEE trust model which we also adopt in this paper admits that the code outside IEEs is potentially malicious and that the machine is under the control of an untrusted party. We believe that there are two aspects to this problem that should be considered separately. The first aspect is related to the production of the IEE-enabled hardware/firmware itself and the protection of the long-term secrets that are used by the attestation security module. If the computations performed by the attestation infrastructure itself are vulnerable to side-channel attacks, then there is nothing that can be done at the protocol design/implementation level. This aspect of trust is within the remit of the equipment manufacturers.

An orthogonal issue is the possibility that software running inside an IEE leaks part of its state or short-term secrets via side channels. Here one should distinguish between software observations and hardware/physical observations. In the former, software co-located in the machine observes timing channels based on memory access patterns, control flow, branch prediction, cache-based based attacks [16], page-fault side channels [39], etc. Protection against this type of side-channel attacks has been widely studied in the practical crypto community, where a consensus exists that writing so-called *constant-time* software is the most effective countermeasure [7,31]. As mentioned above, constant-time software has the property that the entire sequence of memory addresses (in both data and code memory) accessed by a program can be predicted in advance from public inputs, e.g., the length of messages. When it comes to hardware/physical side-channel attacks such as those relying on temperature measurements, power analysis, or electromagnetic radiation, the effectiveness of software countermeasures is very limited, and improving hardware defenses again implies obtaining additional guarantees from the equipment manufacturer.

For clarity, we recall that our two implementations differ in the way they deal with timing channels significantly: while `sgx-mpc-nacl` enforces a strict constant-time policy that is consistent with the IEE trust model, `sgx-mpc-mbed` relies on a standard TLS implementation that was not designed to deal with attacks by an adversary co-located in the same machine. None of the implementations deploy countermeasures against hardware/physical side-channel attacks.

COUNTERMEASURES AGAINST TIMING ATTACKS. To assess the practical impact of assuming a stronger attack model in which the remote party may launch timing attacks against implementations running inside enclaves, we have compared the two implementations of our protocol. We evaluated the performance of the individual steps in the protocol, namely the key exchange between IEE with an input party, and the `Box` component which decrypts all inputs and encrypts all outputs in stage two of the protocol.¹⁷

Table 2 lists the measurements for the individual components of our two implementations, `sgx-mpc-nacl` and `sgx-mpc-mbed`. The reported time comprises the key exchange with *one* input party, and so the overhead will accumulate linearly with an increasing number of input parties. Interestingly, the `Box` components of our `sgx-mpc-nacl` implementation is faster than in our `sgx-mpc-mbed` implementation, which shows that highly optimized constant-time software (deploying cryptographic primitives designed specifically for this purpose) can be faster than their non-constant-time counterparts, as argued for example in [8]. In the key exchange stage, the better performance of RSA public-key operations (encryption and signature verification) gives `sgx-mpc-mbed` an advantage.

In the remainder of the section, when we compare our protocol with previous solutions, we will report the evaluation results for our faster implementation `sgx-mpc-mbed`. The overhead of `sgx-mpc-nacl` over `sgx-mpc-mbed`

¹⁶ SGX does not require the measurement of the entire initial memory state of an enclave. However, the SDK’s default behaviour is to measure the entire enclave (code and data) which is what is required for our protocol.

¹⁷ We note that, when such an attacker is considered, not only the cryptographic components must be implemented following the constant-time coding policies, but also the code that implements the functionality itself (!), and so an additional penalty may be paid in addition to the overhead we report here.

Table 2. LAC Components Benchmarks

Component	sgx-mpc-mbed (<i>ms</i>)	sgx-mpc-nacl (<i>ms</i>)
Key exchange (Stage 1)	35.17	127.6
Box (Stage 2)	0.036	0.012

occurs mostly in the key exchange phase, this means that from the execution times of `sgx-mpc-nacl` can easily be inferred from our `sgx-mpc-mbed` results.

8.3 Comparative analysis with state-of-the-art MPC protocols

We compare our implementation with measurements we performed using the ABY framework [18]. We chose ABY for comparison, as we could evaluate it on the same platform we used for assessing our protocol, therefore avoiding differences due to performance disparities of heterogeneous evaluation platforms. Although it is specific to the two-party secure computation setting, ABY is representative of state-of-the-art MPC implementations and we expect results for other frameworks such as Sharemind¹⁸ and SPDZ [17] to lead to similar conclusions; indeed, the crux of our performance gains resides in the fact that our solution does not require encoding the computation in circuit form, which happens in one form or another for all of the aforementioned protocols.¹⁹

We evaluated the performance of four different secure two-party computation use cases: minimum, Hamming distance, private set intersection, and AES. Like our protocol, the ABY protocol also has two phases: a *preparation phase* and an *online phase*. The preparation phase comprises the key exchange between the input parties by means of oblivious transfer (OT), and the generation of the garbled circuit (GC) representing the desired function. In the online phase the GC gets evaluated and the result are send back to the output party. In our protocol, the preparation phase is used to establish a secure channel between the IEE and the input parties. The online phase of our protocol comprises the decryption of inputs in the `Box` component, the evaluation of the payload function, and the encryption of the results, again by the `Box` component.

Determination of minimum Table 3 shows the performance of the two stages of ABY and `sgx-mpc-mbed` for determining the minimum of two inputs, 32 *bits* each. For both phases the runtime of `sgx-mpc-mbed` is shorter than ABY’s runtime; in both phases `sgx-mpc-mbed` is about 5.6 times faster than ABY.

Table 3. Minimum of two inputs

Component	ABY (<i>ms</i>)	Ours (<i>ms</i>)
Preparation	196.3	35.17
Online	0.404	0.071

Hamming Distance Next, we compared ABY and `sgx-mpc-mbed` in computing the Hamming Distance of two inputs. We evaluated the performance for different input sizes to demonstrate the scaling behaviour for the different solutions, detailed in Table 4.

Table 4. Hamming distance with different input sizes

	Phase	Preparation (<i>ms</i>)		Online (<i>ms</i>)		Total (<i>ms</i>)	
		Protocol	ABY	Ours	ABY	Ours	ABY
<i>Input size (bits)</i>	160	196.3	35.15	0.752	0.072	197.1	35.22
	1600	196.7	35.12	1.819	0.080	198.5	35.20
	16000	201.6	35.20	13.14	0.165	214.7	35.37
	160000	226.2	35.12	144.4	1.037	370.6	36.16

¹⁸ <https://sharemind.cyber.ee/>

¹⁹ We also note that ABY assumes a semi-honest adversary, which is weaker than the one we consider; but still our performance gains are significant.

Again, the preparation phase as well as the online phase of `sgx-mpc-mbed` is faster, compared to ABY. Additionally, the performance in ABY degrades faster with increasing input sizes than that of our protocol.

Private set intersection Table 5 lists our results for private set intersection. In contrast to the previous use cases, the runtime of ABY’s preparation phase is increasing with the size of the input data, while `sgx-mpc-mbed`’s preparation phase remains constant (modulo some measurement inaccuracies). Additionally, the growth in runtime of ABY’s online phase is much stronger compared to `sgx-mpc-mbed`.

Table 5. Private set intersection with different set sizes

	Phase	Preparation		Online		Total	
	Protocol	ABY	Ours	ABY	Ours	ABY	Ours
<i>Set size</i>	100	224.8	35.22	1.084	0.098	225.9	35.32
	1000	368.1	35.23	2.168	0.368	370.3	35.60
	10,000	1442.2	35.19	12.88	3.454	1455.1	38.64
	100,000	10,698.7	35.19	109.5	36.14	10,808.2	71.33
	1,000,000	84,096.6	35.22	1616.0	385.4	85,712.6	420.6

AES As for the previous three use cases, we evaluated the secure multi-party computation of AES with ABY as well as `sgx-mpc-mbed`; the results are in Table 6. AES has become a standard for evaluating MPC protocols, hence, evaluation results for AES are available in a number of related works.

Table 6. AES – 128 bit key and 128 bit block size

Phase	ABY	Ours
<i>Preparation (ms)</i>	197.9	35.10
<i>Online (ms)</i>	3.249	0.093
<i>Total (ms)</i>	201.1	35.19

In comparison to ABY, the preparation phase and online phase are shorter with `sgx-mpc-mbed`, and consequently the overall runtime is faster as well.

9 Conclusion

The main theoretical contribution of this paper is the concept of Labelled Attested Computation, a cryptographic primitive that permits reasoning about the security properties of IEE-enabled systems in higher level protocols. Building on LAC we rigorously prove the security under a strong adversarial model of a highly efficient MPC protocol. Furthermore, we experimentally show that for two-party functionalities our protocol is considerably faster than the state of the art protocol.

In future work, it would be interesting to expand on the current evaluation to explore the limits of SGX both from the perspective of practical resilience to side-channel attacks, scalability in terms of the amount of code and data that can be loaded into an enclave, scalability in terms of the number of parties that are involved, and how to deal with the inherent limitations in a secure way. Achieving security against even stronger adversaries, e.g., considering adaptive corruptions in a way that avoids heavy cryptographic tools like non-committing encryption is also an interesting direction for future work.

References

1. I. Anati, S. Gueron, S. Johnson, and V. Scarlata. Innovative technology for cpu based attestation and sealing. In *HASP 2013*, page 10, 2013.
2. M. Barbosa, B. Portela, G. Scerri, and B. Warinschi. Foundations of hardware-based attested computation and application to SGX. In *EuroS&P*, pages 245–260. IEEE, 2016.

3. M. Barbosa, B. Portela, G. Scerri, and B. Warinschi. Foundations of hardware-based attested computation and application to SGX. *IACR Cryptology ePrint Archive*, 2016:14, 2016.
4. A. Baumann, M. Peinado, and G. C. Hunt. Shielding applications from an untrusted cloud with haven. In *OSDI*, pages 267–283. USENIX Association, 2014.
5. M. Bellare, J. Kilian, and P. Rogaway. The security of cipher block chaining. In *CRYPTO*, volume 839 of *Lecture Notes in Computer Science*, pages 341–358. Springer, 1994.
6. A. Ben-David, N. Nisan, and B. Pinkas. Fairplaymp: a system for secure multi-party computation. In *ACM Conference on Computer and Communications Security*, pages 257–266. ACM, 2008.
7. D. J. Bernstein. Cache-timing attacks on aes, 2005. <http://cr.yp.to/antiforgery/cachetiming-20050414.pdf>.
8. D. J. Bernstein, T. Lange, and P. Schwabe. The security impact of a new cryptographic library. In *LATINCRYPT*, volume 7533 of *Lecture Notes in Computer Science*, pages 159–176. Springer, 2012.
9. D. Bogdanov, S. Laur, and J. Willemson. Sharemind: A framework for fast privacy-preserving computations. In *ESORICS*, volume 5283 of *Lecture Notes in Computer Science*, pages 192–206. Springer, 2008.
10. E. Brickell, L. Chen, and J. Li. A new direct anonymous attestation scheme from bilinear maps. In *TRUST*, volume 4968 of *Lecture Notes in Computer Science*, pages 166–178. Springer, 2008.
11. E. F. Brickell, J. Camenisch, and L. Chen. Direct anonymous attestation. In *ACM Conference on Computer and Communications Security*, pages 132–145. ACM, 2004.
12. R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally composable security with global setup. In *TCC*, volume 4392 of *Lecture Notes in Computer Science*, pages 61–85. Springer, 2007.
13. R. Canetti and M. Fischlin. Universally composable commitments. In *CRYPTO*, volume 2139 of *Lecture Notes in Computer Science*, pages 19–40. Springer, 2001.
14. R. Canetti, Y. Lindell, R. Ostrovsky, and A. Sahai. Universally composable two-party and multi-party secure computation. In *STOC*, pages 494–503. ACM, 2002.
15. L. Catuogno, A. Dmitrienko, K. Eriksson, D. Kuhlmann, G. Ramunno, A. Sadeghi, S. Schulz, M. Schunter, M. Winandy, and J. Zhan. Trusted virtual domains - design, implementation and lessons learned. In *INTRUST*, volume 6163 of *Lecture Notes in Computer Science*, pages 156–179. Springer, 2009.
16. V. Costan and S. Devadas. Intel SGX explained, 2016.
17. I. Damgård, V. Pastro, N. P. Smart, and S. Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, volume 7417 of *Lecture Notes in Computer Science*, pages 643–662. Springer, 2012.
18. D. Demmler, T. Schneider, and M. Zohner. ABY - A framework for efficient mixed-protocol secure two-party computation. In *NDSS*. The Internet Society, 2015.
19. A. Francillon, Q. Nguyen, K. B. Rasmussen, and G. Tsudik. A minimalist approach to remote attestation. In *DATE*, pages 1–6. European Design and Automation Association, 2014.
20. H. Ge and S. R. Tate. A direct anonymous attestation scheme for embedded devices. In *Public Key Cryptography*, volume 4450 of *Lecture Notes in Computer Science*, pages 16–30. Springer, 2007.
21. C. Gebhardt and A. Tomlinson. Secure virtual disk images for grid computing. In *APTC '08*, pages 19–29. IEEE Computer Society, 2008.
22. S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.*, 17(2):281–308, 1988.
23. D. Gupta, B. Mood, J. Feigenbaum, K. R. B. Butler, and P. Traynor. Using intel software guard extensions for efficient two-party secure function evaluation. In *FC Workshop on Encrypted Computing and Applied Homomorphic Cryptography*, volume 9604 of *Lecture Notes in Computer Science*, pages 302–318. Springer, 2016.
24. S. Halevi, Y. Lindell, and B. Pinkas. Secure computation on the web: Computing without simultaneous interaction. In *CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 132–150. Springer, 2011.
25. W. Henecka, S. Kögl, A. Sadeghi, T. Schneider, and I. Wehrenberg. TASTY: tool for automating secure two-party computations. In *ACM Conference on Computer and Communications Security*, pages 451–462. ACM, 2010.
26. M. Hoekstra, R. Lal, P. Pappachan, V. Phegade, and J. del Cuvillo. Using innovative instructions to create trustworthy software solutions. In *HASP@ISCA*, page 11. ACM, 2013.
27. Intel. *Software Guard Extensions Programming Reference*, 2014. <https://software.intel.com/sites/default/files/managed/48/88/329298-002.pdf>.
28. J. Katz. Universally composable multi-party computation using tamper-proof hardware. In *EUROCRYPT*, volume 4515 of *Lecture Notes in Computer Science*, pages 115–128. Springer, 2007.
29. J. Katz and M. Yung. Unforgeable encryption and chosen ciphertext secure modes of operation. In *FSE*, volume 1978 of *Lecture Notes in Computer Science*, pages 284–299. Springer, 2000.
30. P. Koeberl, S. Schulz, A. Sadeghi, and V. Varadharajan. Trustlite: a security architecture for tiny embedded devices. In *EuroSys*, pages 10:1–10:14. ACM, 2014.
31. A. Langley. Lucky thirteen attack on TLS CBC. Imperial Violet, Feb. 2013. <https://www.imperialviolet.org/2013/02/04/luckythirteen.html>, Accessed October 25th, 2015.
32. D. Malkhi, N. Nisan, B. Pinkas, and Y. Sella. Fairplay - secure two-party computation system. In *USENIX Security Symposium*, pages 287–302. USENIX, 2004.
33. J. M. McCune, B. Parno, A. Perrig, M. K. Reiter, and H. Isozaki. Flicker: an execution infrastructure for tcb minimization. In *EuroSys*, pages 315–328. ACM, 2008.

34. Microsoft. *BitLocker Drive Encryption: Data Encryption Toolkit for Mobile PCs: Security Analysis*, 2007. <https://technet.microsoft.com/en-us/library/cc162804.aspx>.
35. J. Noorman, P. Agten, W. Daniels, R. Strackx, A. V. Herrewewege, C. Huygens, B. Preneel, I. Verbauwhede, and F. Piessens. Sancus: Low-cost trustworthy extensible networked devices with a zero-software trusted computing base. In *USENIX Security Symposium*, pages 479–494. USENIX Association, 2013.
36. R. Pass, E. Shi, and F. Tramer. Formal abstractions for attested execution secure processors. Cryptology ePrint Archive, Report 2016/1027, 2016. <http://eprint.iacr.org/2016/1027>.
37. F. Schuster, M. Costa, C. Fournet, C. Gkantsidis, M. Peinado, G. Mainar-Ruiz, and M. Russinovich. VC3: trustworthy data analytics in the cloud using SGX. In *IEEE Symposium on Security and Privacy*, pages 38–54. IEEE Computer Society, 2015.
38. B. Smyth, M. Ryan, and L. Chen. Direct anonymous attestation (DAA): ensuring privacy with corrupt administrators. In *ESAS*, volume 4572 of *Lecture Notes in Computer Science*, pages 218–231. Springer, 2007.
39. Y. Xu, W. Cui, and M. Peinado. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *IEEE Symposium on Security and Privacy*, pages 640–656. IEEE Computer Society, 2015.

A Proof of Theorem 1

The proof is a sequence of three games presented in Figure 13 and Figure 14. The first game is simply the LAC security game instantiated with our protocol.

In game $G_1^{\text{LAC}, \mathcal{A}}(1^\lambda)$, the adversary loses whenever a `sforge` event occurs. Intuitively, this event corresponds to the adversary producing a signature that was not computed by the signing process with handle 0, and hence constitutes a forgery with respect to Σ . Given that the two games are identical until this event occurs, we have that

$$\Pr[\text{Att}^{\text{LAC}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] - \Pr[G_1^{\text{LAC}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] \leq \Pr[\text{sforge}].$$

<p>Game $G_0^{\text{LAC}, \mathcal{A}}(1^\lambda)$:</p> <pre> prms \leftarrow $\mathcal{M}.\text{Init}(1^\lambda)$ $(P, L^*, l, n, \text{st}_A) \leftarrow$ $\mathcal{A}_1(\text{prms})$ $P^* \leftarrow \text{Compile}(\text{prms}, P, L^*)$ For $k \in [1..n]$: $(i_k, o_k^*, \text{st}_A) \leftarrow$ $\mathcal{A}_2^{\mathcal{M}}(\text{st}_V, \text{st}_A)$ Parse $(o_k, \sigma) \leftarrow o_k^*$ If $\Sigma.\text{Vrfy}(\text{prms}, \sigma, (P^*, \text{filter}[l]((l, i_k, o_k) : \text{ios})))$: ios $\leftarrow ((l, i_k, o_k) : \text{ios})$ Else: Return F $T \leftarrow (l, i_1, o_1, \dots, l, i_n, o_n)$ For hdl* s.t. $\text{Program}_{\mathcal{M}}(\text{hdl}^*) = P^*$ $(l'_1, i'_1, o'_1, \dots, l'_m, i'_m, o'_m) \leftarrow \text{Trace}_{\mathcal{M}_R}(\text{hdl}^*)$ $T' \leftarrow \text{filter}[l](\text{Trace}_{P[\text{st}; \text{Coins}_{\mathcal{M}}]}(\text{hdl}^*)) (l'_1, i'_1, \dots, l'_m, i'_m)$ If $T \sqsubseteq T'$ Return F Return T</pre>	<p>Game $G_1^{\text{LAC}, \mathcal{A}}(1^\lambda)$:</p> <pre> prms \leftarrow $\mathcal{M}.\text{Init}(1^\lambda)$ $(P, L^*, l, n, \text{st}_A) \leftarrow$ $\mathcal{A}_1(\text{prms})$ sforge \leftarrow F $P^* \leftarrow \text{Compile}(\text{prms}, P, L^*)$ For $k \in [1..n]$: $(i_k, o_k^*, \text{st}_A) \leftarrow$ $\mathcal{A}_2^{\mathcal{M}}(\text{st}_V, \text{st}_A)$ Parse $(o_k, \sigma) \leftarrow o_k^*$ If $\Sigma.\text{Vrfy}(\text{prms}, \sigma, (P^*, \text{filter}[l]((l, i_k, o_k) : \text{ios})))$: ios $\leftarrow ((l, i_k, o_k) : \text{ios})$ Else: Return F If $((P^*, (l, i_1, o_1, \dots, l, i_k, o_k), \star), \sigma') \notin \text{Trace}_{\mathcal{M}}(0)$: sforge \leftarrow T; Return F $T \leftarrow (l, i_1, o_1, \dots, l, i_n, o_n)$ For hdl* s.t. $\text{Program}_{\mathcal{M}}(\text{hdl}^*) = P^*$ $(l'_1, i'_1, o'_1, \dots, l'_m, i'_m, o'_m) \leftarrow \text{Trace}_{\mathcal{M}_R}(\text{hdl}^*)$ $T' \leftarrow \text{filter}[l](\text{Trace}_{P[\text{st}; \text{Coins}_{\mathcal{M}}]}(\text{hdl}^*)) (l'_1, i'_1, \dots, l'_m, i'_m)$ If $T \sqsubseteq T'$ Return F Return T</pre>
---	--

Fig. 13. First game hop for the proof of security of our AC protocol.

We upper bound the distance between these two games, by constructing an adversary \mathcal{B} against the existential unforgeability of signature scheme Σ in S^* such that

$$\Pr[\text{sforge}] \leq \text{Adv}_{\Sigma, \mathcal{B}}^{\text{UF}}(\lambda)$$

Adversary \mathcal{B} simulates the environment of $G_1^{\text{LAC}, \mathcal{A}}$ as follows: the operation of machine \mathcal{M} is simulated exactly with the caveat that the signing operations performed within the process loaded by the security module are replaced with calls to the `Sign` oracle provided in the existential unforgeability game. More precisely, whenever process 0 in the remote machine is expected to compute a signature on message m , algorithm \mathcal{B} calls its own oracle on (P^*, m) to obtain σ .

When `sforge` is set, according to the rules of game $G_1^{\text{LAC}, \mathcal{A}}$, algorithm \mathcal{B} outputs message $(P^*, \text{filter}[l](\text{ios}))$ and candidate signature σ . It remains to show that this is a valid forgery. To see this, first observe that this is indeed a valid signature, as signature verification is performed on these values immediately before `sforge` occurs. It suffices to establish that message $(P^*, (l, i_1, o_1, \dots, l, i_k, o_k))$ could not have been queried from the `Sign` oracle. Access

to the signing key that allows signatures to be performed is only permitted to the special process with handle 0. From the construction of S^* , we know that producing such a signature would only occur via the inclusion of $(P^*, (l, i_1, o_1, \dots, l, i_k, o_k))$ in its trace. Since we know that this is not the case, $(P^*, \text{filter}[l](\text{ios}))$ could not have been queried from the signature oracle. We conclude therefore that \mathcal{B} outputs a valid forgery whenever **sforge** occurs.

In game $G_2^{\text{LAC}, \mathcal{A}}(1^\lambda)$, the adversary loses whenever a **mforge** event occurs. Intuitively, this event corresponds to the adversary producing a tag that was not computed by the security module, and hence constitutes a forgery with respect to Π . Given that the two games are identical until this event occurs, we have that

$$\Pr[G_1^{\text{LAC}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] - \Pr[G_2^{\text{LAC}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] \leq \Pr[\text{mforge}].$$

Game $G_1^{\text{LAC}, \mathcal{A}}(1^\lambda)$:	Game $G_2^{\text{LAC}, \mathcal{A}}(1^\lambda)$:
<pre> prms \leftarrow $\mathcal{M}.\text{Init}(1^\lambda)$ $(P, L^*, l, n, \text{st}_A) \leftarrow$ $\mathcal{A}_1(\text{prms})$ sforge \leftarrow F $P^* \leftarrow$ Compile(prms, P, L^*) For $k \in [1..n]$: $(i_k, o_k^*, \text{st}_A) \leftarrow$ $\mathcal{A}_2^{\mathcal{M}}(\text{stv}, \text{st}_A)$ Parse $(o_k, \sigma) \leftarrow o_k^*$ If $\Sigma.\text{Vrfy}(\text{prms}, \sigma, (P^*, \text{filter}[l]((l, i_k, o_k) : \text{ios})))$: ios \leftarrow $((l, i_k, o_k) : \text{ios})$ Else: Return F If $((P^*, (l, i_1, o_1, \dots, l, i_k, o_k), \star), \sigma') \notin \text{Trace}_{\mathcal{M}}(0)$: sforge \leftarrow T; Return F T \leftarrow $(l, i_1, o_1, \dots, l, i_n, o_n)$ For hdl* s.t. Program$_{\mathcal{M}}(\text{hdl}^*) = P^*$ $(l'_1, i'_1, o'_1, \dots, l'_m, i'_m, o'_m) \leftarrow$ Trace$_{\mathcal{M}_R}(\text{hdl}^*)$ $T' \leftarrow$ filter$[l](\text{Trace}_{P[\text{st}; \text{Coins}_{\mathcal{M}}(\text{hdl}^*)]}(l'_1, i'_1, \dots, l'_m, i'_m))$ If $T \sqsubseteq T'$ Return F Return $\bar{\text{T}}$ </pre>	<pre> prms \leftarrow $\mathcal{M}.\text{Init}(1^\lambda)$ $(P, L^*, l, n, \text{st}_A) \leftarrow$ $\mathcal{A}_1(\text{prms})$ sforge \leftarrow F; mforge \leftarrow F $P^* \leftarrow$ Compile(prms, P, L^*) For $k \in [1..n]$: $(i_k, o_k^*, \text{st}_A) \leftarrow$ $\mathcal{A}_2^{\mathcal{M}}(\text{stv}, \text{st}_A)$ Parse $(o_k, \sigma) \leftarrow o_k^*$ If $\Sigma.\text{Vrfy}(\text{prms}, \sigma, (P^*, \text{filter}[l]((l, i_k, o_k) : \text{ios})))$: ios \leftarrow $((l, i_k, o_k) : \text{ios})$ Else: Return F If $((P^*, (l, i_1, o_1, \dots, l, i_k, o_k), \star), \sigma') \notin \text{Trace}_{\mathcal{M}}(0)$: sforge \leftarrow T; Return F If $\exists \text{hdl}^*. \text{Program}_{\mathcal{M}}(\text{hdl}^*) = P^* \wedge$ $(l, i_1, o_1, \dots, l, i_k, o_k) \sqsubseteq \text{filter}[l](\text{Trace}_{P[\text{st}; \text{Coins}_{\mathcal{M}}(\text{hdl}^*)]}(\text{Trace}_{\mathcal{M}}(\text{hdl}^*)))$: Then mforge \leftarrow T; Return F T \leftarrow $(l, i_1, o_1, \dots, l, i_n, o_n)$ For hdl* s.t. Program$_{\mathcal{M}}(\text{hdl}^*) = P^*$ $(l'_1, i'_1, o'_1, \dots, l'_m, i'_m, o'_m) \leftarrow$ Trace$_{\mathcal{M}_R}(\text{hdl}^*)$ $T' \leftarrow$ filter$[l](\text{Trace}_{P[\text{st}; \text{Coins}_{\mathcal{M}}(\text{hdl}^*)]}(l'_1, i'_1, \dots, l'_m, i'_m))$ If $T \sqsubseteq T'$ Return F Return $\bar{\text{T}}$ </pre>

Fig. 14. Second game hop for the proof of security of our AC protocol.

We upper bound the distance between these two games, by constructing an adversary \mathcal{C} against the existential unforgeability of MAC scheme Π in the security module such that

$$\Pr[\text{mforge}] \leq \text{Adv}_{\Pi, \mathcal{C}}^{\text{Auth}}(\lambda)$$

Adversary \mathcal{C} simulates the environment of $G_2^{\text{LAC}, \mathcal{A}}$ as follows: the operation of machine \mathcal{M} is simulated exactly with the caveat that the MAC operations computed inside the internal security module are replaced with calls to the Auth oracle provided in the existential unforgeability game. More precisely, whenever a process running code R^* within an IEE in the remote machine requests a MAC on message m from the security module, algorithm \mathcal{C} calls its own oracle on (P^*, m) to obtain t .

Let $T \leftarrow (l, i_1, o_1, \dots, l, i_k, o_k)$. When **mforge** is set according to the rules of game $G_2^{\text{LAC}, \mathcal{A}}$, algorithm \mathcal{C} retrieves the trace of the process with handle 0 running S^* , locates the input/output pair $((P^*, T), t, \sigma')$ and outputs message (P^*, T) and candidate tag t . To see this is a valid forgery, first observe that, having failed the **sforge** check, we know that $((P^*, T), t, \sigma')$ is in the trace of the process with handle 0, so by its construction we also know that the corresponding input $((P^*, T), t)$ must contain a valid tag. It suffices to establish that message (P^*, T) could not have been queried from the Auth oracle. Suppose that the first part of the **mforge** check failed, i.e., that $\nexists \text{hdl}^*. \text{Program}_{\mathcal{M}}(\text{hdl}^*) = P^*$. Then, because the security module signs the code of the processes requesting the signatures, we are sure that such a query was never placed to the Auth oracle. Furthermore, any MAC query for a message starting with P^* must have been caused by the execution of an instance of P^* . Now suppose some instances of P^* were indeed running in the remote machine, but that none of them displayed the property $(l, i_1, o_1, \dots, l, i_k, o_k) \sqsubseteq \text{filter}[l](\text{Translate}(\text{ATrace}_{\mathcal{M}}(\text{hdl}^*)))$. Then, by the construction of P^* , we can also exclude that (P^*, T) was queried from the MAC oracle. As such, we conclude that \mathcal{C} outputs a valid forgery whenever **mforge** occurs.

To complete the proof, we argue that the adversary never wins in game $G_2^{\text{LAC}, \mathcal{A}}$. To see this, observe that when the game reaches the final check, we have the guarantee that

$$\exists \text{hdl}^*. \text{Program}_{\mathcal{M}}(\text{hdl}^*) = P^* \wedge \\ (l, i_1, o_1, \dots, l, i_n, o_n) \sqsubseteq \text{filter}[l](\text{Trace}_{P[\text{st}; \text{Coins}_{\mathcal{M}}]}(\text{Trace}_{\mathcal{M}}(\text{hdl}^*)))$$

Which exactly matches the final criteria of $T \sqsubseteq T'$.

To finish the proof, we must now show that this scheme also provides security with minimum leakage. This implies defining a ppt simulator \mathcal{S} that provides identical distributions with respect to experiment in Figure 5. This is easy to ascertain given the simulator behaviour described in Figure 15: \mathcal{S}_1 and \mathcal{S}_3 follow the exact description of the actual machine, modulo the generation of (pk, sk) and key . \mathcal{S}_2 takes an external output produced by $P[\text{st}](l, i)$ and returns an output in accordance to the behaviour of \mathcal{M} , which given our language \mathcal{L} may differ from a real output only by the random coins. As such, the distribution provided by the simulator is indistinguishable to the one provided by a real machine, and our claim follows. \square

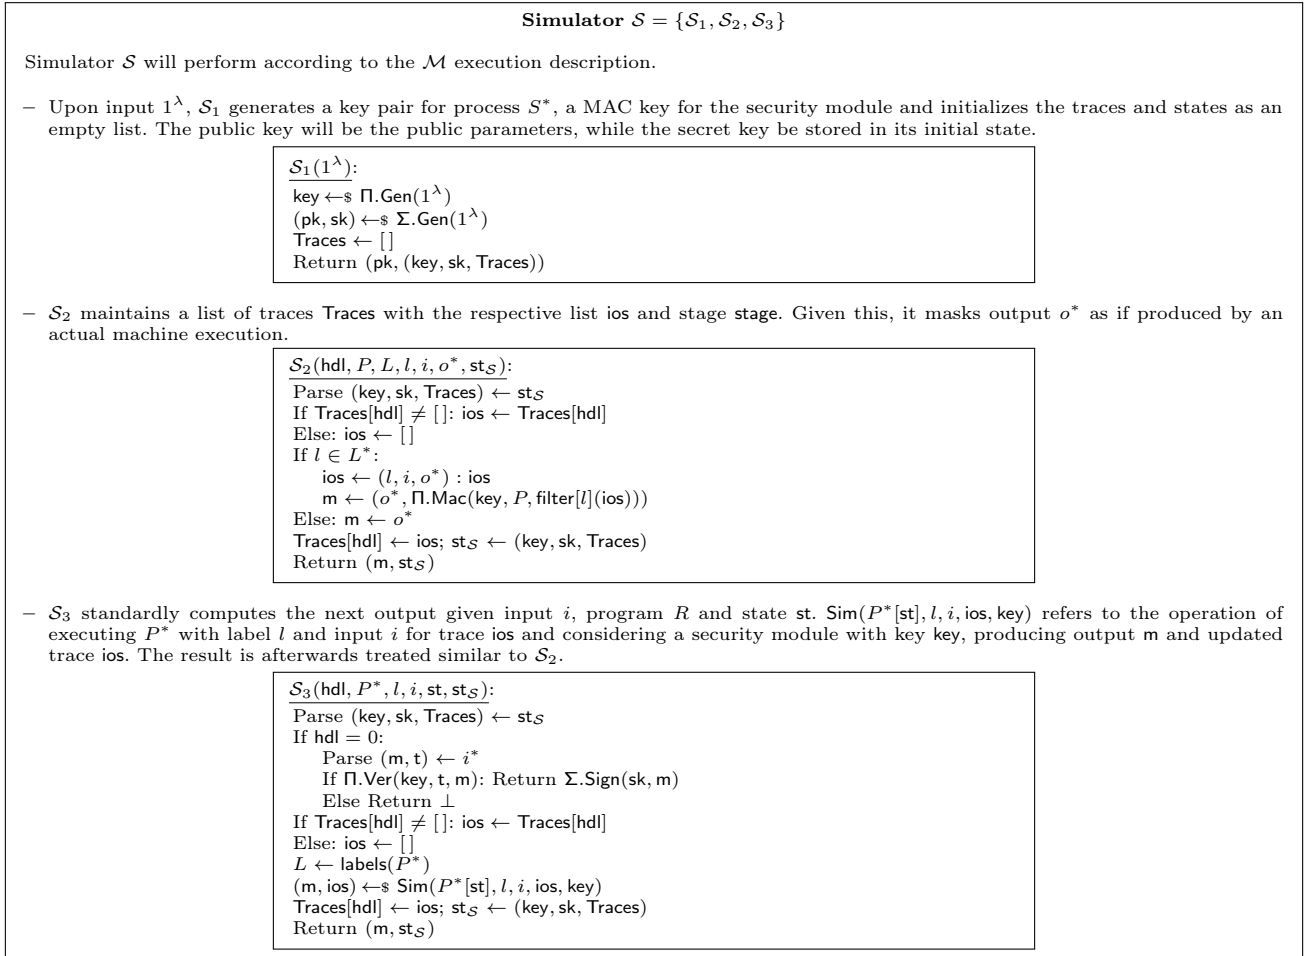


Fig. 15. Description of simulator \mathcal{S}

B Proof of the Utility Theorem

In [3] the proof of utility, presented in Appendix C consists in a sequence of 4 games $G0_{\text{AttKE}, \mathcal{A}}$ to $G3_{\text{AttKE}, \mathcal{A}}$. The first hop removes the possibility of an AC forgery, thus ensuring that messages are delivered properly from the remote program executing the key exchange to the local party. The second hop, using minimal leakage, replaces

execution in the remote machine by a simulated execution based on a real execution of the compiled program. The final game hop replaces the execution of the remote key exchange by call to the AttKE security oracles, and conclude immediately by security of the AttKE. Due to the similarities of both proofs, we do not rewrite the whole proof, instead we highlight the differences coming from our different notion of attestation and our more general composition pattern.

Our proof is a sequence of four games $G0'_{\text{AttKE},\mathcal{A}}$ to $G3'_{\text{AttKE},\mathcal{A}}$. We highlight here how the sequence of games is constructed in relation with the original utility proof from [2].

FIRST GAME HOP. The first game $G0'_{\text{AttKE},\mathcal{A}}$ is the utility game presented in Figure 8. In [2] the first hop consists in adding a `forgeAC` event in the `Send` oracle to ensure that the initial segment of the trace witnessed by a local party matches the initial segment of a valid execution of the distant protocol. Here, similarly we add a `forgeLAC` event, making sure that the subtrace corresponding to the appropriate label matches a remote execution, the `Send` oracle is replaced by the one described in Figure 16.

Oracle $\text{Send}(o^*, i)$:

$o \leftarrow \text{LAC.Verify}[\text{st}_V^i](\text{prms}, l^i, \text{in}_{\text{last}}^i, o^*)$
 If $o = \perp$: Return \perp
 If $b = 0$ then $\mathcal{M} \leftarrow \mathcal{M}$ else $\mathcal{M} \leftarrow \mathcal{M}'$
 If $o \neq \perp \wedge \exists \text{hdl s.t. Program}_{\mathcal{M}}(\text{hdl}) = R_i^*$.
 $\text{Rev}(o : T_L^i) \sqsubseteq \text{filter}[l^i]\text{Trace}_{\mathcal{M}}(\text{hdl})$: `forgeLAC` $\leftarrow \text{T}$
 $m^* \leftarrow_{\$} \text{Loc}_{\text{KE}}[\text{st}_{\text{KE}}^i](o)$
 $\text{in}_{\text{last}}^i \leftarrow m^*; T_L^i \leftarrow m^* : o : T_L^i$
 If $\text{st}_{\text{KE}}^i.\delta \in \{\text{derived}, \text{accept}\} \wedge \text{st}_{\text{KE}}^i.\text{key} \notin \text{fake}$:
 $\text{key}^* \leftarrow_{\$} \{0, 1\}^\lambda$
 $\text{fake} \leftarrow (\text{key}, \text{key}^*) : \text{fake}$
 Return m^*

Fig. 16. Send oracle from $G1'_{\text{AttKE},\mathcal{A}}$

The correctness of this game hop follows from the same arguments as the proof of Theorem 3 in [3].

SECOND GAME HOP. The second game hop in [3] consist in replacing the remote machine by the simulator, provided by the minimum leakage property. The minimal leakage property is exactly the same in AC and LAC, and this second game hop is exactly the same.

This second game allows us to reason on the semantics of the original code instead of the compiled code executed in an IEE. This lets us take advantage of the semantics of parallel and sequential composition in the next game hop.

THIRD GAME HOP. In the third game hop, the crucial point is emulating a run of the protocol using the oracles from the AttKE security game. As in the proof of Theorem 3 from [3], we keep a list of instances of key exchanges related to the various programs, updated in the `Load` oracle. The `NewSession` oracle creates a new instance of `RemKE` using the `NewLoc` oracle and the `Send` oracle uses the `SendLoc` oracle, exactly as in the original utility proof. The crucial modifications appear in the `Run` oracle and are presented in Figure 17.

We remark that for this last game hop to be valid, we crucially need both the fact that only the key and relevant parts of the key exchange state are passed through ϕ , which is ensured by the fact that the mapping function in the sequential composition is $(\phi_{\text{key}}|\phi_1|\dots|\phi_n)$. Additionally, the state of the key exchange has to be completely independent from the state of the other programs composed in parallel in order for us to be able to emulate it using the `SendRem` oracle. This property is ensured by the semantics of the parallel composition. With these remarks, we observe that the semantics of this third game is exactly the same as the semantics of $G2'_{\text{AttKE},\mathcal{A}}$, in a similar way to the utility proof in [3]. Finally, we observe that the adversary wins $G3'_{\text{AttKE},\mathcal{A}}$ if it wins the AttKE security game (modulo the reduction simulating all non-AttKE oracles), which concludes the proof.

C Proof of Theorem 3

Proof. This proof is made by simulation. First, we present the construction of simulator \mathcal{S} , with the task of interacting with \mathcal{A} on behalf of honest participants of the protocol, i.e., $\mathcal{M}.\text{Load}$, $\mathcal{M}.\text{Run}$ and Send for parties 1 to k . We then present arguments for why adversary \mathcal{A} cannot distinguish between this displayed interaction and the real world protocol execution.

```

Oracle Run(hdl, (l, in)):
   $(R_i^*, \text{st}, j, \text{stage}) \leftarrow \text{List}[\text{hdl}]$ 
  If  $(R_i^*, l_0) \in \text{PrgList}$  and  $R_i^* = \text{LAC.Compile}(\langle\langle P|P_1 \dots |P_l \rangle_{(l_0, l_1, \dots, l_n)}; Q \rangle_{(\phi_{\text{key}}|\phi_1|\dots|\phi_n), p, q})$ :
    If stage = 1:
      If  $l = (p, l_0, \epsilon)$ :
        If st.finished.l: Return (F,  $\epsilon$ )
         $o \leftarrow_s \text{SendRem}(\text{in}, i, j)$ 
        Parse  $(o, \text{sid}, \delta, \text{pid}) \leftarrow o$ :
          If  $\delta = \text{accept}$ :
            st.finished.l  $\leftarrow$  T
            st.l.key  $\leftarrow$  TestRem( $i, j$ )
          Else If  $l = (p, l_k, l')$ 
             $o \leftarrow P_k[\text{st}.l_k](l', \text{in})$ 
            st.finished.lk  $\leftarrow$  o.finished
            If st.finished.lk: st.lk  $\leftarrow$   $\phi_k(\text{st}.l_k)$ 
             $o \leftarrow (\wedge_{i=0}^n \text{st.finished}.l_i, o)$ 
             $(o^*, \text{st}_S) \leftarrow_s \mathcal{S}_2(\text{hdl}, R^*, l, \text{in}, o, \text{st}_S)$ 
            If  $(\wedge_{i=0}^n \text{st.finished}.l_i, o)$ : stage = 2
             $T_R^{\text{hdl}} \leftarrow o^* : \text{in} : T_R^{\text{hdl}}$ 
          Else:
             $o \leftarrow_s Q[\text{st}](l, \text{in})$ 
             $(o^*, \text{st}_S) \leftarrow_s \mathcal{S}_2(\text{hdl}, P, \phi, Q, R^*, \text{in}, o, \text{st}_S)$ 
      Else:
         $(o, \text{st}, \text{st}_S) \leftarrow_s \mathcal{S}_3(\text{hdl}, R^*, \text{in}, \text{st}, \text{st}_S)$ 
    List[hdl]  $\leftarrow (R^*, \text{st}, j, \text{stage})$ 
    Return  $o^*$ 

```

Fig. 17. Run oracle from $\text{G3}'_{\text{AttKE}, \mathcal{A}}$

Observe that, according to the experiment in Figure 7, despite being used in different contexts (e.g. the same \mathcal{S} for emulating the machine and the presentation of outputs), the simulator can always distinguish to which call it is responding to. This is because it receives different inputs in different occasions, with exception of honest party initialization and output retrieval, whose orders are predictable (`GetOutput` will always provide \mathcal{S} with an already initialized id). As such, for clarity of presentation, we describe our simulator in Figure 18 (local participants) and Figure 19 (remote machine) with different behavior for different calls. Notice that in this scenario there is no \mathcal{M} , however the simulator perfectly follows the description of \mathcal{M} to emulate its behaviour. Following its description in Section 3, let $\text{SMLnit}(1^\lambda)$ be the initialization function of the security module, producing public parameters `prms` and internal state `sk`, and let $P^*[\text{hdl}_{\text{st}}, \text{sk}](l, i)$ be the execution of compiled P^* given the internal state `hdlst` and private parameters `sk`, according to the description of the security module, producing (possibly attested) output o^* .

The behaviour detailed in \mathcal{S} does not trivially entail indistinguishability from the real world on all cases. The two main differences between how the simulator handles calls and how the same instructions would be executed in the real world are highlighted in the presented figures, and are now further detailed.

- The simulator is replacing the exchanged keys associated with honest participants with randomly generated ones (fake), and using them throughout the second stage of the protocol.
- Instead of the honest participant’s inputs and outputs, the simulator is encrypting strings of 0s with the same length as the real-world values (obtained by `Lin` and `Lout`).

We now argue that, nevertheless, this provides an indistinguishable view for any \mathcal{A} . We prove this in three game hops from the real world, from Figure 20 to Figure 23. The first hop will replace \mathcal{M} with the slightly different \mathcal{M}' , which replaces keys exchanged by honest participants by freshly generated keys (in exactly the same way the simulator is doing it). The correctness of this hop follows from the utility theorem, using a hybrid argument to replace keys of all k honest parties. Afterwards, we replace the encrypted inputs/outputs of honest parties, by encrypting dummy payloads of the correct length. The correctness of this hop follows from the indistinguishability of the underlying authenticated encryption scheme. Finally, we restrict the possibility of \mathcal{A} to produce a forged encryption, by accordingly establishing a *bad* event. The correctness of this final hop follows from the unforgeability of the underlying authenticated encryption scheme.

<pre> $\mathcal{S}(1^\lambda)$: // parameter initialization prms, st.sk \leftarrow SMLnit(1^λ) st.λ \leftarrow 1^λ; st.hdl \leftarrow 0; st.fake \leftarrow [] Return (st, st.prms) $\mathcal{S}(st, id)$: // party setup (st$_L$, Rem$_{KE}$) \leftarrow Setup$_{KE}(st.\lambda, id)$ st.id.st$_L$ \leftarrow st$_L$ st.stage \leftarrow 0 Return (st, Rem$_{KE}$) $\mathcal{S}(st, id)$: // party initialization st.id.InList \leftarrow []; st.id.stage \leftarrow 0; st.id.seq$_in$ \leftarrow 0; st.id.seq$_out$ \leftarrow 1; st.id.in$_{last}$ \leftarrow ϵ (Rem$_{KE}^1, \dots, Rem_{KE}^n$) \leftarrow Pub st.P \leftarrow $\langle (Rem_{KE}^1, \dots, Rem_{KE}^n)_{1, \dots, n}; \text{Box}(\mathcal{F}, \mathcal{A}) \rangle_{\phi_{key}, P, q}$ st.L \leftarrow $\{(p, (id, \epsilon)), (q, id)\}$ st.id.st$_V$ \leftarrow (st.P, st.L) L^* \leftarrow $\{(p, (1, \epsilon)), \dots, (p, (n, \epsilon))\}$ st.P^* \leftarrow LAC.Compile(prms, st.P, L^*) st.id.InLeak \leftarrow []; st.id.InList \leftarrow [] Return st $\mathcal{S}(st, l, id)$: // add inputs st$_id$.InLeak[id] \leftarrow st$_id$.InLeak[id] + [l] Return st $\mathcal{S}(st, id)$: // output retrieval Return (st.id.seq$_out$/2) + 1 </pre>	<pre> $\mathcal{S}^{Fun}(st, id, m)$: // emulate local participant id If st.id.stage = 0 : (i, st.id.st$_V$) \leftarrow LAC.Verify(st.prms, (p, (id, \epsilon)), st.id.in$_{last}$, m, st.id.st$_V$) If i = \perp: Return \perp (o, st.id.st$_L$) \leftarrow S$_{LocKE}(st.id.st_L, i)$ st.id.in$_{last}$ \leftarrow o If st.id.st$_L$.key \notin st.fake \wedge st.id.st$_L$.$\delta \in$ {derived, accept}: key* \leftarrow S{0, 1}$^{st.\lambda}$ st.fake \leftarrow (st.id.st$_L$.key, key*) : fake If st.id.st$_L$.$\delta =$ accept: stage \leftarrow 1 m' \leftarrow (st.id.stage, id, o) Return (st, m') If st.id.stage = 1 : If m = ϵ : l \leftarrow st.id.InLeak[0] (in$_1, \dots, in_k$) \leftarrow st.id.InLeak st.id.InLeak \leftarrow (in$_1, \dots, in_{k-1}$) in \leftarrow {0}l o \leftarrow SΛ.Enc(fake(st.id.st$_L$.key), (st.id.seq$_in$, in)) st.id.InList[st.id.seq$_in$] \leftarrow o st.id.in$_{last}$ \leftarrow o st.id.seq$_in$ \leftarrow st.id.seq$_in$ + 2 m' \leftarrow (st.id.stage, id, o) Return (st, m') Else: m' \leftarrow SΛ.Dec(fake(st.id.st$_L$.key), m) If m' = (st.id.seq$_out$, out') : st.id.seq$_out$ \leftarrow st.id.seq$_out$ + 2 m' \leftarrow (st.id.stage, id, \epsilon) Return (st, m') Else: Return \perp </pre>
---	--

Fig. 18. Description of simulator \mathcal{S} with respect to emulating local participants.

The first game (Figure 20) is simply the real game expanded with the protocol instantiation. In this setting, whenever the adversary sets k as 0, i.e. corrupts all participants, the simulator already produces an indistinguishable view. In this case there are no honest inputs/outputs, so the simulator has access to all information and can therefore execute the protocol without replacing any keys and without encrypting any dummy payloads ($st.id = \epsilon$ for all id), executing Fun whenever a corrupt input is provided to produce the corresponding output. As such, the following steps will only refer to situations in which $k \neq 0$, where indistinguishability is not yet established.

In the second game $\mathbf{G1}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}'}(1^\lambda)$ (Figure 21), we replace the machine in the ideal world with the machine M' of the Utility game for which $b = 1$. This machine performs exactly what the simulator is doing with the list $fake$, i.e., replacing keys for the first k participants whenever they finish the first stage of the protocol (the key exchange). This is possible via two steps.

Fix identity $id = 1$. We can replace the behaviour of \mathcal{M} in $\mathbf{G0}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}}(1^\lambda)$ regarding this participant using the Utility Theorem 2, for which $l^* = id$. In this scenario, the key (both in \mathcal{M}' .Run and in \mathcal{M}' .Send) for that particular participant will be replaced by a fake one and stored in $fake$ (as described in $\mathbf{G1}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}'}(1^\lambda)$, but only for $id = 1$). The advantage gained by the adversary in this intermediate step is bound by its advantage in winning the experiment in Figure 8, by providing

$$\text{Rem}_{KE}^2, (p, (2, \epsilon)), (q, 2), \dots, \text{Rem}_{KE}^n, (p, (n, \epsilon)), (q, n), (p, (1, \epsilon)), \text{Box}(\mathcal{F}, \mathcal{A}), \phi_{key}$$

To NewSession on every call.

Now observe that, for any scenario in which m participants have had their keys replaced by fake ones, it is possible to apply the same Utility theorem for replacing the keys of $m + 1$ participants. In order to replace all k keys, we are therefore required to apply the same Utility theorem k times, and thus

$$\Pr[\mathbf{G0}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}}(1^\lambda) \Rightarrow \mathbf{T}] - \Pr[\mathbf{G1}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}'}(1^\lambda) \Rightarrow \mathbf{T}] \leq \text{Adv}_{\text{AttKE}, \mathcal{A}}^{\text{UT}}(\lambda) * k.$$

In the third game $\mathbf{G2}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)$ (Figure 22), we open the machine M' and change its behaviour for instances of running program P^* on the second stage as follows:

- Upon receiving an honest participant input, instead of using it for computing Fun instead uses the first element of $ListIn$.
- When producing an honest participant output, instead of returning an encryption of the value received from F , it stores the value from F on an output list for this identity OutList and returns an encryption of zeros of the same length as the output.

Similarly, on the local side for instances running the second stage:

```

 $\mathcal{S}(\text{st}, \text{Pub}, P): // \mathcal{M}.\text{Load}$ 
 $\text{st}.\text{hdl} \leftarrow \text{st}.\text{hdl} + 1$ 
For  $i \in L$ :  $\text{seq}[i] \leftarrow 0$ 
 $\text{st}.\text{HdlList} \leftarrow (\text{st}.\text{hdl}, \text{seq}, \epsilon)$ 
Return  $\text{st}.\text{hdl}$ 

 $\mathcal{S}^{\text{Fun}}(\text{st}, \text{hdl}, m): // \mathcal{M}.\text{Run}$ 
 $(P^*, \text{seq}, \text{st}_{\text{hdl}}) \leftarrow \text{st}.\text{HdlList}[\text{hdl}]$ 
If  $P^* = \text{st}.\text{P}^*$ : // The agreed protocol.
  If  $(p, (\text{id}, \epsilon)) \notin \text{st}.\text{L}$ : Return  $\perp$ 
  If  $\text{st}_{\text{hdl}}[\text{id}].\text{stage} = 0$  :
     $(\text{id}, \text{in}) \leftarrow m$ 
     $m' \leftarrow \text{P}^*[\text{st}_{\text{hdl}}, \text{st}.\text{sk}](\text{id}, m)$ 
    If  $\text{st}.\text{id} \neq \epsilon \wedge \text{st}_{\text{hdl}}[\text{id}].\text{key} \notin \text{st}.\text{fake} \wedge \text{st}_{\text{hdl}}[\text{id}].\delta \in \{\text{derived}, \text{accept}\}$ :
       $\text{key}^* \leftarrow \{0, 1\}^{\text{st}.\lambda}$ 
       $\text{st}.\text{fake} \leftarrow (\text{st}_{\text{hdl}}[\text{id}].\text{key}, \text{key}^*) : \text{fake}$ 
  Else If  $\text{st}_{\text{hdl}}[\text{id}].\text{stage} = 1$  :
     $(\text{seq}_{\text{in}}, \text{id}, \text{in}) \leftarrow m$ 
    If  $(\text{seq}[\text{id}] \neq \text{seq}_{\text{in}})$ : Return  $\perp$ 
    If  $\text{st}.\text{id} \neq \epsilon$ : // Honest participant
      If  $\text{st}.\text{id}.\text{InList}[\text{seq}[\text{id}]] \neq \text{in}$ : Return  $\perp$ 
       $l \leftarrow \text{Fun}(\text{honest}, \text{id}, \epsilon)$ 
       $\text{out} \leftarrow \{0\}^l$ 
       $m' \leftarrow \Lambda.\text{Enc}(\text{fake}(\text{st}_{\text{hdl}}.\text{key}[\text{id}]), (\text{seq}[\text{id}] + 1, \text{in}))$ 
    Else: // Corrupt participant
       $\text{in}^* \leftarrow \Lambda.\text{Dec}(\text{fake}(\text{st}_{\text{hdl}}.\text{key}[\text{id}]), (\text{seq}[\text{id}] + 1, \text{in}))$ 
       $\text{out} \leftarrow \text{Fun}(\text{corrupt}, \text{id}, \text{in}^*)$ 
       $m' \leftarrow \Lambda.\text{Enc}(\text{fake}(\text{st}_{\text{hdl}}.\text{key}[\text{id}]), (\text{seq}[\text{id}] + 1, \text{out}))$ 
     $\text{seq}[\text{id}] \leftarrow \text{seq}[\text{id}] + 2$ 
  Else: // Any other program on  $\mathcal{M}$ .
     $(\text{id}, \text{in}) \leftarrow m$ 
     $m' \leftarrow P^*[\text{st}_{\text{hdl}}, \text{st}.\text{sk}](\text{id}, m)$ 
 $\text{st}.\text{HdlList}[\text{hdl}] \leftarrow (P, \text{seq}, \text{st}_{\text{hdl}})$ 
Return  $(\text{st}, m')$ 

```

Fig. 19. Description of simulator \mathcal{S} with respect to emulating the remote machine.

- When called for presenting the input, instead of encrypting the actual input, it stores it on a list of inputs `InList` and encrypts a string of zeros of the same length.
- Upon receiving an output, instead of decrypting and storing it on `ListOut`, it retrieves the value of `OutList` and stores it on `ListOut`.

We upper bound the distance between these two games, by constructing an adversary \mathcal{B} against the indistinguishability of encryption scheme Λ such that

$$\Pr[\text{G1}_{\mathcal{F}, \pi, \mathcal{M}'}(1^\lambda) \Rightarrow \text{T}] - \Pr[\text{G2}_{\mathcal{F}, \pi}(1^\lambda) \Rightarrow \text{T}] \leq \text{Adv}_{\Lambda, \mathcal{B}}^{\text{IND}}(\lambda) * k * 2I$$

Adversary \mathcal{B} simulates the environment of $\text{G2}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)$ as follows: it first has to try and guess which message will be used to distinguish. Let I be the maximum number of inputs adversary \mathcal{A} chooses to input for any participant. It samples uniformly from $[1..k]$ a participant p , and from $[1..(I * 2)]$ a message m . Since every input produces an output, we establish that

- If $m \in [1..I]$, \mathcal{B} picked the m -th input.
- If $m \in [I + 1, \dots, (I * 2)]$, \mathcal{B} picked the $\frac{m}{I}$ -th output.

and proceed accordingly. \mathcal{B} replaces all calls for encryption/decryption for inputs/outputs of participant p with similar calls to $\Lambda.\text{Enc}$ and $\Lambda.\text{Dec}$, with exception of the following. If $m \in [1..I]$, whenever $\text{Send}(\text{id}, m)$ for $\text{id} = p$ is called for the m -th time on the second stage, \mathcal{B} challenges $\text{IND}^{\Lambda, \mathcal{B}}(1^\lambda)$ with message

$$((\text{st}_{\text{id}}.\text{seq}_{\text{in}}, \text{in}), (\text{st}_{\text{id}}.\text{seq}_{\text{in}}, \{0\}^{|\text{in}|}))$$

Otherwise, whenever $\text{Run}(\text{hdl}, l, m)$ for $l = p$ and $P = P^*$ (the agreed protocol) is called for the $\frac{m}{I}$ -th time on the second stage, \mathcal{B} challenges $\text{IND}^{\Lambda, \mathcal{B}}(1^\lambda)$ with message

$$((\text{seq}[\text{id}] + 1, \text{out}), (\text{seq}[\text{id}] + 1, \{0\}^{|\text{out}|}))$$

Observe that any advantage \mathcal{B} acquires in this transformation can be effectively used to distinguish between $\text{G1}_{\mathcal{F}, \pi, \mathcal{M}'}(1^\lambda)$ and $\text{G2}_{\mathcal{F}, \pi}(1^\lambda)$, since the only difference between the two games is the encryption of either the first message (the real value) or the second (the dummy payload with the same length). The two games are identical modulo this difference.

<p>$\mathbf{G0}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}}(1^\lambda)$: $(n, \mathcal{F}, \text{Lin}, \text{Lout}) \leftarrow \mathcal{F}$ $\text{prms} \leftarrow \mathcal{M}.\text{Init}(1^\lambda)$ $(\text{st}_{\mathcal{A}}, k) \leftarrow \mathcal{A}(\text{prms})$ For $\text{id} \in [1..k]$: $(\text{st}_{\text{id}}.\text{st}_L, \text{st}_{\text{id}}.\text{pub}) \leftarrow \text{Setup}_{\text{KE}}(1^\lambda, \text{id})$ $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_k)$ For $\text{id} \in [k+1..n]$: $(\text{st}_{\mathcal{A}}, \text{pub}_{\text{id}}) \leftarrow \mathcal{A}(\text{st}_{\mathcal{A}}, \text{id}, \text{Pub})$ $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_n)$ For $\text{id} \in [1..k]$: $\text{st}_{\text{id}}.\text{InList} \leftarrow []$; $\text{st}_{\text{id}}.\text{stage} \leftarrow 0$; $\text{st}_{\text{id}}.\text{seq}_{\text{in}} \leftarrow 0$; $\text{st}_{\text{id}}.\text{seq}_{\text{out}} \leftarrow 1$; $\text{st}_{\text{id}}.\text{in}_{\text{last}} \leftarrow \epsilon$ $(\text{Rem}_{\text{KE}}^1, \dots, \text{Rem}_{\text{KE}}^n) \leftarrow \text{Pub}$ $P \leftarrow \langle (\text{Rem}_{\text{KE}}^1, \dots, \text{Rem}_{\text{KE}}^n)_{1, \dots, n}; \text{Box}(\mathcal{F}, \mathcal{A}) \rangle_{\phi_{\text{key}, P, q}}$ $L \leftarrow \{(p, (\text{st}_{\text{id}}.\text{id}, \epsilon)), (q, \text{st}_{\text{id}}.\text{id})\}$ $\text{st}_{\text{id}}.\text{st}_V \leftarrow (P, L)$ $b \leftarrow \mathcal{A}^O(\text{st}_{\mathcal{A}})$</p> <p>Oracle SetInput(in, id): If $\text{id} \notin [1..k]$ Return \perp $\text{st}_{\text{id}}.\text{InList} \leftarrow \text{st}_{\text{id}}.\text{InList} + [\text{in}]$</p> <p>Oracle Load(P): Return $\mathcal{M}.\text{Load}(P)$</p> <p>Oracle Run(hdl, l, m): Return $\mathcal{M}.\text{Run}(\text{hdl}, l, m)$</p>	<p>Oracle Send(id, m): If $\text{id} \notin [1..k]$ Return \perp If $\text{st}_{\text{id}}.\text{stage} = 0$: $(i, \text{st}_{\text{id}}.\text{st}_V) \leftarrow \text{LAC}.\text{Verify}(\text{st}_{\text{id}}.\text{prms}, (p, (\text{st}_{\text{id}}.\text{id}, \epsilon)), \text{in}_{\text{last}}, m, \text{st}_V)$ If $i = \perp$: Return \perp $(o, \text{st}_{\text{id}}.\text{st}_L) \leftarrow \text{Loc}_{\text{KE}}(\text{st}_{\text{id}}.\text{st}_L, i)$ $\text{st}_{\text{id}}.\text{in}_{\text{last}} \leftarrow o$ If $(\text{st}_{\text{id}}.\text{st}_L.\text{st}_{\text{KE}}.\delta) = \text{accept}$: Then $\text{stage} \leftarrow 1$ $m' \leftarrow (\text{st}_{\text{id}}.\text{stage}, \text{st}_{\text{id}}.\text{id}, o)$ Return m' If $\text{st}_{\text{id}}.\text{stage} = 1$: If $m = \epsilon$: $\text{in} \leftarrow \text{st}_{\text{id}}.\text{InList}[0]$ $(\text{in}_1, \dots, \text{in}_k) \leftarrow \text{st}_{\text{id}}.\text{ListInid}$ $\text{st}_{\text{id}}.\text{ListInid} \leftarrow (\text{in}_1, \dots, \text{in}_{k-1})$ $o \leftarrow \mathcal{A}.\text{Enc}(\text{st}_{\text{id}}.\text{st}_L.\text{key}, (\text{st}_{\text{id}}.\text{seq}_{\text{in}}, \text{in}))$ $\text{st}_{\text{id}}.\text{in}_{\text{last}} \leftarrow o$ $\text{st}_{\text{id}}.\text{seq}_{\text{in}} \leftarrow \text{st}_{\text{id}}.\text{seq}_{\text{in}} + 2$ $m' \leftarrow (\text{st}_{\text{id}}.\text{stage}, \text{st}_{\text{id}}.\text{id}, o)$ Return m' Else: $m' \leftarrow \mathcal{A}.\text{Dec}(\text{st}_{\text{id}}.\text{st}_L.\text{key}, m)$ If $m' = (\text{st}_{\text{id}}.\text{seq}_{\text{out}}, \text{out}')$: $\text{st}_{\text{id}}.\text{seq}_{\text{out}} \leftarrow \text{st}_{\text{id}}.\text{seq}_{\text{out}} + 2$ $\text{st}_{\text{id}}.\text{out} \leftarrow \text{out}'$ $m' \leftarrow (\text{st}_{\text{id}}.\text{stage}, \text{st}_{\text{id}}.\text{id}, \epsilon)$ Return m' Else: Return \perp</p> <p>Oracle GetOutput(id): If $\text{id} \notin [1..k]$ Return \perp Return $\text{st}_{\text{id}}.\text{out}$</p>
---	---

Fig. 20. Real world expanded.

In the fourth game $\mathbf{G3}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)$ (Figure 23), the adversary loses whenever `authForge` event occurs. Intuitively, this event corresponds to the adversary producing an encryption that was not produced by either `Send` or `M.Run` for $\text{id} \in [1..k]$ (honest participant), and hence constitutes a forgery with respect to \mathcal{A} . Given that the two games are identical until this event occurs, we have that

$$\Pr[\mathbf{G2}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda) \Rightarrow \mathsf{T}] - \Pr[\mathbf{G3}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda) \Rightarrow \mathsf{T}] \leq \Pr[\text{authForge}].$$

We upper bound the distance between these two games, by constructing an adversary \mathcal{C} against the existential unforgeability of encryption scheme \mathcal{A} such that

$$\Pr[\text{authForge}] \leq \text{Adv}_{\mathcal{A}, \mathcal{C}}^{\text{UF}}(\lambda) * k$$

Adversary \mathcal{C} simulates the environment of $\mathbf{G3}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)$ as follows: it first has to try and guess which session will produce the forgery. As such, it samples uniformly from $[1..k]$ a participant p and replaces the key generated for honest participant p (before adding to fake) with the key generated by $\mathcal{A}.\text{Gen}$. From there on, every time an encryption/decryption is requested for p , the same operation will be requested to $\mathcal{A}.\text{Enc}$ and $\mathcal{A}.\text{Dec}$, respectively.

When `authForge` is set, according to the rules of $\mathbf{G3}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)$, algorithm \mathcal{C} outputs candidate encryption m . It remains to show that this is a valid forgery. To see this, first observe that this is indeed a valid encryption, as decryption is performed on this value immediately before `authForge` occurs. It suffices to establish that message m could not have been queried from the \mathcal{A} oracle. Access to this oracle is only permitted on the encryption of inputs for this participant, and on outputs to this participant (when executing `Run`). From the construction of these operations and the sequence numbers they entail, we know that producing such an encryption would only occur via the inclusion of m in `authList`. Since we know this is not the case, m could not have been queried to the encryption oracle. We conclude therefore that \mathcal{C} outputs a valid forgery whenever `authForge` occurs.

Finally, we argue that the behavior displayed by the simulator is indistinguishable to what adversary \mathcal{A} observes in game $\mathbf{G3}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)$. This is the case because the simulator no longer has private information to which he has no access to. In $\mathbf{G3}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)$, only the length of honest inputs and outputs is required to emulate their private inputs and outputs, to which \mathcal{S} has access to via `Lin` and `Lout`. Additionally, the simulator uses the same message sequence numbers to prevent \mathcal{A} from forcing an execution that deviates from the order in which inputs are provided and outputs are retrieved. Since the executions are the same for all other aspects (including key replacements to fake and exclusion of forged encryptions), \mathcal{A} is provided the same view in both worlds.

<p>G1$_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}'}(1^\lambda)$:</p> <p>$(n, F, \text{Lin}, \text{Lout}) \leftarrow \mathcal{F}$ $\text{prms} \leftarrow \mathcal{M}'.\text{Init}(1^\lambda)$ $\text{fake} \leftarrow []$ $(\text{st}_{\mathcal{A}}, k) \leftarrow \mathcal{A}(\text{prms})$ For $\text{id} \in [1..k]$: $(\text{st}_{\text{id}}.\text{st}_L, \text{st}_{\text{id}}.\text{pub}) \leftarrow \text{Setup}_{\text{KKE}}(1^\lambda, \text{id})$ $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_k)$ For $\text{id} \in [k+1..n]$: $(\text{st}_{\mathcal{A}}, \text{pub}_{\text{id}}) \leftarrow \mathcal{A}(\text{st}_{\mathcal{A}}, \text{id}, \text{Pub})$ $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_n)$ For $\text{id} \in [1..k]$: $\text{st}_{\text{id}}.\text{InList} \leftarrow []$; $\text{st}_{\text{id}}.\text{stage} \leftarrow 0$; $\text{st}_{\text{id}}.\text{seq}_{\text{in}} \leftarrow 0$; $\text{st}_{\text{id}}.\text{seq}_{\text{out}} \leftarrow 1$; $\text{st}_{\text{id}}.\text{in}_{\text{last}} \leftarrow \epsilon$ $(\text{Rem}_{\text{KKE}}^1, \dots, \text{Rem}_{\text{KKE}}^n) \leftarrow \text{Pub}$ $P \leftarrow \langle (\text{Rem}_{\text{KKE}}^1, \dots, \text{Rem}_{\text{KKE}}^n)_{1, \dots, n}; \text{Box}(\mathcal{F}, \Lambda) \rangle_{\phi_{\text{key}, P, q}}$ $L \leftarrow \{(p, (\text{st}_{\text{id}}.\text{id}, \epsilon)), (q, \text{st}_{\text{id}}.\text{id})\}$ $\text{st}_{\text{id}}.\text{st}_V \leftarrow (P, L)$ $P^* \leftarrow \text{LAC.Compile}(\text{prms}, P, L)$ $b \leftarrow \mathcal{A}^O(\text{st}_{\mathcal{A}})$</p> <p>Oracle SetInput(in, id): If $\text{id} \notin [1..k]$ Return \perp $\text{st}_{\text{id}}.\text{InList} \leftarrow \text{st}_{\text{id}}.\text{InList} + [\text{in}]$</p> <p>Oracle Load(P): Return $\mathcal{M}'.\text{Load}(P)$</p> <p>Oracle Run(hdl, l, m): $\text{flag} \leftarrow F$ If $\text{Program}_{\mathcal{M}}(\text{hdl}) = P^*$: $\text{flag} \leftarrow T$ Return $\mathcal{M}'.\text{Run}(\text{hdl}, l, m, \text{flag}, \text{fake})$</p>	<p>Oracle Send(id, m): If $\text{id} \notin [1..k]$ Return \perp If $\text{st}_{\text{id}}.\text{stage} = 0$: $(i, \text{st}_{\text{id}}.\text{st}_V) \leftarrow \text{LAC.Verify}(\text{st}_{\text{id}}.\text{prms}, (p, (\text{st}_{\text{id}}.\text{id}, \epsilon)), \text{in}_{\text{last}}, m, \text{st}_V)$ If $i = \perp$: Return \perp $(o, \text{st}_{\text{id}}.\text{st}_L) \leftarrow \text{Loc}_{\text{KKE}}(\text{st}_{\text{id}}.\text{st}_L, i)$ $\text{st}_{\text{id}}.\text{in}_{\text{last}} \leftarrow o$ If $\text{st}_{\text{id}}.\text{st}_L.\text{key} \notin \text{fake} \wedge \text{st}_{\text{id}}.\text{st}_L.\delta \in \{\text{derived}, \text{accept}\}$: $\text{key}^* \leftarrow \{0, 1\}^\lambda$ $\text{st}.\text{fake} \leftarrow (\text{st}_{\text{id}}.\text{st}_L.\text{key}, \text{key}^*) : \text{fake}$ If $(\text{st}_{\text{id}}.\text{st}_L.\text{st}_{\text{KKE}}.\delta) = \text{accept}$: Then $\text{stage} \leftarrow 1$ $m' \leftarrow (\text{st}_{\text{id}}.\text{stage}, \text{st}_{\text{id}}.\text{id}, o)$ Return m'</p> <p>If $\text{st}_{\text{id}}.\text{stage} = 1$: If $m = \epsilon$: $\text{in} \leftarrow \text{st}_{\text{id}}.\text{InList}[0]$ $(\text{in}_1, \dots, \text{in}_k) \leftarrow \text{st}_{\text{id}}.\text{ListIn}_{\text{id}}$ $\text{st}_{\text{id}}.\text{ListIn}_{\text{id}} \leftarrow (\text{in}_1, \dots, \text{in}_{k-1})$ $o \leftarrow \Lambda.\text{Enc}(\text{fake}(\text{st}_{\text{id}}.\text{st}_L.\text{key}), (\text{st}_{\text{id}}.\text{seq}_{\text{in}}, \text{in}))$ $\text{st}_{\text{id}}.\text{in}_{\text{last}} \leftarrow o$ $\text{st}_{\text{id}}.\text{seq}_{\text{in}} \leftarrow \text{st}_{\text{id}}.\text{seq}_{\text{in}} + 2$ $m' \leftarrow (\text{st}_{\text{id}}.\text{stage}, \text{st}_{\text{id}}.\text{id}, o)$ Return m'</p> <p>Else: $m' \leftarrow \Lambda.\text{Dec}(\text{fake}(\text{st}_{\text{id}}.\text{st}_L.\text{key}), m)$ If $m' = (\text{st}_{\text{id}}.\text{seq}_{\text{out}}, \text{out}')$: $\text{st}_{\text{id}}.\text{seq}_{\text{out}} \leftarrow \text{st}_{\text{id}}.\text{seq}_{\text{out}} + 2$ $\text{st}_{\text{id}}.\text{out} \leftarrow \text{out}'$ $m' \leftarrow (\text{st}_{\text{id}}.\text{stage}, \text{st}_{\text{id}}.\text{id}, \epsilon)$ Return m'</p> <p>Else: Return \perp</p> <p>Oracle GetOutput(id): If $\text{id} \notin [1..k]$ Return \perp Return $\text{st}_{\text{id}}.\text{out}$</p>
--	---

Fig. 21. First hop of the proof.

Let

$$\text{Adv}_{\mathcal{F}, \mathcal{A}}^{\text{Distinguish}} = \Pr[\text{Real}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}}(1^\lambda) \Rightarrow T] - \Pr[\text{Ideal}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{S}}(1^\lambda) \Rightarrow T]$$

To conclude, we have that

$$\begin{aligned} \text{Adv}_{\mathcal{F}, \mathcal{A}}^{\text{Distinguish}} &= \Pr[\text{G0}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}}(1^\lambda)] - \Pr[\text{G3}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)] \\ &= (\Pr[\text{G0}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}}(1^\lambda)] - \Pr[\text{G1}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}'}(1^\lambda)]) + (\Pr[\text{G1}_{\mathcal{F}, \pi, \mathcal{A}, \mathcal{M}'}(1^\lambda)] - \Pr[\text{G2}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)]) + \\ &\quad (\Pr[\text{G2}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)] - \Pr[\text{G3}_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)]) \\ &\leq \text{Adv}_{\text{UT}, \mathcal{A}}^{\text{Att}}(\lambda) * k + \text{Adv}_{\Lambda, \mathcal{B}}^{\text{IND}}(\lambda) * k * 2I + \Pr[\text{forgeAuth}] \\ &\leq \text{Adv}_{\text{UT}, \mathcal{A}}^{\text{Att}}(\lambda) * k + \text{Adv}_{\Lambda, \mathcal{B}}^{\text{IND}}(\lambda) * k * 2I + \text{Adv}_{\Lambda, \mathcal{C}}^{\text{UF}}(\lambda) * k \end{aligned}$$

and Theorem 3 follows.

```

G2 $\mathcal{F}, \pi, \mathcal{A}(1^\lambda)$ :
( $n, F, \text{Lin}, \text{Lout}$ )  $\leftarrow \mathcal{F}$ 
( $\text{prms}, \text{sk}$ )  $\leftarrow \mathcal{S} \text{ SMLnit}(1^\lambda)$ 
 $\text{hdl} \leftarrow 0$ 
 $\text{fake} \leftarrow []$ 
( $\text{st}_A, k$ )  $\leftarrow \mathcal{S} \mathcal{A}(\text{prms})$ 
For  $\text{id} \in [1..k]$ :
  ( $\text{st}_{\text{id}}.\text{st}_L, \text{st}_{\text{id}}.\text{pub}$ )  $\leftarrow \text{Setup}_{\text{KE}}(1^\lambda, \text{id})$ 
 $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_k)$ 
For  $\text{id} \in [k+1..n]$ :
  ( $\text{st}_A, \text{pub}_{\text{id}}$ )  $\leftarrow \mathcal{S} \mathcal{A}(\text{st}_A, \text{id}, \text{Pub})$ 
 $\text{Pub} \leftarrow (\text{pub}_1, \dots, \text{pub}_n)$ 
For  $\text{id} \in [1..k]$ :
   $\text{st}_{\text{id}}.\text{ListIn} \leftarrow []$ ;  $\text{st}_{\text{id}}.\text{ListOut} \leftarrow []$ ;  $\text{st}_{\text{id}}.\text{stage} \leftarrow 0$ 
   $\text{st}_{\text{id}}.\text{seq}_{\text{in}} \leftarrow 0$ ;  $\text{st}_{\text{id}}.\text{seq}_{\text{out}} \leftarrow 1$ ;  $\text{st}_{\text{id}}.\text{in}_{\text{last}} \leftarrow \epsilon$ 
  ( $\text{Rem}_{\text{KE}}^1, \dots, \text{Rem}_{\text{KE}}^n$ )  $\leftarrow \text{Pub}$ 
   $P \leftarrow \langle \langle \text{Rem}_{\text{KE}}^1, \dots, \text{Rem}_{\text{KE}}^n \rangle_{1, \dots, n}; \text{Box}(\mathcal{F}, A) \rangle_{\phi_{\text{key}, P, q}}$ 
   $L \leftarrow \{(p, (\text{st}_{\text{id}}.\text{id}, \epsilon)), (q, \text{st}_{\text{id}}.\text{id})\}$ 
   $\text{st}_{\text{id}}.\text{st}_V \leftarrow (P, L)$ 
   $P^* \leftarrow \text{LAC.Compile}(\text{prms}, P, L)$ 
 $b \leftarrow \mathcal{A}^\mathcal{O}(\text{st}_A)$ 

Oracle Run( $\text{hdl}, l, m$ ):
( $P, \text{seq}, \text{st}_{\text{hdl}}$ )  $\leftarrow \text{HdlList}[\text{hdl}]$ 
If  $P = P^*$ : // The agreed protocol.
  If  $(p, (\text{id}, \epsilon)) \notin L$ : Return  $\perp$ 
  If  $\text{st}_{\text{hdl}}[\text{id}].\text{stage} = 0$ :
    ( $\text{id}, \text{in}$ )  $\leftarrow m$ 
     $m' \leftarrow P^*[\text{st}_{\text{hdl}}, \text{sk}](\text{id}, m)$ 
    If  $\text{st}_{\text{id}} \neq \epsilon \wedge \text{st}_{\text{hdl}}[\text{id}].\text{key} \notin \text{fake} \wedge \text{st}_{\text{hdl}}[\text{id}].\delta \in \{\text{derived}, \text{accept}\}$ :
       $\text{key}^* \leftarrow \{0, 1\}^{1^\lambda}$ 
       $\text{st}_{\text{id}}.\text{fake} \leftarrow (\text{st}_{\text{hdl}}[\text{id}].\text{key}, \text{key}^*) : \text{fake}$ 
    Else If  $\text{st}_{\text{hdl}}[\text{id}].\text{stage} = 1$ :
      ( $\text{seq}_{\text{in}}, \text{id}, \text{in}$ )  $\leftarrow m$ 
      If  $(\text{seq}[\text{id}] \neq \text{seq}_{\text{in}})$ : Return  $\perp$ 
      If  $\text{st}_{\text{id}} \neq \epsilon$ : // Honest participant
         $m' \leftarrow \Lambda.\text{Dec}(\text{fake}(\text{st}_{\text{id}}.\text{st}_L.\text{key}), \text{in})$ 
        If  $m' = (\text{seq}[\text{id}], \text{out}')$ :
           $\text{out} \leftarrow F[\text{st}_{\text{id}}](\text{id}, \text{InList}[\text{seq}[\text{id}]])$ 
           $\text{st}_{\text{id}}.\text{OutList}[\text{seq}[\text{id}] + 1] \leftarrow \text{out}$ 
           $m' \leftarrow \Lambda.\text{Enc}(\text{fake}(\text{st}_{\text{hdl}}.\text{key}[\text{id}]), (\text{seq}[\text{id}] + 1, \{0\}^{\text{out}}))$ 
        Else: // Corrupt participant
           $\text{in}^* \leftarrow \Lambda.\text{Dec}(\text{st}_{\text{hdl}}.\text{key}[\text{id}], (\text{seq}[\text{id}] + 1, \text{in}))$ 
           $\text{out} \leftarrow F[\text{st}_{\text{id}}](\text{id}, \text{in})$ 
           $m' \leftarrow \Lambda.\text{Enc}(\text{st}_{\text{hdl}}.\text{key}[\text{id}], (\text{seq}[\text{id}] + 1, \text{out}))$ 
           $\text{seq}[\text{id}] \leftarrow \text{seq}[\text{id}] + 2$ 
      Else: // Any other program on  $\mathcal{M}$ .
        ( $\text{id}, \text{in}$ )  $\leftarrow m$ 
         $m' \leftarrow P^*[\text{st}_{\text{hdl}}, \text{st}.\text{sk}](\text{id}, m)$ 
   $\text{HdlList}[\text{hdl}] \leftarrow (P, \text{seq}, \text{st}_{\text{hdl}})$ 
  Return  $m'$ 

Oracle Load( $P$ ):
 $\text{hdl} \leftarrow \text{hdl} + 1$ 
For  $i \in L$ :  $\text{seq}[i] \leftarrow 0$ 
 $\text{HdlList} \leftarrow (\text{hdl}, \text{seq}, \epsilon)$ 
Return  $\text{hdl}$ 

Oracle SetInput( $\text{in}, \text{id}$ ):
If  $\text{id} \notin [1..k]$  Return  $\perp$ 
 $\text{st}_{\text{id}}.\text{InList} \leftarrow \text{st}_{\text{id}}.\text{InList} + [\text{in}]$ 

Oracle Send( $\text{id}, m$ ):
If  $\text{id} \notin [1..k]$  Return  $\perp$ 
If  $\text{st}_{\text{id}}.\text{stage} = 0$ :
  ( $i, \text{st}_{\text{id}}.\text{st}_V$ )  $\leftarrow \text{LAC.Verify}(\text{st}_{\text{id}}.\text{prms}, (p, (\text{st}_{\text{id}}.\text{id}, \epsilon)), \text{in}_{\text{last}}, m, \text{st}_V)$ 
  If  $i = \perp$ : Return  $\perp$ 
  ( $o, \text{st}_{\text{id}}.\text{st}_L$ )  $\leftarrow \mathcal{S} \text{Loc}_{\text{KE}}(\text{st}_{\text{id}}.\text{st}_L, i)$ 
   $\text{st}_{\text{id}}.\text{in}_{\text{last}} \leftarrow o$ 
  If  $\text{st}_{\text{id}}.\text{st}_L.\text{key} \notin \text{fake} \wedge \text{st}_{\text{id}}.\text{st}_L.\delta \in \{\text{derived}, \text{accept}\}$ :
     $\text{key}^* \leftarrow \{0, 1\}^{1^\lambda}$ 
     $\text{st}_{\text{id}}.\text{fake} \leftarrow (\text{st}_{\text{id}}.\text{st}_L.\text{key}, \text{key}^*) : \text{fake}$ 
  If  $(\text{st}_{\text{id}}.\text{st}_L.\text{st}_{\text{KE}}.\delta) = \text{accept}$ : Then  $\text{stage} \leftarrow 1$ 
   $m' \leftarrow (\text{st}_{\text{id}}.\text{stage}, \text{st}_{\text{id}}.\text{id}, o)$ 
  Return  $m'$ 
If  $\text{st}_{\text{id}}.\text{stage} = 1$ :
  If  $m = \epsilon$ :
     $\text{in} \leftarrow \text{st}_{\text{id}}.\text{InList}[0]$ 
    ( $\text{in}_1, \dots, \text{in}_k$ )  $\leftarrow \text{st}_{\text{id}}.\text{ListIn}_{\text{id}}$ 
     $\text{st}_{\text{id}}.\text{ListIn}_{\text{id}} \leftarrow (\text{in}_1, \dots, \text{in}_{k-1})$ 
     $\text{st}_{\text{id}}.\text{InList}[\text{st}_{\text{id}}.\text{seq}_{\text{in}}] \leftarrow \text{in}$ 
     $o \leftarrow \mathcal{S} \Lambda.\text{Enc}(\text{fake}(\text{st}_{\text{id}}.\text{st}_L.\text{key}), (\text{st}_{\text{id}}.\text{seq}_{\text{in}}, \{0\}^{\text{in}}))$ 
     $\text{st}_{\text{id}}.\text{in}_{\text{last}} \leftarrow o$ 
     $\text{st}_{\text{id}}.\text{seq}_{\text{in}} \leftarrow \text{st}_{\text{id}}.\text{seq}_{\text{in}} + 2$ 
     $m' \leftarrow (\text{st}_{\text{id}}.\text{stage}, \text{st}_{\text{id}}.\text{id}, o)$ 
    Return  $m'$ 
  Else:
     $m' \leftarrow \Lambda.\text{Dec}(\text{fake}(\text{st}_{\text{id}}.\text{st}_L.\text{key}), m)$ 
    If  $m' = (\text{st}_{\text{id}}.\text{seq}_{\text{out}}, \text{out}')$ :
       $\text{st}_{\text{id}}.\text{ListOut} \leftarrow \text{st}_{\text{id}}.\text{OutList}[\text{st}_{\text{id}}.\text{seq}_{\text{out}}] : \text{st}_{\text{id}}.\text{ListOut}$ 
       $\text{st}_{\text{id}}.\text{seq}_{\text{out}} \leftarrow \text{st}_{\text{id}}.\text{seq}_{\text{out}} + 2$ 
       $m' \leftarrow (\text{st}_{\text{id}}.\text{stage}, \text{st}_{\text{id}}.\text{id}, \epsilon)$ 
      Return  $m'$ 
    Else: Return  $\perp$ 

Oracle GetOutput( $\text{id}$ ):
If  $\text{id} \notin [1..k]$  Return  $\perp$ 
( $\text{out}_1, \dots, \text{out}_k$ )  $\leftarrow \text{st}_{\text{id}}.\text{ListOut}$ 
Return  $\text{out}_1 || \dots || \text{out}_k$ 

```

Fig. 22. Second hop of the proof.

G3 $_{\mathcal{F}, \pi, \mathcal{A}}(1^\lambda)$:

```

(n, F, Lin, Lout) ←  $\mathcal{F}$ 
(prms, sk) ←  $\mathcal{S}$  SMInit( $1^\lambda$ )
hdl ← 0
fake ← []
forgeAuth ← F
authList ← []
(st $_{\mathcal{A}}, k$ ) ←  $\mathcal{A}$ (prms)
For id ∈ [1..k]:
  (st $_{id}$ .st $_L$ , st $_{id}$ .pub) ← Setup $_{\text{KKE}}$ ( $1^\lambda$ , id)
  Pub ← (pub $_1$ , ..., pub $_k$ )
  For id ∈ [k + 1..n]:
    (st $_{\mathcal{A}}$ , pub $_{id}$ ) ←  $\mathcal{A}$ (st $_{\mathcal{A}}$ , id, Pub)
  Pub ← (pub $_1$ , ..., pub $_n$ )
  For id ∈ [1..k]:
    st $_{id}$ .ListIn ← []; st $_{id}$ .ListOut ← []; st $_{id}$ .stage ← 0
    st $_{id}$ .seq $_{in}$  ← 0; st $_{id}$ .seq $_{out}$  ← 1; st $_{id}$ .in $_{last}$  ←  $\epsilon$ 
    (Rem $_{\text{KKE}}^1, \dots, \text{Rem}_{\text{KKE}}^n$ ) ← Pub
    P ←  $\langle (\text{Rem}_{\text{KKE}}^1, \dots, \text{Rem}_{\text{KKE}}^n)_{1, \dots, n}; \text{Box}(\mathcal{F}, \mathcal{A}) \rangle_{\phi_{\text{key}, P, q}}$ 
    L ← {(p, (st $_{id}$ .id,  $\epsilon$ )), (q, st $_{id}$ .id)}
    st $_{id}$ .st $_V$  ← (P, L)
    P* ← LAC.Compile(prms, P, L)
  b ←  $\mathcal{A}^{\mathcal{O}}$ (st $_{\mathcal{A}}$ )
  If forgeAuth = T: b ←  $\mathcal{S}$  {0, 1}

```

Oracle Run(hdl, l, m):

```

(P, seq, st $_{hdl}$ ) ← HdList[hdl]
If P = P*: // The agreed protocol.
  If (p, (id,  $\epsilon$ ))  $\notin$  L: Return  $\perp$ 
  If st $_{hdl}$ [id].stage = 0 :
    (id, in) ← m
    m' ←  $\mathcal{S}$  P* [st $_{hdl}$ , sk](id, m)
    If st $_{id}$   $\neq$   $\epsilon$   $\wedge$  st $_{hdl}$ [id].key  $\notin$  fake  $\wedge$  st $_{hdl}$ [id]. $\delta$  ∈ {derived, accept}:
      key* ←  $\mathcal{S}$  {0, 1} $^{1^\lambda}$ 
      st.fake ← (st $_{hdl}$ [id].key, key*) : fake
    Else If st $_{hdl}$ [id].stage = 1 :
      (seq $_{in}$ , id, in) ← m
      If (seq[id]  $\neq$  seq $_{in}$ ): Return  $\perp$ 
      If st $_{id}$   $\neq$   $\epsilon$ : // Honest participant
        If m' = (seq[id], out') :
          out ← F[st $_F$ ](id, InList[seq[id]])
          st $_{id}$ .OutList[seq[id] + 1] ← out
          m' ←  $\mathcal{A}$ .Enc(fake(st $_{hdl}$ .key[id]), (seq[id] + 1, {0} $^{|\text{out}'|}$ ))
          authList ← (fake(st $_{hdl}$ .key[id]), (seq[id] + 1, {0} $^{|\text{out}'|}$ )) : authList
        Else: // Corrupt participant
          in* ←  $\mathcal{A}$ .Dec(st $_{hdl}$ .key[id], (seq[id] + 1, in))
          out ← F[st $_F$ ](id, in)
          m' ←  $\mathcal{A}$ .Enc(st $_{hdl}$ .key[id], (seq[id] + 1, out))
          seq[id] ← seq[id] + 2
      Else: // Any other program on  $\mathcal{M}$ .
        (id, in) ← m
        m' ←  $\mathcal{S}$  P* [st $_{hdl}$ , st.sk](id, m)
  HdList[hdl] ← (P, seq, st $_{hdl}$ )
  Return m'

```

Oracle Load(P):

```

hdl ← hdl + 1
For i ∈ L: seq[i] ← 0
HdList ← (hdl, seq,  $\epsilon$ )
Return hdl

```

Oracle SetInput(in, id):

```

If id  $\notin$  [1..k] Return  $\perp$ 
st $_{id}$ .InList ← st $_{id}$ .InList + [in]

```

Oracle Send(id, m):

```

If id  $\notin$  [1..k] Return  $\perp$ 
If st $_{id}$ .stage = 0 :
  (i, st $_{id}$ .st $_V$ ) ← LAC.Verify(st $_{id}$ .prms, (p, (st $_{id}$ .id,  $\epsilon$ )), in $_{last}$ , m, st $_V$ )
  If i =  $\perp$ : Return  $\perp$ 
  (o, st $_{id}$ .st $_L$ ) ←  $\mathcal{S}$  Loc $_{\text{KKE}}$ (st $_{id}$ .st $_L$ , i)
  st $_{id}$ .in $_{last}$  ← o
  If st $_{id}$ .st $_L$ .key  $\notin$  fake  $\wedge$  st $_{id}$ .st $_L$ . $\delta$  ∈ {derived, accept}:
    key* ←  $\mathcal{S}$  {0, 1} $^{1^\lambda}$ 
    st.fake ← (st $_{id}$ .st $_L$ .key, key*) : fake
  If (st $_{id}$ .st $_L$ .st $_{\text{KKE}}$ . $\delta$ ) = accept : Then stage ← 1
  m' ← (st $_{id}$ .stage, st $_{id}$ .id, o)
  Return m'
If st $_{id}$ .stage = 1 :
  If m =  $\epsilon$  :
    in ← st $_{id}$ .InList[0]
    (in $_1, \dots, \text{in}_k$ ) ← st $_{id}$ .ListIn $_{id}$ 
    st $_{id}$ .ListIn $_{id}$  ← (in $_1, \dots, \text{in}_{k-1}$ )
    st $_{id}$ .InList[st $_{id}$ .seq $_{in}$ ] ← in
    o ←  $\mathcal{A}$ .Enc(fake(st $_{id}$ .st $_L$ .key), (st $_{id}$ .seq $_{in}$ , {0} $^{|\text{in}|}$ ))
    authList ← (fake(st $_{hdl}$ .key[id]), (st $_{id}$ .seq $_{in}$ , {0} $^{|\text{in}|}$ )) : authList
    st $_{id}$ .in $_{last}$  ← o
    st $_{id}$ .seq $_{in}$  ← st $_{id}$ .seq $_{in}$  + 2
    m' ← (st $_{id}$ .stage, st $_{id}$ .id, o)
    Return m'
  Else:
    m' ←  $\mathcal{A}$ .Dec(fake(st $_{id}$ .st $_L$ .key), m)
    If m'  $\neq$   $\perp$   $\wedge$  (fake(st $_{id}$ .st $_L$ .key), m)  $\notin$  authList:
      forgeAuth ← T
    If m' = (st $_{id}$ .seq $_{out}$ , out') :
      st $_{id}$ .ListOut ← st $_{id}$ .OutList[st $_{id}$ .seq $_{out}$ ] : st $_{id}$ .ListOut
      st $_{id}$ .seq $_{out}$  ← st $_{id}$ .seq $_{out}$  + 2
      st $_{id}$ .out ← out'
      m' ← (st $_{id}$ .stage, st $_{id}$ .id,  $\epsilon$ )
      Return m'
  Else: Return  $\perp$ 

```

Oracle GetOutput(id):

```

If id  $\notin$  [1..k] Return  $\perp$ 
(out $_1, \dots, \text{out}_k$ ) ← ListOut $_{id}$ 
Return out $_1$  || ... || out $_k$ 

```

Fig. 23. Third hop of the proof.