

# Foundations of Hardware-Based Attested Computation and Application to SGX\*

Manuel Barbosa  
HASLab, INESC TEC, FCUP  
mbb@dcc.fc.up.pt

Bernardo Portela  
HASLab, INESC TEC, UMinho  
blfportela@gmail.com

Guillaume Scerri  
University of Bristol  
guillaume.scerri@bris.ac.uk

Bogdan Warinschi  
University of Bristol  
csxbw@bristol.ac.uk

## Abstract

Exciting new capabilities of modern trusted hardware technologies allow for the execution of arbitrary code within environments completely isolated from the rest of the system and provide cryptographic mechanisms for securely reporting on these executions to remote parties.

Rigorously proving security of protocols that rely on this type of hardware faces two obstacles. The first is to develop models appropriate for the induced trust assumptions (e.g., what is the correct notion of a *party* when the peer one wishes to communicate with is a specific instance of an outsourced program). The second is to develop scalable analysis methods, as the inherent stateful nature of the platforms precludes the application of existing modular analysis techniques that require high degrees of independence between the components.

We give the first steps in this direction by studying three cryptographic tools which have been commonly associated with this new generation of trusted hardware solutions. Specifically, we provide formal security definitions, generic constructions and security analysis for *attested computation*, *key-exchange for attestation* and *secure outsourced computation*. Our approach is incremental: each of the concepts relies on the previous ones according to an approach that is quasi-modular. For example we show how to build a secure outsourced computation scheme from an arbitrary attestation protocol combined together with a key-exchange and an encryption scheme.

## 1 Introduction

**BACKGROUND.** The many applications that routinely manipulate sensitive data require strong guarantees which ensure i. that adversaries cannot tamper with their execution; and ii. that no sensitive information is leaked. Yet, satisfying these guarantees on modern execution platforms rife with vulnerabilities (e.g. in mobile devices, PCs) or inherently not trustworthy (e.g., cloud infrastructures) is a major challenge.

---

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

A promising starting point for solutions are the *remote attestation* capabilities offered by modern trusted hardware: computational platforms equipped with this technology can guarantee to a remote party various degrees of integrity for the software that it runs. For example the Trusted Platform Module (TPM) can provide certified measurements on the state of the platform and can be used to guarantee integrity of BIOS and boot code right before it is executed. More recent technologies (e.g., ARM’s TrustZone and Intel’s Software Guard Extension (SGX) [18]) have significantly expanded the scope and guarantees of trusted hardware. They offer the ability to run applications in “clean-slate” isolated execution environments (IEE) completely independent of anything else running on the processor; the desired attestation guarantees come from reports that are authenticated cryptographically.

A second major challenge is to provide security guarantees that go beyond heuristic arguments. Here, the established methodology is the “provable security” approach, which advocates carrying out the analysis of systems with respect to rigorously specified models that clarify the trust relations, the powers of the adversary and what constitutes a security breach. The approach offers well-established definitional paradigms for all basic primitives and some of the more used protocols. On the other hand, the success of applying this approach to new and more complex scenarios fundamentally hinges on one’s ability to tame scalability problems, as models and proofs, even for moderate size systems, tend to be unwieldy. Some solutions to this issue exist in the form of compositional principles that, when incorporated as a native feature in the security abstractions, allow to establish the guarantees of larger systems from the guarantees on its components [8].

In this paper we take a provable security approach to protocols that rely on the IEE-capabilities of modern trusted hardware. It may be tempting to assume that, for the analysis of such protocols, designing security models is a simple matter of overlaying/merging the trust model induced by the use of such hardware over well-established security abstractions. If this were true, one could rely on established models and methodologies to perform the analysis. Unfortunately, this is not the case.

Consider the problem of secure outsourced computation of a program  $P$  to a remote machine. The owner of  $P$  wants to ensure that the (potentially malicious) remote platform does not tamper with the execution of the program and that it learns no information about the input/output (I/O) behavior of  $P$ . For remote machines with IEE capabilities, the following straightforward design has been informally proposed in the literature and should, intuitively, provide the desired guarantees. First, execute a key-exchange with a remote instance of an IEE; after the key exchange finishes, use the key thus established to send encrypted inputs to the IEE who can decrypt and pass them to  $P$ . The output returned by  $P$  is encrypted within the IEE and sent to the user. The construction relies on standard building blocks (key-exchange, authenticated encryption) so the security of the overall design should be reducible to that of the key exchange and authenticated encryption, for which we already have widely-accepted security models and constructions.

We highlight two important issues that show that neither existent models nor existent techniques are immediately suitable for the analysis of IEE-based protocols in general, and for the protocol above in particular. The first is the concept of a *party* which is a key notion in specifying and reasoning about the security of distributed systems. Traditionally, one considers security (e.g., of a key-exchange) in a setting where there is a PKI and at least some of the parties (e.g., the servers) have associated public keys. Parties and their cryptographic material are then essentially the same thing for the purpose of the security analysis. In the context that we study, users (who wish to use trusted hardware) are not expected to have

long term keys; furthermore, privacy considerations require that the cryptographic operations performed by the trusted hardware on remote machines should not allow one to track different instances – which makes long term cryptographic material inadequate as a technical anchor of a party’s participation in such a protocol. Indeed, the desirable functionality for many usages of such systems is that the cryptographic material associated with a computation outsourcing protocol (both for local and remote *parties*) can be arbitrary and fixed on-the-fly, when the protocol is executed. An interesting problem is therefore how to define the security of outsourced computation in this setting, and how to rely on the asymmetry afforded by the trust model specific to IEE systems to realise it.

The second issue is *composability*. In the protocol for outsourced computation that we present above, one might be led to think that security simply follows if the key-exchange is secure and the channel between the user and the IEE uses authenticated encryption (with appropriate replay protection): if the only information passed from the key exchange to the channel is the encryption key, then one can design and analyze the two parts separately. While in more standard scenarios this may be true, reliance on the IEE breaks the independence assumption that allows for composability results: what the specification above hides is that the code run by the IEE (i.e., the program for key-exchange, the one for the secure channels and the program that they protect) needs to be loaded at once, or else no isolation guarantees are given by the trusted hardware.<sup>1</sup> This means that the execution of the different parts of the program is not necessarily independent, as they unavoidably share the state of the IEE. An important question is therefore whether the above intuitive construction is sound, and under which conditions can one use it to perform IEE-enabled outsourcing of computation.

To summarise the above discussion, two remarks help motivate the work in this paper. Protocols relying on IEE are likely to be deployed in applications with stringent security requirements, so ensuring that they fall under the scope of the provable security approach is important. Moreover, such applications are likely to be complex: inherently, they involve communication between remote parties, incorporate diverse code executed at different levels of trust, and rely on multiple cryptographic primitives and protocols as building blocks (see, for example, the One-time password protocol based on SGX[18]). However, we have concluded that key aspects of existing cryptographic models do not naturally translate to this new setting and, perhaps more worryingly, that the type of compositional reasoning enabled by such cryptographic models is clearly unsuitable for protocols that rely on IEE.

**OUR APPROACH.** We will first outline the high-level decisions that underly our approach, and then describe our technical contributions more in detail. We use SGX and TrustZone as inspiration, but we do not hardwire our models to a specific platform, in order to ensure that the scope of our work encompasses other similar technologies. Instead, we use an abstract notion of a machine which captures the relevant aspects that such platforms offer: their capability to run processes with isolation guarantees and the ability to directly use secure

---

<sup>1</sup>Intuitively, mechanisms such as SGX and TrustZone are designed to protect and provide attestation guarantees over monolithic pieces of software, which must be fixed when an IEE is created and are identified using a fingerprint of the code. To go around this restriction and compose multiple programs, one has two options: i. to build a single composed program and load it in its entirety into an IEE (this is the approach we follow in this paper); or ii. to use multiple IEEs to host the various programs, and employ cryptographic protocols to protect the interactions between them using the (potentially malicious) host operating system as a communications channel. We note that this latter option would again lead to loading composed programs into each IEE, since cryptographic code would need to be added to each of the individual programs, and this would then lead to the same problem that we intend to solve with the framework we propose in this paper.

cryptography without going through potentially untrusted software. Additionally, we target our effort on the combination of remote attestation and key-exchange protocols. The former is *the* raison d’être for trusted hardware, while the latter is the natural building block towards adding secrecy guarantees for code running under the protection of IEEs (per our example above).

As explained above, one challenge we set out to address is to incorporate into our approach a new form of compositional reasoning that permits dealing with potentially shared state between all the code that is loaded into an IEE, which means that a-priori there are no guarantees of independence between the executions of different cryptographic primitives that one could incorporate into the same program. Indeed, at the very least, the reporting mechanisms for the IEE will refer to the code of the full program, which immediately constrains modular reasoning — a crucial tool to enable scalability. We sidestep this problem by providing definitions (for both syntax and security) that are *composition-aware*: they explicitly assume that the code loaded into an IEE may result from the composition of multiple programs.

## Contributions

We focus on the interplay between composition and attestation in hardware-based settings. In particular, we concentrate on a pervasive use case: attestation is often used to protect a security-critical part of the code ( $P^*$  in our setting) whereas the remaining code  $Q$  does not need attestation: it can rely on guarantees established by  $P^*$  (e.g. an authenticated secret key), and can therefore be much more efficient. As explained above and in footnote 1 the stateful nature of the execution environment does not allow for independent analysis of these two components and new techniques for rigorous validation are needed.

Our starting point is therefore a program  $Q$  that a local user wishes to outsource to a remote machine. Such a program will need to be transformed (*compiled* in the cryptographic sense of the word) into another program that, intuitively, will result from the composition of a handshake/bootstrapping procedure  $P^*$  that will establish a secure channel with the IEE in which program  $Q$  will be executed, and an instrumented version of  $Q$ , say  $Q^*$ , that uses the aforementioned secure channel to ensure that  $Q$  is indeed securely executed.

Our approach to formalising and realising this composition pattern has three main stepping-stones: i. we introduce *attested computation* as the formalisation of the raw guarantees provided by IEEs with cryptographic functionalities; ii. we show how a passively secure key exchange can be efficiently combined with an *attested computation* scheme to obtain the bootstrapping procedure  $P^*$  referred above; and iii. we rely on our composition-aware formalisation to show that by instrumenting program  $Q^*$  using standard cryptographic techniques one achieves secure outsourced computation. Details follow.

**ATTESTED COMPUTATION.** Our first contribution is a formal treatment of IEE-based remote attested computation. We consider a setting where a user wishes to remotely execute a program  $P$  and rely on the cryptographic infrastructure available within the IEE to attest that some incarnation  $P^*$  of this program is indeed executing within an IEE of a specific remote machine (or group of machines). We provide a general solution to this problem in the form of a new cryptographic concept called an *attested computation*. We formalize two core guarantees.

First, we demand that the user’s local view of the execution is “as expected”, i.e., that the I/O behavior that is reconstructed locally corresponds to an honest execution of  $P$ . The

second guarantee is more subtle and requires that such an *execution has actually occurred in an IEE within a specific remote platform*; in other words, the attested computation client is given the assurance that its code is being run in isolation (and displays a given I/O behaviour) within a prescribed remote physical machine (or group of machines) associated with some authenticated public parameters. This latter guarantee is crucial for bootstrapping the secure outsourcing of code: consider for example  $P$  to be a key-exchange protocol, which will be followed by some other program that relies on the derived key. It must be the case that, at the end of the key exchange, the remote state of the key-exchange is protected by an IEE.

The second guarantee we demand from attested computation follows easily from the previous observation. If the attested program keeps sensitive information in its internal state (which is not revealed by its I/O behaviour as in the case of a key exchange protocol) then execution within the remote IEE should safeguard its internal state. If this were not the case, we would again run into problems when trying to compose  $P^*$  with some program that relies on the security properties of  $P$ . We therefore exclude attested computation schemes where the instrumented program  $P^*$  might leak more information in its I/O behaviour than  $P$  itself by introducing the notion of *minimal leakage*, which essentially states that the I/O of an attested program does not leak any information beyond what is unavoidably leaked by an honest execution.

Finally, we provide a scheme for attested computation that relies on a remote machine offering a combination of symmetric authentication and digital signatures (a capability similar to what SGX provides) and show that our scheme is secure in the sense that we define.

**KEY-EXCHANGE FOR ATTESTED COMPUTATION.** On its own, attested computation only provides integrity guarantees: the I/O behavior of the outsourced code is exposed to untrusted code in the remote machine on which it is run. The natural solution to the problem is to establish a secure communication channel with the IEE via a key-exchange protocol. It is unclear, however, how the standard security models for key exchange protocols map into the attested computation scenario, and how existing constructions for secure key-exchange fare in the novel scenario that we study. Indeed, for efficiency reasons one should use a key exchange protocol that is *just strong enough* to achieve this goal.

To clarify this issue we formalize the notion of key-exchange *for* attested computation. The name that we propose is intentional: key-exchange protocols as used in the our context differ significantly in the syntax and security models from their more traditional counterparts. For example, our syntax reflects that the code of the key-exchange is not fixed a-priori: a user can set parameters both for the component to be run locally and for the one to be executed within the IEE. This allows a user to hardwire in the code to be run remotely a new nonce (or as in our examples some cryptographic public key for which it knows the secret key).

As explained above, the notion of *party* in the context of attested computation needs to be different from that adopted by traditional notions of secure key-exchange. Our solution is to rely on the trust model specific to IEE settings: we can assign some arbitrary strings as identifiers for the users of the local machine, and we allow these users to specify arbitrary strings as identifiers for the remote code (a secure instantiation would require that this identifier corresponds to some cryptographic material possibly generated on the fly as explained above). We then adapt the execution model and definitions for key-exchange for the modified syntax and the new notion of communicating parties to reflect the expected guarantees: different local and remote sessions agree on each other's identifiers, derive the same key and the key is unknown to the adversary. One crucial aspect of our security model for key-exchange

is that it explicitly accounts for the fact that the remote process will be run under attestation guarantees, which maps to a *semi-active* adversarial environment.

To improve usability of our notion of secure key-exchange for attested computation we provide two results. The first result simplifies the design of such protocols. Here, we show a generic construction that combines a key-exchange protocol that is passively secure and a standard signature scheme to derive a (potentially very efficient) key exchange protocol for attestation. The second result simplifies reasoning about the composition of a key-exchange for attestation with an arbitrary protocol that relies on the agreed key. Specifically, we provide a *utility theorem* which specifically states what composition guarantees one gets for an arbitrary program  $Q$  that is run within a remote IEE and relies on a shared key that was established via the attested computation of a key exchange protocol that satisfies our tailored definition.

**SECURE OUTSOURCED COMPUTATION.** The last layer in our framework is a formalisation of secure outsourced computation, the principal motivating use-case for our approach. We provide syntax and two security notions for a secure outsourced computation protocol, one for authenticity and the second one for the privacy of the I/O of the outsourced program. We then prove that the construction that combines a key-exchange for attested computation with an authenticated symmetric encryption scheme and replay protection gives rise to a scheme for secure outsourced computation. We present our result as a general formalisation (i.e., not application specific) of the intuition that by relying on more powerful hardware assumption such as those offered by SGX, one can indeed efficiently achieve a well-defined notion of secure outsourced computation that simultaneously offers verifiability *and* privacy. The proof of this result crucially relies on the utility theorem we defined for the combination of attested computation with key exchange.

## 2 Other Related work

Work that looks at provable security of realistic protocols that use trusted hardware-based protocols has developed around the protocols offered by the Trusted Platform Module (TPM) [5, 28, 6, 12, 11]. However, the functionality and efficiency of the protocols offered by the TPM makes them more suitable for static attestation (i.e., ensuring integrity of programs right before they are executed). Run-time guarantees, like those that we study here are in principle possible but cumbersome to obtain. Linking attestation guarantees provided by the TPM with those offered by a secure channel onto a remote machine had been studied before but only informally [16]. Although more rigorous approaches used to analyze attestation guarantees for protocols based on the TPM exist, they use more abstract models (with weaker guarantees) [27, 10].

Another related line of research leverages the trusted hardware to bootstrap entire platforms for secure software execution (e.g. Flicker [22], Trusted Virtual Domains [9], Haven [2]). These are large systems that are currently outside the scope of provable-security techniques. Smaller protocols which solve specific problems (secure disk encryption [23], one-time password authentication [18] outsourced Map-Reduce computations [26], Secure Virtual Disk Images [13], secure embedded devices [24, 21]) are more susceptible to rigorous analysis. Although some protocols (e.g., those of Hoekstra et al. [18]) come only with intuition regarding their security, others – most notably those by Schuster et. al [26] which uses SGX platforms to outsource map-reduce computation – come with a proof of security. The constructions in

<b>Game <math>\text{Auth}^{\Pi, \mathcal{A}}(1^\lambda)</math>:</b> List $\leftarrow \square$ key $\leftarrow_{\$} \text{Gen}(1^\lambda)$ $(m, t) \leftarrow_{\$} \mathcal{A}^{\text{Auth}}(1^\lambda)$ Return $\text{Ver}(\text{key}, m, t) = \text{T} \wedge m \notin \text{List}$	<b>Oracle <math>\text{Auth}(m)</math>:</b> List $\leftarrow (m : \text{List})$ $t \leftarrow \text{Mac}(\text{key}, m)$ Return $t$
--	---

Figure 1: Game defining the security of a MAC scheme  $\Pi$ .

that paper are close to those that we abstract and analyze here.

Our construction of secure outsourced computation can be seen as a solution to the problem of verifiable computation as specified by [14]. That rich line of work concentrates on the much harder problem of providing crypto-only solutions and usually gives up privacy for efficiency and verifiability (e.g., [4], [15], [25]). Here, we show how to obtain a reasonably efficient and secure system with technology that is likely to be deployed in the near future.

### 3 Preliminaries

#### 3.1 Message Authentication Codes

**SYNTAX.** A message authentication code scheme  $\Pi$  is a triple of PPT algorithms  $(\text{Gen}, \text{Auth}, \text{Ver})$ . On input  $1^\lambda$ , where  $\lambda$  is the security parameter, the randomized key generation algorithm returns a fresh key. On input key and message  $m$ , the deterministic MAC algorithm  $\text{Auth}$  returns a tag  $t$ . On input key,  $m$  and  $t$ , the deterministic verification algorithm  $\text{Ver}$  returns  $\text{T}$  or  $\text{F}$  indicating whether  $t$  is a valid MAC for  $m$  relative to key. We require that, for all  $\lambda \in \mathbb{N}$ , all  $\text{key} \in [\text{Gen}(1^\lambda)]$  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 [3]. We say that  $\Pi$  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 returns  $\text{T}$ .

#### 3.2 Digital Signature Schemes

**SYNTAX.** A signature scheme  $\Sigma$  is a triple of PPT algorithms  $(\text{Gen}, \text{Sign}, \text{Vrfy})$ . On input  $1^\lambda$ , where  $\lambda$  is the security parameter, the randomized key generation algorithm returns a fresh key pair  $(\text{pk}, \text{sk})$ . On input secret key  $\text{sk}$  and message  $m$ , the possibly randomized signing algorithm  $\text{Sign}$  returns a signature  $\sigma$ . On input public key  $\text{pk}$ ,  $m$  and  $\sigma$ , the deterministic verification algorithm  $\text{Vrfy}$  returns  $\text{T}$  or  $\text{F}$  indicating whether  $\sigma$  is a valid signature for  $m$  relative to  $\text{pk}$ . We require that, for all  $\lambda \in \mathbb{N}$ , all  $(\text{pk}, \text{sk}) \in [\text{Gen}(1^\lambda)]$  and all  $m$ , it is the case that  $\text{Vrfy}(\text{pk}, m, (\text{Sign}(\text{sk}, m))) = \text{T}$ .

**SECURITY.** We use the standard notion of existential unforgeability for signature schemes [17]. We say that  $\Sigma$  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 2 returns  $\text{T}$ .

#### 3.3 Passively secure key exchange

We define a form of key exchange protocol that does not rely on long term secret/state or global setup and for which we require weak security guarantees (essentially security against a passive adversary); the classical Diffie-Hellman key exchange is a standard example. Later we

<p><b>Game <math>\text{UF}^{\Sigma, \mathcal{A}}(1^\lambda)</math>:</b></p> <p>List <math>\leftarrow \square</math></p> <p><math>(\text{pk}, \text{sk}) \leftarrow_{\\$} \text{Gen}(1^\lambda)</math></p> <p><math>(\text{m}, \sigma) \leftarrow_{\\$} \mathcal{A}^{\text{Sign}}(1^\lambda, \text{pk})</math></p> <p>Return <math>\text{Vrfy}(\text{pk}, \text{m}, \sigma) = \top \wedge \text{m} \notin \text{List}</math></p>	<p><b>Oracle <math>\text{Sign}(\text{m})</math>:</b></p> <p>List <math>\leftarrow (\text{m} : \text{List})</math></p> <p><math>\sigma \leftarrow \text{Sign}(\text{sk}, \text{m})</math></p> <p>Return <math>\sigma</math></p>
---	--

Figure 2: Game defining the security of a signature scheme  $\Sigma$ .

<p><b>Game <math>\text{Corr}_{\Pi}(1^\lambda)</math>:</b></p> <p><math>\text{st}_j \leftarrow \epsilon; \text{t} \leftarrow j</math></p> <p><math>(\text{m}, \text{st}_i) \leftarrow_{\\$} \Pi(\epsilon, i, \text{initiator}, \epsilon)</math></p> <p>While <math>\text{m} \neq \epsilon</math>:</p> <p style="padding-left: 20px;">If <math>\text{t} = j</math>: <math>\text{t} \leftarrow i; (\text{m}, \text{st}_j) \leftarrow_{\\$} \Pi(\text{m}, j, \text{responder}, \text{st}_j)</math></p> <p style="padding-left: 20px;">Else: <math>\text{t} \leftarrow j; (\text{m}, \text{st}_i) \leftarrow_{\\$} \Pi(\text{m}, i, \text{initiator}, \text{st}_i)</math></p> <p>Return <math>\delta_i = \delta_j = \text{accept} \wedge \text{key}_i = \text{key}_j \wedge \text{sid}_i = \text{sid}_j \wedge \text{pid}_j = i \wedge \text{pid}_i = j</math></p>
---

Figure 3: Game defining the correctness of protocol  $\Pi$ .

show how, combined with attestation, such protocols yield secure key-exchange when facing active adversaries. A key exchange protocol is therefore defined by a single ppt algorithm  $\Pi$  used by communicating parties. When analysing the security and correctness of the protocol we will consider that each party with identifier  $\text{id}$  can execute several instances of the protocol with different parties. Throughout the paper we let identifiers be arbitrary strings, which will be given meaning by the higher-level application relying on the protocol under analysis. For  $s \in \mathbb{N}$ , we write  $\Pi_{\text{id}}^s$  for the  $s$  instance of party  $\text{id}$ . We assume that each instance maintains variables  $\text{st}, \delta, \text{key}$  which record respectively, local state information, the state of the key (derived, accept, reject or  $\perp$ ) and the value of the key. In addition, we assume variables for the role  $\rho$  of the session (initiator or responder), the party identifier of the owner of the session  $\text{oid}$  and that of its partner  $\text{pid}$  and a session identifier  $\text{sid}$ . We require that  $\text{key} = \perp$  unless  $\delta \in \{\text{derived}, \text{accept}\}$ , that  $\text{oid}, \rho, \text{sid}, \text{key}$  are only assigned once during the entire execution of the protocol (the first two when the session is initialized). We denote running  $\Pi$  with message  $\text{m}$ , role  $\rho$  and state  $\text{st}$  to produce  $\text{m}'$  and the updated state  $\text{st}'$  by  $(\text{m}', \text{st}') \leftarrow_{\$} \Pi(1^\lambda, \text{m}, \text{id}, \rho, \text{st})$ , and will omit the security parameter input throughout the paper for the sake of compactness. We will use message  $\epsilon$  to refer to the empty string. This will be passed to both parties as the initial state; it will be passed to the initiator as first input message; and it will be returned as output message by the party that executes last, to denote that no further interaction is needed. Long term secrets and/or shared initial state between several instances run by the same identity can be captured by setting their initial state accordingly. **CORRECTNESS.** A key exchange protocol is correct if, after a complete (honest) run between two participants with complementary roles, both reach the `accept` state, both derive the same key and session identifier, and both obtain correct partner identities. More formally, a protocol  $\Pi$  is correct if, for any distinct party identities  $i$  and  $j$ , the experiment in Figure 3 always returns  $\top$ .

**EXECUTION MODEL.** We will consider key exchange schemes that are secure against passive adversaries. Our adopted security notion is a restriction of the scenario considered in [20] that excludes corruptions.<sup>2</sup> The execution model considers an adversary, which is run on

<sup>2</sup>The trust model we consider for attested computation excludes corruptions for the sake of simplicity, but all our results can be extended to consider that possibility. Having said that, it seems reasonable to exclude the possibility of isolated executed environment breach based on a hardware assumption. The possibility of

the security parameter, and which can interact with the following oracles whose behaviour depends on a secret sampled bit  $b$  and a shared list of pairs of keys `fake`, which is initially empty:

- `Execute( $i, j$ )` runs a new instance of the protocol between distinct parties  $i$  and  $j$ . It then checks if the key derived for the executed session exists in list `fake`. If not, it generates a new `key*` uniformly at random, and adds `(key, key*)` to the list. Finally, it outputs the transcript of the protocol execution and the session identifier associated with it.
- `Reveal( $i, s$ )` outputs the session key `key` of  $\Pi_i^s$ .
- `Test( $i, s$ )` will return  $\perp$  if  $\delta_i^s \neq \text{accept}$  (i.e. if no execute query actually created such a session). Otherwise, if  $b = 0$  it outputs the key associated with  $\Pi_i^s$ . If  $b = 1$ , it searches for the key associated with  $\Pi_i^s$  in list `fake` and returns the associated `key*`.

When the adversary terminates interacting with the oracles, it will eventually output a bit  $b'$  which represents his guess on what the challenge bit  $b$  is.

**PARTNERING RELATIONS.** We define entity authentication following [7], and observe that in the case of passively secure key exchange, this is essentially a correctness property. First we introduce a notion of partnering, which informally states that two oracles which have derived keys are partners if they share the same session identifier. The definition makes use of the following predicate on two instances  $\Pi_i^s$  and  $\Pi_j^t$  holding states  $(\text{st}_i^s, \delta_i^s, \rho_i, \text{sid}_i^s, \text{pid}_i^s, \text{key}_i^s)$  and  $(\text{st}_j^t, \delta_j^t, \rho_j, \text{sid}_j^t, \text{pid}_j^t, \text{key}_j^t)$ , respectively:

$$P(\Pi_i^s, \Pi_j^t) = \begin{cases} \text{T} & \text{if } \text{sid}_i^s = \text{sid}_j^t \wedge \delta_i^s, \delta_j^t \in \{\text{derived}, \text{accept}\} \\ \text{F} & \text{otherwise.} \end{cases}$$

**Definition 1.** [Partner] Two players  $\Pi_i^s$  and  $\Pi_j^t$  are partnered if  $P(\Pi_i^s, \Pi_j^t) = \text{T}$ .

Our authentication notion relies on three further definitions, which demand that partnerings will need to be *valid*, *confirmed* and *unique*. In short, these three requirements ensure that any instance that accepts has a partner, that this partner is unique and that partners share the same key.

**Definition 2** (Valid Partners). A protocol  $\Pi$  ensures valid partners if the bad event `notval` does not occur, where `notval` is defined as follows:

$$\begin{aligned} \exists \Pi_i^s, \Pi_j^t \text{ s.t. } P(\Pi_i^s, \Pi_j^t) = \text{T} \wedge \\ (\text{pid}_i^s \neq \text{oid}_j^t \vee \text{pid}_j^t \neq \text{oid}_i^s \vee \rho_i = \rho_j \vee \text{key}_i^s \neq \text{key}_j^t). \end{aligned}$$

**Definition 3** (Confirmed Partners). A protocol  $\Pi$  ensures confirmed partners if the bad event `notconf` does not occur, where `notconf` is defined as follows:

$$\exists \Pi_i^s \text{ s.t. } \delta_i^s = \text{accept} \wedge \forall \Pi_j^t, P(\Pi_i^s, \Pi_j^t) = \text{F}.$$

**Definition 4** (Unique Partners). A protocol  $\Pi$  ensures unique partners if the bad event `notuni` does not occur, where `notuni` is defined as follows:

$$\begin{aligned} \exists \Pi_i^s, \Pi_j^t, \Pi_k^r \text{ s.t.} \\ (j, t) \neq (k, r) \wedge P(\Pi_i^s, \Pi_j^t) = \text{T} \wedge P(\Pi_i^s, \Pi_k^r) = \text{T} \end{aligned}$$

---

local machine corruption should, however, be considered.

Intuitively, we will consider that an adversary violates two-sided entity authentication if he can lead an instance of an honest party running the protocol to accept, and in doing that cause one of the bad events `notval`, `notconf`, `notuni`.

**SECURITY DEFINITION.** We are now ready to present the security notion we will be using for key exchange protocols. To exclude breaks via trivial attacks, we define legitimate adversaries as those who ensure the following freshness criteria is satisfied for his `Test`( $i, s$ ) queries: i. `Reveal`( $i, s$ ) was not queried; and ii. for all  $\Pi_j^t$  such that  $P(\Pi_i^s, \Pi_j^t) = \mathsf{T}$ , `Reveal`( $j, t$ ) was not queried. We only consider experiments in which the adversary is found to be legitimate and define the winning event `guess` to be  $b = b'$  at the end of the experiment.

**Definition 5.** [Passive AKE security] A protocol  $\Pi$  is passively secure if, for any legitimate ppt adversary interacting with the execution environment described above: 1. the adversary violates two-sided entity authentication with negligible probability  $\Pr[\text{notval} \vee \text{notconf} \vee \text{notuni}]$ ; and 2. its key secrecy advantage  $2 \cdot \Pr[\text{guess}] - 1$  is negligible.

**ONE-SIDED AUTHENTICATION.** We will also consider a weaker form of entity authentication guarantee, where only some parties are authenticated. To this end, we will distinguish between `Loc` and `Rem` parties, identify the former with responders, and the latter with initiators. In one-sided authentication, only initiators (i.e., remote parties) are authenticated, which means that responders (i.e., local parties) do not need to keep any long term secrets. Again, we follow [7] and give one-sided versions of the definitions for valid partners and confirmed partners. Intuitively, on acceptance, a local party will be assured that a valid unique partnering session exists. Security will still need to ensure that each local (or remote) session has at most one remote (respectively local) partnered session. However, we no longer confirmation for the remote machine: we allow the remote party to accept even if there is no local matching session that has accepted. This weaker guarantee is sufficient for many applications and has been used in the context of attested computation, for example, in [18].

**Definition 6** (One-Sided Valid Partners). A protocol  $\Pi$  ensures one-sided valid partners if the bad event `os-notval` does not occur, where `os-notval` is defined as follows:

$$\begin{aligned} \exists \Pi_i^s, \Pi_j^t \text{ s.t. } & i \in \text{Loc} \wedge P(\Pi_i^s, \Pi_j^t) = \mathsf{T} \wedge \\ & (\text{pid}_i^s \neq j \vee \rho_i^s \neq \text{initiator} \vee \rho_j \neq \text{responder} \vee \text{key}_i^s \neq \text{key}_j^t). \end{aligned}$$

**Definition 7** (One-Sided Confirmed Partners). A protocol  $\Pi$  ensures one-sided confirmed partners if the bad event `os-notconf` does not occur, where `os-notconf` is defined as follows:

$$\begin{aligned} \exists \Pi_i^s \text{ s.t. } & i \in \text{Loc} \wedge \\ & \delta_i^s = \text{accept} \wedge \forall \Pi_j^t, P(\Pi_i^s, \Pi_j^t) = \mathsf{F}. \end{aligned}$$

When one-sided authentication suffices, then the definition of key exchanged security is weakened by relaxing condition 1 in Definition 5, which is modified to refer to the following event.

$$\Pr[\text{os-notval} \vee \text{os-notconf} \vee \text{notuni}]$$

## 4 IEEs, Programs, and Machines

**ISOLATED EXECUTION ENVIRONMENTS.** At the high-level, an IEE can be seen as an idealised random access machine running some fixed program  $P$ , whose behaviour can only be influenced via a well-specified interface that permits passing inputs to the program, and receiving

its outputs. Intuitively, an IEE gives the following security guarantees, which we will formalise later in this section. 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. This means, in particular, that there is strict isolation between processes running in different IEEs (and any other program running on the machine). Furthermore, the only information that is revealed about a program running within an IEE is contained in its input-output behaviour (which in most hardware systems is simply shared memory between the protected code and the untrusted software outside).

We emphasize that our notion of a machine is intended to be inclusive of any hardware platform that supports some form of isolated execution. For this reason, the syntax of this abstraction is minimalistic, so that it can be restricted/extended to capture the specific guarantees awarded by different concrete hardware architectures, including TPM, TrustZone, SGX, etc. As an example, our “vanilla” machine supports an arbitrary number of IEEs, where programs can be loaded only once, and where multiple input/output interactions are allowed with the protected code. This is a close match to the SGX/TrustZone functionalities. However, for something like TPM, one could consider a restricted machine where a limited number of IEEs exist, with constrained input/output capabilities, and running specific code (e.g., to provide key storage). Similarly, we consider IEE environments where the underlying hardware is assumed to only keep *benevolent* state, i.e., state that cannot be used to introduce destructive correlations between multiple interactions with an IEE. Again, this closely matches what happens in SGX/Trustzone, but different types of state keeping could be allowed for scenarios where such correlations are not a problem or where they must be dealt with explicitly.

PROGRAMS. Implicit throughout the paper will be a programming language  $\mathcal{L}$  in which programs are written. We assume that this language is used by all computational platforms, but we admit IEE-specific system calls giving access to different cryptographic functionalities. These are referred as the *security module* interface. An additional system call `rand` is also assumed to be present in all platforms, giving access to fresh random coins sampled uniformly at random. Language  $\mathcal{L}$  is assumed to be deterministic modulo the operation of system calls. 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.

A program  $P$  must be written as a transition function, mapping bit-strings to bit-strings. Such functions take a current state `st` and an input  $i$ , and they will produce a new output  $o$  and an updated state. We will refer to this as an *activation* and express it as  $o \leftarrow P[\text{st}](i)$ . Unless otherwise stated, `st` will be 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. The transition function may return arbitrary output until it produces an output where `finished` = `T`, at which point it can return no further output or change its state. We extend our notation as  $o \leftarrow P[\text{st}; r](i)$  to account for the randomness

obtained via the `rand` system call as extra input  $r$ ; and as  $(o_1, \dots, o_n) \leftarrow P[\text{st}; r](i_1, \dots, i_n)$  to represent a sequence of activations. We write  $\text{Trace}_{P[\text{st}; r]}(i_1, \dots, i_n)$  for the corresponding I/O trace  $(i_1, o_1, \dots, i_n, o_n)$ .

**PROGRAM COMPOSITION.** Given two programs  $P$  and  $Q$ , and a projection function between the internal states of the two programs  $\phi$ , we will refer to the sequential composition of the two programs as  $\text{Compose}_\phi(P, Q)$ . This is defined as a transition function  $R$  that has two execution stages, which are signaled in its output via an additional `stage` bit. In the first stage, every input to  $R$  will activate 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  $R$  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 (termination being counted as a third stage) is appended to any output of a composition. When dealing with such a composed program, we will denote by  $\text{ATrace}_{R[\text{st}; r]}(i_1, \dots, i_n)$  the prefix of the trace that corresponds to the execution of  $P$ . Intuitively, this denotes the *attested trace* where only the initial part of the program must be protected via attestation.

**MACHINES.** A *machine*  $\mathcal{M}$  is an abstract computational device that captures the resources offered by a real world computer or group of computers, whose hardware security functionalities are initialised by a specific manufacturer before being deployed, possibly in 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.<sup>3</sup> In this paper we will restrict ourselves to the simplest case where standard public key signatures are used; but all our results easily extend to more complex group management schemes.

We will 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. Loosely speaking, this interface can be thought of as the ideal functionality that captures a system such as SGX [19]. 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. We emphasize that the global parameters of machines are the only pieces of information that are assumed to be authenticated using external mechanisms such as a PKI in the entire paper.
- $\text{Load}(P)$  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

---

<sup>3</sup>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.

the only information that is leaked about internal state and randomness input is that revealed (indirectly) via the outputs of the program.

- $\text{Run}(\text{hdl}, i)$  is the process activation procedure. On input a handle  $\text{hdl}$  and an input  $i$ , it will activate process running in isolated execution environment with handle  $\text{hdl}$  with  $i$  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  $\text{hdl}$  running in some machine  $\mathcal{M}$  as the tuple  $(i_1, o_1, \dots, i_n, o_n)$  that includes the entire sequence of  $n$  inputs/outputs resulting from all invocations of the  $\text{Run}$  procedure on  $\text{hdl}$ ;  $\text{Program}_{\mathcal{M}}(\text{hdl})$  is the code (program) running inside the process with handle  $\text{hdl}$ ;  $\text{Coins}_{\mathcal{M}}(\text{hdl})$  represents the coins given to the program by the  $\text{rand}$  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  $\text{Load}$  and  $\text{Run}$  oracles defined above.

REMARK. Here, we use  $\text{hdl}$  as a convenient identifier for the secure environment executing process  $P$ ; in some incarnation the handle could be defined as a tuple containing the identity of the machine, some identifier for the secure environment and, say, the hash of the program  $P$ . More detailed formalisms are possible. We may consider, for example, different entry/exit points related to  $P$ . We may also explicitly refine  $P$  as a program and some initial associated data.

## 5 Attested Computation

We now formalise a cryptographic primitive that aims to address the remote execution, i.e., outsourcing, of programs as illustrated in Figure 4. In this setting, a user running software in a trusted local machine wishes to use an untrusted network to access a pool of remote machines with IEE facilities. The remote machines will be running general-purpose operating systems and other untrusted software. The goal of the user is to run a specific program  $P$  within an IEE in one of the remote machines, and to obtain assurance that, not only the program is indeed executing there, but also that it is displaying a particular I/O behaviour. We call this *attested computation*, and introduce it as the cryptographic primitive that formalises the simplest cryptographic application of trusted hardware systems offering IEE functionality, such as Intel’s SGX architecture. Attested computation will be the lowest of a series of cryptographic layers that we will be building on top of each other throughout the paper, until eventually we show how such hardware assumptions permit efficiently realising a strong form of secure verifiable computation.

SYNTAX. An *Attested Computation* (AC) scheme is defined by the following algorithms:

- $\text{Compile}(\text{prms}, P, \phi, Q)$  is the program compilation algorithm. On input global parameters for some machine  $\mathcal{M}_R$ , and programs  $P$  and  $Q$ , whose composition under projection function  $\phi$  will be outsourced, it will output program  $R^*$ , together with an initial (possibly empty) state  $\text{st}$  for the verification algorithm. This algorithm is run locally.  $R^*$  is the code to be run as an isolated process in the remote machine. Intuitively,  $P$  is the initial part of the remote code that requires attestation guarantees, whereas  $Q$  is any subsequent code that may be remotely executed (generally leveraging the security guarantees that have been bootstrapped using the initial attested execution).

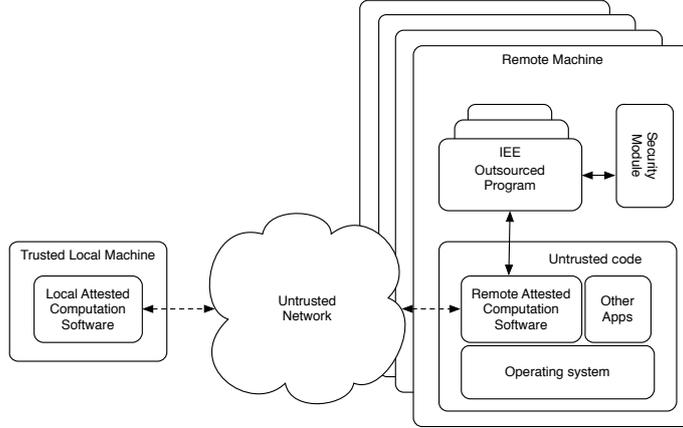


Figure 4: Attested Computation scenario.

- $\text{Attest}(\text{prms}, \text{hdl}, i)$  is the attestation algorithm. On input global parameters for  $\mathcal{M}_R$ , a process handle  $\text{hdl}$  and an input  $i$ , it will use the interface of  $\mathcal{M}_R$  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  $R^*$ , providing it with inputs and recovering the (possibly attested) outputs that should be returned to the local machine.
- $\text{Verify}(\text{prms}, i, o^*, \text{st})$  is the (stateful) output verification algorithm. On input global parameters for  $\mathcal{M}_R$ , an input  $i$ , a (possibly attested) output  $o^*$  and some state  $\text{st}$ , 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  $\text{Attest}$  algorithm.

In Figure 4, the local attested computation software block corresponds to  $\text{Compile}$  (one initial usage per program) and  $\text{Verify}$  (one usage per incoming attested output), whereas the remote attested computation software block corresponds to  $\text{Attest}$  (one usage per remote program activation, i.e. per I/O transition). The above syntax can be naturally extended to accommodate the simultaneous compilation of multiple input programs and/or the possibility that  $\text{Compile}$  may generate multiple output programs. This would allow us to capture, e.g., map/reduce applications such as those described in [26].

**CORRECTNESS.** Intuitively, an AC scheme is correct if, for any given programs  $P$  and  $Q$  and assuming an honest execution of all components in the scheme, both locally and remotely, the local user is able to accurately reconstruct a view of the I/O sequence that took place in the remote environment. Furthermore, this I/O sequence must be consistent with the semantics of  $\text{Compose}_\phi(P; Q)$ . In other words, suppose the compiled program is run under handle  $\text{hdl}^*$  in remote machine  $\mathcal{M}_R$ , and the local user uses  $\text{Verify}$  to reconstruct the remote I/O behaviour  $(i_1, o_1, \dots, i_n, o_n)$ . Then, if we define  $R := \text{Compose}_\phi(P; Q)$ , we must have

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

The following definition formalizes the notion of a local user *correctly remotely executing program*  $P$  using attested computation.

**Definition 8.** [Correctness] An Attested Computation scheme AC is correct if, for all  $\lambda$ , and all adversaries  $\mathcal{A}$ , the experiment in Figure 5 (top) always returns T.

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. Since we are dealing with composed programs, we extend the correctness requirements on attested computation schemes to preserve the structure of the input program  $(P, \phi, Q)$ , and to modify only the part of the code that will be attested. Formally, we impose that, given any program  $P$ , there exists a (unique) compiled program  $P^*$ , such that, for any mapping function  $\phi$  and any program  $Q$ , we have that  $\text{Compose}_\phi\langle P^*; Q \rangle = \text{Compile}(P, \phi, Q)$ .

SECURITY. Security of an attested computation scheme imposes that an adversary with absolute control of the remote machine cannot convince the local user that some arbitrary remote execution of a program  $P$  has occurred, when it has not (nothing is said about the subsequent remote execution of program  $Q$ ). Formally, we allow the adversary to freely interact with the remote machine, whilst providing a sequence of (potentially forged) attested outputs. 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. the execution trace that is validated by `Verify` is inconsistent with the semantics of  $P$  (in which case an adversary would be able to convince the local user of an I/O sequence that could not possibly have occurred!); or ii. there does not exist a remote process `hdl*` exhibiting a consistent execution trace (in which case, the adversary would be able to convince the local user that a process running  $P$  was executing in the remote machine, when it was not).

Since the adversary is free to interact with the remote machine as it pleases, we can not hope to prevent it from appending arbitrary inputs to the trace of any remote process, while refusing to deliver all of the resulting attested outputs to the local user. This justifies the winning condition in our security game referring to a prefix of the trace in the remote machine, rather than imposing trace equality. Indeed, the definition’s essence is to impose that the locally recovered trace and the remote trace share a common prefix ( $\sqsubseteq$ ), which exactly corresponds to the part of the source program’s behaviour that should be protected by attestation.

Formally, we need to account for the fact that the actual I/O sequence of the remote program includes more information than that of  $R$ , e.g., to allow for the cryptographic enforcement of security guarantees. Our definition is parametrised by a `Translate` algorithm that permits formalising this notion of *semantic consistency*. Another way to see  $\text{Translate}(\text{prms}, \text{ATrace}_{\mathcal{M}_R}(\text{hdl}^*))$  is as a trace translation procedure associated with a given AC scheme, which maps remote traces into traces at the source level.

**Definition 9.** [Security] An attested computation scheme is secure if there exists an efficient deterministic algorithm `Translate` s.t., for all ppt adversaries  $\mathcal{A}$ , the probability that experiment in Figure 5 (bottom) returns T is negligible.

<p><b>Game</b> <math>\text{Corr}_{\text{AC}, \mathcal{A}}(1^\lambda)</math>:</p> <pre> prms <math>\leftarrow</math> <math>\mathcal{M}_R.\text{Init}(1^\lambda)</math> <math>(P, \phi, Q, n, \text{st}_A) \leftarrow</math> <math>\mathcal{A}_1(\text{prms})</math> <math>(R^*, \text{st}_V) \leftarrow \text{Compile}(\text{prms}, P, \phi, Q)</math> <math>\text{hdl}^* \leftarrow \mathcal{M}_R.\text{Load}(R^*)</math> For <math>k \in [1..n]</math>:   <math>(i_k, \text{st}_A) \leftarrow</math> <math>\mathcal{A}_2(o_1^*, \dots, o_{k-1}^*, \text{st}_A)</math>   <math>o_k^* \leftarrow \text{Attest}^{\mathcal{M}_R}(\text{prms}, \text{hdl}^*, i_k)</math>   <math>(o_{R,k}, \text{st}_V) \leftarrow \text{Verify}(\text{prms}, i_k, o_k^*, \text{st}_V)</math>   If <math>o_{R,k} = \perp</math>:     Return F Define <math>R := \text{Compose}_\phi(P; Q)</math> <math>T \leftarrow \text{Trace}_{R[\text{st}; \text{Coins}_{\mathcal{M}_R}(\text{hdl}^*)]}(i_1, \dots, i_n)</math> <math>T' \leftarrow (i_1, o_{R,1}, \dots, i_n, o_{R,n})</math> Return <math>T = T'</math> </pre>	<p><b>Game</b> <math>\text{Att}_{\text{AC}, \mathcal{A}}(1^\lambda)</math>:</p> <pre> prms <math>\leftarrow</math> <math>\mathcal{M}_R.\text{Init}(1^\lambda)</math> <math>(P, \phi, Q, n, \text{st}_A) \leftarrow</math> <math>\mathcal{A}_1(\text{prms})</math> <math>(R^*, \text{st}_V) \leftarrow \text{Compile}(\text{prms}, P, \phi, Q)</math> For <math>k \in [1..n]</math>:   <math>(i_k, o_k^*, \text{st}_A) \leftarrow</math> <math>\mathcal{A}_2^{\mathcal{M}_R}(\text{st}_A)</math>   <math>(o_{R,k}, \text{st}_V) \leftarrow \text{Verify}(\text{prms}, i_k, o_k^*, \text{st}_V)</math>   If <math>o_{R,k} = \perp</math> Return F <math>T' \leftarrow (i_1, o_{R,1}, \dots, i_n, o_{R,n})</math> Define <math>R := \text{Compose}_\phi(P; Q)</math> For <math>\text{hdl}^*</math> s.t. <math>\text{Program}_{\mathcal{M}_R}(\text{hdl}^*) = R^*</math>:   <math>T \leftarrow \text{ATrace}_{R[\text{st}; \text{Coins}_{\mathcal{M}_R}(\text{hdl}^*)]}(i_1, \dots, i_n)</math>   If <math>T \subseteq T' \wedge T \subseteq \text{Translate}(\text{prms}, \text{ATrace}_{\mathcal{M}_R}(\text{hdl}^*))</math>:     Return F Return T </pre>
--	--

Figure 5: Games defining the correctness (left) and security (right) of an AC scheme.

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

**MINIMUM LEAKAGE.** From the discussion above, an AC scheme should guarantee 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. However, it is important to establish an additional restriction on what AC compilation actually does to a source program, to ensure that we are able to take advantage of this primitive to achieve more ambitious goals, namely to perform attestation of the remote execution of cryptographic code.

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 when it is remotely executed.

**Definition 10.** [Minimal leakage] Attested Computation scheme AC ensures security with minimal leakage if it is secure according to Definition 9 and there exists a ppt simulator  $\mathcal{S}$  that, for every adversary  $\mathcal{A}$ , the following distributions are identical:

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

where games  $\text{Leak-Real}_{\text{AC}, \mathcal{A}}$  and  $\text{Leak-Ideal}_{\text{AC}, \mathcal{A}, \mathcal{S}}$  are shown in Figure 6.

Notice that we allow the simulator to replace the global parameters of the machine with some value  $\text{prms}$  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

<b>Game Leak-Real<math>_{AC, \mathcal{A}}(1^\lambda)</math>:</b> $\text{PrgList} \leftarrow []$ $\text{prms} \leftarrow \mathcal{M}_R.\text{Init}(1^\lambda)$ $b \leftarrow \mathcal{A}^{\mathcal{O}}(\text{prms})$ Return $b$	<b>Oracle Compile<math>(P, \phi, Q)</math>:</b> $(R, \text{st}_V) \leftarrow \text{Compile}(\text{prms}, P, \phi, Q)$ $\text{PrgList} \leftarrow R : \text{PrgList}$ Return $R$  <b>Oracle Load<math>(R)</math>:</b> Return $\mathcal{M}_R.\text{Load}(R)$	<b>Oracle Run<math>(\text{hdl}, i)</math>:</b> Return $\mathcal{M}_R.\text{Run}(\text{hdl}, i)$
<b>Game Leak-Ideal<math>_{AC, \mathcal{A}, \mathcal{S}}(1^\lambda)</math>:</b> $\text{PrgList} \leftarrow []$ $\text{List} \leftarrow []$ $\text{hdl} \leftarrow 0$ $(\text{prms}, \text{st}_S) \leftarrow \mathcal{S}_1(1^\lambda)$ $b \leftarrow \mathcal{A}^{\mathcal{O}}(\text{prms})$ Return $b$	<b>Oracle Compile<math>(P, \phi, Q)</math>:</b> $(R, \text{st}_V) \leftarrow \text{Compile}(\text{prms}, P, \phi, Q)$ $\text{PrgList} \leftarrow (P, \phi, Q, R) : \text{PrgList}$ Return $R$  <b>Oracle Load<math>(R)</math>:</b> $\text{hdl} \leftarrow \text{hdl} + 1$ $\text{List}[\text{hdl}] \leftarrow (R, \epsilon)$ Return $\text{hdl}$	<b>Oracle Run<math>(\text{hdl}, i)</math>:</b> $(R, \text{st}) \leftarrow \text{List}[\text{hdl}]$ If $(P, \phi, Q, R) \in \text{PrgList}$ : $R^* \leftarrow \text{Compose}_\phi(P, Q)$ $o^* \leftarrow R^*[\text{st}](i)$ $(o, \text{st}_S) \leftarrow \mathcal{S}_2(\text{hdl}, P, \phi, Q, R, i, o^*, \text{st}_S)$ Else: $(o, \text{st}, \text{st}_S) \leftarrow \mathcal{S}_3(\text{hdl}, R, i, \text{st}, \text{st}_S)$ $\text{List}[\text{hdl}] \leftarrow (R, \text{st})$ Return $o$

Figure 6: Games defining minimum leakage of an AC scheme.

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).

## 6 Attested Computation à la SGX

The remote attestation protocol we will consider is inspired in the Secure Guard Extensions (SGX) architecture proposed by Intel [1]. The main feature of this system is that the remote machine is equipped with a security module that manages both short-term and long-term cryptographic keys, with which it is capable of producing MACs that enable authenticated communication between various IEEs and digital signatures that can be publicly verified by anyone holding the (long-term) public key for that machine (or group of machines). We first formalise the operation of (a simplified version of) this security module.

**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  $(\text{pk}, \text{sk})$  using  $\Sigma.\text{Gen}$  and a symmetric key  $\text{key}$  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  $\text{sk}$  and  $\text{key}$ .
- The code of  $S^*$  is dedicated to transforming messages authenticated with  $\text{key}$  into messages signed with  $\text{sk}$ . On each activation, it expects an input  $(\text{m}, \text{t})$ . It obtains  $\text{key}$  from the security module and verifies the tag using  $\Pi.\text{Ver}(\text{key}, \text{t}, \text{m})$ . If the previous oper-

ation was successful, it obtains  $\text{sk}$  from the security module, signs the message using  $\sigma \leftarrow_{\S} \Sigma.\text{Sign}(\text{sk}, \text{m})$  and writes  $\sigma$  to the output. Otherwise, it writes  $\perp$  in the output.

- The security module exposes a single system call  $\text{mac}(\text{m})$  to code running in all other IEEs. On such a request from a process running program  $P$ , the security module returns a MAC tag  $\text{t}$  computed using  $\text{key}$  over both the code of  $P$  and the input message ( $\text{m}$ ).

We note that the operation of the security module allows any process to produce an authenticated message that can be validated by the special process running  $S^*$  as coming from within another IEE in the same machine.

We will assume that the message authentication code scheme  $\Pi$  and the signature scheme  $\Sigma$  satisfy the standard notions of correctness and existential unforgeability, and that the machine's public key is authenticated by some external PKI.

**ATTESTED COMPUTATION SCHEME.** We now define an AC scheme that relies on a remote machine supporting a security module with the above functionality. The operation of the various algorithms is intuitive, except for the fact that basic replay protection using a sequence number does not suffice to bind a remote process to a full trace, since the adversary could then run multiple copies of the same process and *mix and match* outputs from various traces. Instead, the remote process must commit to its entire trace whenever an attested output is produced. Details follow:

- $\text{Compile}(\text{prms}, P, \phi, Q)$  will generate a new program  $R^* = \text{Compose}_{\phi}(P^*, Q)$  and output it along with the initial state of the verification algorithm  $(R^*, [], 1)$ , where 1 is an indicator of the stage in which remote program  $R^*$  is supposed to be executing. Program  $P^*$  is instrumented as follows: it keeps a list  $\text{ios}$  of all the I/O pairs it has previously received and computed, i.e, its own trace; on each activation with input  $i$ ,  $P^*$  first computes  $o \leftarrow_{\S} P[\text{st}_P](i)$  and updates the list by adding a new  $(i, o)$  pair; it then requests from the security module a MAC of the updated  $\text{ios}$ . Due to the operation of the security module, this will correspond to a tag  $\text{t}$  on the tuple  $(R^*, \text{ios})$ ; it finally outputs  $(o, \text{t}, R^*, \text{ios})$ . We note that we include  $(R^*, \text{ios})$  explicitly in the outputs of  $R^*$  for clarity of presentation only. This value would be kept in an insecure environment by a stateful **Attest** program.
- $\text{Attest}(\text{prms}, \text{hdl}, i)$  invokes  $\mathcal{M}_R.\text{Run}(\text{hdl}, i)$  using the handle and input value it has received. When the process produces an output  $o$ , **Attest** parses it into  $(o', \text{t}, R^*, \text{ios})$ . It may happen that parsing fails, e.g., if  $Q$  is already executing, in which case **Attest** simply produces  $o$  as its own output. Otherwise, it uses  $\mathcal{M}_R.\text{Run}(0, (R^*, \text{ios}, \text{t}))$  to convert the tag into a signature  $\sigma$  on the same message. If this conversion fails, then **Attest** produces the original output  $o$  as its own output. Otherwise, it outputs  $(o', \sigma)$ .
- $\text{Verify}(\text{prms}, i, o^*, (R^*, \text{ios}, \text{stage}))$  returns  $o^*$  if  $\text{stage} = 2$ . Otherwise, it first parses  $o^*$  into  $(o, \sigma)$ , appends  $(i, o)$  to  $\text{ios}$ , and verifies the digital signature  $\sigma$  using  $\text{prms}$  and  $(R^*, \text{ios})$ . If parsing or verification fails, **Verify** outputs  $\perp$ . If not, then **Verify** will check if output  $o$  indicates that program  $P^*$  has finished. If so, it will update  $\text{stage}$  to value 2. In any case, it terminates outputting  $o$ .

**CORRECTNESS.** It is easy to see that our AC scheme is correct, provided that the underlying signature and message authentication code schemes are themselves correct. To see this, first note that, during the execution of  $P^*$ , unless a MAC or signature verification fails, the I/O

sequence produced by `Verify` exactly matches that of  $\text{Compose}_\phi\langle P; Q \rangle$ , and therefore  $T = T'$  is always  $\top$ . This follows from the construction of  $R^*$ , the operation of `Attest`, and the fact that the associated randomness tapes are established by  $\text{Coins}_{\mathcal{M}_R}(\text{hdl}^*)$  as identical. Furthermore, 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 because the combined actions of  $R^*$ , `Attest`, the signing process running  $S^*$  and the security module lead to tags and signatures on tuples  $(R^*, \text{ios})$  that exactly match those input to the verification algorithms  $\Pi.\text{Ver}$  and  $\Sigma.\text{Verify}$ . Finally, after executing  $P^*$ , given that the associated randomness tapes are established by  $\text{Coins}_{\mathcal{M}_R}(\text{hdl}^*)$  are identical and that traces are identical up to that point, so will be  $\phi(\text{st}_P)$  in both sides, and all subsequent calls to  $Q$  will display a similar behaviour.

**SECURITY.** Let `Translate` be the deterministic function that receives the machine parameters and a list of tuples of the form  $(i, (o, \text{t}, R^*, \text{ios}))$  and returns a list of pairs of the form  $(i, o)$ .

**Theorem 1.** *The AC scheme presented above provides secure attestation if the underlying MAC scheme  $\Pi$  and signature scheme  $\Sigma$  are existentially unforgeable. Furthermore, it unconditionally ensures minimum leakage.*

The proof of the following theorem can be found in Appendix A. The intuition behind the proof of secure attestation is straightforward: since all attested outputs (i.e. those processed by `Verify` until  $\text{st}_P.\text{finished} = \top$ ) are bound to a full trace of the execution, all accepted messages that pass `AC.Verify` must terminate a prefix of a remote trace for some instance of  $R^*$ . The only case in which the adversary could win would be if the signature verification performed by `Verify` accepts a message that was never authenticated by an IEE running  $R^*$ . However, in this case, the adversary is either breaking the MAC (to dishonestly execute `Attest`) or breaking the signature (and forging attested outputs directly). The minimum leakage property can be proven by constructing the trivial simulator that generates the machine parameters itself, simulates the entire machine for non-attested processes, and attaches MACs to the source I/O traces of attested programs.

## 7 AKE for Attested Computation

An intermediate step in constructing high-level applications that rely on attested computation is the establishment of a secure communications channel with a process running a particular program inside an IEE in the remote machine. After such a channel has been established, standard cryptographic techniques can be used to ensure (in combination with the isolation provided by IEEs) the integrity and confidentiality of subsequent computations. In this section we will see how attested computation, in combination with a specific flavour of key exchange protocol can be seen as a bootstrapping process for this scenario.

We first formalize the precise requirements for a key exchange protocol that can be used in this setting (we call this *authenticated key exchange for attested computation*) and show how a simple transformation can be used to construct such protocols from any passively secure key exchange protocol. Later on we present a utility theorem that precisely describes what it means to use attested computation and a suitable key exchange protocol to establish a secure channel with an arbitrary remote program.

## Definitions

SYNTAX. A *Key Exchange for Attested Computation* (AttKE) protocol is defined by the following pair of algorithms.

- $\text{Setup}(1^\lambda, \text{id})$  is the remote program generation algorithm, which is run on the local machine to initialise a fresh instance of the AttKE protocol under party identifier  $\text{id}$ . On input the security parameter and  $\text{id}$ , it will output the code for a program  $\text{Rem}_{\text{KE}}$  and the initial state  $\text{st}_L$  of the  $\text{Loc}_{\text{KE}}$  algorithm. This algorithm is run locally.
- $\text{Rem}_{\text{KE}}$  (which is generated dynamically by  $\text{Setup}$ ) is a program that will be run as a part of an IEE process in the remote machine, and it will keep the entire remote state of the key exchange protocol in that protected environment.
- $\text{Loc}_{\text{KE}}(\text{st}_L, \text{m})$  is the algorithm that runs the local end of the AttKE protocol, interacting with  $\text{Rem}_{\text{KE}}$ . On input its current state, and an incoming message  $\text{m}$ , it will output an updated state and an outgoing message.

When analysing the security of such a protocol we will impose that the  $\text{Loc}_{\text{KE}}$  algorithm and all  $\text{Rem}_{\text{KE}}$  programs that may be produced by  $\text{Setup}$  keep in their state the same information that was imposed on general key exchange algorithms in Section 3.3. We will refer to the instances of local key exchange executions as  $\text{Loc}_{\text{KE}}^s$ , for  $s \in \mathbb{N}$ . The local identity will be implicit in our notation since, in the following discussion we will concentrate our attention on the case where a single local identity  $\text{id}$  is considered. We do this for the sake of rigour and clarity of presentation: by looking at this simplified case we can present our security models in game-based form, whilst taming the complexity of the resulting games. The extension of these results to the more general case where several local identities are considered is straightforward. On the remote side, the identity of the remote process will actually be generated on the fly by the combined actions of the  $\text{Setup}$  algorithm and possibly the protocol execution itself, as it may depend for example on the code of the remote program. For this reason we will enumerate over remote instances as  $\text{Rem}_{\text{KE}}^{i,j}$  for  $i, j \in \mathbb{N}$ , and observe that the value of variable  $\text{oid}$  in this case will be set during the execution of the program itself, rather than passed explicit as an input to one of the algorithms. CORRECTNESS. An AttKE is correct if, after a complete (honest) run between two participants, one local and one remote, and where the remote program is always the one to initiate the communication, both reach the `accept` state, both derive the same key and session identifier and have matching partner identities. More formally, a protocol  $P = \{\text{Setup}, \text{Loc}_{\text{KE}}\}$  is correct if, for any arbitrary identity  $\text{id}$ , the experiment in Figure 7 always returns `T`. We note that our definition of correctness imposes that remote programs always operate as initiators and local machines as the responders in the key exchange.

EXECUTION ENVIRONMENT. The specific flavour of key exchange that we will be considering is clarified by the execution environment in Figure 8. This follows the standard modelling of active attackers, e.g. [20], when one excludes the possibility of corruption (which we do only for the sake of simplicity). There are, however, two modifications that attend to the fact that AttKE remote programs are designed to be executed under attested computation guarantees. On one hand, the adversary is given the power to create as many remote AttKE programs as it may need, by using the `NewLocal` oracle. This reveals the entire code of the remote AttKE program (and implicitly all of its initial internal state, which is assumed to be

<p><b>Game</b> <math>\text{Corr}_{\text{AttKE}}(1^\lambda)</math>:</p> <p><math>(\text{st}_L, \text{Rem}_{\text{KE}}) \leftarrow \text{Setup}(1^\lambda, \text{id})</math>  <math>\text{st}_R \leftarrow \epsilon</math>  <math>m \leftarrow \text{Rem}_{\text{KE}}[\text{st}_R](\epsilon)</math>  <math>t \leftarrow \top</math>  <b>While</b> <math>m \neq \epsilon</math>:    <b>If</b> <math>t</math>: <math>(\text{st}_L, m) \leftarrow \text{Loc}_{\text{KE}}(\text{st}_L, m)</math>    <b>Else</b>: <math>m \leftarrow \text{Rem}_{\text{KE}}[\text{st}_R](m)</math>    <math>t \leftarrow \neg t</math>  <b>Return</b> <math>\text{st}_L.\delta = \text{st}_R.\delta = \text{accept} \wedge \text{st}_L.\text{key} = \text{st}_R.\text{key} \wedge \text{st}_L.\text{sid} = \text{st}_R.\text{sid} \wedge</math>  <math>\text{st}_L.\text{pid} = \text{st}_R.\text{oid} \wedge \text{st}_L.\text{oid} = \text{st}_R.\text{pid} \wedge \text{st}_L.\text{oid} = \text{id}</math></p>
---

Figure 7: Game defining the correctness of an AttKE scheme.

empty) to the adversary. This captures the fact that remote AttKE programs will be loaded into IEE execution environments in an otherwise untrusted remote machine, and it implies that remote AttKE programs cannot keep *any* long term secret information. Intuitively, this limitation will be compensated by the attested computation protocol. On the other hand, the adversary is able to freely interact with remote processes, but it is constrained in its interaction with the local machine. Indeed, the `SendLocal` oracle filters which messages the adversary can deliver to the local machine by checking that these are consistent with at least one remote process that the adversary is interacting with. This captures the fact that AttKE is designed to interact over a partially authenticated channel from the remote machine to the local machine, which will be provided by an attested computation protocol.

<p><b>Game</b> <math>\text{Att}_{\text{AttKE}, \mathcal{A}}(1^\lambda)</math>:</p> <p><math>\text{InsList} \leftarrow []</math>; <math>\text{fake} \leftarrow []</math>  <math>i \leftarrow 0</math>  <math>b \leftarrow \{0, 1\}</math>  <math>b' \leftarrow \mathcal{A}^{\mathcal{O}}(1^\lambda, \text{id})</math>  <b>Return</b> <math>b = b'</math></p> <p><b>Oracle</b> <math>\text{NewLoc}()</math>:</p> <p><math>i \leftarrow i + 1</math>; <math>T_L^i \leftarrow []</math>  <math>(\text{Rem}_{\text{KE}}^i, \text{st}_L^i) \leftarrow \text{Setup}(1^\lambda, \text{id})</math>  <math>\text{InsList}[i] \leftarrow 0</math>  <b>Return</b> <math>\text{Rem}_{\text{KE}}^i</math></p> <p><b>Oracle</b> <math>\text{TestLoc}(i)</math>:</p> <p><b>If</b> <math>\text{st}_L^i.\delta \neq \text{accept}</math> <b>return</b> <math>\perp</math>  <b>If</b> <math>b = 0</math> <b>return</b> <math>\text{st}_L^i.\text{key}</math>  <b>Return</b> <math>\text{fake}(\text{st}_L^i.\text{key})</math></p> <p><b>Oracle</b> <math>\text{SendLoc}(m, i)</math>:</p> <p><b>If</b> <math>\nexists j, (m : T_L^i) \sqsubseteq T_R^{i,j}</math> <b>return</b> <math>\perp</math>  <math>(m', \text{st}_L^i) \leftarrow \text{Loc}_{\text{KE}}^i(\text{st}_L^i, m)</math>  <math>T_L^i \leftarrow m' : m : T_L^i</math>  <b>If</b> <math>\text{st}_L^i.\delta \in \{\text{accept}, \text{derived}\}</math>:    <b>If</b> <math>(\text{st}_L^i.\text{key}, \text{key}^*) \notin \text{fake}</math>:      <math>\text{key}^* \leftarrow \{0, 1\}^\lambda</math>      <math>\text{fake} \leftarrow (\text{st}_L^i.\text{key}, \text{key}^*) : \text{fake}</math>    <b>Return</b> <math>(m', \text{st}_L^i.\text{sid}, \text{st}_L^i.\delta, \text{st}_L^i.\text{pid})</math></p>	<p><b>Oracle</b> <math>\text{RevealLoc}(i)</math>:</p> <p><b>Return</b> <math>\text{st}_L^i.\text{key}</math></p> <p><b>Oracle</b> <math>\text{RevealRem}(i, j)</math>:</p> <p><b>Return</b> <math>\text{st}_R^{i,j}.\text{key}</math></p> <p><b>Oracle</b> <math>\text{NewRem}(i)</math>:</p> <p><math>\text{InsList}[i] \leftarrow \text{InsList}[i] + 1</math>  <math>j \leftarrow \text{InsList}[i]</math>  <math>T_R^{i,j} \leftarrow []</math>; <math>\text{st}_R^{i,j} \leftarrow \epsilon</math>  <b>Return</b> <math>\epsilon</math></p> <p><b>Oracle</b> <math>\text{TestRem}(i, j)</math>:</p> <p><b>If</b> <math>\text{st}_R^{i,j}.\delta \neq \text{accept}</math> <b>return</b> <math>\perp</math>  <b>If</b> <math>b = 0</math> <b>return</b> <math>\text{st}_R^{i,j}.\text{key}</math>  <b>Return</b> <math>\text{fake}(\text{st}_R^{i,j}.\text{key})</math></p> <p><b>Oracle</b> <math>\text{SendRem}(m, i, j)</math>:</p> <p>// No restriction  <math>m' \leftarrow \text{Rem}_{\text{KE}}[\text{st}_R^{i,j}](m)</math>  <math>T_R^{i,j} \leftarrow m' : m : T_R^{i,j}</math>  <b>If</b> <math>\text{st}_R^{i,j}.\delta \in \{\text{accept}, \text{derived}\}</math>:    <b>If</b> <math>(\text{st}_R^{i,j}.\text{key}, \text{key}^*) \notin \text{fake}</math>:      <math>\text{key}^* \leftarrow \{0, 1\}^\lambda</math>      <math>\text{fake} \leftarrow (\text{st}_R^{i,j}.\text{key}, \text{key}^*) : \text{fake}</math>    <b>Return</b> <math>(m', \text{st}_R^{i,j}.\text{sid}, \text{st}_R^{i,j}.\delta, \text{st}_R^{i,j}.\text{pid})</math></p>
---	--

Figure 8: Execution environment for AttKEs.

PARTNERING. We will consider the natural extension of the partnering properties introduced for passive key exchange in Section 3.3 to the AttKE setting. In addition to the

syntactic modifications that result from referring to  $\text{Loc}_{\text{KE}}^s$  and  $\text{Rem}_{\text{KE}}^{i,j}$ , we further restrict validity so that partnering is only valid when it occurs between local and remote instances, in which the latter is the initiator. To this end, we will use the following predicate on two instances  $\text{Loc}_{\text{KE}}^s$  and  $\text{Rem}_{\text{KE}}^{i,j}$  holding  $\text{st}_L^s = (\text{st}^s, \delta^s, \rho^s, \text{sid}^s, \text{pid}^s, \text{oid}^s, \text{key}^s)$  and  $\text{st}_R^{i,j} = (\text{st}^{i,j}, \delta^{i,j}, \rho^{i,j}, \text{sid}^{i,j}, \text{pid}^{i,j}, \text{oid}^{i,j}, \text{key}^{i,j})$ , respectively:

$$P(\text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i,j}) = \begin{cases} \text{T} & \text{if } \text{sid}^s = \text{sid}^{i,j} \wedge \delta^s, \delta^{i,j} \in \{\text{derived}, \text{accept}\} \\ \text{F} & \text{otherwise.} \end{cases}$$

The definition of partner is the obvious one, whereas invalid partners now includes an extra possibility.

**Definition 11** (Partner). Two instances  $\text{Loc}_{\text{KE}}^s$  and  $\text{Rem}_{\text{KE}}^{i,j}$  are partnered if

$$P(\text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i,j}) = \text{T}.$$

**Definition 12** (Valid Partners). A protocol  $\text{AttKE}$  ensures valid partners if the bad event  $\text{notval}$  does not occur, where  $\text{notval}$  is defined as one of the following events occurring:

$$\begin{aligned} & \exists \text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i,j} \text{ s.t. } P(\text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i,j}) = \text{T} \wedge (\text{pid}^s \neq \text{oid}^{i,j} \vee \text{oid}^s \neq \text{pid}^{i,j} \vee \\ & \quad \text{key}^s \neq \text{key}^{i,j} \vee \rho^s \neq \text{responder} \vee \rho^{i,j} \neq \text{initiator}) \\ & \exists \text{Loc}_{\text{KE}}^r, \text{Loc}_{\text{KE}}^s \text{ s.t. } r \neq s \wedge P(\text{Loc}_{\text{KE}}^r, \text{Loc}_{\text{KE}}^s) = \text{T} \\ & \exists \text{Rem}_{\text{KE}}^{i,j}, \text{Rem}_{\text{KE}}^{k,l} \text{ s.t. } (i, j) \neq (k, l) \wedge P(\text{Rem}_{\text{KE}}^{i,j}, \text{Rem}_{\text{KE}}^{k,l}) \end{aligned}$$

For completeness, we present also the adapted definitions of confirmed and unique partners.

**Definition 13** (Confirmed Partners). A protocol  $\text{AttKE}$  ensures confirmed partners if the bad event  $\text{notconf}$  does not occur, where  $\text{notconf}$  is defined as at least one of the following two events occurring:

$$\begin{aligned} & \exists \text{Loc}_{\text{KE}}^s \text{ s.t. } \delta^s = \text{accept} \wedge \forall \text{Rem}_{\text{KE}}^{i,j}, P(\text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i,j}) = \text{F} \\ & \exists \text{Rem}_{\text{KE}}^{i,j} \text{ s.t. } \delta^{i,j} = \text{accept} \wedge \forall \text{Loc}_{\text{KE}}^s, P(\text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i,j}) = \text{F}. \end{aligned}$$

**Definition 14** (Unique Partners). A protocol  $\text{AttKE}$  ensures unique partners if the bad event  $\text{notuni}$  does not occur, where  $\text{notuni}$  is defined as at least one of the following two events occurring:

$$\begin{aligned} & \exists \text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i,j}, \text{Rem}_{\text{KE}}^{i',j'} \text{ s.t.} \\ & \quad (i, j) \neq (i', j') \wedge P(\text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i,j}) = \text{T} \wedge P(\text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i',j'}) = \text{T} \\ & \exists \text{Rem}_{\text{KE}}^{i,j}, \text{Loc}_{\text{KE}}^s, \text{Loc}_{\text{KE}}^{s'} \text{ s.t.} \\ & \quad s \neq s' \wedge P(\text{Loc}_{\text{KE}}^s, \text{Rem}_{\text{KE}}^{i,j}) = \text{T} \wedge P(\text{Loc}_{\text{KE}}^{s'}, \text{Rem}_{\text{KE}}^{i,j}) = \text{T}. \end{aligned}$$

As before, we will consider that an adversary violates entity authentication if he can get a session to accept, but there is no unique and confirmed valid session in its intended partner. More formally, we wish to verify that none of the bad events  $\text{notval}$ ,  $\text{notconf}$ ,  $\text{notuni}$  occurs. In the attested computation scenario, it is common to use one-sided authentication where only the local party receives authentication guarantee. Such definitions can be easily derived from the ones we have presented above, analogously to what was done in Section 3.3. SECURITY. Again, the set of  $\text{TestLoc}$  and  $\text{TestRem}$  queries must be restricted in order to exclude trivial attacks. An adversary is legitimate if it respects the following freshness criteria:

- For all  $\text{TestLoc}(i)$  queries, the following holds:
  1.  $\text{RevealLoc}(i)$  was not queried; and
  2. for all  $\text{Rem}_{\text{KE}}^{j,k}$  s.t.  $\text{P}(\text{Rem}_{\text{KE}}^{j,k}, \text{Loc}_{\text{KE}}^s) = \text{T}$ ,  $\text{RevealRem}(j, k)$  was not queried.
- For all  $\text{TestRem}(i, j)$  queries, the following holds:
  1.  $\text{RevealRem}(i, j)$  was not queried; and
  2. for all  $\text{Loc}_{\text{KE}}^k$  s.t.  $\text{P}(\text{Loc}_{\text{KE}}^k, \text{Rem}_{\text{KE}}^{i,j}) = \text{T}$ ,  $\text{RevealLoc}(i)$  was not queried.

We only consider legitimate adversaries, and say that the winning event  $\text{guess}$  occurs if  $b = b'$  at the end of the experiment. We define  $\text{AttKE}$  security by requiring both mutual authentication of parties and key secrecy.

**Definition 15** (*AttKE security*). An  $\text{AttKE}$  protocol is secure if, for any ppt adversary in Figure 8, and for any local party identifier string  $\text{id}$ : 1. the adversary violates entity authentication with negligible probability  $\text{Pr}[\text{notval} \vee \text{notconf} \vee \text{notuni}]$ ; and 2. its key secrecy advantage  $2 \cdot \text{Pr}[\text{guess}] - 1$  is negligible.

## Generic Construction

We now present a construction of an  $\text{AttKE}$  scheme from any passively secure key exchange protocol, relying additionally on a existentially unforgeable signature scheme. The intuition here is that the attested computation protocol guarantees correct remote execution of a program, but does not ensure uniqueness, i.e., it does not exclude that potentially many replicas of the same key exchange protocol instance could be running in the remote machine. By binding a fresh signature verification key with the identifier for the remote party associated with the key exchange protocol and generating a fresh nonce at the start of every execution, we can remotely execute the key exchange code whilst ensuring one-to-one authentication at the process level. This transformation can be seen as a weaker version of the well-known passive-to-active compilation process by Katz et al. [20], since our target security model is not fully active. We now present the details.

Consider a passively-secure authenticated key exchange protocol  $\Pi$  and a signature scheme  $\Sigma = (\text{Gen}, \text{Sign}, \text{Vrfy})$ . Our construction splits the execution of  $\Pi$  between the local machine and a remote isolated execution environment: the responder will run locally and the initiator will run remotely within a program  $\text{Rem}_{\text{KE}}$ .<sup>4</sup> The code of the remote program will have hardwired into it a unique verification key for the signature scheme. The first activation of  $\text{Rem}_{\text{KE}}$  initialises an internal state and computes a nonce, together with the first message in the key exchange protocol. The party identifier string of the remote process will then be defined to comprise the verification key and the nonce. The local part of the protocol signs the full communication trace so far. Subsequent activations of remote program  $\text{Rem}_{\text{KE}}$  will simply respond according to the key exchange protocol description, rejecting all inputs that fail signature verification. The details of our construction are shown in Figure 9.

- **Setup** first generates a fresh key pair for the signature scheme and constructs program  $\text{Rem}_{\text{KE}}$ , parametrised by algorithm  $\Pi$  and verification key  $\text{pk}$ , as described in Figure 9 (top). In this program state variables  $\delta$ ,  $\rho$ ,  $\text{key}$ ,  $\text{sid}$  and  $\text{pid}$  are all shared with  $\Pi$  (this

<sup>4</sup>Setting the remote machine as the initiator of the protocol is the most common scenario. We considered it for simplicity; the converse can be treated analogously.

<b>Program <math>\text{Rem}_{\text{KE}}(\Pi, \text{pk})</math>:</b>	
Upon activation with input $m$ and state $\text{st}$ :	
If $\text{st} = \epsilon$ :	
$\delta \leftarrow \perp$ ; if $m \neq \epsilon$ then $\delta \leftarrow \text{reject}$	
$t \leftarrow []$ ; $r \leftarrow_{\$} \{0, 1\}^k$ ; $\text{oid} \leftarrow \text{pk}    r$ ; $m' \leftarrow \epsilon$	
Else:	
Parse $(m', \sigma) \leftarrow m$	
If $\Sigma.\text{Vrfy}(\text{pk}, \sigma, m' : t) = \perp$ then $\delta \leftarrow \text{reject}$	
$(m^*, \text{st}) \leftarrow_{\$} \Pi(m', \text{oid}, \text{initiator}, \text{st})$	
$m \leftarrow (m^*, r)$ ; $t \leftarrow m : m' : t$	
If $\delta = \text{reject}$ return $\epsilon$	
Return $m$	
<b>Algorithm <math>\text{Setup}(1^\lambda, \text{id})</math>:</b>	<b>Algorithm <math>\text{Loc}_{\text{KE}}(\text{st}_L, m)</math>:</b>
$(\text{pk}, \text{sk}) \leftarrow_{\$} \Sigma.\text{Gen}(1^k)$	$(\text{id}, \text{st}_{\text{KE}}, \text{sk}, t) \leftarrow \text{st}_L$
$R^* := \text{Rem}_{\text{KE}}(\Pi, \text{pk})$	Parse $(m^*, r) \leftarrow m$
$t \leftarrow []$	$(m', \text{st}_{\text{KE}}) \leftarrow_{\$} \Pi(m^*, \text{id}, \text{responder}, \text{st}_{\text{KE}})$
$\text{st}_{\text{KE}} \leftarrow \epsilon$	$t \leftarrow m' : m : t$
$\text{st}_L \leftarrow (\text{id}, \text{st}_{\text{KE}}, \text{sk}, t)$	$\sigma \leftarrow \Sigma.\text{Sign}(\text{sk}, t)$
Return $(\text{st}_L, R^*)$	$\text{st}_L \leftarrow (\text{id}, \text{st}_{\text{KE}}, \text{sk}, t)$
	Return $((m', \sigma), \text{st}_L)$

Figure 9: Details of the AttKE construction.

is implicit in the figure). The initial value of  $\text{st}_L$  will store  $\text{id}$ , along with the initially empty state for the key exchange  $\text{st}_{\text{KE}}$ , the signing key for the signature scheme and an initially empty trace  $t$  log.

- $\text{Loc}_{\text{KE}}$  takes  $(\text{st}_L, m)$  runs  $\Pi(m, \text{id}, \text{responder}, \text{st}_{\text{KE}})$  to compute the next message  $o$ , produces signature  $\sigma$  of the entire updated protocol trace, and returns the updated state  $\text{st}_L$  and message  $(o, \sigma)$ .

The following theorem establishes the correctness and security of the generic construction.

**Theorem 2.** *Given a correct passively secure key exchange protocol  $\Pi$  and an existentially unforgeable signature scheme  $\Sigma$ , the generic construction above yields a correct and secure AttKE protocol.*

The full proof is given in Appendix B. The intuition behind the proof is that each local instance of the key exchange protocol is bound to a verification key for the signature scheme which is hardwired into the party identifier of the associated remote program code. Each remote instance of the code initialises its own party identifier by attaching a nonce to the verification key. This nonce is also transmitted along with every remote-to-local message. These facts combined with the restriction on the  $\text{SendLoc}$  oracle imply that the adversary is essentially restricted to function in a passive way, by passing around messages between the two  $\text{Send}$  oracles. Therefore the security guarantees can be reduced to the security of the underlying key exchange protocol.

## Utility

As an intermediate result building up to the construction of full-fledged authenticated and private remote attested computation, we will now present a utility theorem that describes precisely the guarantees one obtains when combining an attested computation protocol with an AttKE. Intuitively, this theorem states that attested computation guarantees that the authentication and secrecy assurance offered by AttKE are retained when we use it to establish

session keys with remote IEEs, in the presence of fully active adversaries that control the remote machine, and when the key exchange is composed with arbitrary programs.

Figure 10 shows an idealised game where an adversary must distinguish between two remote machines where an AttKE scheme is executed in combination with an AC scheme. Machine  $\mathcal{M}_R$  is any standard remote machine that is supported by the attested computation protocol, whereas  $\mathcal{M}'_R$  represents a modification of  $\mathcal{M}_R$  where one can tweak the operation of Rem<sub>KE</sub> programs. The differences of  $\mathcal{M}'_R$  with respect to  $\mathcal{M}_R$  are concentrated on the Run interface, which now operates as follows:

- It takes as additional parameters a list `fake` of pairs of keys and Boolean flag `tweak` that, when activated, identifies a process that is running an instance of Rem<sub>KE</sub> composed with some program  $Q$ . This flag triggers the following modifications with respect to the operations of  $\mathcal{M}_R$ .
- When it detects that Rem<sub>KE</sub> has transitioned into `derived` or `accept` state, it will check if the derived key exists in list `fake`. If not, it generates a new random `key*`, and  $(\text{key}, \text{key}^*)$  is added to the list.
- When it detects that program  $Q$  is set to start executing, rather than using the key as an input to  $\phi$ , it uses `fake(key)` instead.

The environment presented to the adversary models a standard attested computation interaction, where it is given total control over the remote machine using oracles `Load` and `Run` (these oracles will either give access to  $\mathcal{M}_R$  or to  $\mathcal{M}'_R$ , depending on a secret bit  $b$  generated in the beginning of the game). The adversary is also able to obtain challenge remote programs using a `NewSession(Q)` oracle that uses the attested computation scheme to compile Rem<sub>KE</sub> composed with arbitrary program  $Q$  of its choice under a mapping function  $\phi_{\text{key}}$  that reveals the relevant parts of the key exchange state (namely the secret key `key`, the party identifiers `oid` and `pid`, the state  $\delta$  and the session identifier `sid`). We observe that such arbitrary programs can leak all of the information revealed by  $\phi_{\text{key}}$  to the attacker. If the adversary chooses to `Load` a challenge program, and if  $\mathcal{M}'_R$  is being used in the game, then it will be tweaked as described above. Whenever `NewSession(Q)` is called, the environment creates a new local session  $i$  that the adversary can interact with using a `Send(i, m)` oracle. The `Send` oracle uses the `Verify` algorithm of the attested computation scheme to validate attested outputs and, if they are accepted, feeds them to the `LocKE` instance (and also ensures that list `fake` is updated). Finally, the adversary can explicitly choose to be tested (as opposed to the implicit testing it may trigger using arbitrary programs  $Q$ ) by calling `Test` on a local instance. This oracle will either return the true key, if  $b = 0$ , or the associated random key that is kept in the `fake` list. As before, we define the winning event `guess` to occur when  $b = b'$  in the end of the game.

The proof of the following theorem can be found in Appendix C.

**Theorem 3** (AttKE utility). *If AttKE is correct and secure and the AC protocol is correct, secure and ensures minimum leakage, then for all ppt adversaries in the utility experiment: 1. the probability that the adversary violates AttKE two-sided entity authentication is negligible; and the key secrecy advantage  $2 \cdot \Pr[\text{guess}] - 1$  is negligible.*

The intuition behind the proof is that one can use the security of attested computation to exclude `Send` queries that do not match a legitimate remote trace. This essentially maps

<p><b>Game</b> <math>\text{Att}_{\text{AttKE}, \mathcal{A}}(1^\lambda)</math>:</p> <p><math>\text{prms}_0 \leftarrow \mathcal{M}_R.\text{Init}(1^\lambda)</math>  <math>\text{prms}_1 \leftarrow \mathcal{M}'_R.\text{Init}(1^\lambda)</math>  <math>\text{PrgList} \leftarrow []</math>  <math>\text{fake} \leftarrow []</math>  <math>i \leftarrow 0</math>  <math>b \leftarrow \{0, 1\}</math>  <math>b' \leftarrow \mathcal{A}^{\mathcal{O}}(\text{prms}_b, \text{id})</math>  Return <math>b = b'</math></p> <p><b>Oracle</b> <math>\text{Load}(R^*)</math>:</p> <p><math>\text{hdl}_0 \leftarrow \mathcal{M}_R.\text{Load}(R^*)</math>  <math>\text{hdl}_1 \leftarrow \mathcal{M}'_R.\text{Load}(R^*)</math>  Return <math>\text{hdl}_b</math></p> <p><b>Oracle</b> <math>\text{Run}(\text{hdl}, \text{in})</math>:</p> <p><math>o_0 \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, \text{in})</math>  <math>\text{tweak} \leftarrow \text{F}</math>  If <math>\text{Program}_{\mathcal{M}'_R}(\text{hdl}) \in \text{PrgList}</math> then <math>\text{flag} \leftarrow \text{T}</math>  <math>(o_1, \text{fake}) \leftarrow \mathcal{M}'_R.\text{Run}(\text{hdl}, \text{in}, \text{tweak}, \text{fake})</math>  Return <math>o_b</math></p>	<p><b>Oracle</b> <math>\text{NewSession}(Q)</math>:</p> <p><math>i \leftarrow i + 1</math>  <math>(\text{Rem}_{\text{KE}}^i, \text{st}_{\text{KE}}^i) \leftarrow \mathcal{S}.\text{Setup}(1^\lambda, \text{id})</math>  <math>(R_i^*, \text{st}_L^i) \leftarrow \mathcal{S}.\text{AC.Compile}(\text{prms}_b, \text{Rem}_{\text{KE}}^i, \phi_{\text{key}}, Q)</math>  <math>\text{in}_{\text{last}}^i \leftarrow \epsilon</math>  <math>\text{PrgList} \leftarrow R_i^* : \text{PrgList}</math>  Return <math>R_i^*</math></p> <p><b>Oracle</b> <math>\text{Send}(m', i)</math>:</p> <p><math>(m, \text{st}_L^i) \leftarrow \mathcal{S}.\text{AC.Verify}(\text{prms}_b, \text{in}_{\text{last}}^i, m', \text{st}_L^i)</math>  If <math>m = \perp</math> then return <math>\perp</math>  <math>(m^*, \text{st}_{\text{KE}}^i) \leftarrow \mathcal{S}.\text{Loc}_{\text{KE}}^i(\text{st}_{\text{KE}}^i, m)</math>  <math>\text{in}_{\text{last}}^i \leftarrow m^*</math>  If <math>\text{st}_{\text{KE}}^i.\delta \in \{\text{derived}, \text{accept}\} \wedge \text{st}_{\text{KE}}^i.\text{key} \notin \text{fake}</math>:  <math>\text{key}^* \leftarrow \{0, 1\}^\lambda</math>  <math>\text{fake} \leftarrow (\text{key}, \text{key}^*) : \text{fake}</math>  Return <math>m^*</math></p> <p><b>Oracle</b> <math>\text{Test}(i)</math>:</p> <p>If <math>\text{st}_{\text{KE}}^i.\delta \neq \text{accept}</math> return <math>\perp</math>  If <math>b = 0</math> then return <math>\text{st}_{\text{KE}}^i.\text{key}</math>  Return <math>\text{fake}(\text{st}_{\text{KE}}^i.\text{key})</math></p>
--	--

Figure 10: Game defining the utility of an AttKE scheme when used in the context of attested computation.

to the restriction in the AttKE security game imposed on the `SendLoc` oracle. The proof is concluded by applying the minimum leakage property of the attested computation protocol to show that any attack by the adversary against AttKE when it is run inside a remote machine can be transformed (via trace simulation) to an attack against the original AttKE when it is run in source code form.

## 8 Secure Outsourced Computation

In this section we build on the results in previous sections to design and analyze a protocol for secure outsourced computation. Informally, we require two properties i) that only the legitimate local user can pass inputs to the outsourced program and ii) that the I/O of the remote program is secret from any observer (even an actively malicious one).

We first give syntax for the protocols that solve this problem, then propose formal definitions for the properties that we outlined above and conclude with a generic construction that combines a key-exchange for attestation, a scheme for attested computation and an authenticated encryption scheme.

**SYNTAX.** A *Secure Outsourced Computation* scheme (SOC) for a remote machine  $\mathcal{M}_R$  is defined by the following algorithms:

- $\text{Compile}(\text{prms}, P, \text{id})$  is the program compilation algorithm. On input a program  $P$  and a party identifier  $\text{id}$ , it outputs a compiled program  $P^*$ , together with an initial state  $\text{st}_l$  for the local side algorithms. We assume that initially  $\text{st}_l.\text{accept} = \perp$ . Note that unlike the AC compilation algorithm, this algorithm only takes one program as input, as this scheme is intended for providing guarantees for the whole trace and not only an initial segment.
- $\text{BootStrap}(\text{prms}, o, \text{st}_l)$  is the client side initialization algorithm. On input  $o$  (presumably

the last message from the remote machine) it returns the next message  $i$  to be delivered to the remote machine in the bootstrapping step, together with the updated local state. We assume `BootStrap` sets an `accept` flag to `T` when the initialization process successfully terminates.

- `Verify(prms, o*, stl)` is the verification algorithm. It fulfills the same function as the AC verification algorithm. Note that, as all the inputs are provided by the local machine, we do not need to feed it the last input as it can be stored in the state. It is expected to return  $\perp$  if `stl.accept`  $\neq$  `T`.
- `Encode(prms, i, stl)` is the encoding algorithm. On input the local state and the next intended input for  $P$ , it returns the next input  $i^*$  for  $P^*$  together with the updated local state. It is expected to return  $\perp$  if `stl.accept`  $\neq$  `T`.
- `Attest(prms, hdl, i)` is, as in an AC scheme, the (untrusted) attestation algorithm.

A party  $A$  with identifier `id` who wants to outsource program  $P$  to the remote machine first compiles  $P$  with his `id`, thus obtaining  $P^*$  and some secret data `stl`. He then loads  $P^*$  on the remote machine using some untrusted protocol. As it is, the program  $P^*$  is not ready to receive inputs intended for  $P$ : an initial bootstrapping phase (until `BootStrap` sets the `accept` flag) is necessary to establish some shared secrets between the IEE in which  $P^*$  is executed and  $A$ . Then when  $A$  wants to send an input to the remote execution, he encodes it using `Encode`, sends it (using `Attest`) and verifies the output provided by `Attest` using `Verify`.

In this section, for simplicity reasons, we assume that the program  $P$  is deterministic. However, as for an AC scheme it would be easy to extend all the definitions to a non-deterministic program.

**INPUT INTEGRITY.** While security of attested computation aims at ensuring that a trace was honestly produced on the remote side, it does nothing to restrict the provenance of the inputs received.

We provide a stronger notion which we call *input integrity* which, intuitively, ensures that if a program is compiled by a party with identifier `id`, then only that party may use the remote compiled program. We ensure this property by making sure that the local and remote views coincide (up to the last message exchanged, which may not have yet been delivered). The following formula  $\Psi$  which relates two input/output traces captures this intuition.

$$\Psi(T, T') := T = T' \vee \exists o. (T = o :: T') \exists i. (T' = i :: T)$$

The formalization that we provide in Figure 11 is as follows. The adversary chooses a program  $P$  that is compiled with an honest party's `id` yielding  $P^*$  (which is given to the adversary). The adversary is given access to two oracles. A bootstrapping oracle that simply executes `BootStrap` honestly; and a send oracle that verifies the last (presumed) output of the remote program and encodes the next input (which is provided by the adversary), while keeping track of the local view of the trace. The goal of the adversary is then create a mismatch between the local and remote view of the trace.

**Definition 16** (Input Integrity). We say that a SOC scheme satisfies *input integrity* if there exists a polynomial time algorithm `Translate` such that for all ppt  $\mathcal{A}$  the experiment described in Figure 11 returns true with probability negligible in the security parameter.

<p><b>Game</b> <math>\text{Int}_{\text{SOC}, \mathcal{A}}(1^\lambda)</math>:</p> <p><math>\text{prms} \leftarrow \mathcal{M}_R.\text{Init}(1^\lambda)</math>  <math>(P, \text{st}_A) \leftarrow \mathcal{A}_1(\text{prms})</math>  <math>(P^*, \text{st}_I) \leftarrow \mathcal{C}(\text{prms}, P, \text{id})</math>  <math>\text{tr} \leftarrow \square</math>  Run <math>\mathcal{A}_2^{\mathcal{O}, \mathcal{M}_R}(\text{st}_A, P^*)</math>  If <math>\nexists \text{hdl}</math> such that      <math>\text{Program}_{\mathcal{M}_R}(\text{hdl}) = P^* \wedge</math>      <math>\text{Translate}(\text{prms}, \text{Trace}_{\mathcal{M}_R}(\text{hdl})) \neq \square</math>  Return <math>\perp</math>  <math>\text{hdl} \leftarrow \text{Program}_{\mathcal{M}_R}^{-1}(P^*)</math>  <math>T \leftarrow \text{Translate}(\text{prms}, \text{Trace}_{\mathcal{M}_R}(\text{hdl}))</math>  <math>T' \leftarrow \text{tr}</math>  Return <math>\neg\Psi(T, T')</math></p>	<p><b>Oracle</b> <math>\text{Send}(o^*, i)</math>:</p> <p><math>o, \text{st}_I \leftarrow \text{Verify}(\text{prms}, o^*, \text{st}_I)</math>  If <math>o = \perp</math> Return <math>\perp</math>  <math>i^*, \text{st}_I \leftarrow \text{Encode}(\text{prms}, i, \text{st}_I)</math>  <math>\text{tr} \leftarrow i : o : \text{tr}</math>  Return <math>o, i^*</math></p> <p><b>Oracle</b> <math>\text{BootStrap}(o)</math>:</p> <p>If <math>\text{st}_I.\text{accept}</math>  Return <math>\perp</math>  <math>i, \text{st}_I \leftarrow \text{BootStrap}(\text{prms}, o, \text{st}_I)</math>  Return <math>i</math></p>
--	--

Figure 11: Input integrity of a SOC scheme

**INPUT PRIVACY.** We define the privacy of I/O with an indistinguishability game. One important point here is that we chose to restrict the class of programs we consider to *length-uniform* (written lu) programs. A program is length uniform if the length of its outputs depends only on the length of its inputs. Intuitively, this is because the encryption scheme is allowed to leak the length of the messages, which in turn would leak information about the inputs for a non lu program.

The formalization described in Figure 12 is as follows. We start by choosing a bit  $b$  that will determine whether the adversary will be talking with the left send oracle or the right send oracle (described later). As for input integrity, the adversary then chooses a program  $P$ . We compile it for an honest party's identifier and give the resulting  $P^*$  to the adversary. The adversary is also given access to the bootstrapping oracle. In addition, he is given access to a left or right send oracle. This oracle, on a request with the last candidate output of the remote machine and two inputs  $i_0$  and  $i_1$ , verifies the last candidate output and, depending on the bit  $b$ , encodes either  $i_0$  or  $i_1$  and returns the result. The goal of the adversary is to guess the bit  $b$  with non-negligible bias from  $1/2$ .

<p><b>Game</b> <math>\text{Priv}_{\text{SOC}, \mathcal{A}}(1^\lambda)</math>:</p> <p><math>b \leftarrow \{0, 1\}</math>  <math>\text{prms} \leftarrow \mathcal{M}_R.\text{Init}(1^\lambda)</math>  <math>(P, \text{st}_A) \leftarrow \mathcal{A}_1(\text{prms})</math>  If <math>\neg \text{lu}(P)</math> Return <math>b' \leftarrow \{0, 1\}</math>  <math>(P^*, \text{st}_I) \leftarrow \mathcal{C}(\text{prms}, P, \text{id})</math>  <math>b' \leftarrow \mathcal{A}_2(\text{st}_A, P^*)^{\mathcal{O}, \mathcal{M}_R}</math>  Return <math>b = b'</math></p>	<p><b>Oracle</b> <math>\text{Send}_b(o^*, i_0, i_1)</math>:</p> <p><math>o, \text{st}_I \leftarrow \text{Verify}(\text{prms}, o^*, \text{st}_I)</math>  If <math> m_0  \neq  m_1 </math> Return <math>\perp</math>  <math>i^*, \text{st}_I \leftarrow \text{Encode}(\text{prms}, i_b, \text{st}_I)</math>  Return <math>i^*</math></p> <p><b>Oracle</b> <math>\text{BootStrap}(o)</math>:</p> <p>If <math>\text{st}_I.\text{accept}</math>  Return <math>\perp</math> // (1 init max)  <math>i, \text{st}_I \leftarrow \text{BootStrap}(\text{prms}, o, \text{st}_I)</math>  Return <math>i</math></p>
--	--

Figure 12: Input privacy of a SOC scheme

**Definition 17.** We say that a SOC scheme satisfies *input privacy* if, for all ppt  $\mathcal{A}$ , the experiment in Figure 12 returns true with probability  $1/2$  up to a negligible function.

This definition ensures that there exist no two traces (with messages of the same length) played by an honest party over a SOC protocol that are distinguishable for an (active) adversary. This means that no adversary can gain information on the inputs sent out by a local machine using a SOC scheme, besides the length of the messages exchanged, achieving our goal of hiding the honest party's inputs.

**Definition 18.** We say that a SOC scheme is *secure* if it satisfies both input privacy and input integrity.

AN IMPLEMENTATION OF A SECURE SOC SCHEME. Having defined what security we expect from a SOC scheme, we now define a scheme that satisfies these requirements. We base our construction on an AttKE, and an AC scheme. The main idea is using the AttKE to establish a key between the party agent and the IEE, and then communicate with the IEE over the secure channel established with this key.

Formally, let  $(\text{Compile}, \text{Attest}, \text{Verify})$  be an AC scheme,  $(\text{Setup}, \text{Loc}_{\text{KE}})$  be an AttKE and  $(E, D, K)$  be an authenticated encryption scheme. Figure 13 defines a SOC scheme. The most important part is the compilation part, which uses the AC scheme compilation to compile the composition of the  $\text{Rem}_{\text{KE}}$  program generated by  $\text{Setup}$  together with program  $P$  running over a secure channel (denoted by  $C(P)$ ). The initial local state is the union of the state provided by the AC compilation and the AttKE setup. The program  $C(P)$  simply decrypts the message it receives checks that the sequence number of the message matches its view the passes the decrypted message to  $P$ . It then retrieves the output of  $P$ , appends the corresponding next sequence number and outputs it. This mechanism ensures that all messages received (resp. sent out) by  $P^*$  after the bootstrapping phase have the form  $E(i\#m, k)$  where  $i$  is the position of the message in the trace,  $m$  is the message intended to (resp. produced by)  $P$ , and  $k$  is the key established by the AttKE.

On the local side, the bootstrapping mechanism simply consists of running the local KE over the AC protocol as already described in the utility definition. Once the key has been established, the local state keeps track of the local view of the sequence number. Verifying an output consists in decrypting it and checking that the sequence number against the local view of it. Encoding an input, is just appending the correct sequence number and encrypting it with the shared key.

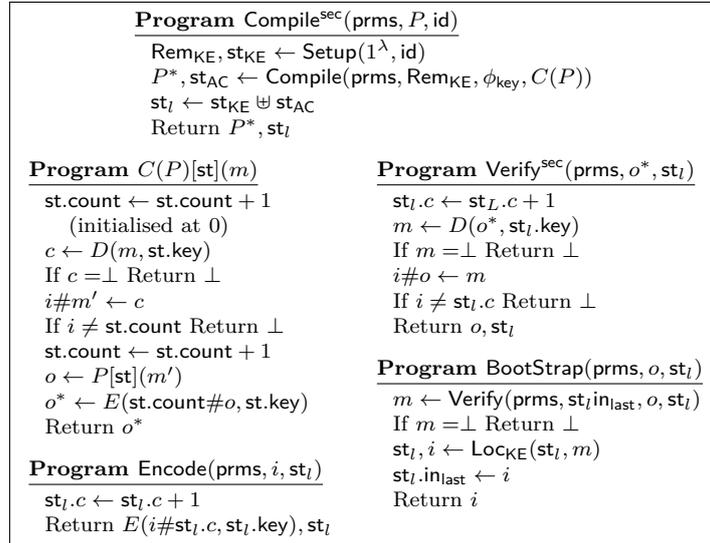


Figure 13: SOC algorithms

**Theorem 4.** *If  $(\text{Compile}, \text{Attest}, \text{Verify})$  is a correct and secure AC scheme,  $\{\text{Setup}, \text{Loc}_{\text{KE}}, \text{Rem}_{\text{KE}}\}$  is a secure AttKE and  $(E, D, KG)$  is an secure authenticated encryption scheme, then the SOC presented in Figure 13 is secure.*

The proof is provided in Appendix D, we provide here a sketch of proof. We first do a game hop that consists in replacing the key established by  $P^*$  and the party using the AttKE by a “magically” shared fresh key. The utility property of the AttKE provides us with the fact that we can replace the key shared by the remote machine and the local agent by a freshly generated key. We are left with showing that this key is shared with an IEE which is indeed running  $P^*$  and not  $\text{Compile}(\text{Rem}_{\text{KE}}, \phi_{\text{key}}, Q)$  for some other  $Q$ , this is provided by the security of the AC scheme.

We then prove input integrity by remarking that injecting new messages in the trace would contradict the unforgeability of the authenticated encryption scheme. The sequence number ensures that the messages are delivered in the right order and that replays are impossible.

We remark that the input integrity property ensures that we know that the only meaningful action the adversary can take is to forward messages between the remote and local machines. Taking advantage of that fact we can reduce the input privacy game to the IND-CPA property of the authenticated encryption.

## 9 Conclusion

This paper offers a set of building blocks for constructing protocols that leverage the guarantees of IEEs. First, we define and construct attested computation based on IEEs. In the process we identify and formalize two key properties that such protocols need to satisfy to be useful: composition awareness and minimal leakage. Our instantiation of attested computation relies on SGX.

The next component that we provide are key-exchange protocols between a remote party and an IEE. Such protocols are components which need to be used in any protocol where secrecy of data communicated to and from the IEE is important. Our contribution is to adapt existing models of security for key-exchange to the novel setting where protocol participants do not necessarily have an a-priori identity – IEEs cannot be uniquely identified before actively communicating with them. For constructions, we present a modular approach where by combining a passively secure key exchange protocol with an arbitrary attested computation protocol we obtain a fully secure key-exchange protocol. As an application, we show how to use attested computation, key-exchange and symmetric authenticated encryption to generically construct a secure outsourced computation scheme for arbitrary functionalities.

In terms of follow-up work, the natural next step is to use the building blocks that we provide to construct other protocols. One interesting target (which generalizes the application in this paper) is to construct secure multi-party computation based on IEEs and experimentally compare their efficiency with the existent software-only alternatives.

## References

- [1] I. Anati, S. Gueron, S. Johnson, and V. Scarlata. Innovative technology for cpu based attestation and sealing. In *Workshop on Hardware and Architectural Support for Security and Privacy*, page 10, 2013.
- [2] Andrew Baumann, Marcus Peinado, and Galen C. Hunt. Shielding applications from an untrusted cloud with haven. In *11th USENIX Symposium on Operating Systems*

- Design and Implementation, OSDI '14, Broomfield, CO, USA, October 6-8, 2014.*, pages 267–283. USENIX Association, 2014.
- [3] Mihir Bellare, Joe Kilian, and Phillip Rogaway. The security of cipher block chaining. In *Advances in Cryptology - CRYPTO '94, 14th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1994, Proceedings*, volume 839 of *Lecture Notes in Computer Science*, pages 341–358. Springer, 1994.
  - [4] Siavosh Benabbas, Rosario Gennaro, and Yevgeniy Vahlis. Verifiable delegation of computation over large datasets. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 111–131. Springer, 2011.
  - [5] Ernest F. Brickell, Jan Camenisch, and Liqun Chen. Direct anonymous attestation. In *Proceedings of the 11th ACM Conference on Computer and Communications Security, CCS 2004, Washington, DC, USA, October 25-29, 2004*, pages 132–145. ACM, 2004.
  - [6] Ernie Brickell, Liqun Chen, and Jiangtao Li. A new direct anonymous attestation scheme from bilinear maps. In *Trusted Computing - Challenges and Applications, First International Conference on Trusted Computing and Trust in Information Technologies, Trust 2008, Villach, Austria, March 11-12, 2008, Proceedings*, volume 4968 of *Lecture Notes in Computer Science*, pages 166–178. Springer, 2008.
  - [7] Christina Brzuska, Nigel P. Smart, Bogdan Warinschi, and Gaven J. Watson. An analysis of the EMV channel establishment protocol. In *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 373–386. ACM, 2013.
  - [8] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 136–145. IEEE Computer Society, 2001.
  - [9] Luigi Catuogno, Alexandra Dmitrienko, Konrad Eriksson, Dirk Kuhlmann, Gianluca Ramunno, Ahmad-Reza Sadeghi, Steffen Schulz, Matthias Schunter, Marcel Winandy, and Jing Zhan. Trusted virtual domains - design, implementation and lessons learned. In *Trusted Systems, First International Conference, INTRUST 2009, Beijing, China, December 17-19, 2009. Revised Selected Papers*, volume 6163 of *Lecture Notes in Computer Science*, pages 156–179. Springer, 2009.
  - [10] Anupam Datta, Jason Franklin, Deepak Garg, and Dilsun Kirli Kaynar. A logic of secure systems and its application to trusted computing. In *30th IEEE Symposium on Security and Privacy, SP 2009, 17-20 May 2009, Oakland, California, USA*, pages 221–236. IEEE Computer Society, 2009.
  - [11] Aurélien Francillon, Quan Nguyen, Kasper Bonne Rasmussen, and Gene Tsudik. A minimalist approach to remote attestation. In *Design, Automation & Test in Europe Conference & Exhibition, DATE 2014, Dresden, Germany, March 24-28, 2014*, pages 1–6. European Design and Automation Association, 2014.

- [12] He Ge and Stephen R. Tate. A direct anonymous attestation scheme for embedded devices. In *Public Key Cryptography - PKC 2007, 10th International Conference on Practice and Theory in Public-Key Cryptography, Beijing, China, April 16-20, 2007, Proceedings*, volume 4450 of *Lecture Notes in Computer Science*, pages 16–30. Springer, 2007.
- [13] C. Gebhardt and A. Tomlinson. Secure virtual disk images for grid computing. In *Third Asia-Pacific Trusted Infrastructure Technologies Conference, APTC '08*, pages 19–29, Washington, DC, USA, 2008. IEEE Computer Society.
- [14] Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, volume 6223 of *Lecture Notes in Computer Science*, pages 465–482. Springer, 2010.
- [15] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct nizks without pcps. In *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 626–645. Springer, 2013.
- [16] Kenneth A. Goldman, Ronald Perez, and Reiner Sailer. Linking remote attestation to secure tunnel endpoints. In *Proceedings of the 1st ACM Workshop on Scalable Trusted Computing, STC 2006, Alexandria, VA, USA, November 3, 2006*, pages 21–24. ACM, 2006.
- [17] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.*, 17(2):281–308, 1988.
- [18] Matthew Hoekstra, Reshma Lal, Pradeep Pappachan, Vinay Phegade, and Juan del Cuvillo. Using innovative instructions to create trustworthy software solutions. In *HASP 2013, The Second Workshop on Hardware and Architectural Support for Security and Privacy, Tel-Aviv, Israel, June 23-24, 2013*, page 11. ACM, 2013.
- [19] Intel. *Software Guard Extensions Programming Reference*, 2014. <https://software.intel.com/sites/default/files/managed/48/88/329298-002.pdf>.
- [20] Jonathan Katz and Moti Yung. Scalable protocols for authenticated group key exchange. In *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 110–125. Springer, 2003.
- [21] Patrick Koeberl, Steffen Schulz, Ahmad-Reza Sadeghi, and Vijay Varadharajan. Trustlite: a security architecture for tiny embedded devices. In *Ninth EuroSys Conference 2014, EuroSys 2014, Amsterdam, The Netherlands, April 13-16, 2014*, pages 10:1–10:14. ACM, 2014.
- [22] J. M. McCune, B. J. Parno, A. Perrig, M. Reiter, and H. Isozaki. Flicker: An execution infrastructure for tcb minimization. *SIGOPS Oper. Syst. Rev.*, 42(4):315–328, April 2008.

- [23] Microsoft. *BitLocker Drive Encryption: Data Encryption Toolkit for Mobile PCs: Security Analysis*, 2007. <https://technet.microsoft.com/en-us/library/cc162804.aspx>.
- [24] Job Noorman, Pieter Agten, Wilfried Daniels, Raoul Strackx, Anthony Van Herrewege, Christophe Huygens, Bart Preneel, Ingrid Verbauwhede, and Frank Piessens. Sancus: Low-cost trustworthy extensible networked devices with a zero-software trusted computing base. In *Proceedings of the 22th USENIX Security Symposium, Washington, DC, USA, August 14-16, 2013*, pages 479–494. USENIX Association, 2013.
- [25] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*, pages 238–252. IEEE Computer Society, 2013.
- [26] Felix Schuster, Manuel Costa, Cédric Fournet, Christos Gkantsidis, Marcus Peinado, Gloria Mainar-Ruiz, and Mark Russinovich. VC3: trustworthy data analytics in the cloud using SGX. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 38–54. IEEE Computer Society, 2015.
- [27] Sean W. Smith. Outbound authentication for programmable secure coprocessors. *Int. J. Inf. Sec.*, 3(1):28–41, 2004.
- [28] Ben Smyth, Mark Ryan, and Liqun Chen. Direct anonymous attestation (DAA): ensuring privacy with corrupt administrators. In *Security and Privacy in Ad-hoc and Sensor Networks, 4th European Workshop, ESAS 2007, Cambridge, UK, July 2-3, 2007, Proceedings*, volume 4572 of *Lecture Notes in Computer Science*, pages 218–231. Springer, 2007.

## A Proof of Theorem 1

The proof is a sequence of three games presented in Figure 14 and Figure 15. The first game is simply the AC security game instantiated with our protocol. In game  $G_1^{\text{AC},A}(1^\lambda)$ , 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  $\Sigma$ . Given that the two games are identical until this event occurs, we have that

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

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

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

Adversary  $\mathcal{B}$  simulates the environment of  $G_1^{\text{AC},A}$  as follows: the operation of machine  $\mathcal{M}_R$  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

<p><b>Game G0<sub>AC,A</sub>(1<sup>λ</sup>):</b></p> <pre> prms ← \$ M<sub>R</sub>.Init(1<sup>λ</sup>) (P, φ, Q, n, st<sub>A</sub>) ← \$ A<sub>1</sub>(prms)  (R*, (R*, ios, stage')) ← Compile(prms, P, φ, Q) For k ∈ [1..n]:   (i<sub>k</sub>, o<sub>k</sub><sup>*</sup>, st<sub>A</sub>) ← \$ A<sub>2</sub><sup>M<sub>R</sub></sup>(st<sub>A</sub>)   If stage' = 1:     Parse (o'<sub>k</sub>, σ) ← o<sub>k</sub><sup>*</sup>; (o<sub>k</sub>, finished, stage) ← o'<sub>k</sub>     If Σ.Vrfy(prms, σ, (R*, (i<sub>k</sub>, o'<sub>k</sub>) : ios)):       (o<sub>R,k</sub>, ios) ← (o'<sub>k</sub>, (i<sub>k</sub>, o'<sub>k</sub>) : ios)       If finished = T: stage' ← 2     Else: Return F    Else o<sub>R,k</sub> ← o<sub>k</sub> T' ← (i<sub>1</sub>, o<sub>R,1</sub>, ..., i<sub>n</sub>, o<sub>R,n</sub>) Define R := Compose<sub>φ</sub>(P; Q) For hdl* s.t. Program<sub>M<sub>R</sub></sub>(hdl*) = R*:   T ← ATrace<sub>R[st;Coins<sub>M<sub>R</sub></sub>](hdl*)</sub>(i<sub>1</sub>, ..., i<sub>n</sub>)   If T ⊆ T' ∧ T ⊆ Translate(prms, ATrace<sub>M<sub>R</sub></sub>(hdl*)):     Return F Return T </pre>	<p><b>Game G1<sub>AC,A</sub>(1<sup>λ</sup>):</b></p> <pre> prms ← \$ M<sub>R</sub>.Init(1<sup>λ</sup>) (P, φ, Q, n, st<sub>A</sub>) ← \$ A<sub>1</sub>(prms) sforge ← F (R*, (R*, ios, stage')) ← Compile(prms, P, φ, Q) For k ∈ [1..n]:   (i<sub>k</sub>, o<sub>k</sub><sup>*</sup>, st<sub>A</sub>) ← \$ A<sub>2</sub><sup>M<sub>R</sub></sup>(st<sub>A</sub>)   If stage' = 1:     Parse (o'<sub>k</sub>, σ) ← o<sub>k</sub><sup>*</sup>; (o<sub>k</sub>, finished, stage) ← o'<sub>k</sub>     If Σ.Vrfy(prms, σ, (R*, (i<sub>k</sub>, o'<sub>k</sub>) : ios)):       (o<sub>R,k</sub>, ios) ← (o'<sub>k</sub>, (i<sub>k</sub>, o'<sub>k</sub>) : ios)       If finished = T: stage' ← 2     Else: Return F   If (((R*, (i<sub>1</sub>, o<sub>R,1</sub>, ..., i<sub>k</sub>, o<sub>k</sub>)), *, σ') ∉ Trace<sub>M<sub>R</sub></sub>(0)):     sforge ← T; Return F   Else o<sub>R,k</sub> ← o<sub>k</sub> T' ← (i<sub>1</sub>, o<sub>R,1</sub>, ..., i<sub>n</sub>, o<sub>R,n</sub>) Define R := Compose<sub>φ</sub>(P; Q) For hdl* s.t. Program<sub>M<sub>R</sub></sub>(hdl*) = R*:   T ← ATrace<sub>R[st;Coins<sub>M<sub>R</sub></sub>](hdl*)</sub>(i<sub>1</sub>, ..., i<sub>n</sub>)   If T ⊆ T' ∧ T ⊆ Translate(prms, ATrace<sub>M<sub>R</sub></sub>(hdl*)):     Return F Return T </pre>
--	--

Figure 14: First game hop for the proof of security of our AC protocol.

expected to compute a signature on message  $m$ , algorithm  $\mathcal{B}$  calls its own oracle on  $(R^*, m)$  to obtain  $\sigma$ .

When `sforge` is set, according to the rules of game  $G_1^{AC,A}$ , algorithm  $\mathcal{B}$  outputs message  $(R^*, ios)$  and candidate signature  $\sigma$ . 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  $(R^*, (i_1, o'_1, \dots, i_k, o'_k))$  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  $(R^*, (i_1, o'_1, \dots, i_k, o'_k))$  in its trace. Since we know that this is not the case,  $(R^*, ios)$  could not have been queried from the signature oracle. We conclude therefore that  $\mathcal{B}$  outputs a valid forgery whenever `sforge` occurs.

In game  $G_2^{AC,A}(1^\lambda)$ , 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  $\Pi$ . Given that the two games are identical until this event occurs, we have that

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

We upper bound the distance between these two games, by constructing an adversary  $\mathcal{C}$  against the existential unforgeability of MAC scheme  $\Pi$  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  $G_2^{AC,A}$  as follows: the operation of machine  $\mathcal{M}_R$  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

<p><b>Game G1<sub>AC,A</sub>(1<sup>λ</sup>):</b></p> <pre> prms ← \$ M<sub>R</sub>.Init(1<sup>λ</sup>) (P, φ, Q, n, st<sub>A</sub>) ← \$ A<sub>1</sub>(prms) sforge ← F (R*, (R*, ios, stage')) ← Compile(prms, P, φ, Q) For k ∈ [1..n]:   (i<sub>k</sub>, o<sub>k</sub><sup>*</sup>, st<sub>A</sub>) ← \$ A<sub>2</sub><sup>M<sub>R</sub></sup>(st<sub>A</sub>)   If stage' = 1:     Parse (o'<sub>k</sub>, σ) ← o<sub>k</sub><sup>*</sup>; (o<sub>k</sub>, finished, stage) ← o'<sub>k</sub>     If Σ.Vrfy(prms, σ, (R*, (i<sub>k</sub>, o'<sub>k</sub>)): ios):       (o<sub>R,k</sub>, ios) ← (o'<sub>k</sub>, (i<sub>k</sub>, o'<sub>k</sub>)): ios       If finished = T: stage' ← 2     Else: Return F   If (((R*, (i<sub>1</sub>, o<sub>R,1</sub>, ..., i<sub>k</sub>, o<sub>k</sub>)), *, σ') ∉ Trace<sub>M<sub>R</sub></sub>(0):     sforge ← T; Return F    Else o<sub>R,k</sub> ← o<sub>k</sub>   T' ← (i<sub>1</sub>, o<sub>R,1</sub>, ..., i<sub>n</sub>, o<sub>R,n</sub>)   Define R := Compose<sub>φ</sub>(P; Q)   For hdl* s.t. Program<sub>M<sub>R</sub></sub>(hdl*) = R*:     T ← ATrace<sub>R[st;Coins<sub>M<sub>R</sub></sub>](hdl*)</sub>(i<sub>1</sub>, ..., i<sub>n</sub>)     If T ⊆ T' ∧ T ⊆ Translate(prms, ATrace<sub>M<sub>R</sub></sub>(hdl*)):       Return F   Return T </pre>	<p><b>Game G2<sub>AC,A</sub>(1<sup>λ</sup>):</b></p> <pre> prms ← \$ M<sub>R</sub>.Init(1<sup>λ</sup>) (P, φ, Q, n, st<sub>A</sub>) ← \$ A<sub>1</sub>(prms) sforge ← F; mforge ← F (R*, (R*, ios, stage')) ← Compile(prms, P, φ, Q) For k ∈ [1..n]:   (i<sub>k</sub>, o<sub>k</sub><sup>*</sup>, st<sub>A</sub>) ← \$ A<sub>2</sub><sup>M<sub>R</sub></sup>(st<sub>A</sub>)   If stage' = 1:     Parse (o'<sub>k</sub>, σ) ← o<sub>k</sub><sup>*</sup>; (o<sub>k</sub>, finished, stage) ← o'<sub>k</sub>     If Σ.Vrfy(prms, σ, (R*, (i<sub>k</sub>, o'<sub>k</sub>)): ios):       (o<sub>R,k</sub>, ios) ← (o'<sub>k</sub>, (i<sub>k</sub>, o'<sub>k</sub>)): ios       If finished = T: stage' ← 2     Else: Return F   If (((R*, (i<sub>1</sub>, o<sub>R,1</sub>, ..., i<sub>k</sub>, o<sub>k</sub>)), *, σ') ∉ Trace<sub>M<sub>R</sub></sub>(0):     sforge ← T; Return F   If ∄ hdl*. Program<sub>M<sub>R</sub></sub>(hdl*) = R* ∧   (i<sub>1</sub>, o<sub>R,1</sub>, ..., i<sub>k</sub>, o<sub>R,k</sub>) ⊆ Translate(prms, ATrace<sub>M<sub>R</sub></sub>(hdl*)):     Then mforge ← T; Return F   Else o<sub>R,k</sub> ← o<sub>k</sub>   T' ← (i<sub>1</sub>, o<sub>R,1</sub>, ..., i<sub>n</sub>, o<sub>R,n</sub>)   Define R := Compose<sub>φ</sub>(P; Q)   For hdl* s.t. Program<sub>M<sub>R</sub></sub>(hdl*) = R*:     T ← ATrace<sub>R[st;Coins<sub>M<sub>R</sub></sub>](hdl*)</sub>(i<sub>1</sub>, ..., i<sub>n</sub>)     If T ⊆ T' ∧ T ⊆ Translate(prms, ATrace<sub>M<sub>R</sub></sub>(hdl*)):       Return F   Return T </pre>
--	--

Figure 15: Second game hop for the proof of security of our AC protocol.

remote machine requests a MAC on message  $m$  from the security module, algorithm  $\mathcal{C}$  calls its own oracle on  $(R^*, m)$  to obtain  $t$ .

Let  $T \leftarrow (i_1, o_{R,1}, \dots, i_k, o_{R,k})$ . When  $mforge$  is set according to the rules of game  $G_2^{AC,A}$ , algorithm  $\mathcal{C}$  retrieves the trace of the process with handle 0 running  $S^*$ , locates the input/output pair  $((R^*, T), t, \sigma')$  and outputs message  $(R^*, T)$  and candidate tag  $t$ . To see this is a valid forgery, first observe that, having failed the  $sforge$  check, we know that  $((R^*, T), t, \sigma')$  is in the trace of the process with handle 0, so by its construction we also know that the corresponding input  $((R^*, T), t)$  must contain a valid tag. It suffices to establish that message  $(R^*, T)$  could not have been queried from the Auth oracle. Suppose that the first part of the  $mforge$  check failed, i.e., that  $\nexists hdl^*. Program_{M_R}(hdl^*) = R^*$ . 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  $R^*$  must have been caused by the execution of an instance of  $R^*$ . Now suppose some instances of  $R^*$  were indeed running in the remote machine, but that none of them displayed the property  $(i_1, o_{R,1}, \dots, i_k, o_{R,k}) \subseteq Translate(prms, ATrace_{M_R}(hdl^*))$ . Then, by the construction of  $R^*$ , we can also exclude that  $(R^*, T)$  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  $G_2^{AC,A}$ . To see this, observe that when the game reaches the final check, we have the guarantee that

$$\exists hdl^*. Program_{M_R}(hdl^*) = R^* \wedge (i_1, o_{R,1}, i_2, o_{R,2}, \dots, i_k, o_{R,k}) \subseteq Translate(prms, Trace_{M_R}(hdl^*))$$

By the construction of  $R^*$ , it immediately follows that the final check in the game will always cause the adversary to lose:

- $T$  is fixed by the input sequence, the value of the randomness tape and the semantics of  $R$ , which determines the sequence of outputs  $(o_{1,L}, \dots, o_{n,L})$ .
- The above existential guarantee for  $\text{hdl}^*$  implies that an instance of  $R^*$  in the remote machine received the same initial  $k$  input sequence as that fixed by  $T$ .
- Since we have

$$T \leftarrow \text{ATrace}_{R[\text{st}; \text{Coins}, \mathcal{M}_R(\text{hdl}^*)]}(i_1, \dots, i_n)$$

we also know that the randomness tape used to produce values for  $T$  is identical to the one used in  $\text{hdl}^*$ .

- One can therefore inductively deduce, by the semantics of  $R^*$ , that the same process has produced an initial sequence of  $k$  outputs that (modulo the action of `Translate`) is identical to that included in  $T$ .
- Subsequent inputs after  $\text{stage}' \leftarrow 2$  may produce  $n - k$  additional non-attested outputs that are appended to  $T'$ , and that differ from  $T$ . However, the above observation implies that  $T \sqsubseteq T'$ , and our security claim follows.

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 6. This is easy to ascertain given the simulator behaviour described in Figure 16:  $\mathcal{S}_1$  and  $\mathcal{S}_3$  follow the exact description of the actual machine, modulo the generation of  $(\text{pk}, \text{sk})$  and `key`.  $\mathcal{S}_2$  takes an external output produced by  $R[\text{st}](i)$  and returns an output in accordance to the behaviour of  $\mathcal{M}_R$ , 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.  $\square$

## B Proof for Theorem 2

Our proof will follow the intuition of Katz and Yung for Theorem 1 in [20]. We will start by bounding the probabilities for the occurrence of bad events ( $\text{G0}^{\text{AttKE}, \mathcal{A}}$  to  $\text{G3}^{\text{AttKE}, \mathcal{A}}$ ), and then argue that the behaviour of  $\text{G3}^{\text{AttKE}, \mathcal{A}}$  towards an adversary is the same as the one of  $\text{G4}^{\Pi, \mathcal{A}}$  using a passive adversary for the original protocol  $\Pi$ . At this final stage we show that, given correctness and security guarantees of  $\Pi$ , we have a correct and secure `AttKE` protocol. The proof consists in a sequence of five games presented in Figures 17 to 21.

The first game is simply the `AttKE` security game in Figure 8 instantiated with our construction in Figure 9.

In the second game  $\text{G1}^{\text{AttKE}, \mathcal{A}}$ , the adversary loses whenever a `repeat` event occurs. Intuitively, this event corresponds to the adversary generating two sessions with the same key pair, and hence constitutes a forgery with respect to  $\Sigma$ . Given that the two games are identical until this event occurs, we have that

$$\Pr[\text{G0}^{\text{AttKE}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] - \Pr[\text{G1}^{\text{AttKE}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] \leq \Pr[\text{repeat}].$$

Let  $q$  be the maximum number of calls to `NewLocal` allowed. We upper bound the distance between these two games, by constructing an adversary  $\mathcal{B}$  against the existential unforgeability of signature scheme  $\Sigma$  such that

**Simulator  $\mathcal{S} = \{\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3\}$**

Simulator  $\mathcal{S}$  will perform according to the  $\mathcal{M}_R$  execution description.

- Upon input  $1^\lambda$ ,  $\mathcal{S}_1$  generates a key pair for process  $S^*$ , a MAC key for the security module and initializes the traces as an empty list. The public key will be the public parameters, while the secret key be stored in its initial state.

$\mathcal{S}_1(1^\lambda)$ :

```

key  $\leftarrow$   $\Pi$ .Gen( $1^\lambda$ )
(pk, sk)  $\leftarrow$   $\Sigma$ .Gen( $1^\lambda$ )
Traces  $\leftarrow$  []
Return (pk, (key, sk, Traces))

```

- $\mathcal{S}_2$  maintains a list of traces Traces with the respective list ios and stage stage. Given this, it masks output  $o^*$  as if produced by an actual machine execution.

$\mathcal{S}_2(\text{hdl}, P, \phi, Q, R, i, o^*, \text{st}_{\mathcal{S}})$ :

```

Parse (key, sk, Traces)  $\leftarrow$   $\text{st}_{\mathcal{S}}$ 
If  $\nexists$  (ios, stage)  $\in$  Traces[hdl]: ios  $\leftarrow$  []; stage  $\leftarrow$  1
ios  $\leftarrow$  ( $i, o^*$ ) : ios
If stage = 1: m  $\leftarrow$  ( $o^*$ ,  $\Pi$ .Mac(key, R, ios))
Else m  $\leftarrow$   $o^*$ 
Parse (o, finished, stage)  $\leftarrow$   $o^*$ 
If finished = T: stage  $\leftarrow$  2
Traces[hdl]  $\leftarrow$  (ios, stage);  $\text{st}_{\mathcal{S}} \leftarrow$  (key, sk, Traces)
Return (m,  $\text{st}_{\mathcal{S}}$ )

```

- $\mathcal{S}_3$  standardly computes the next output given input  $i$ , program  $R$  and state  $\text{st}$ . The result is afterwards treated similar to  $\mathcal{S}_2$ .

$\mathcal{S}_3(\text{hdl}, R, i, \text{st}, \text{st}_{\mathcal{S}})$ :

```

Parse (key, sk, Traces)  $\leftarrow$   $\text{st}_{\mathcal{S}}$ 
If hdl = 0:
  Parse ( $i, t$ )  $\leftarrow$   $i^*$ 
  If  $\Pi$ .Ver(key, t,  $i$ ): Return  $\Sigma$ .Sign(sk,  $i$ )
  Else Return  $\perp$ 
If  $\nexists$  (ios, stage)  $\in$  Traces[hdl]: ios  $\leftarrow$  []; stage  $\leftarrow$  1
 $o^* \leftarrow$   $R[\text{st}](i)$ 
If stage = 1: m  $\leftarrow$  ( $o^*$ ,  $\Pi$ .Mac(key, R, st.ios))
Else m  $\leftarrow$   $o^*$ 
Parse (o, finished, stage)  $\leftarrow$   $o^*$ 
If finished = T: stage  $\leftarrow$  2
Traces[hdl]  $\leftarrow$  (ios, stage);  $\text{st}_{\mathcal{S}} \leftarrow$  (key, sk, Traces)
Return (m,  $\text{st}_{\mathcal{S}}$ )

```

Figure 16: Description of simulator  $\mathcal{S}$

$$\Pr[\text{repeat}] \leq \frac{\text{Adv}_{\Sigma, \mathcal{B}}^{\text{UF}}(\lambda) * q}{2}.$$

Adversary  $\mathcal{B}$  simulates the environment of  $\text{G1}^{\text{AttKE}, \mathcal{A}}$  as follows: at the beginning of the game,  $\mathcal{B}$  has to try and guess which session will have a duplicate key. As such, it samples uniformly from  $[1..q]$  a session  $s$  and replaces the public key generated by  $\Sigma$ .Gen in `NewLocal` for instance  $i = s$  with the public key  $\text{pk}_{\text{UF}}$  provided by  $\text{UF}_{\Sigma, \mathcal{B}}^{\text{UF}}$ . Every time instance  $s$  has to produce a signature in `SendLocal`( $m, s$ ), instead of  $\Sigma$ .Sign(sk,  $t$ ),  $\mathcal{B}$  calls Oracle `Sign`( $t$ ). Additionally,  $\mathcal{B}$  will store all key pairs and session ids generated by `NewLocal` s.t.  $i \neq s$  in a list `keys`. When

<p><b>G0<sub>AttKE, A</sub>(1<sup>λ</sup>):</b>  <math>\text{InsList} \leftarrow []</math>  <math>\text{fake} \leftarrow []</math>  <math>i \leftarrow 0</math>  <math>b \leftarrow \{0, 1\}</math>  <math>b' \leftarrow \mathcal{A}^{\mathcal{O}}(1^\lambda, \text{id})</math>  Return <math>b = b'</math></p>	<p><b>Oracle NewLoc():</b>  <math>i \leftarrow i + 1</math>  <math>(\text{pk}, \text{sk}) \leftarrow \Sigma.\text{Gen}(1^k)</math>  <math>R^i \leftarrow \text{Rem}_{\text{KE}} &lt; \Pi, \text{pk} &gt;</math>  <math>\text{st}_L^i \leftarrow (\text{id}, \epsilon, \text{sk}, [])</math>  <math>\text{InsList}[i] \leftarrow 0</math>  <math>T_L^i \leftarrow []</math>  Return <math>R^i</math></p> <p><b>Oracle SendLoc(m, i):</b>  If <math>\nexists j, (m : T_L^i) \sqsubseteq T_R^{i,j}</math> then return <math>\perp</math>  <math>(\text{id}, \text{st}_{\text{KE}}, \text{sk}, t) \leftarrow \text{st}_L^i</math>  Parse <math>(m^*, r) \leftarrow m</math>  <math>(m', \text{st}_{\text{KE}}) \leftarrow \Pi(1^\lambda, m^*, \text{id}, \text{responder}, \text{st}_{\text{KE}})</math>  <math>t \leftarrow m' : m : t</math>  <math>\sigma = \Sigma.\text{Sign}(\text{sk}, t)</math>  <math>\text{st}_L^i \leftarrow (\text{id}, \text{st}_{\text{KE}}, \text{sk}, t)</math>  <math>T_L^i \leftarrow m' : m : T_L^i</math>  If <math>\text{st}_L^i.\delta \in \{\text{accept}, \text{derived}\}</math>:    If <math>(\text{st}_L^i.\text{key}, \text{key}^*) \notin \text{fake}</math>:      <math>\text{key}^* \leftarrow \{0, 1\}^\lambda</math>      <math>\text{fake} \leftarrow (\text{st}_L^i.\text{key}, \text{key}^*) : \text{fake}</math>  Return <math>(m', \text{st}_L^i.\text{sid}, \text{st}_L^i.\delta, \text{st}_L^i.\text{pid})</math></p> <p><b>Oracle TestLoc(i):</b>  If <math>\text{st}_L^i.\delta \neq \text{accept}</math> return <math>\perp</math>  If <math>b = 0</math> return <math>\text{st}_L^i.\text{key}</math>  Return <math>\text{fake}(\text{st}_L^i.\text{key})</math></p> <p><b>Oracle RevealLoc(i):</b>  Return <math>\text{st}_L^i.\text{key}</math></p>	<p><b>Oracle NewRem(i):</b>  <math>\text{InsList}[i] \leftarrow \text{InsList}[i] + 1</math>  <math>j \leftarrow \text{InsList}[i]; \text{st}_R^{i,j} \leftarrow \epsilon</math>  <math>T_R^{i,j} \leftarrow []</math>  Return <math>\epsilon</math></p> <p><b>Oracle SendRem(m, i, j):</b>  <math>m' \leftarrow R^i[\text{st}_R^{i,j}](m)</math>  <math>T_R^{i,j} \leftarrow m' : m : T_R^{i,j}</math>  If <math>\text{st}_R^{i,j}.\delta \in \{\text{accept}, \text{derived}\}</math>:    If <math>(\text{st}_R^{i,j}.\text{key}, \text{key}^*) \notin \text{fake}</math>:      <math>\text{key}^* \leftarrow \{0, 1\}^\lambda</math>      <math>\text{fake} \leftarrow (\text{st}_R^{i,j}.\text{key}, \text{key}^*) : \text{fake}</math>  Return <math>(m', \text{st}_R^{i,j}.\text{sid}, \text{st}_R^{i,j}.\delta, \text{st}_R^{i,j}.\text{pid})</math></p> <p><b>Oracle TestRem(i, j):</b>  If <math>\text{st}_R^{i,j}.\delta \neq \text{accept}</math> return <math>\perp</math>  If <math>b = 0</math> return <math>\text{st}_R^{i,j}.\text{key}</math>  Return <math>\text{fake}(\text{st}_R^{i,j}.\text{key})</math></p> <p><b>Oracle RevealRem(i, j):</b>  Return <math>\text{st}_R^{i,j}.\text{key}</math></p>
---	--	---

Figure 17: Game defining the execution environment for the security analysis of an AttKE scheme for an arbitrary local identity  $\text{id}$ .  $\mathcal{O}$  denotes all oracles associated with the game.

the game terminates, the adversary looks up on list keys for a key pair with public key  $\text{pk}_{\text{UF}}$ , extracts the associated secret key  $\text{sk}$ , executes  $\sigma \leftarrow \Sigma.\text{Sign}(\text{sk}, m')$  for any  $m'$  not yet queried to Oracle Sign and presents  $(m', \sigma)$  as a challenge for  $\text{UF}^{\Sigma, \mathcal{B}}$ . It remains to show that, when  $\text{repeat}$  is set,  $\mathcal{B}$  wins  $\text{UF}^{\Sigma, \mathcal{B}}$  with probability  $2/q$

When the game ends and  $\text{repeat} = \text{T}$ , we have at least two duplicate  $R$ . This implies that  $\exists i, j$  s.t.  $(\text{pk}_i, \text{sk}_i) = (\text{pk}_j, \text{sk}_j), i \neq j$ . If it is the case that  $i = s \vee j = s$ ,  $\mathcal{B}$  has either  $\text{sk}_i$  or  $\text{sk}_j$  in keys, and can use that to generate the signature that wins the  $\text{UF}^{\Sigma, \mathcal{B}}$  game. Since  $i, j \in [1..q]$  and  $s$  is sampled uniformly from  $[1..q]$ , we have that this happens with probability  $2/q$ , and we therefore conclude that  $\mathcal{B}$  outputs a valid forgery with the same probability whenever  $\text{repeat}$  occurs.

In game  $\text{G2}^{\text{AttKE}, \mathcal{A}}$ , the adversary loses whenever a  $\text{rnonce}$  event occurs. Intuitively, this event corresponds to the adversary generating two duplicate nonces. Given that the two games are identical until this event occurs, we have that

$$\Pr[\text{G1}^{\text{AttKE}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] - \Pr[\text{G2}^{\text{AttKE}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] \leq \Pr[\text{rnonce}].$$

Let  $q_R$  be the number of calls to  $\text{NewRemote}$  allowed to the adversary. We upper bound the distance between these two games such that

<p><b>G1<sub>AttKE, A</sub>(1<sup>λ</sup>):</b>  <math>\text{InsList} \leftarrow []</math>  <math>\text{fake} \leftarrow []</math>  <math>\text{PrgList} \leftarrow []</math>  <b>repeat</b> <math>\leftarrow</math> <b>F</b>  <math>i \leftarrow 0</math>  <math>b \leftarrow_{\\$} \{0, 1\}</math>  <math>b' \leftarrow_{\\$} \mathcal{A}^{\mathcal{O}}(1^\lambda, \text{id})</math>  <b>If repeat = T:</b>  <math>b' \leftarrow_{\\$} \{0, 1\}</math>  <b>Return</b> <math>b = b'</math></p>	<p><b>Oracle NewLoc():</b>  <math>i \leftarrow i + 1</math>  <math>(\text{pk}, \text{sk}) \leftarrow_{\\$} \Sigma.\text{Gen}(1^k)</math>  <math>R^i \leftarrow \text{Rem}_{\text{KE}} &lt; \Pi, \text{pk} &gt;</math>  <b>If</b> <math>R^i \in \text{PrgList}</math>:  <b>repeat</b> <math>\leftarrow</math> <b>T</b>  <math>\text{PrgList} \leftarrow (R^i : \text{PrgList})</math>  <math>\text{st}_L^i \leftarrow (\text{id}, \epsilon, \text{sk}, [])</math>  <math>\text{InsList}[i] \leftarrow 0</math>  <math>T_L^i \leftarrow []</math>  <b>Return</b> <math>R^i</math></p> <p><b>Oracle SendLoc(m, i):</b>  <b>If</b> <math>\nexists j, (m : T_L^i) \sqsubseteq T_R^{i,j}</math> <b>then return</b> <math>\perp</math>  <math>(\text{id}, \text{st}_{\text{KE}}, \text{sk}, t) \leftarrow \text{st}_L^i</math>  <math>\text{Parse}(m^*, r) \leftarrow m</math>  <math>(m', \text{st}_{\text{KE}}) \leftarrow_{\\$} \Pi(1^\lambda, m^*, \text{id}, \text{responder}, \text{st}_{\text{KE}})</math>  <math>t \leftarrow m' : m : t</math>  <math>\sigma = \Sigma.\text{Sign}(\text{sk}, t)</math>  <math>\text{st}_L^i \leftarrow (\text{id}, \text{st}_{\text{KE}}, \text{sk}, t)</math>  <math>T_L^i \leftarrow m' : m : T_L^i</math>  <b>If</b> <math>\text{st}_L^i.\delta \in \{\text{accept}, \text{derived}\}</math>:  <b>If</b> <math>(\text{st}_L^i.\text{key}, \text{key}^*) \notin \text{fake}</math>:  <math>\text{key}^* \leftarrow_{\\$} \{0, 1\}^\lambda</math>  <math>\text{fake} \leftarrow (\text{st}_L^i.\text{key}, \text{key}^*) : \text{fake}</math>  <b>Return</b> <math>(m', \text{st}_L^i.\text{sid}, \text{st}_L^i.\delta, \text{st}_L^i.\text{pid})</math></p> <p><b>Oracle TestLoc(i):</b>  <b>If</b> <math>\text{st}_L^i.\delta \neq \text{accept}</math> <b>return</b> <math>\perp</math>  <b>If</b> <math>b = 0</math> <b>return</b> <math>\text{st}_L^i.\text{key}</math>  <b>Return</b> <math>\text{fake}(\text{st}_L^i.\text{key})</math></p> <p><b>Oracle RevealLoc(i):</b>  <b>Return</b> <math>\text{st}_L^i.\text{key}</math></p>	<p><b>Oracle NewRem(i):</b>  <math>\text{InsList}[i] \leftarrow \text{InsList}[i] + 1</math>  <math>j \leftarrow \text{InsList}[i]; \text{st}_R^{i,j} \leftarrow \epsilon</math>  <math>T_R^{i,j} \leftarrow []</math>  <b>Return</b> <math>\epsilon</math></p> <p><b>Oracle SendRem(m, i, j):</b>  <math>m' \leftarrow_{\\$} R^i[\text{st}_R^{i,j}](m)</math>  <math>T_R^{i,j} \leftarrow m' : m : T_R^{i,j}</math>  <b>If</b> <math>\text{st}_R^{i,j}.\delta \in \{\text{accept}, \text{derived}\}</math>:  <b>If</b> <math>(\text{st}_R^{i,j}.\text{key}, \text{key}^*) \notin \text{fake}</math>:  <math>\text{key}^* \leftarrow_{\\$} \{0, 1\}^\lambda</math>  <math>\text{fake} \leftarrow (\text{st}_R^{i,j}.\text{key}, \text{key}^*) : \text{fake}</math>  <b>Return</b> <math>(m', \text{st}_R^{i,j}.\text{sid}, \text{st}_R^{i,j}.\delta, \text{st}_R^{i,j}.\text{pid})</math></p> <p><b>Oracle TestRem(i, j):</b>  <b>If</b> <math>\text{st}_R^{i,j}.\delta \neq \text{accept}</math> <b>return</b> <math>\perp</math>  <b>If</b> <math>b = 0</math> <b>return</b> <math>\text{st}_R^{i,j}.\text{key}</math>  <b>Return</b> <math>\text{fake}(\text{st}_R^{i,j}.\text{key})</math></p> <p><b>Oracle RevealRem(i, j):</b>  <b>Return</b> <math>\text{st}_R^{i,j}.\text{key}</math></p>
---	---	---

Figure 18: Game defining the execution environment for the security analysis of an AttKE scheme for an arbitrary local identity  $\text{id}$ .  $\mathcal{O}$  denotes all oracles associated with the game.

$$\Pr[\text{rnonce}] \leq \frac{q_R^2}{2^k}.$$

From the rules of  $\text{G2}^{\text{AttKE}, \mathcal{A}}$  and the construction of  $R$ , we know that a new  $r$  is generated at every query of  $\text{SendRemote}$  such that  $T_R^{i,j} = []$ . This only happens at most once for every new  $T_R^{i,j}$ , i.e., at most once for every call of  $\text{NewRemote}$ . Since  $r$  is sampled uniformly from a subset of  $\{0, 1\}^k$ , we conclude that the probability of  $\text{rnonce}$  is  $q_R^2/2^k$ .

In game  $\text{G3}^{\text{AttKE}, \mathcal{A}}$ , the adversary loses whenever a  $\text{forge}$  event occurs. Intuitively, this event corresponds to the adversary producing a signature that was not computed by  $\text{SendLocal}$ , and hence constitutes a forgery with respect to  $\Sigma$ . Given that the two games are identical until this event occurs, we have that

$$\Pr[\text{G2}^{\text{AttKE}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] - \Pr[\text{G3}^{\text{AttKE}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] \leq \Pr[\text{forge}].$$

Let  $q$  be the maximum number of calls to  $\text{NewLocal}$  allowed. We upper bound the distance between these two games, by constructing an adversary  $\mathcal{C}$  against the existential unforgeability of signature scheme  $\Sigma$  such that

<p><b>G2</b><sub>AttKE, A</sub>(1<sup>λ</sup>):</p> <pre> InsList ← [] fake ← [] PrgList ← [] NonList ← [] repeat ← F rnonce ← F i ← 0 b ← \$ {0, 1} b' ← \$ A<sup>O</sup>(1<sup>λ</sup>, id) If repeat = T:   b' ← \$ {0, 1} If rnonce = T:   b' ← \$ {0, 1} Return b = b'</pre>	<p><b>Oracle NewLoc</b>(<i>i</i>):</p> <pre> i ← i + 1 (pk, sk) ← \$ Σ.Gen(1<sup>k</sup>) R<sup>i</sup> ← Rem<sub>KE</sub> &lt; Π, pk &gt; If R<sup>i</sup> ∈ PrgList:   repeat ← T PrgList ← (R<sup>i</sup> : PrgList) st<sub>L</sub><sup>i</sup> ← (id, ε, sk, []) InsList[<i>i</i>] ← 0 T<sub>L</sub><sup>i</sup> ← [] Return R<sup>i</sup></pre> <p><b>Oracle SendLoc</b>(<i>m</i>, <i>i</i>):</p> <pre> If ∄j, (m : T<sub>L</sub><sup>i</sup>) ⊆ T<sub>R</sub><sup>i,j</sup> then return ⊥ (id, st<sub>KE</sub>, sk, t) ← st<sub>L</sub><sup>i</sup> Parse (m*, r) ← m (m', st<sub>KE</sub>) ← \$ Π(1<sup>λ</sup>, m*, id, responder, st<sub>KE</sub>) t ← m' : m : t σ = Σ.Sign(sk, t) st<sub>L</sub><sup>i</sup> ← (id, st<sub>KE</sub>, sk, t) T<sub>L</sub><sup>i</sup> ← m' : m : T<sub>L</sub><sup>i</sup> If st<sub>L</sub><sup>i</sup>.δ ∈ {accept, derived}:   If (st<sub>L</sub><sup>i</sup>.key, key*) ∉ fake:     key* ← \$ {0, 1}<sup>λ</sup>     fake ← (st<sub>L</sub><sup>i</sup>.key, key*) : fake Return (m', st<sub>L</sub><sup>i</sup>.sid, st<sub>L</sub><sup>i</sup>.δ, st<sub>L</sub><sup>i</sup>.pid)</pre> <p><b>Oracle TestLoc</b>(<i>i</i>):</p> <pre> If st<sub>L</sub><sup>i</sup>.δ ≠ accept return ⊥ If b = 0 return st<sub>L</sub><sup>i</sup>.key Return fake(st<sub>L</sub><sup>i</sup>.key)</pre> <p><b>Oracle RevealLoc</b>(<i>i</i>):</p> <pre> Return st<sub>L</sub><sup>i</sup>.key</pre>	<p><b>Oracle NewRem</b>(<i>i</i>):</p> <pre> InsList[<i>i</i>] ← InsList[<i>i</i>] + 1 j ← InsList[<i>i</i>]; st<sub>R</sub><sup>i,j</sup> ← ε T<sub>R</sub><sup>i,j</sup> ← [] Return ε</pre> <p><b>Oracle SendRem</b>(<i>m</i>, <i>i</i>, <i>j</i>):</p> <pre> m' ← \$ R<sup>i</sup>[st<sub>R</sub><sup>i,j</sup>](m) If T<sub>R</sub><sup>i,j</sup> = []:   Parse (o, r) ← m'   If r ∈ nonList:     rnonce ← T     nonList ← (r : nonList)   T<sub>R</sub><sup>i,j</sup> ← m' : m : T<sub>R</sub><sup>i,j</sup> If st<sub>R</sub><sup>i,j</sup>.δ ∈ {accept, derived}:   If (st<sub>R</sub><sup>i,j</sup>.key, key*) ∉ fake:     key* ← \$ {0, 1}<sup>λ</sup>     fake ← (st<sub>R</sub><sup>i,j</sup>.key, key*) : fake Return (m', st<sub>R</sub><sup>i,j</sup>.sid, st<sub>R</sub><sup>i,j</sup>.δ, st<sub>R</sub><sup>i,j</sup>.pid)</pre> <p><b>Oracle TestRem</b>(<i>i</i>, <i>j</i>):</p> <pre> If st<sub>R</sub><sup>i,j</sup>.δ ≠ accept return ⊥ If b = 0 return st<sub>R</sub><sup>i,j</sup>.key Return fake(st<sub>R</sub><sup>i,j</sup>.key)</pre> <p><b>Oracle RevealRem</b>(<i>i</i>, <i>j</i>):</p> <pre> Return st<sub>R</sub><sup>i,j</sup>.key</pre>
---	---	--

Figure 19: Game defining the execution environment for the security analysis of an AttKE scheme for an arbitrary local identity *id*.  $\mathcal{O}$  denotes all oracles associated with the game.

$$\Pr[\text{forge}] \leq \text{Adv}_{\Sigma, \mathcal{C}}^{\text{UF}}(\lambda) * q.$$

Adversary  $\mathcal{C}$  simulates the environment of  $\text{G3}^{\text{AttKE}, \mathcal{A}}$  as follows: at the beginning of the game,  $\mathcal{C}$  has to try and guess which session will have a duplicate key. As such, it samples uniformly from  $[1..q]$  a session  $s$  and replaces the public key generated by  $\Sigma.\text{Gen}$  in **NewLocal** for instance  $i = s$  with the public key  $\text{pk}_{\text{UF}}$  provided by  $\text{UF}^{\Sigma, \mathcal{C}}$ . Every time instance  $s$  has to produce a signature in **SendLocal**( $m, s$ ), instead of  $\Sigma.\text{Sign}(\text{sk}, t)$ ,  $\mathcal{C}$  calls **Oracle Sign**( $t$ ). During the execution of **SendRemote**( $m, i, j$ ) that sets  $\text{forge} = \text{T}$ ,  $\mathcal{C}$  presents  $((o : \text{st}_R^{i,j}.t), \sigma)$  as a challenge for  $\text{UF}^{\Sigma, \mathcal{C}}$ . It remains to show that, when  $\text{forge}$  is set,  $\mathcal{C}$  wins  $\text{UF}^{\Sigma, \mathcal{C}}$  with probability  $1/q$ .

When  $\text{forge} = \text{T}$  we have that, for some execution of **SendRemote**( $m, i, j$ ),  $\text{st}_R^{i,j}.\delta \neq \text{reject} \wedge (o : \text{st}_R^{i,j}) \notin \text{sigList}$ . From the construction of  $R$ , we know that  $\text{st}_R^{i,j}.\delta \neq \text{reject}$  implies that the provided  $((o : \text{st}_R^{i,j}.t), \sigma)$  is a valid message/signature pair for session  $i$ . If  $i = s$ , this is also a valid message/signature pair for  $\text{UF}^{\Sigma, \mathcal{C}}$ . Now observe that  $(o : \text{st}_R^{s,j}.t) \notin \text{sigList}$  assures us that  $(o : \text{st}_R^{s,j}.t)$  was not queried to oracle **Sign**: **rnonce** establishes that every nonce is unique, and every execution of **SendLocal** adds the nonce  $r$  to the signed message, so every call to oracle **Sign** is unique. From

If  $\nexists j, (m : T_L^s) \sqsubseteq T_R^{s,j}$  then return  $\perp$

and repeat, we know that if oracle **Sign** was called for some  $(m' : \text{st}_R^{s,j}.t)$ , then **SendLocal** of session  $s$  would be responding to the sequence of messages exchanged with the unique instance  $j$ , matching  $T_R^{s,j}$  that coincides (modulo signatures) with  $\text{st}_R^{s,j}.t$ . However, by the construction of  $\text{G3}^{\text{AttKE}, \mathcal{A}}$ , that would imply  $\text{sigList} \leftarrow (\text{st}_R^{s,j}.t : \text{sigList})$ , and we know that  $\text{st}_R^{s,j}.t \notin \text{sigList}$ . We therefore conclude that  $(\text{st}_R^{i,j}.t, \sigma)$  is a winning output for game  $\text{UF}^{\Sigma, \mathcal{C}}$  if  $i = s$ . That probability is  $1/q$ .

<p><b>G3<sup>AttKE, A</sup>(1<sup>λ</sup>):</b>  <math>\text{InsList} \leftarrow []</math>  <math>\text{fake} \leftarrow []</math>  <math>\text{PrgList} \leftarrow []</math>  <math>\text{NonList} \leftarrow []</math>  <math>\text{sigList} \leftarrow []</math>  <math>\text{repeat} \leftarrow \text{F}</math>  <math>\text{rnonce} \leftarrow \text{F}</math>  <math>\text{forge} \leftarrow \text{F}</math>  <math>i \leftarrow 0</math>  <math>b \leftarrow \{0, 1\}</math>  <math>b' \leftarrow \mathcal{A}^{\mathcal{O}}(1^\lambda, \text{id})</math>            If <math>\text{repeat} = \text{T}</math>:  <math>b' \leftarrow \{0, 1\}</math>            If <math>\text{rnonce} = \text{T}</math>:  <math>b' \leftarrow \{0, 1\}</math>            If <math>\text{forge} = \text{T}</math>:  <math>b' \leftarrow \{0, 1\}</math>            Return <math>b = b'</math></p>	<p><b>Oracle NewLoc():</b>  <math>i \leftarrow i + 1</math>  <math>(\text{pk}, \text{sk}) \leftarrow \Sigma.\text{Gen}(1^k)</math>  <math>R^i \leftarrow \text{Rem}_{\text{KE}} &lt; \Pi, \text{pk} &gt;</math>            If <math>R^i \in \text{PrgList}</math>:  <math>\text{repeat} \leftarrow \text{T}</math>  <math>\text{PrgList} \leftarrow (R^i : \text{PrgList})</math>  <math>\text{st}_L^i \leftarrow (\text{id}, \epsilon, \text{sk}, [])</math>  <math>\text{InsList}[i] \leftarrow 0</math>  <math>T_L^i \leftarrow []</math>            Return <math>R^i</math></p> <p><b>Oracle SendLoc(m, i):</b>            If <math>\nexists j, (m : T_L^i) \sqsubseteq T_R^{i,j}</math> then return <math>\perp</math>  <math>(\text{id}, \text{st}_{\text{KE}}, \text{sk}, t) \leftarrow \text{st}_L^i</math>  <math>\text{Parse}(m^*, r) \leftarrow m</math>  <math>(m', \text{st}_{\text{KE}}) \leftarrow \Pi(1^\lambda, m^*, \text{id}, \text{responder}, \text{st}_{\text{KE}})</math>  <math>t \leftarrow m' : m : t</math>  <math>\sigma = \Sigma.\text{Sign}(\text{sk}, t)</math>  <math>\text{sigList} \leftarrow t : \text{sigList}</math>  <math>\text{st}_L^i \leftarrow (\text{id}, \text{st}_{\text{KE}}, \text{sk}, t)</math>  <math>T_L^i \leftarrow m' : m : T_L^i</math>            If <math>\text{st}_L^i.\delta \in \{\text{accept}, \text{derived}\}</math>:              If <math>(\text{st}_L^i.\text{key}, \text{key}^*) \notin \text{fake}</math>:                <math>\text{key}^* \leftarrow \{0, 1\}^\lambda</math>                <math>\text{fake} \leftarrow (\text{st}_L^i.\text{key}, \text{key}^*) : \text{fake}</math>            Return <math>(m', \text{st}_L^i.\text{sid}, \text{st}_L^i.\delta, \text{st}_L^i.\text{pid})</math></p> <p><b>Oracle TestLoc(i):</b>            If <math>\text{st}_L^i.\delta \neq \text{accept}</math> return <math>\perp</math>            If <math>b = 0</math> return <math>\text{st}_L^i.\text{key}</math>            Return <math>\text{fake}(\text{st}_L^i.\text{key})</math></p> <p><b>Oracle RevealLoc(i):</b>            Return <math>\text{st}_L^i.\text{key}</math></p>	<p><b>Oracle NewRem(i):</b>  <math>\text{InsList}[i] \leftarrow \text{InsList}[i] + 1</math>  <math>j \leftarrow \text{InsList}[i]; \text{st}_R^{i,j} \leftarrow \epsilon</math>  <math>T_R^{i,j} \leftarrow []</math>            Return <math>\epsilon</math></p> <p><b>Oracle SendRem(m, i, j):</b>  <math>m' \leftarrow R^i[\text{st}_R^{i,j}](m)</math>            If <math>T_R^{i,j} = []</math>:  <math>\text{Parse}(o, r) \leftarrow m'</math>            If <math>r \in \text{nonList}</math>:  <math>\text{rnonce} \leftarrow \text{T}</math>  <math>\text{nonList} \leftarrow (r : \text{nonList})</math>            Else:  <math>\text{Parse}(o, \sigma) \leftarrow m</math>            If <math>\text{st}_R^{i,j}.\delta \neq \text{reject} \wedge (o : \text{st}_R^{i,j}.t) \notin \text{sigList}</math>  <math>\text{forge} \leftarrow \text{T}</math>  <math>T_R^{i,j} \leftarrow m' : m : T_R^{i,j}</math>            If <math>\text{st}_R^{i,j}.\delta \in \{\text{accept}, \text{derived}\}</math>:              If <math>(\text{st}_R^{i,j}.\text{key}, \text{key}^*) \notin \text{fake}</math>:                <math>\text{key}^* \leftarrow \{0, 1\}^\lambda</math>                <math>\text{fake} \leftarrow (\text{st}_R^{i,j}.\text{key}, \text{key}^*) : \text{fake}</math>            Return <math>(m', \text{st}_R^{i,j}.\text{sid}, \text{st}_R^{i,j}.\delta, \text{st}_R^{i,j}.\text{pid})</math></p> <p><b>Oracle TestRem(i, j):</b>            If <math>\text{st}_R^{i,j}.\delta \neq \text{accept}</math> return <math>\perp</math>            If <math>b = 0</math> return <math>\text{st}_R^{i,j}.\text{key}</math>            Return <math>\text{fake}(\text{st}_R^{i,j}.\text{key})</math></p> <p><b>Oracle RevealRem(i, j):</b>            Return <math>\text{st}_R^{i,j}.\text{key}</math></p>
--	---	--

Figure 20: Game defining the execution environment for the security analysis of an AttKE scheme for an arbitrary local identity  $\text{id}$ .  $\mathcal{O}$  denotes all oracles associated with the game.

Finally, in game  $\text{G4}^{\Pi, \mathcal{A}}$ , we no longer require **AttKE**, but instead make use of a passive adversary and its corresponding oracles  $\{\text{Execute}(i, j), \text{Reveal}(i, s), \text{Test}(i, s)\}$ . The intuition is that, at this fifth game, all calls to the compiled **AttKE** protocol can be perfectly simulated using the original protocol  $\Pi$ . More formally, we want to show that

$$\Pr[\text{G3}^{\text{AttKE}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] = \Pr[\text{G4}^{\Pi, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}].$$

<p><b>G4<sub>II,A</sub>(1<sup>λ</sup>):</b>  <math>\text{InsList} \leftarrow []</math>  <math>\text{fake} \leftarrow []</math>  <math>\text{PrgList} \leftarrow []</math>  <math>\text{NonList} \leftarrow []</math>  <math>\text{sigList} \leftarrow []</math>  <math>\text{repeat} \leftarrow \text{F}</math>  <math>\text{rnonce} \leftarrow \text{F}</math>  <math>\text{forge} \leftarrow \text{F}</math>  <math>i \leftarrow 0</math>  <math>b \leftarrow_{\\$} \{0, 1\}</math>  <math>b' \leftarrow_{\\$} \mathcal{A}^{\mathcal{O}}(1^\lambda, \text{id})</math>            If <math>\text{repeat} = \text{T}</math>:  <math>b' \leftarrow_{\\$} \{0, 1\}</math>            If <math>\text{rnonce} = \text{T}</math>:  <math>b' \leftarrow_{\\$} \{0, 1\}</math>            If <math>\text{forge} = \text{T}</math>:  <math>b' \leftarrow_{\\$} \{0, 1\}</math>            Return <math>b = b'</math></p>	<p><b>Oracle NewLoc():</b>  <math>i \leftarrow i + 1</math>  <math>(\text{pk}, \text{sk}) \leftarrow_{\\$} \Sigma.\text{Gen}(1^k)</math>  <math>R^i \leftarrow \text{Rem}_{\text{KE}} &lt; \Pi, \text{pk} &gt;</math>            If <math>R^i \in \text{PrgList}</math>:  <math>\text{repeat} \leftarrow \text{T}</math>  <math>\text{PrgList} \leftarrow (R^i : \text{PrgList})</math>  <math>\text{st}_L^i \leftarrow (\text{id}, (\text{pk}, \text{sk}), \perp, [], 1)</math>  <math>\text{InsList}[i] \leftarrow 0</math>  <math>T_L^i \leftarrow []</math>            Return <math>R^i</math></p> <p><b>Oracle SendLoc(m, i):</b>            If <math>\nexists j, (m : T_L^i) \sqsubseteq T_R^{i,j}</math> then return <math>\perp</math>  <math>(\text{id}, (\text{pk}, \text{sk}), r, t, n) \leftarrow \text{st}_L^i</math>  <math>\text{Parse}(m^*, r) \leftarrow m</math>            If <math>t = []</math> then <math>r \leftarrow r'</math>  <math>m^* \leftarrow E^{i,r}[n]</math>  <math>t \leftarrow m' : m : t</math>  <math>\sigma = \Sigma.\text{Sign}(\text{sk}, t)</math>  <math>\text{sigList} \leftarrow t : \text{sigList}</math>  <math>\text{st}_L^i \leftarrow (\text{id}, (\text{pk}, \text{sk}), r, t, n)</math>  <math>T_L^i \leftarrow m' : m : T_L^i</math>  <math>(\text{sid}, \delta, \text{pid}) \leftarrow \text{Local}_i</math>            Return <math>((m^*, \sigma), \text{sid}, \delta, \text{pid})</math></p> <p><b>Oracle TestLoc(i):</b>  <math>(\text{sid}, \delta, \text{pid}) \leftarrow \text{Local}_i</math>            If <math>\delta \neq \text{accept}</math> return <math>\perp</math>            Return <math>\text{Test}(\text{id}, \text{sid})</math></p> <p><b>Oracle RevealLoc(i):</b>  <math>(\text{sid}, \delta, \text{pid}) \leftarrow \text{Local}_i</math>            If <math>\delta \neq \{\text{derived}, \text{accept}\}</math> then Return <math>\perp</math>            Return <math>\text{Reveal}(\text{id}, \text{sid})</math></p>	<p><b>Oracle NewRem(i):</b>  <math>\text{InsList}[i] \leftarrow \text{InsList}[i] + 1</math>  <math>j \leftarrow \text{InsList}[i]; \text{st}_R^{i,j} \leftarrow \epsilon</math>  <math>T_R^{i,j} \leftarrow []</math>            Return <math>\epsilon</math></p> <p><b>Oracle SendRem(m, i, j):</b>            If <math>T_R^{i,j} = []</math>:            If <math>m \neq \epsilon</math> then <math>\delta \leftarrow \text{reject}</math>  <math>t \leftarrow []; n \leftarrow 0; r \leftarrow_{\\$} \{0, 1\}^k</math>            If <math>r \in \text{nonList}</math>:  <math>\text{rnonce} \leftarrow \text{T}</math>  <math>\text{nonList} \leftarrow (r : \text{nonList})</math>  <math>\text{oid} \leftarrow \text{st}_L^i.\text{pk}    r</math>  <math>E^{i,r} \leftarrow_{\\$} \text{Execute}(\text{id}, \text{oid}); o \leftarrow \epsilon</math>            Else:  <math>(\text{pk}, r, t, n) \leftarrow \text{st}_R^{i,j}</math>  <math>\text{Parse}(o, \sigma) \leftarrow m</math>            If <math>\Sigma.\text{Vrfy}(\text{pk}, \sigma, (o : t)) = \perp</math>: <math>\delta \leftarrow \text{reject}</math>            If <math>\delta \neq \text{reject} \wedge (o : t) \notin \text{sigList}</math>  <math>\text{forge} \leftarrow \text{T}</math>  <math>m' \leftarrow E^{i,r}[n]</math>  <math>t \leftarrow (m', r) : o : t</math>  <math>\text{st}_R^{i,j} \leftarrow (\text{pk}, r, t, n + 2)</math>  <math>T_R^{i,j} \leftarrow (m', r) : m : T_R^{i,j}</math>  <math>(\text{sid}, \delta, \text{pid}, \text{oid}) \leftarrow \text{Remote}_{i,j}</math>            Return <math>((m', r), \text{sid}, \delta, \text{pid})</math></p> <p><b>Oracle TestRem(i, j):</b>  <math>(\text{sid}, \delta, \text{pid}, \text{oid}) \leftarrow \text{Remote}_{i,j}</math>            If <math>\delta \neq \text{accept}</math> return <math>\perp</math>            Return <math>\text{Test}(\text{oid}, \text{sid})</math></p> <p><b>Oracle RevealRem(i, j):</b>  <math>(\text{sid}, \delta, \text{pid}, \text{oid}) \leftarrow \text{Remote}_{i,j}</math>            If <math>\delta \neq \{\text{derived}, \text{accept}\}</math> then Return <math>\perp</math>            Return <math>\text{Reveal}(\text{oid}, \text{sid})</math></p>
--	--	---

Figure 21: Game defining the execution environment for the security analysis of an AttKE scheme for an arbitrary local identity  $\text{id}$ .  $\mathcal{O}$  denotes all oracles associated with the game.

Whenever  $\text{G4}^{\Pi, \mathcal{A}}$  is required to respond to  $\text{SendLocal}$  or  $\text{SendRemote}$ , instead of executing  $\Pi$  and  $R$  (respectively), it will follow a unique  $\text{Execute}(\text{id}, \text{oid})$  that is associated with sessions  $i$  and nonce  $r$ . Additionally,  $\text{G4}^{\Pi, \mathcal{A}}$  tracks when any execution should be rejected, and responds to  $\text{Reveal}/\text{Test}$  with either  $\perp$  or with the output of the same oracles from the passive adversary. To help with this,  $\text{G4}^{\Pi, \mathcal{A}}$  stores a list of  $\text{Execute}$  transcripts  $E^{i,r}$ . Messages exchanged from such transcripts will be tracked locally and remotely with a counter  $n$ . From the information in  $(\text{st}_L^i, \text{st}_R^{i,j})$  and transcripts  $E^{i,r}$  we assume that it is possible to infer  $(\text{sid}, \delta, \text{pid})$  and  $(\text{sid}, \delta, \text{pid}, \text{oid})$  via  $\text{Local}_i$  and  $\text{Remote}_{i,j}$ , respectively.

The changes in  $\text{NewLocal}$  merely reflect additional information stored in existing structures. In  $\text{SendLocal}$ ,  $E^{i,r}[n]$  is used instead of  $\Pi$  to provide the next message. In  $\text{SendRemote}$ , instead of running  $R$ , the specification in Figure 9 is followed on the oracle itself, with the

exception of calling  $E^{i,j}[n]$  instead of using  $\Pi$  to provide the next message. In **Reveal/Test**, either  $\perp$  is returned or the result of another oracle **Reveal/Test** is given, instead of directly providing a key. In these scenarios, handling of fake keys is delegated to **Test** queries of the passive adversary. As such, to validate these games as equivalent, we must show that all calls to  $E^{i,r}$  correspond to the  $\Pi$  replaced, and that **Reveal/Test** are responding similarly to  $\text{G3}^{\text{AttKE},\mathcal{A}}$ .

Observe that, by the construction of  $E^{i,r} \leftarrow_{\$} \text{Execute}(\text{id}, \text{oid})$ ,

$$E^{i,r} = [o_1 \leftarrow_{\$} \Pi(1^\lambda, \epsilon, \text{oid}, \text{initiator}, st_r), o_2 \leftarrow_{\$} \Pi(1^\lambda, o_1, \text{id}, \text{responder}, st_i), \dots]$$

the transcript  $E^{i,r}$  contains a list of specific executions of  $\Pi(1^\lambda, m, i, \rho, st)$ . We must show that every  $E^{i,j}[n]$  matches to the output of  $\Pi(1^\lambda, m, i, \rho, st)$  executed in  $\text{G3}^{\text{AttKE},\mathcal{A}}$ . First note that, from **repeat**, we know that every  $i$  is associated with a different  $R$ . This means that every  $\text{SendRemote}(m, i, j)$  will be associated with a unique  $st_L^i$ . Many executions of the same session  $i$  may occur, so the nonce associated cannot be determined at **NewLocal** (hence  $st_L^i \leftarrow (\text{id}, (\text{pk}, \text{sk}), \perp, [], 1)$ ), but we know that every first message of **AttKE** will fix a unique  $r$ , so  $E^{i,r}$  can be established in the first non-reject call of **SendLocal**( $m, i$ ), and a 1-to-1 relation with instance  $j$  is given by

$$\text{If } \nexists j, (m : T_L^i) \sqsubseteq T_R^{i,j} \text{ then return } \perp \quad (1)$$

and forge. Let  $\text{SendLocal}_k$  and  $\text{SendRemote}_k$  be the  $k$ -th execution of these Oracles.

- $\text{SendRemote}_1(m, i, j)$ : If  $\delta \neq \text{reject}$ ,  $r \leftarrow_{\$} \{0, 1\}^k$ ,  $[\Pi(1^\lambda, \epsilon, \text{pk}||r, \text{initiator}, st), \dots] \leftarrow_{\$} \text{Execute}(\text{id}, \text{pk}||r)$  and we have that  $\Pi(1^\lambda, \epsilon, \text{pk}||r, \text{initiator}, st) = E^{i,r}[0] = m_1$ ;  $st_R^{i,j}.t = [(m_1, r), \epsilon]$ ;  $st_R^{i,j}.n = 2$ .
- $\text{SendLocal}_2(m, i)$ : If  $st_L^i.\delta \neq \text{reject}$ , we know that  $((o, r) : st_L^i.t)$  is in the prefix trace of a unique instance  $\text{SendRemote}_1$  (from (1)). Since  $st_L^i.t = []$ ,  $m$  was the first message produced by  $\text{SendRemote}_1$ , so we can set  $r$  as the nonce of the execution, retrieving the second message from  $\text{Execute}(\text{id}, \text{pk}||r)$  in  $E^{i,r}$ , the unique transcript between session  $i$  and the remote  $j$  that produced  $r$ . As such, we have that  $\Pi(1^\lambda, m', \text{id}, \text{responder}, st) = E^{i,r}[1] = m_2$ .  $st_L^i.t = [m_2, (m_1, r'), \epsilon]$ ;  $st_L^i.n \leftarrow 3$ .
- $\text{SendRemote}_n(m, i, j)$ : If  $\delta \neq \text{reject}$ ,  $\text{Parse}(o, \sigma) \leftarrow m$ ;  $\Sigma.\text{Vrfy}(\text{pk}, \sigma, (o : st_R^{i,j}.t)) \neq \perp$  and  $(o : st_R^{i,j}.t) \in \text{sigList}$  means that  $m$  was the  $(n-1)$ -th message produced by instance  $\text{SendLocal}_{n-1}(m_{n-2}, i)$  with  $st_L^i.t = [o, \dots, (m_1, r), \epsilon]$ , following  $(o : st_R^{i,j}.t)$  (from forge).  $(o : st_R^{i,j}.t)$  was constructed with the first  $n-1$  messages in the transcript of  $\text{Execute}(\text{id}, \text{pk}||r) : E^{i,r}$ , so given that  $E^{i,r}[n]$  provides the  $n$ -th message of  $E^{i,r}$ , we have that  $\Pi(1^\lambda, o, \text{pk}||r, \text{initiator}, st) = E^{i,r}[n]$ .  $st_R^{i,j}.n = n + 2$ .
- $\text{SendLocal}_n(m, i)$ : If  $st_L^i.\delta \neq \text{reject}$ , we know that  $m$  was the  $(n-1)$ -th message produced by some instance  $\text{SendRemote}_{n-1}$  with  $([m, \dots, m_2, (m_1, r), \epsilon])$  in its trace  $st_R^{i,j}.t$  (from (1)).  $st_R^{i,j}.t$  was also constructed with the first  $n-1$  messages in the transcript of  $\text{Execute}(\text{id}, \text{pk}||r) : E^{i,r}$ , so given that  $E^{i,r}[n]$  provides the  $n$ -th message of  $E^{i,r}$ , we have that  $\Pi(1^\lambda, m', \text{id}, \text{responder}, st) = E^{i,r}[n]$ .  $st_L^i.n = n + 2$ .

Regarding **Reveal** and **Test** queries, observe that either the instances of  $i$  or  $j$  have  $\delta \neq \{\text{derived}, \text{accept}\}$  or  $\delta \neq \text{accept}$ , respectively, which given the previous is the same in both

scenarios, or  $\mathcal{A}$  gets the response from the Oracle of the passive adversary regarding the unique  $\text{Execute}(\text{id}, \text{oid})$  associated with sessions  $i$  and  $j$ , and the output is also the same in both scenarios.

As such, we have that the behaviour of  $\text{G3}^{\text{AttKE}, \mathcal{A}}$  towards an adversary  $\mathcal{A}$  is indistinguishable to the one provided in  $\text{G4}^{\Pi, \mathcal{A}}$  interacting with a passive adversary for the original protocol. This provides us with the correctness and security guarantees of the passive protocol  $\Pi$  modulo any attacks that would also be possible in the passive scenario. Now observe that, given that  $\text{rnonce}$  assures unique values for  $\text{oid}$ , and given that  $\text{Local}_{\text{KE}}$  and  $\text{Remote}_{\text{KE}}$  have  $\rho = \text{responder}$  and  $\rho = \text{initiator}$ , respectively, it is trivial to infer that this implies the correctness and security of  $\text{AttKE}$ .

To conclude, we have that

$$\begin{aligned}
\text{Adv}_{\text{AttKE}, \mathcal{A}}^{\text{Att}} &= \Pr[\text{G}_0^{\text{AttKE}, \mathcal{A}}(1^\lambda)] - \Pr[\text{G}_4^{\Pi, \mathcal{A}}(1^\lambda)] \\
&= \left( \sum_{i=0}^3 \Pr[\text{G}_i^{\text{AttKE}, \mathcal{A}}(1^\lambda)] - \Pr[\text{G}_{i+1}^{\text{AttKE}, \mathcal{A}}(1^\lambda)] \right) + (\Pr[\text{G}_3^{\text{AttKE}, \mathcal{A}}(1^\lambda)] - \Pr[\text{G}_4^{\Pi, \mathcal{A}}(1^\lambda)]) \\
&\leq \Pr[\text{repeat}] + \Pr[\text{rnonce}] + \Pr[\text{forge}] + \text{Adv}_{\Pi, \mathcal{A}}^{\text{Att}}(\lambda) \\
&\leq \frac{3 * \text{Adv}_{\Sigma, \mathcal{D}}^{\text{UF}}(\lambda) * q}{2} + \frac{q_R^2}{2^k} + \text{Adv}_{\Pi, \mathcal{A}}^{\text{Att}}(\lambda)
\end{aligned}$$

and Theorem 2 follows. □

## C Proof for Theorem 3

<p><b><math>\text{G0}_{\text{AttKE}, \mathcal{A}}(1^\lambda)</math>:</b>  <math>\text{prms}_0 \leftarrow \mathcal{M}_R.\text{Init}(1^\lambda)</math>  <math>\text{prms}_1 \leftarrow \mathcal{M}'_R.\text{Init}(1^\lambda)</math>  <math>\text{PrgList} \leftarrow []</math>  <math>\text{fake} \leftarrow []</math>  <math>i \leftarrow 0</math>  <math>b \leftarrow \{0, 1\}</math>  <math>b' \leftarrow \mathcal{A}^{\mathcal{O}}(\text{prms}_b, \text{id})</math>  Return <math>b = b'</math></p>	<p><b>Oracle <math>\text{NewSession}(Q)</math>:</b>  <math>i \leftarrow i + 1</math>  <math>(\text{Rem}_{\text{KE}}^i, \text{st}_{\text{KE}}^i) \leftarrow \text{Setup}(1^\lambda, \text{id})</math>  <math>(R_i^*, \text{st}_L^i) \leftarrow \text{AC.Compile}(\text{prms}_b, \text{Rem}_{\text{KE}}^i, \phi_{\text{key}}, Q)</math>  <math>\text{in}_{\text{last}}^i \leftarrow \epsilon</math>  <math>\text{PrgList} \leftarrow ((R_i^*, \text{Rem}_{\text{KE}}^i, Q) : \text{PrgList})</math>  Return <math>R_i^*</math></p> <p><b>Oracle <math>\text{Run}(\text{hdl}, \text{in})</math>:</b>  <math>\text{flag} \leftarrow \text{F}</math>  If <math>\text{Program}_{\mathcal{M}'_R}(\text{hdl}) \in \text{PrgList}</math> then <math>\text{flag} \leftarrow \text{T}</math>  If <math>b = 0</math> return <math>\mathcal{M}_R.\text{Run}(\text{hdl}, \text{in})</math>  <math>(o, \text{fake}) \leftarrow \mathcal{M}'_R.\text{Run}(\text{hdl}, \text{in}, \text{flag}, \text{fake})</math>  Return <math>o</math></p>	<p><b>Oracle <math>\text{Send}(m', i)</math>:</b>  <math>(m, \text{st}_L^i) \leftarrow \text{AC.Verify}(\text{prms}_b, \text{in}_{\text{last}}^i, m', \text{st}_L^i)</math>  <math>(m^*, \text{st}_{\text{KE}}^i) \leftarrow \text{Loc}_{\text{KE}}^i(\text{st}_{\text{KE}}^i, m)</math>  <math>\text{in}_{\text{last}}^i \leftarrow m^*</math>  If <math>\text{st}_{\text{KE}}^i.\text{key} \notin \text{fake} \wedge \text{st}_{\text{KE}}^i.\delta \in \{\text{derived}, \text{accept}\}</math>:  <math>\text{key}^* \leftarrow \{0, 1\}^\lambda</math>  <math>\text{fake} \leftarrow (\text{key}, \text{key}^*) : \text{fake}</math>  Return <math>m^*</math></p> <p><b>Oracle <math>\text{Test}(i)</math>:</b>  If <math>\text{st}_{\text{KE}}^i.\delta \neq \text{accept}</math> return <math>\perp</math>  If <math>b = 0</math> then return <math>\text{st}_{\text{KE}}^i.\text{key}</math>  Return <math>\text{fake}(\text{st}_{\text{KE}}^i.\text{key})</math></p>
---	--	--

Figure 22: Game defining the utility of an  $\text{AttKE}$  scheme when used in the context of attested computation.

In this proof, we will start by bounding the possibility for the occurrence of a bad event ( $\text{G0}^{\text{AC}, \mathcal{A}}$  to  $\text{G1}^{\text{AC}, \mathcal{A}}$ ). We will then replace the machine execution with the indistinguishable behavior of the simulator of minimum leakage game in Figure 6 for  $\text{G2}^{\text{AC}, \mathcal{A}}$ . Finally, we argue that, given these circumstances, this scenario is the same as one of  $\text{G3}^{\text{AttKE}, \mathcal{A}}$  using the oracles of key exchange for attested computation in Figure 8 modulo any advantage the adversary may gain from the  $\text{AttKE}$  scheme. The proof consists in a sequence of four games presented in figures 22 to 25. The first game is simply the utility game in Figure 10.

<p><b>G1<sub>AttKE, A</sub>(1<sup>λ</sup>):</b>  prms<sub>0</sub> ← \$ M<sub>R</sub>.Init(1<sup>λ</sup>)  prms<sub>1</sub> ← \$ M'<sub>R</sub>.Init(1<sup>λ</sup>)  PrgList ← []  fake ← []  forgeAC ← F  i ← 0  b ← \$ {0, 1}  b' ← \$ A<sup>O</sup>(prms<sub>b</sub>, id)  <b>If forgeAC = T:</b>  b' ← {0, 1}  Return b = b'</p> <p><b>Oracle Load(R*):</b>  If b = 0  Return M<sub>R</sub>.Load(R*)  Return M'<sub>R</sub>.Load(R*)</p>	<p><b>Oracle NewSession(Q):</b>  i ← i + 1  (Rem<sub>KE</sub><sup>i</sup>, st<sub>KE</sub><sup>i</sup>) ← \$ Setup(1<sup>λ</sup>, id)  (R<sub>i</sub><sup>*</sup>, st<sub>L</sub><sup>i</sup>) ← \$ AC.Compile(prms<sub>b</sub>, Rem<sub>KE</sub><sup>i</sup>, φ<sub>key</sub>, Q)  in<sub>last</sub><sup>i</sup> ← ε  PrgList ← ((R<sub>i</sub><sup>*</sup>, Rem<sub>KE</sub><sup>i</sup>, Q) : PrgList)  T<sub>L</sub><sup>i</sup> = []  Return R<sub>i</sub><sup>*</sup></p> <p><b>Oracle Run(hdl, in):</b>  flag ← F  <b>If</b> Program<sub>M'</sub>(hdl) ∈ PrgList <b>then</b> flag ← T  <b>If</b> b = 0: o ← M<sub>R</sub>.Run(hdl, in)  <b>Else:</b> (o, fake) ← \$ M'<sub>R</sub>.Run(hdl, in, flag, fake)  Return o</p>	<p><b>Oracle Send(m', i):</b>  (m, st<sub>L</sub><sup>i</sup>) ← \$ AC.Verify(prms<sub>b</sub>, in<sub>last</sub><sup>i</sup>, m', st<sub>L</sub><sup>i</sup>)  <b>If</b> b = 0 <b>then</b> M ← M<sub>R</sub> <b>else</b> M ← M'<sub>R</sub>  <b>If</b> m ≠ ⊥ ∧ ∃hdl s.t. Program<sub>M</sub>(hdl) = R<sub>i</sub><sup>*</sup> :  Rev(m' : T<sub>L</sub><sup>i</sup>) ⊆ ATrace<sub>M</sub>(hdl): forgeAC ← T  (m*, st<sub>KE</sub><sup>i</sup>) ← \$ Loc<sub>KE</sub><sup>i</sup>(st<sub>KE</sub><sup>i</sup>, m)  in<sub>last</sub><sup>i</sup> ← m*; T<sub>L</sub><sup>i</sup> ← m* : m' : T<sub>L</sub><sup>i</sup>  <b>If</b> st<sub>KE</sub><sup>i</sup>.key ∉ fake ∧ st<sub>KE</sub><sup>i</sup>.δ ∈ {derived, accept}:  key* ← \$ {0, 1}<sup>λ</sup>  fake ← (key, key*) : fake  Return m*</p> <p><b>Oracle Test(i):</b>  <b>If</b> st<sub>KE</sub><sup>i</sup>.δ ≠ accept <b>return</b> ⊥  <b>If</b> b = 0 <b>then</b> <b>return</b> st<sub>KE</sub><sup>i</sup>.key  Return fake(st<sub>KE</sub><sup>i</sup>.key)</p>
---	--	--

Figure 23: First hop of the utility proof.

<p><b>G2<sub>AttKE, A</sub>(1<sup>λ</sup>):</b>  (prms, st<sub>S</sub>) ← \$ S<sub>1</sub>(1<sup>λ</sup>)  PrgList ← []  fake ← []  List ← []  forgeAC ← F  i ← 0  hdl ← 0  b ← \$ {0, 1}  b' ← \$ A<sup>O</sup>(prms, id)  <b>If</b> forgeAC = T:  b' ← {0, 1}  Return b = b'</p> <p><b>Oracle Load(R*):</b>  hdl ← hdl + 1  List[hdl] ← (R*, ε)  T<sub>R</sub><sup>hdl</sup> ← []  Return hdl</p>	<p><b>Oracle NewSession(Q):</b>  i ← i + 1  (Rem<sub>KE</sub><sup>i</sup>, st<sub>KE</sub><sup>i</sup>) ← \$ Setup(1<sup>λ</sup>, id)  (R<sub>i</sub><sup>*</sup>, st<sub>L</sub><sup>i</sup>) ← \$ Compile(Rem<sub>KE</sub><sup>i</sup>, φ<sub>key</sub>, Q)  in<sub>last</sub><sup>i</sup> ← ε  PrgList ← ((Rem<sub>KE</sub><sup>i</sup>, φ<sub>key</sub>, Q, R<sub>i</sub><sup>*</sup>) : PrgList)  T<sub>L</sub><sup>i</sup> = []  Return R<sub>i</sub><sup>*</sup></p> <p><b>Oracle Run(hdl, in):</b>  (R*, st) ← List[hdl]  <b>If</b> (P, φ, Q, R) ∈ PrgList:  R ← Compose<sub>φ</sub>[P, Q]  o* ← \$ R[st](in)  (o, st<sub>S</sub>) ← \$ S<sub>2</sub>(hdl, P, φ, Q, R*, in, o*, st<sub>S</sub>)  <b>If</b> st<sub>P</sub>.stage = 1:  <b>If</b> st<sub>P</sub>.key ∉ fake ∧ st<sub>P</sub>.δ ∈ {derived, accept}:  key* ← \$ {0, 1}<sup>λ</sup>  fake ← (key, key*) : fake  <b>If</b> δ = accept ∧ b = 1:  st<sub>P</sub>.key ← fake(key)  T<sub>R</sub><sup>hdl</sup> ← o : in : T<sub>R</sub><sup>hdl</sup>  <b>Else:</b>  (o, st, st<sub>S</sub>) ← \$ S<sub>3</sub>(hdl, R*, in, st, st<sub>S</sub>)  List[hdl] ← (R*, st)  Return o</p>	<p><b>Oracle Send(m', i):</b>  (m, st<sub>L</sub><sup>i</sup>) ← \$ AC.Verify(prms, in<sub>last</sub><sup>i</sup>, m', st<sub>L</sub><sup>i</sup>)  <b>If</b> m ≠ ⊥ ∧ ∃hdl s.t. List[hdl] = R<sub>i</sub><sup>*</sup> :  Rev(m' : T<sub>L</sub><sup>i</sup>) ⊆ T<sub>R</sub><sup>hdl</sup>: forgeAC ← T  (m*, st<sub>KE</sub><sup>i</sup>) ← \$ Loc<sub>KE</sub><sup>i</sup>(st<sub>KE</sub><sup>i</sup>, m)  in<sub>last</sub><sup>i</sup> ← m*; T<sub>L</sub><sup>i</sup> ← m* : m' : T<sub>L</sub><sup>i</sup>  <b>If</b> st<sub>KE</sub><sup>i</sup>.key ∉ fake ∧ st<sub>KE</sub><sup>i</sup>.δ ∈ {derived, accept}:  key* ← \$ {0, 1}<sup>λ</sup>  fake ← (key, key*) : fake  Return m*</p> <p><b>Oracle Test(i):</b>  <b>If</b> st<sub>KE</sub><sup>i</sup>.δ ≠ accept <b>return</b> ⊥  <b>If</b> b = 0 <b>then</b> <b>return</b> st<sub>KE</sub><sup>i</sup>.key  Return fake(st<sub>KE</sub><sup>i</sup>.key)</p>
---	--	---

Figure 24: Second hop of the utility proof.

In the second game G1<sup>AttKE, A</sup>, the adversary loses whenever a forgeAC event occurs. Intuitively, this event corresponds to the adversary producing an output that is successfully validated, but was not computed using Run, and hence constitutes a forgery with respect to AC. We establish that Rev reverses a list given as input<sup>5</sup>. Given that the two games are identical until this event occurs, we have that

$$\Pr[\text{G0}^{\text{AttKE, A}}(1^\lambda) \Rightarrow \text{T}] - \Pr[\text{G1}^{\text{AttKE, A}}(1^\lambda) \Rightarrow \text{T}] \leq \Pr[\text{forgeAC}].$$

<sup>5</sup>This is for handling a technicality, in which the trace of M and the constructed T<sub>L</sub><sup>i</sup> are in reverse order.

<p><b>G3</b><sub>AttKE, A</sub>(1<sup>λ</sup>):</p> <pre> (prms, st<sub>S</sub>) ←<sub>S</sub> S<sub>1</sub>(1<sup>λ</sup>) PrgList ← [] fake ← [] List ← [] InsList ← [] forgeAC ← F i ← 0 hdl ← 0 b ←<sub>S</sub> {0, 1} b' ←<sub>S</sub> A<sup>O</sup>(prms, id) If forgeAC = T:   b' ← {0, 1} Return b = b'  Oracle Load(R*): hdl ← hdl + 1 T<sub>R</sub><sup>hdl</sup> = [] If ∄(R*, j) ∈ InsList:   InsList[R*] ← 1 Else: j ← j + 1   InsList[R*] ← j List[hdl] ← (R*, ε, j, 1) NewRem() T<sub>R</sub><sup>hdl</sup> ← [] Return hdl </pre>	<p><b>Oracle NewSession</b>(Q):</p> <pre> i ← i + 1 Rem<sub>KE</sub><sup>i</sup> ←<sub>S</sub> NewLoc() (R<sub>i</sub><sup>*</sup>, st<sub>L</sub><sup>i</sup>) ←<sub>S</sub> Compile(Rem<sub>KE</sub><sup>i</sup>, φ<sub>key</sub>, Q) in<sub>last</sub><sup>i</sup> ← ε PrgList ← ((Rem<sub>KE</sub><sup>i</sup>, φ<sub>key</sub>, Q, R<sub>i</sub><sup>*</sup>) : PrgList) T<sub>L</sub><sup>i</sup> = [] Return R<sub>i</sub><sup>*</sup>  Oracle Run(hdl, in): (R<sub>i</sub><sup>*</sup>, st, j, stage) ← List[hdl] If (P, φ, Q, R) ∈ PrgList:   If stage = 1:     o ←<sub>S</sub> SendRem(in, i, j)     Parse(o, sid, δ, pid) ← o:       (o*, st<sub>S</sub>) ←<sub>S</sub> S<sub>2</sub>(hdl, P, φ, Q, R*, in, o, st<sub>S</sub>)       If δ = accept:         stage ← 2         st ← TestRem(i, j)         T<sub>R</sub><sup>hdl</sup> ← o* : in : T<sub>R</sub><sup>hdl</sup>       Else:         o ←<sub>S</sub> Q[st](in)         (o*, st<sub>S</sub>) ←<sub>S</sub> S<sub>2</sub>(hdl, P, φ, Q, R*, in, o, st<sub>S</sub>)     Else:       (o, st, st<sub>S</sub>) ←<sub>S</sub> S<sub>3</sub>(hdl, R*, in, st, st<sub>S</sub>)   List[hdl] ← (R*, st, j, stage) Return o* </pre>	<p><b>Oracle Send</b>(m', i):</p> <pre> (m, st<sub>L</sub><sup>i</sup>) ←<sub>S</sub> AC.Verify(prms, in<sub>last</sub><sup>i</sup>, m', st<sub>L</sub><sup>i</sup>) If m ≠ ⊥ ∧ ∄hdl s.t. List[hdl]<sub>last</sub> = R<sub>i</sub><sup>*</sup>:   Rev(m' : T<sub>L</sub><sup>i</sup>) ⊆ T<sub>R</sub><sup>hdl</sup>: forgeAC ← T   m* ←<sub>S</sub> SendLoc(m, i)   Parse(o, sid, δ, pid) ← m*:     in<sub>last</sub><sup>i</sup> ← m*; T<sub>L</sub><sup>i</sup> ← o : m' : T<sub>L</sub><sup>i</sup>   Return o  Oracle Test(i): Return TestLoc(i) </pre>
---	--	---

Figure 25: Third hop of the utility proof.

Let  $q$  be the maximum number of calls to `NewSession` allowed, and  $N$  the number of messages exchanged in  $\text{Rem}_{\text{KE}}^i$ . We upper bound the distance between these two games, by constructing an adversary  $\mathcal{B}$  against the security of `AC` such that

$$\Pr[\text{forgeAC}] \leq \text{Adv}_{\text{AC}, \mathcal{B}}^{\text{Att}}(\lambda) * q * \lceil N/2 \rceil.$$

Adversary  $\mathcal{B}$  simulates the environment of  $\text{G1}^{\text{AttKE}, A}$  as follows: at the beginning of the game,  $\mathcal{B}$  has to try and guess which session will have a forged message, and which message of the protocol it will be. As such, it samples uniformly from  $[1..q]$  a session  $s$  and from  $[1..\lceil N/2 \rceil]$  a message  $k$ . During `NewSession` such that  $i = s$ ,  $\mathcal{B}_1$  will output  $(\text{Rem}_{\text{KE}}^i, \phi_{\text{key}}, Q, k, \perp)$  to  $\text{Att}^{\text{AC}, \mathcal{B}}$ , store the produced  $R^*$  and all `hdl` output by `Load`( $R^*$ ) from there on. Afterwards, all calls to  $\mathcal{M}_R.\text{Load}(R^*)$  and  $\mathcal{M}_R.\text{Run}(\text{hdl}, m)$  or  $\mathcal{M}'_R.\text{Load}(R^*)$  and  $\mathcal{M}'_R.\text{Run}(\text{hdl}, m, \text{flag}, \text{fake})$  with the same calls on  $\text{Att}^{\text{AC}, \mathcal{B}}$ . Whenever `Send`( $m', i$ ) is called,  $i = s$  and  $m \neq \perp$ , let  $n$  be the number of messages sent for session  $s$ :

- If  $n = 0$ ,  $\mathcal{B}_2$  outputs  $(\perp, m', m^*)$ .
- If  $n > 0 \wedge n < k$ ,  $\mathcal{B}_2$  outputs  $(\text{st}_{\mathcal{B}}, m', m^*)$

It remains to show that, when `forgeAC` is set,  $\mathcal{B}$  wins  $\text{Att}^{\text{AC}, \mathcal{B}}$  with probability  $1/(q * \lceil N/2 \rceil)$ .

When the game ends and `forgeAC` = T, we have that

$$m \neq \perp \wedge \exists \text{hdl s.t. Program}_{\mathcal{M}}(\text{hdl}) = R_i^* . (m' : T_L^i) \sqsubseteq \text{ATrace}_{\mathcal{M}}(\text{hdl})$$

From the construction of `Send`,  $m \neq \perp$  means that this message has been successfully validated by `AC.Verify`. Furthermore, this matches the verifications in  $\text{Att}^{\text{AC}, \mathcal{B}}$ , considering that  $\text{in}_{\text{last}}^i \leftarrow$

$m^*$  in `Send`, so we know that all verifications that succeed in  $G1^{\text{AttKE}, \mathcal{A}}$  also do so in  $\text{Att}^{\text{AC}, \mathcal{B}}$ . If `forgeAC` is set upon receiving message  $k$ , we reach the final check. All calls that produce handles such that  $\text{Program}_{\mathcal{M}}(\text{hdl}) = R_i^*$  are also performed in  $\text{Att}^{\text{AC}, \mathcal{B}}$ , so it remains to show that  $\text{Rev}(m' : T_L^i)$  matches  $T$ .

We know that  $T \leftarrow \text{ATrace}_{R[\text{st}; \text{Coins}_{\mathcal{M}_R}(\text{hdl}^*)]}(i_1, \dots, i_n)$ . From the construction of `Send` and the behavior of the adversary, we know that this is constructed with the outputs given by  $\mathcal{B}_2$ . We will show, inductively, why  $\text{Rev}(m' : T_L^i)$  matches the  $k$  outputs of  $\mathcal{B}_2$ .

- Initially, given that  $(\perp, m'_0, m_0^*) \leftarrow \mathcal{B}_2$ :

$$\text{Rev}(m'_0 : T_L^i) = \text{Rev}[m'_0, \perp] = [\perp, m'_0]$$

- For all subsequent messages up to  $k$ ,  $(m_{n-1}^*, m'_n, m_n^*) \leftarrow \mathcal{B}_2$ , so

$$\text{Rev}(m' : T_L^i) = \text{Rev}[m'_n, m_{n-1}^*, \dots, m'_0, \perp] = [\perp, m'_0, \dots, m_{n-1}^*, m'_n]$$

As such, when `forgeAC` is set on session  $s$ , and in the  $k$ -th message input on `Send`, we have that

$$\exists \text{hdl s.t. } \text{Program}_{\mathcal{M}}(\text{hdl}) = R_i^* \cdot (m' : T_L^i) \sqsubseteq \text{ATrace}_{\mathcal{M}}(\text{hdl})$$

which results in a winning output for  $\text{Att}^{\text{AC}, \mathcal{B}}$  with probability  $1/(q * \lceil N/2 \rceil)$ .

In the third game  $G2^{\text{AttKE}, \mathcal{A}}$ , when the adversary loads and runs code, it is no longer interacting with machine  $\mathcal{M}$ , but rather with a simulator that is given the trace of a legitimate execution. Furthermore,  $\mathcal{S}$  handles the different behavior of  $\mathcal{M}_R$  and  $\mathcal{M}'_R$  following the description in Section 7. Intuitively, this difference corresponds to the indistinguishable scenario presented in the minimum leakage game of Figure 6. Given the presented differences, we have that

$$\Pr[G1^{\text{AttKE}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] - \Pr[G2^{\text{AttKE}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] \leq \Pr[\text{Leak-Real}^{\text{AC}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}] - \Pr[\text{Leak-Ideal}^{\text{AC}, \mathcal{A}}(1^\lambda) \Rightarrow \text{T}].$$

We must now argue that the difference between the games is bound by the adversary's advantage of breaking minimal leakage. In both possibilities for the bit  $b$ , the transformation in `Load` is exactly the same as the one in the minimum leakage game. Furthermore,  $T_R^{\text{hdl}}$  will always correspond to the  $\text{ATrace}_{\mathcal{M}}(\text{hdl})$ . First, consider  $b = 0$ . The behavior of `Run` is exactly the same as the one in the minimum leakage game, modulo the generation of the fake key, which will not be taken into consideration since it is only used in `Test` when  $b = 1$ . Now consider  $b = 1$ . The behavior of `Run` is exactly the same as the one in the minimum leakage game, modulo generating and setting the fake key. However, these additional operations match the described behavior expected from  $\mathcal{M}'_R$  for establishing the fake key.

As such, the advantage of the adversary in  $G2^{\text{AttKE}, \mathcal{A}}$  with respect to  $G1^{\text{AttKE}, \mathcal{A}}$  is limited by its advantage of breaking the minimum leakage.

Finally, in game  $G3^{\text{AttKE}, \mathcal{A}}$ , we no longer run the key exchange part of the code, but rather make use of `AttKE` and its corresponding oracles  $\{\text{NewRem}, \text{NewLoc}, \text{SendRem}, \text{SendLoc}, \text{TestRem}, \text{TestLoc}\}$ . The intuition is that, at this fourth game, calls to the attested part of the protocol can be provided using a key exchange for attested computation `AttKE`. More formally, we want to show that

$$\Pr[\mathsf{G2}^{\text{AttKE},\mathcal{A}}(1^\lambda) \Rightarrow \mathsf{T}] - \Pr[\mathsf{G3}^{\text{AttKE},\mathcal{A}}(1^\lambda) \Rightarrow \mathsf{T}] \leq \text{Adv}_{\text{AttKE},\mathcal{A}}^{\text{Att}}.$$

Whenever  $\mathsf{G3}^{\text{AttKE},\mathcal{A}}$  creates a local or a remote session (respectively, `NewSession` or `Load`), it will initialize accordingly on the `AttKE` protocol. On the remote side, a new `InsList` will also keep track of how many remote sessions are created (just like oracle `NewRem` in `AttKE`). On the local side, algorithm `Setup` will be replaced by a call to `NewLoc` that by its construction will return the same value as in  $\mathsf{G2}^{\text{AttKE},\mathcal{A}}$ . As such, it must now be argued that, every time the game is required to produce either an output from the local/remote machine or a key, the response given by the corresponding oracle is only distinguishable by an adversary that breaks `AttKE` security.

- For stage = 1  $\wedge \delta \neq \text{accept}$ : Locally, `SendLoc` replaces calling `LociKE`, which by the oracle behavior implies the same result. Remotely, instead of producing output via `R[st](in)`, we now directly execute oracle `SendRem`, which by construction holds and updates `st` in the same way as  $\mathsf{G2}^{\text{AttKE},\mathcal{A}}$ .
- For stage = 1  $\wedge \delta = \text{accept}$ : In  $\mathsf{G2}^{\text{AttKE},\mathcal{A}}$ , `SendLoc` would replace the real key with a fake key according to decision bit  $b$ . In  $\mathsf{G3}^{\text{AttKE},\mathcal{A}}$ , the same thing happens, but according to `AttKE` decision bit. Remotely, the behavior of `st`  $\leftarrow$  `TestRem` also differs in the same manner.
- For stage = 2: We know by `Composeφ[P,Q]`, that the state `st` considered for  $Q$  only maintains the key. As such, from thereon, the only difference between  $\mathsf{G2}^{\text{AttKE},\mathcal{A}}$  and  $\mathsf{G3}^{\text{AttKE},\mathcal{A}}$  is the key depending on the decision bit taken into consideration.

In this final game, the adversary's decision is based on a bit that is unrelated to any information obtainable via the oracles. Let

$$\text{Adv}_{\text{AC},\mathcal{A}}^{\text{Leak}} = \Pr[\text{Leak-Real}^{\text{AC},\mathcal{A}}(1^\lambda) \Rightarrow \mathsf{T}] - \Pr[\text{Leak-Ideal}^{\text{AC},\mathcal{A}}(1^\lambda) \Rightarrow \mathsf{T}]$$

we have that

$$\begin{aligned} \text{Adv}_{\text{AttKE},\mathcal{A}}^{\text{Ut}} &= \Pr[\mathsf{G}_0^{\text{AttKE},\mathcal{A}}(1^\lambda)] - \Pr[\mathsf{G}_3^{\text{AttKE},\mathcal{A}}(1^\lambda)] \\ &= \sum_{i=0}^3 \Pr[\mathsf{G}_i^{\text{AttKE},\mathcal{A}}(1^\lambda)] - \Pr[\mathsf{G}_{i+1}^{\text{AttKE},\mathcal{A}}(1^\lambda)] \\ &\leq \Pr[\text{forgeAC}] + \text{Adv}_{\text{AC},\mathcal{A}}^{\text{Leak}} + \text{Adv}_{\text{AttKE},\mathcal{A}}^{\text{Att}} \\ &\leq \text{Adv}_{\text{AC},\mathcal{B}}^{\text{Att}}(\lambda) * q * \lceil N/2 \rceil + \text{Adv}_{\text{AC},\mathcal{A}}^{\text{Leak}} + \text{Adv}_{\text{AttKE},\mathcal{A}}^{\text{Att}} \end{aligned}$$

and Theorem 3 follows. □

## D Proof of Theorem 4

In this section we provide the proof of security for the scheme defined in Section 8. In all the Section we assume that the encryption scheme  $(E, D, K)$  used is a secure authenticated encryption scheme. Therefore this encryption scheme satisfies both `IND-CPA` and `INT-CTXT`. In this section, let `AttKE` be the `AttKE` scheme used in the construction of our `SOC` scheme and `AC` be the `AC` scheme used.

## D.1 Integrity

The proof consists of two game hops described in Figure 26. The first hop from  $G_0$  to  $G_1$  consists in using the utility of the key exchange to replace the shared key by a magically shared fresh key. The second game hop from  $G_1$  to  $G_2$  is simply using sequence numbers and integrity of the encryption scheme to ensure that the local and remote traces actually match. Let us now give the details of each game hop.

$G_0$  to  $G_1$ :

Let  $\mathcal{A}$  be an adversary against  $G_0$  or  $G_1$ , let us build an adversary  $\mathcal{B}$  against  $\text{Att}_{\text{AttKE}}$ . The machine  $\mathcal{B}$  simulates  $\mathcal{A}$  giving  $\text{NewSession}(P)$  as input to  $\mathcal{A}_2$  instead of  $\text{Compile}^{\text{sec}}(\text{prms}, P, \text{id})$ , answering the oracle calls as follows:

**BootStrap( $o$ ):** Return  $\text{Send}(o, 0)$

**Send( $o^*, i$ ):**  $k \leftarrow \text{Test}(0)$   
 $o, c \leftarrow D(o^*, k)$   
 $\text{count}_l \leftarrow \text{count}_l + 1$   
 If  $c \neq \text{count}_l$  Return  $\perp$   
 $\text{count}_l \leftarrow \text{count}_l + 1$   
 $i^* \leftarrow E(\text{count}_l \# i, k)$   
 $\text{tr} \leftarrow i : o : \text{tr}$   
 Return  $o, i^*$

**Load( $R^*$ ):** Return  $\text{Load}(R^*)$

**Run( $\text{hdl}, i$ ):** Return  $\text{Run}(\text{hdl}, i)$

We now remark that if the bit  $b$  chosen in the utility game is 0, the game being played by  $\mathcal{B}$  is exactly the integrity game. On the other hand, if the bit chosen is 1, either  $\mathcal{B}$  violates entity authentication or  $\mathcal{B}$  behaves as  $\mathcal{A}$  playing against  $G_1$ . We conclude that

$$|G_0\text{-Int}_{\text{SOC}, \mathcal{A}}(1^\lambda) - G_1\text{-Int}_{\text{SOC}, \mathcal{A}}(1^\lambda)| \leq \text{Adv}_{\text{AttKE}, \mathcal{B}}^{\text{utility}}(1^\lambda)$$

**D.1.1  $G_1$  to  $G_2$ :**

First note that the only difference between  $G_1$  and  $G_2$  occurs when the **forge** event is raised (the other differences are simple rewritings). Now let  $\mathcal{A}$  be an adversary against  $G_2$ . We build an adversary  $\mathcal{B}$  against the INT-CTXT game by simulating  $\mathcal{A}$  playing  $G_2$  using the encryption/decryption oracles provided by the INT-CTXT game. Note that the event **forge** is raised if and only if  $\mathcal{B}$  wins the INT-CTXT game. Indeed, remarking that the counters are strictly increasing ensure that no encryption can be accepted twice as input to **Run** or **Send**. This entails the fact that if a message is accepted, it was the last message produced. We conclude that

$$|G_1\text{-Int}_{\text{SOC}, \mathcal{A}}(1^\lambda) - G_2\text{-Int}_{\text{SOC}, \mathcal{A}}(1^\lambda)| \leq \text{Adv}_{(E, D, K), \mathcal{B}}^{\text{INT-CTXT}}(1^\lambda)$$

It is now enough to remark that in  $G_2$ , the **Run** oracle and the **Send** oracle agree on inputs and outputs (unless **forge** is raised). As  $G_2$  always returns true, we conclude

$$\text{Adv}_{\text{SOC}}^{\text{Int}}(1^\lambda) \leq \text{Adv}_{\text{AttKE}}^{\text{utility}}(1^\lambda) + \text{Adv}_{(E, D, K)}^{\text{INT-CTXT}}(1^\lambda)$$

<p><b>Game <math>G_0\text{-IntSOC}, \mathcal{A}(1^\lambda)</math>:</b></p> <pre> prms <math>\leftarrow</math> <math>\mathcal{M}_R.\text{Init}(1^\lambda)</math> <math>(P, \text{st}_A) \leftarrow</math> <math>\mathcal{A}_1(\text{prms})</math> <math>P^*, \text{st}_I \leftarrow \text{Compile}^{\text{sec}}(\text{prms}, P, \text{id})</math> Run <math>\mathcal{A}_2^{\text{BootStrap, Send, Run, Load}}(\text{st}_A, P^*)</math> If <math>\nexists_{i=1} \text{hdl}</math> such that   Program<math>_{\mathcal{M}_R}(\text{hdl}) = P^* \wedge</math>   Translate(<math>\text{prms}, \text{Trace}_{\mathcal{M}_R}(\text{hdl})</math>) <math>\neq \square</math>   Return F hdl <math>\leftarrow</math> Program<math>_{\mathcal{M}_R}^{-1}(P^*)</math> <math>T \leftarrow \text{Translate}(\text{prms}, \text{Trace}_{\mathcal{M}_R}(\text{hdl}))</math> <math>T' \leftarrow \text{tr}</math> Return <math>\neg\Psi(T, T')</math> </pre>	<p><b>Oracle Bootstrap(<math>i</math>):</b></p> <pre> If <math>\text{st}_I.\text{accept}</math> Return <math>\perp</math> <math>i, \text{st}_I \leftarrow \text{BootStrap}(\text{prms}, o, \text{st}_I)</math> Return <math>i</math> </pre> <p><b>Oracle Send(<math>o^*, i</math>):</b></p> <pre> <math>o, c \leftarrow D(o^*, \text{st}_I.\text{key})</math> <math>\text{count}_I \leftarrow \text{count}_I + 1</math> If <math>c \neq \text{count}_I</math> Return <math>\perp</math> <math>\text{count}_I \leftarrow \text{count}_I + 1</math> <math>i^* \leftarrow E(\text{count}_I \# i, \text{st}_I.\text{key})</math> <math>\text{tr} \leftarrow i : o : \text{tr}</math> Return <math>o, i^*</math> </pre>	<p><b>Oracle Load(<math>R^*</math>):</b></p> <pre> hdl <math>\leftarrow \mathcal{M}_R.\text{Load}(R^*)</math> Return hdl </pre> <p><b>Oracle Run(hdl, <math>i</math>):</b></p> <pre> <math>o \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i)</math> Return <math>o</math> </pre>
<p><b>Game <math>G_1\text{-IntSOC}, \mathcal{A}(1^\lambda)</math>:</b></p> <pre> <math>k \leftarrow K(1^\lambda)</math> <math>\text{st}_P \leftarrow \emptyset</math> <math>\text{count}_R \leftarrow 0</math> <math>\text{count}_I \leftarrow 0</math> prms <math>\leftarrow</math> <math>\mathcal{M}_R.\text{Init}(1^\lambda)</math> <math>(P, \text{st}_A) \leftarrow</math> <math>\mathcal{A}_1(\text{prms})</math> <math>P^*, \text{st}_I \leftarrow \text{Compile}^{\text{sec}}(\text{prms}, P, \text{id})</math> Run <math>\mathcal{A}_2^{\text{BootStrap, Send, Run, Load}}(\text{st}_A, P^*)</math> If <math>\nexists_{i=1} \text{hdl}</math> such that   Program<math>_{\mathcal{M}_R}(\text{hdl}) = P^* \wedge</math>   Translate(<math>\text{prms}, \text{Trace}_{\mathcal{M}_R}(\text{hdl})</math>) <math>\neq \square</math>   Return F hdl <math>\leftarrow</math> Program<math>_{\mathcal{M}_R}^{-1}(P^*)</math> <math>T \leftarrow \text{Translate}(\text{prms}, \text{Trace}_{\mathcal{M}_R}(\text{hdl}))</math> <math>T' \leftarrow \text{tr}</math> Return <math>\neg\Psi(T, T')</math> </pre>	<p><b>Oracle Bootstrap(<math>o</math>):</b></p> <pre> If <math>\text{st}_I.\text{accept}</math> Return <math>\perp</math> <math>i, \text{st}_I \leftarrow \text{BootStrap}(\text{prms}, o, \text{st}_I)</math> Return <math>i</math> </pre> <p><b>Oracle Send(<math>o^*, i</math>):</b></p> <pre> If <math>\neg \text{st}_I.\text{accept}</math> Return <math>\perp</math> <math>o, c \leftarrow D(o^*, k)</math> <math>\text{count}_I \leftarrow \text{count}_I + 1</math> If <math>c \neq \text{count}_I</math> Return <math>\perp</math> <math>\text{count}_I \leftarrow \text{count}_I + 1</math> <math>i^* \leftarrow E(\text{count}_I \# i, k)</math> <math>\text{tr} \leftarrow i : o : \text{tr}</math> Return <math>o, i^*</math> </pre> <p><b>Oracle Load(<math>R^*</math>):</b></p> <pre> hdl <math>\leftarrow \mathcal{M}_R.\text{Load}(R^*)</math> Return hdl </pre>	<p><b>Oracle Run(hdl, <math>i^*</math>):</b></p> <pre> If Program<math>_{\mathcal{M}_R}(\text{hdl}) \neq P^*</math> <math>o^* \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i^*)</math> Return <math>o^*</math> If <math>\neg \text{st}_{\mathcal{M}_R}(\text{hdl}).\text{accept}</math> <math>o \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i^*)</math> Return <math>o^*</math> If <math>\exists \text{hdl}_2 \neq \text{hdl}.</math>Program<math>_{\mathcal{M}_R}(\text{hdl}) = P^*</math> <math>\wedge \text{st}_{\mathcal{M}_R}(\text{hdl}_2).\text{accept}</math>   raise twoPartners <math>i \# c \leftarrow D(i^*, k)</math> <math>\text{count}_R \leftarrow \text{count}_R + 1</math> If <math>c \neq \text{count}_R</math> Return <math>\perp</math> <math>o \leftarrow P[\text{st}_P](i)</math> <math>\text{count}_R \leftarrow \text{count}_R + 1</math> Return <math>E(\text{count} \# o, k)</math> </pre>
<p><b>Game <math>G_2\text{-IntSOC}, \mathcal{A}(1^\lambda)</math>:</b></p> <pre> <math>k \leftarrow K(1^\lambda)</math> <math>\text{st}_P \leftarrow \emptyset</math> <math>\text{count}_R \leftarrow 0</math> <math>\text{count}_I \leftarrow 0</math> <math>\text{tr}_R \leftarrow \square</math> prms <math>\leftarrow</math> <math>\mathcal{M}_R.\text{Init}(1^\lambda)</math> <math>(P, \text{st}_A) \leftarrow</math> <math>\mathcal{A}_1(\text{prms})</math> <math>P^*, \text{st}_I \leftarrow \text{Compile}^{\text{sec}}(\text{prms}, P, \text{id})</math> Run <math>\mathcal{A}_2^{\text{BootStrap, Send, Run, Load}}(\text{st}_A, P^*)</math> If <math>\nexists_{i=1} \text{hdl}</math> such that   Program<math>_{\mathcal{M}_R}(\text{hdl}) = P^* \wedge</math>   Translate(<math>\text{prms}, \text{Trace}_{\mathcal{M}_R}(\text{hdl})</math>) <math>\neq \square</math>   Return F hdl <math>\leftarrow</math> Program<math>_{\mathcal{M}_R}^{-1}(P^*)</math> <math>T \leftarrow \text{tr}_R</math> <math>T' \leftarrow \text{tr}</math> Return <math>\neg\Psi(T, T')</math> </pre>	<p><b>Oracle Bootstrap(<math>o</math>):</b></p> <pre> If <math>\text{st}_I.\text{accept}</math> Return <math>\perp</math> <math>i, \text{st}_I \leftarrow \text{BootStrap}(\text{prms}, o, \text{st}_I)</math> Return <math>i</math> </pre> <p><b>Oracle Send(<math>o^*, i</math>):</b></p> <pre> If <math>\neg \text{st}_I.\text{accept}</math> Return <math>\perp</math> <math>o, c \leftarrow D(o^*, k)</math> <math>\text{count}_I \leftarrow \text{count}_I + 1</math> If <math>c \neq \text{count}_I</math> Return <math>\perp</math> If <math>o \neq \text{last}_o</math> Raise forge <math>\text{count}_I \leftarrow \text{count}_I + 1</math> <math>i^* \leftarrow E(\text{count}_I \# i, k)</math> <math>\text{last}_i \leftarrow i</math> <math>\text{tr} \leftarrow i : o : \text{tr}</math> Return <math>o, i^*</math> </pre> <p><b>Oracle Load(<math>R^*</math>):</b></p> <pre> hdl <math>\leftarrow \mathcal{M}_R.\text{Load}(R^*)</math> Return hdl </pre>	<p><b>Oracle Run(hdl, <math>i^*</math>):</b></p> <pre> If Program<math>_{\mathcal{M}_R}(\text{hdl}) \neq P^*</math> <math>o^* \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i^*)</math> Return <math>o^*</math> If <math>\neg \text{st}_{\mathcal{M}_R}(\text{hdl}).\text{accept}</math> <math>o \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i^*)</math> Return <math>o^*</math> If <math>\exists \text{hdl}_2 \neq \text{hdl}.</math>Program<math>_{\mathcal{M}_R}(\text{hdl}) = P^*</math> <math>\wedge \text{st}_{\mathcal{M}_R}(\text{hdl}_2).\text{accept}</math>   raise twoPartners <math>i \# c \leftarrow D(i^*, k)</math> <math>\text{count}_R \leftarrow \text{count}_R + 1</math> If <math>c \neq \text{count}_R</math> Return <math>\perp</math> If <math>i \neq \text{last}_i</math> Raise forge <math>o \leftarrow P[\text{st}_P](i)</math> <math>\text{count}_R \leftarrow \text{count}_R + 1</math> <math>\text{last}_o \leftarrow o</math> <math>\text{tr}_R \leftarrow o : i : \text{tr}_R</math> Return <math>E(\text{count} \# o, k)</math> </pre>

Figure 26: Game hops for integrity of the SOC scheme

## D.2 Privacy

The proof of consists, as previously in replacing the derived key by a magically shared key. The second game hop, as in the integrity property, makes sure that the inputs received on both sides coincide. The game  $G_2$  obtained reduces then quite simply to IND-CPA of the encryption scheme. The game hop is presented in Figure 27. The reduction from  $G_0$  to  $G_1$

is exactly the same as the one in the integrity proof, we do not detail it further. We get as previously (for some PPT  $\mathcal{B}$ )

$$\left| G_0\text{-Priv}_{\text{SOC},\mathcal{A}}(1^\lambda) - G_1\text{-Priv}_{\text{SOC},\mathcal{A}}(1^\lambda) \right| \leq \text{Adv}_{\text{AttKE},\mathcal{B}}^{\text{utility}}(1^\lambda)$$

**$G_1$  to  $G_2$ :**

Let  $\mathcal{A}$  be an adversary against  $G_1$  or  $G_2$ . Let us build an adversary  $\mathcal{B}$  against the INT-CTXT game for  $(E, D, K)$ . The machine  $\mathcal{B}$  simulates  $\mathcal{A}$  playing the game  $G_1$ , with the exception of not drawing  $k$  and using the encryption and decryption oracle of the INT-CTXT game. Remark that  $G_1$  and  $G_2$  behave differently only if  $\mathcal{A}$  is able to submit a forged encryption to either the send or the run oracle. Indeed, if the encryption submitted is neither the last produced encryption nor a forgery, the sequence number makes sure that the corresponding oracle return  $\perp$ . We get

$$\left| G_1\text{-Priv}_{\text{SOC},\mathcal{A}}(1^\lambda) - G_2\text{-Priv}_{\text{SOC},\mathcal{A}}(1^\lambda) \right| \leq \text{Adv}_{(E,D,K),\mathcal{B}}^{\text{INT-CTXT}}(1^\lambda)$$

**Reducing  $G_2$  to IND-CPA:**

The key point in the proof is reducing  $G_2$  to the security of the authenticated encryption scheme. Let us call  $\text{Chal}$  the IND-CPA challenge oracle and  $\text{Enc}$  the encryption oracle in the IND-CPA game for  $(E, D, K)$ . Let  $\mathcal{A}$  be an adversary against  $G_2\text{-Priv}$ . We build the following adversary  $\mathcal{B}$  against the IND-CPA game for  $(E, D, K)$ . The machine  $\mathcal{B}$  simulates the game  $G_2$ , without drawing the key  $k$  or the bit  $b$  and using the IND-CPA oracles to perform encryptions. In the  $\text{Send}$  oracle, instead of computing  $E(\text{count}_i \# i_b, k)$ ,  $\mathcal{B}$  calls  $\text{Chal}(\text{count}_i \# i_0, \text{count}_i \# i_1)$ . Similarly, in the  $\text{Run}$  oracle, instead of computing  $E(\text{count} \# o_b, k)$ ,  $\mathcal{B}$  calls  $\text{Chal}(\text{count} \# o_0, \text{count} \# o_1)$ . All calls are well formed as at each call of  $\text{Send}$ , the length of the two inputs is required to be identical and, in the  $\text{Run}$  oracle,  $P$  is assumed length-uniform. It is now enough to remark that  $\mathcal{B}$  wins the IND-CPA game if and only if  $\mathcal{A}$  wins  $G_0$  to conclude

$$\left| G_2\text{-Priv}_{\text{SOC},\mathcal{A}}(1^\lambda) - \frac{1}{2} \right| \leq \text{Adv}_{(E,D,K),\mathcal{B}}^{\text{IND-CPA}}(1^\lambda)$$

From this result, we can sum everything up and obtain the advantage of an adversary against the privacy game:

$$\text{Adv}_{\text{SOC}}^{\text{Priv}}(1^\lambda) \leq \text{Adv}_{(E,D,K),\mathcal{B}}^{\text{IND-CPA}}(1^\lambda) + \text{Adv}_{(E,D,K),\mathcal{B}}^{\text{INT-CTXT}}(1^\lambda) + \text{Adv}_{\text{AttKE},\mathcal{B}}^{\text{utility}}(1^\lambda)$$

<p><b>Game <math>G_0\text{-Priv}_{\text{SOC},\mathcal{A}}(1^\lambda)</math>:</b></p> <p><math>k \leftarrow \mathcal{K}(1^\lambda)</math>  <math>\text{st}_P \leftarrow \emptyset</math>  <math>\text{count}_R \leftarrow 0</math>  <math>\text{count}_I \leftarrow 0</math>  <math>\text{prms} \leftarrow \mathcal{M}_R.\text{Init}(1^\lambda)</math>  <math>(P, \text{st}_A) \leftarrow \mathcal{A}_1(\text{prms})</math>  <math>P^*, \text{st}_I \leftarrow \text{Compile}^{\text{sec}}(\text{prms}, P, \text{id})</math>  <math>b' \leftarrow \mathcal{A}_2^{\mathcal{O}}(\text{st}_A, P^*)</math>  Return <math>b = b'</math></p>	<p><b>Oracle <math>\text{BootStrap}(o)</math>:</b></p> <p>If <math>\text{st}_I.\text{accept}</math> Return <math>\perp</math>  <math>i, \text{st}_I \leftarrow \text{BootStrap}(\text{prms}, o, \text{st}_I)</math>  Return <math>i</math></p> <p><b>Oracle <math>\text{Send}_b(o^*, i_0, i_1)</math>:</b></p> <p>If <math> i_0  \neq  i_1 </math> Return <math>\perp</math>  <math>o, c \leftarrow D(o^*, \text{st}_I.\text{key})</math>  <math>\text{count}_I \leftarrow \text{count}_I + 1</math>  If <math>c \neq \text{count}_I</math> Return <math>\perp</math>  <math>\text{count}_I \leftarrow \text{count}_I + 1</math>  <math>i^* \leftarrow E(\text{count}_I \# i_b, \text{st}_I.\text{key})</math>  <math>\text{tr} \leftarrow i : o : \text{tr}</math>  Return <math>o, i^*</math></p>	<p><b>Oracle <math>\text{Load}(R^*)</math>:</b></p> <p><math>\text{hdl} \leftarrow \mathcal{M}_R.\text{Load}(R^*)</math>  Return <math>\text{hdl}</math></p> <p><b>Oracle <math>\text{Run}(\text{hdl}, i^*)</math>:</b></p> <p><math>o^* \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i^*)</math>  Return <math>o^*</math></p>
<p><b>Game <math>G_1\text{-Priv}_{\text{SOC},\mathcal{A}}(1^\lambda)</math>:</b></p> <p><math>k \leftarrow \mathcal{K}(1^\lambda)</math>  <math>\text{st}_P \leftarrow \emptyset</math>  <math>\text{count}_R \leftarrow 0</math>  <math>\text{count}_I \leftarrow 0</math>  <math>\text{prms} \leftarrow \mathcal{M}_R.\text{Init}(1^\lambda)</math>  <math>(P, \text{st}_A) \leftarrow \mathcal{A}_1(\text{prms})</math>  <math>P^*, \text{st}_I \leftarrow \text{Compile}^{\text{sec}}(\text{prms}, P, \text{id})</math>  <math>b' \leftarrow \mathcal{A}_2^{\mathcal{O}}(\text{st}_A, P^*)</math>  Return <math>b = b'</math></p>	<p><b>Oracle <math>\text{BootStrap}(o)</math>:</b></p> <p>If <math>\text{st}_I.\text{accept}</math> Return <math>\perp</math>  <math>i, \text{st}_I \leftarrow \text{BootStrap}(\text{prms}, o, \text{st}_I)</math>  Return <math>i</math></p> <p><b>Oracle <math>\text{Send}_b(o^*, i_0, i_1)</math>:</b></p> <p>If <math> i_0  \neq  i_1 </math> Return <math>\perp</math>  If <math>\neg \text{st}_I.\text{accept}</math> Return <math>\perp</math>  <math>o, c \leftarrow D(o^*, k)</math>  <math>\text{count}_I \leftarrow \text{count}_I + 1</math>  If <math>c \neq \text{count}_I</math> Return <math>\perp</math>  <math>\text{count}_I \leftarrow \text{count}_I + 1</math>  <math>i^* \leftarrow E(\text{count}_I \# i_b, k)</math>  <math>\text{tr} \leftarrow i : o : \text{tr}</math>  Return <math>o, i^*</math></p> <p><b>Oracle <math>\text{Load}(R^*)</math>:</b></p> <p><math>\text{hdl} \leftarrow \mathcal{M}_R.\text{Load}(R^*)</math>  Return <math>\text{hdl}</math></p>	<p><b>Oracle <math>\text{Run}(\text{hdl}, i^*)</math>:</b></p> <p>If <math>\text{Program}_{\mathcal{M}_R}(\text{hdl}) \neq P^*</math>  <math>o^* \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i^*)</math>  Return <math>o^*</math></p> <p>If <math>\neg \text{st}_{\mathcal{M}_R}(\text{hdl}).\text{accept}</math>  <math>o \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i^*)</math>  Return <math>o^*</math></p> <p>If <math>\exists \text{hdl}_2 \neq \text{hdl}.\text{Program}_{\mathcal{M}_R}(\text{hdl}) = P^*</math>  <math>\wedge \text{st}_{\mathcal{M}_R}(\text{hdl}).\text{accept}</math>  raise twoPartners  <math>i \# c \leftarrow D(i^*, k)</math>  <math>\text{count}_R \leftarrow \text{count}_R + 1</math>  If <math>c \neq \text{count}_R</math> Return <math>\perp</math>  <math>o \leftarrow P[\text{st}_P](i)</math>  <math>\text{count}_R \leftarrow \text{count}_R + 1</math>  Return <math>E(\text{count} \# o, k)</math></p>
<p><b>Game <math>G_2\text{-Priv}_{\text{SOC},\mathcal{A}}(1^\lambda)</math>:</b></p> <p><math>k \leftarrow \mathcal{K}(1^\lambda)</math>  <math>\text{st}_P \leftarrow \emptyset</math>  <math>\text{count}_R \leftarrow 0</math>  <math>\text{count}_I \leftarrow 0</math>  <math>\text{prms} \leftarrow \mathcal{M}_R.\text{Init}(1^\lambda)</math>  <math>(P, \text{st}_A) \leftarrow \mathcal{A}_1(\text{prms})</math>  <math>P^*, \text{st}_I \leftarrow \text{Compile}^{\text{sec}}(\text{prms}, P, \text{id})</math>  <math>b' \leftarrow \mathcal{A}_2^{\mathcal{O}}(\text{st}_A, P^*)</math>  Return <math>b = b'</math></p>	<p><b>Oracle <math>\text{BootStrap}(o)</math>:</b></p> <p>If <math>\text{st}_I.\text{accept}</math> Return <math>\perp</math>  <math>i, \text{st}_I \leftarrow \text{BootStrap}(\text{prms}, o, \text{st}_I)</math>  Return <math>i</math></p> <p><b>Oracle <math>\text{Send}_b(o^*, i_0, i_1)</math>:</b></p> <p>If <math> i_0  \neq  i_1 </math> Return <math>\perp</math>  <b>If <math>o^* \neq \text{last}_{o^*}</math> Return <math>\perp</math></b>  <math>i^* \leftarrow E(\text{count}_I \# i_b, k)</math>  <math>\text{last}_{i_0} \leftarrow i_0</math>  <math>\text{last}_{i_1} \leftarrow i_1</math>  <math>\text{last}_{i^*} \leftarrow i^*</math>  <math>\text{tr} \leftarrow i : o : \text{tr}</math>  Return <math>o, i^*</math></p> <p><b>Oracle <math>\text{Load}(R^*)</math>:</b></p> <p><math>\text{hdl} \leftarrow \mathcal{M}_R.\text{Load}(R^*)</math>  Return <math>\text{hdl}</math></p>	<p><b>Oracle <math>\text{Run}_b(\text{hdl}, i^*)</math>:</b></p> <p>If <math>\text{Program}_{\mathcal{M}_R}(\text{hdl}) \neq P^*</math>  <math>o^* \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i)</math>  Return <math>o^*</math></p> <p>If <math>\neg \text{st}_{\mathcal{M}_R}(\text{hdl}).\text{accept}</math>  <math>o \leftarrow \mathcal{M}_R.\text{Run}(\text{hdl}, i)</math>  Return <math>o^*</math></p> <p>If <math>\exists \text{hdl}_2 \neq \text{hdl}.\text{Program}_{\mathcal{M}_R}(\text{hdl}) = P^*</math>  <math>\wedge \text{st}_{\mathcal{M}_R}(\text{hdl}).\text{accept}</math>  raise twoPartners  <b>If <math>i \neq \text{last}_{i^*}</math> Return <math>\perp</math></b>  <math>\text{count}_R \leftarrow \text{count}_R + 1</math>  <math>o_0 \leftarrow P[\text{st}_P^0](\text{last}_{i_0})</math>  <math>o_1 \leftarrow P[\text{st}_P^1](\text{last}_{i_1})</math>  <math>\text{count}_R \leftarrow \text{count}_R + 1</math>  <math>o^* \leftarrow E(\text{count} \# o_b, k)</math>  <math>\text{last}_{o^*} \leftarrow o^*</math>  Return <math>o^*</math></p>

Figure 27: Game hops in the privacy proof