

# Steel: Composable Hardware-based Stateful and Randomised Functional Encryption

Pramod Bhatotia<sup>1,2</sup>, Markulf Kohlweiss<sup>1</sup>, Lorenzo Martinico<sup>1</sup>, and Yiannis Tselekounis<sup>1</sup>

<sup>1</sup> School of Informatics, University of Edinburgh

<sup>2</sup> Department of Informatics, TU Munich

**Abstract.** Trusted execution environments (TEEs) enable secure execution of programs on untrusted hosts and cryptographically attest the correctness of outputs. As these are complex systems, it is essential to formally capture the exact security achieved by protocols employing TEEs, and ultimately, prove their security under composition, as TEEs are typically employed in multiple protocols, simultaneously.

Our contribution is twofold. On the one hand, we show that under existing definitions of attested execution setup, we can realise cryptographic functionalities that are unrealisable in the standard model. On the other hand, we extend the adversarial model to capture a broader class of *realistic adversaries*, we demonstrate weaknesses of existing security definitions this class, and we propose stronger ones.

Specifically, we first define a generalization of *Functional Encryption* that captures *Stateful and Randomised* functionalities (FESR). Then, assuming the ideal functionality for attested execution of Pass et al. (Eurocrypt '2017), we construct the associated protocol, **Steel**, and we prove that **Steel** UC-realises FESR in the *universal composition with global subroutines* model by Badertscher et al. (TCC '2020). Our work is also a validation of the compositionality of the **Iron** protocol by Fisch et al. (CCS '2017), capturing (non-stateful) hardware-based functional encryption.

As the existing functionality for attested execution of Pass et al. is too strong for real world use, we propose a weaker functionality that allows the adversary to conduct *rollback and forking attacks*. We demonstrate that **Steel** (realising stateful functionalities), contrary to the stateless variant corresponding to **Iron**, is not secure in this setting and discuss possible mitigation techniques.

## 1 Introduction

Due to the rise of cloud computing, most people living in countries with active digital economies can expect a significant amount of information about them to be stored on cloud platforms. Cloud computing offers economies of scale for computational resources with ease of management, elasticity, and fault tolerance driving further centralization. While cloud computing is ubiquitously employed for building modern online service, it also poses security and privacy risks. Cloud

storage and computation are outside the control of the data owner and users currently have no mechanism to verify whether the third-party operator, even with good intentions, can handle their data with confidentiality and integrity guarantees.

*Hardware-based solutions.* To overcome these limitations, trusted execution environments (TEEs), such as Intel SGX [27], ARM Trustzone [45], RISC-V Keystone [29, 38], AMD-SEV [33] provide an appealing way to build secure systems. TEEs provide a hardware-protected secure memory region called a *secure enclave* whose residing code and data are isolated from any layers in the software stack including the operating system and/or the hypervisor. In addition, TEEs offer remote attestation for proving their trustworthiness to third-parties. In particular, the remote attestation enables a remote party to verify that an enclave has a specific identity and is indeed running on a genuine TEE hardware platform. Given they promise a hardware-assisted secure abstraction, TEEs are now commercially offered by major cloud computing providers including Microsoft Azure [47], Google Cloud [46], and Alibaba Cloud [26].

*Modeling challenges.* While TEEs provide a promising building block, it is not straightforward to design secure applications on top of TEEs. In particular applications face the following three challenges: (1) Most practical applications require combining trusted and untrusted components for improved performance and a low trusted computing base; (2) TEEs are designed to protect only the volatile, in-memory, “stateless” computations and data. Unfortunately, this abstraction is insufficient for most practical applications, which rely on stateful computation on untrusted storage mediums (SSDs, disks). Ensuring security for such untrusted storage mediums is challenging because TEEs are prone to rollback attacks; and lastly, (3) TEE hardware designs are prone to numerous side channel attacks exploiting memory access patterns, cache timing channels, etc. These side channel attacks have the potential to completely compromise the confidentiality, integrity, and authenticity (remote attestation) of enclaves.

Therefore, it is important to carefully model the security achieved by the protocols of such systems as well as the assumptions in the cryptography and the hardware, and the trust afforded in protocol participants. Ideally such modelling must be compositional to facilitate the construction of larger systems based on smaller hardware and cryptography components. Given a sufficiently expressive model of TEEs, they can be used as a powerful setup assumption to realise many protocols.

The model of Pass, Shi, and Tramer (PST) [44] takes an initial step towards modelling protocols employing TEEs. The PST model provides a compositional functionality for attested execution and shows how to instantiate various primitives impossible in the standard model, as well as some limitations of TEEs. The PST model was first weakened in [52], which provides a compelling example of how an excessively weak enclave, susceptible to side channel attacks that break confidentiality (but not integrity and authenticity), can still be used as setup for useful cryptographic primitives. Both models, however, live at two opposite

extremes, and thus fail to capture realistic instantiations of real world trusted execution.

*Functional encryption & limitations.* One of the core primitives that enables privacy preserving computation and storage is *Functional Encryption* (FE), introduced by [15]. FE is a generalisation of Attribute/Identify Based Encryption [49, 48], that enables authorized entities to compute over encrypted data, and learn the results in the clear. In particular, parties possessing the so-called functional key,  $sk_f$ , for the function  $f$ , can compute  $f(x)$ , where  $x$  is the plaintext, by applying the decryption algorithm on  $sk_f$  and an encryption of  $x$ . Access to the functional key is regulated by a trusted third party. While out of scope for our work, identifying such a party is an interesting question that requires establishing metrics for the trustworthiness of entities we might want to be able to decrypt functions, and the kind of functions that should be authorised for a given level of trust. An obvious option for the role of trusted authority would be that of a data protection authority, who can investigate the data protection practices of organisations and levy fines in case these are violated. Another approach could be decentralising this role, by allowing the functional key to be generated collectively by a number of data owners [22, 1].

FE is a very powerful primitive but in practice highly non-trivial to construct. Motivated by the inefficiency of existing instantiations of FE for arbitrary functions, the work of [28] introduces *Iron*, which is a practically efficient protocol that realises FE based on Intel’s SGX. In [28] the authors formally prove security of the proposed protocol, however their proof is in the standalone setting. In a related work, Matt and Maurer [41] show (building on [3]) that composable functional encryption (CFE) is impossible to achieve in the standard model, but achievable in the random oracle model. For another important variant of the primitive, namely, *randomized functional encryption*, existing constructions [2, 30, 37], are limited in the sense that they require a new functional key for each invocation of the function, i.e., decryptions with the same functional key always return the same output. Finally, existing notions of FE only capture *stateless* functionalities, which we believe further restricts the usefulness and applicability of the primitive. For instance, imagine a financial institution that sets its global lending rate based on the total liquidity of its members. Financial statements can be sent, encrypted, by each member, with each of these transactions updating the global view for the decryptor, who can then compute the function’s result in real time.

Given the above limitations, in this work we leverage the power of hardware assisted computation to construct FE for a *broader class of functionalities* under the strongest notion of *composable security*.

## 1.1 Our Contributions

We consider a generalization of FE to arbitrary *stateful and probabilistic functionalities* (FESR), that subsumes multi-client FE [22] and enables crypto-

graphic computations in a natural way, due to the availability of internal randomness. Our contributions are as follows:

- We formally define functional encryption for stateful and randomized functionalities (FESR), in the Universal Composition (UC) setting [28].
- We construct the protocol **Steel** and prove that it realizes FESR in the newly introduced Universal Composition with Global Subroutines (UCGS) model [8]. Our main building blocks are: (1) the functional encryption scheme of [28] and (2) the global attestation functionality of PST. Our treatment lifts the PST model to the UCGS setting, and by easily adapting our proofs one can also establish the UCGS-security of [28].
- Finally, we introduce a weaker functionality for attested execution in the UCGS model to allow rollback and forking attacks, and use it to demonstrate that **Steel** does not protect against these. Finally, we sketch possible mitigation techniques.

## 1.2 Technical Overview

*Attested execution via the global attestation functionality  $G_{\text{att}}$  of PST [44].* Our UC protocols assume access to the *global attestation functionality*,  $G_{\text{att}}$ , that captures the core abstraction provided by a broad class of attested execution processors, such as Intel SGX [27]. It models multiple hardware-protected memory regions of a TEE, called *secure enclaves*. Each enclave contains trusted code and data. In combination with a call-gate mechanism to control entry and exit into the trusted execution environment, this guarantees that this memory can only be accessed by the enclave it belongs to, i.e., the enclave memory is protected from concurrent enclaves and other (privileged) code on the platform. TEE processing environments guarantee the authenticity, the integrity and the confidentiality of their executing code, data and runtime states, e.g. CPU registers, memory and others.

$G_{\text{att}}$  is parametrised by a signature scheme and a registry that captures all the platforms that are equipped with an attested execution processor. At a high level,  $G_{\text{att}}$  allows parties to register programs and ask for evaluations over arbitrary inputs, while also receiving signatures that ensure correctness of the computation. Since the manufacturer’s signing key pair can be used in multiple protocols simultaneously,  $G_{\text{att}}$  is defined as a *global functionality* that uses the same key pair across sessions.

*Universal composition with global subroutines [9].* In our work we model global information using the newly introduced UCGS framework, which resolves inconsistencies in GUC [18], an earlier work that aims to model executions in the presence of global functionalities. UCGS handles such executions via a *management* protocol, that combines the target protocol and one or more instances of the global functionality, and creates an embedding within the standard UC framework. In our work,  $G_{\text{att}}$  (cf. Section 2.2) is modeled as a global functionality in the UCGS framework (updating the original PST formulation in GUC).

*Setting, adversarial model & security.* Our treatment considers three types of parties namely, encryptors, denoted by A, decryptors, denoted by B, as well as a single party that corresponds to the trusted authority, denoted by C. The adversary is allowed to corrupt parties in B and request for evaluations of functions of its choice over messages encrypted by parties in A. We then require *correctness* of the computation, meaning that the state for each function has not been tampered with by the adversary, as well as *confidentiality* of the encrypted message, which ensures that the adversary learns only the output of the computation (and any information implied by it) and nothing more. Our treatment covers both stateful and randomized functionalities.

*Steel: UCGS-secure FE for stateful and randomized functionalities.* Steel is executed by the sets of parties discussed above, where besides encryptors, all other parties receive access to  $G_{\text{att}}$ , abstracting an execution in the presence of secure hardware enclaves. Our protocol is based on Iron [28], so we briefly revisit the main protocol operations: (1) **Setup**, executed by the trusted party C, installs a *key management enclave* (KME), running a program to generate *public-key encryption* and *digital signature*, key pairs. The public keys are published, while the equivalent secrets are kept encrypted in storage (using SGX’s terminology, the memory is sealed). Each of the decryptors installs a *decryption enclave* (DE), and attests its authenticity to the KME to receive the secret key for the encryption scheme over a secure channel. (2) **KeyGen**, on input function F, calls KME, where the latter produces a signature on the measurement of an instantiated enclave that computes F. (3) When **Encrypt** is called by an encryptor, it uses the published public encryption key to encrypt a message and sends the ciphertext to the intended recipients. (4) **Decrypt** is executed by a decrypting party seeking to compute some function F on a ciphertext. This operation instantiates a matching function enclave (or resume an existing one), whose role is that of computing the functional decryption, if an authorised functional key is provided.

Steel consists of the above operations, with the appropriate modifications to enable stateful functionalities. In addition, Steel provides some simplifications over the Iron protocol. In particular, we repurpose attestation’s signature capabilities to supplant the need for a separate signature scheme to generate functional keys, and thus minimise the trusted computing base. In practice, a functional key for a function F can be produced by just letting the key generation process return F; as part of  $G_{\text{att}}$ ’s execution, this produces an attestation signature  $\sigma$  over F, which becomes the functional key  $\text{sk}_F$  for that function, provided the generating enclave id is also made public (a requirement for verification, due to the signature syntax of attestation in  $G_{\text{att}}$ ).

The statefulness of functional encryption is simply enabled by adding a state array to each functional enclave. The array is also stored locally by the corresponding decryption enclave, and is updated for every decryption of a given function. Similar to [44], a curious artefact in the protocol’s modeling is the addition of a “backdoor” that programs the output of the function evaluation subroutine, such that, if a specific argument is set on the input, the function evaluation returns the value of that argument. The reason for this addition is

to enable simulation of signatures over function evaluations that have already been computed using the ideal functionality. We note that this addition does not impact correctness, as the state array is not modified if the backdoor is used, nor confidentiality, since the output of this subroutine is never passed to any other party besides the caller **B**. Finally, a further addition is that our protocol requires the addition of a proof of plaintext knowledge on top of the underlying encryption scheme. The **Steel** protocol definition is presented in Section 4.

*Security of Steel.* Our protocol uses an existentially unforgeable under chosen message attacks (EU-CMA) signature scheme,  $\Sigma$ , a CCA-secure public-key encryption scheme,  $\text{PKE}$ , and a non-interactive zero knowledge scheme,  $\text{N}$ . Informally,  $\Sigma$  provides the guarantees required for realizing attested computation (as discussed above),  $\text{PKE}$  is used to protect the communication between enclaves, and for protecting the encryptors’ inputs. For the latter usage, it is possible to reduce the security requirement to CPA-security as we additionally compute a simulation-extractable NIZK proof of well-formedness of the ciphertext that guarantees non-malleability.

Our proof is via a sequence of hybrids in which we prove that the real world protocol execution w.r.t. **Steel** is indistinguishable from the ideal execution, in the presence of an ideal functionality that captures FE for stateful and randomized functionalities. The goal is to prove that the decryptor learns nothing more than an authorized function of the private input plaintext, thus our hybrids gradually fake all relevant information accessed by the adversary. In the first hybrid,<sup>3</sup> all signature verifications w.r.t. the attestation key are replaced by an idealized verification process, that only accepts message/signature pairs that have been computed honestly (i.e., we omit verification via  $\Sigma$ ). Indistinguishability is proven via reduction to the EU-CMA security of  $\Sigma$ . Next we fake all ciphertexts exchanged between enclaves that carry the decryption key for the target ciphertext, over which the function is evaluated (those hybrids require reductions to the CCA security of  $\text{PKE}$ ).<sup>4</sup> The next hybrid substitutes ZK proofs over the target plaintexts with simulated ones, and indistinguishability with the previous one reduces to the zero knowledge property of  $\text{N}$ . Then, for maliciously generated ciphertexts under  $\text{PKE}$  – which might result via tampering with honestly generated encryptors’ ciphertexts – instead of using the decryption operation of  $\text{PKE}$ , our simulator recovers the corresponding plaintext using the extractability property of  $\text{N}$ . Finally, we fake all ciphertexts of  $\text{PKE}$ , that encrypt the inputs to the functions (this reduces to CPA security). Note that, in [28], the adversary outputs the target message, which is then being encrypted and used as a parameter to the ideal world functionality that is accessed by the simulator in a black box way. In this work, we consider a stronger setting in which the adversary directly outputs ciphertexts of it’s choice. While in the classic setting for Functional Encryption (where **Iron** lives) simulation security is easily achieved

<sup>3</sup> Here we omit some standard UC-related hybrids.

<sup>4</sup> Here CCA security is a requirement as the adversary is allowed to tamper with honestly generated ciphertexts.

by asking the adversarial enclave to produce an evaluation for the challenge ciphertext, in FESR the simulator is required to conduct all decryptions through the ideal functionality, so that the decryptor’s state for that function can be updated. We address the above challenge by using the extractability property of NIZKs: for maliciously generated ciphertexts our simulator extracts the original plaintext and ask the ideal FESR functionality for it’s evaluation. Simulation-extractable NIZK can be efficiently instantiated, e.g., using zk-SNARKs [11]. Security of our protocol is formally proven in Section 5. The simulator therein provided could be easily adapted to show that the *Iron* protocol UCGS-realises Functional Encryption, by replacing the NIZK operations for maliciously generated ciphertexts with a decryption from the enclave, as described above.

*Rollback & forking attacks.* Modeling attested execution via  $G_{\text{att}}$  facilitates composable protocol design, however, such a functionality cannot be easily realized since real world adversaries can perform highly non-trivial *rollback* and *forking* attacks against hardware components. In Section 6, we define a weaker functionality for attested execution, called  $G_{\text{att}}^{\text{rollback}}$ , that aims to capture rollback and forking attacks. To achieve this, we replace the enclave storage array in  $G_{\text{att}}$  with a tree data structure. While the honest party only ever accesses the last leaf of the tree (equivalent to a linked list), a corrupt party is able to provide an arbitrary path within the tree. This allows them to rollback the enclave, by re-executing a previous (non-leaf) state, and to support multiple forks of the program by interactively selecting different sibling branches. We give an example FESR function where we can show that correctness does not hold if  $G_{\text{att}}^{\text{rollback}}$  is used instead of  $G_{\text{att}}$  within *Steel*, and discuss how countermeasures from the rollback protection literature can be adopted to address these attacks, with a consideration on efficiency.

### 1.3 Related work

Hardware is frequently used to improve performance or circumvent impossibility results, e.g. [42, 4, 25]. As a relevant example, Chung et al. [24] show how to use of stateless hardware tokens to implement functional encryption

The use of attestation has been widely adopted in the design of computer systems to bootstrap security [43]. In addition to formalising attested execution, Pass, Shi and Tramer (PST) [44] show that two-party computation is realisable in UC only if both parties have access to attested execution, and fair two-party computation is also possible if additionally both secure processors have access to a trusted clock. The PST model is the first work to formalise attested execution in the UC framework. The compositional aspect of UC allows for the reused of the model in several successive works [52, 54, 23, 21]. Other attempts at providing a formal model for attested execution include the game-based models of Barbosa et al. [14], Bahmani et al. [12], Fisch et al. [28]. The latter model arises from the need to evaluate the security of *Iron*, a hardware-based realisation of functional encryption, which was later extended to verifiable functional encryption in Suzuki et al. [51].

Rollback attacks (also known as reset attacks in the cryptographic literature) are a common attack vectors against third-party untrusted computing infrastructure. An attacker who is in control of the underlying infrastructure can at times simply restart the system to restore a previous system state. Yilek [55] presents a general attack that is applicable to both virtual machine and enclave executions: it shows that an adversary capable of executing multiple rollback attacks on IND-CCA or IND-CPA secure encryption schemes might learn information about encrypted messages by running the encryption algorithm on multiple messages with the same randomness. In the absence of true hardware-based randomness that cannot be rolled back, these kinds of attacks can be mitigated using hedged encryption, a type of key-wrap scheme [32], such that for each encryption round, the original random coin and the plaintext are passed through a pseudorandom function to generate the randomness for the ciphertext.

The area of rollback attacks on TEEs is well studied. Platforms like SGX [20], TPMs [39], etc. provide trusted monotonic counters, from which it is possible to bootstrap rollback-resilient storage. However, trusted counters are too slow for most practical applications. Furthermore, they wear out after a short period of time. As their lifetime is limited, they are unreliable for applications that require frequent updates [40]. Moreover, an adversary that is aware of this vulnerability can attack protocols that rely exclusively on counters, by instantiating a malicious enclave on the same platform that artificially damages the counters.

To overcome the limitation of SGX counters, ROTE [40] uses a consensus protocol to build a distributed trusted counter service, with performance necessarily reduced through several rounds of network communication. In the same spirit, Ariadne [50] is an optimized (local) synchronous technique to increment the counter by a single bit flip for deterministic enclaves.

Speicher [13] and Palaemon [31] proposed an asynchronous trusted counter interface, which provide a systematic trade-off between performance and rollback protection, addressing some limitations of synchronous counters. The asynchronous counter is backed up by a synchronous counter interface with a period of vulnerability, where an adversary can rollback the state of a TEE-equipped storage server in a system until the last stable synchronous point. To protect against such attacks, these systems rely on the clients to keep the changes in their local cache until the counter stabilizes to the next synchronisation point.

Lightweight Collective Memory (Brandenburger et al. [16]) is a proposed framework that claims to achieve fork-linearizability: each honest client that communicates with a TEE (on an untrusted server that might be rolled back) can detect if the server is being inconsistent in their responses to any of the protocol clients (i.e. if they introduce any forks or non-linearity in their responses). Finally, [35, 36, 53], protect hardware memory against active attacks, while [34, 5], protect cryptographic hardware against tampering and Trojan injection attacks, respectively.

## 2 Preliminaries

### 2.1 UC background

Universal Composability (UC), introduced by Canetti [17], is a security framework that enables the security analysis of cryptographic protocols. It supports the setting where multiple instances of the *same*, or *different protocols*, can be executed concurrently. Many extensions and variants of the framework have been proposed over the years; our treatment is based on the recently released Universal Composability with Global Subroutines framework (UCGS) [9] and the 2020 version of UC [17]. We briefly summarise the aspects of UC and UCGS necessary to understand our work .

**Universal Composability** Consider two systems of PPT interactive Turing machine instances  $(\pi, \mathcal{A}, \mathcal{Z})$  and  $(\phi, \mathcal{S}, \mathcal{Z})$ , where  $\mathcal{Z}$  is the initial instance, and  $\pi, \mathcal{A}$  (and respectively  $\phi, \mathcal{S}$ ) have comparable runtime balanced by the inputs of  $\mathcal{Z}$ . We say that the two systems are indistinguishable if  $\mathcal{Z}$  making calls to  $\pi, \mathcal{A}$  (resp.  $\phi, \mathcal{S}$ ) cannot distinguish which system it is located in. The two systems are commonly referred to as the *real* and *ideal world* (respectively).  $\mathcal{Z}$  can make calls to instances within the protocol by assuming the (external) identity of arbitrary instances (as defined by the control function). Depending on the protocol settings, it might be necessary to restrict the external identities available to the environment. A  $\xi$ -identity-bounded environment is limited to assume external identities as specified by  $\xi$ , a polynomial time boolean predicate on the current system configuration.

We now recall a few definitions. Please consult [17, 9] or our full version for the formal definitions of terms such as *balanced*, *respecting*, *exposing*, *compliant*.

**Definition 1 (UC emulation [17]).** *Given two PPT protocols  $\pi, \phi$  and some predicate  $\xi$ , we say that  $\pi$  UC-emulates  $\phi$  with respect to  $\xi$ -identity bound environments (or  $\pi$   $\xi$ -UC-emulates  $\phi$ ) if for any balanced  $\xi$ -identity-bounded environment and any PPT adversary, there exists a PPT simulator  $\mathcal{S}$  such that the systems  $(\phi, \mathcal{S}, \mathcal{Z})$  and  $(\pi, \mathcal{A}, \mathcal{Z})$  are indistinguishable.*

Given a protocol  $\pi$  which UC-emulates a protocol  $\phi$ , and a third protocol  $\rho$ , which calls  $\phi$  as a subroutine, we can construct a protocol where all calls to  $\phi$  are replaced with calls to  $\pi$ , denoted as  $\rho^{\phi \rightarrow \pi}$ .

**Theorem 1 (Universal Composition [17]).** *Given PPT protocols  $\pi, \phi, \rho$  and predicate  $\xi$ , if  $\pi, \phi$  are both subroutine respecting and subroutine exposing,  $\rho$  is  $(\pi, \phi, \xi)$ -compliant and  $\pi$   $\xi$ -UC-emulates  $\phi$ , then protocol  $\rho^{\phi \rightarrow \pi}$  UC-emulates  $\rho$*

By the composition theorem, any protocol that leverages subroutine  $\phi$  in its execution can now be instantiated using protocol  $\pi$ .

**UCGS** As the name suggests, generalised UC (GUC) [18] is an important generalization of the UC model. It accounts for the existence of a shared subroutine  $\gamma$ , such that both  $\rho$  and its subroutine  $\pi$  (regardless of how many instances of  $\pi$  are called by  $\rho$ ) can have  $\gamma$  as a subroutine. The presence of the *global subroutine* allows proving protocols that rely on some powerful functionality that needs to be globally accessible, such as a public key infrastructure (PKI) [19], a global clock [7], or trusted hardware [44].

Unfortunately GUC has inconsistencies and has not been updated from the 2007 to the 2020 version of UC.<sup>5</sup> Universal Composability with Global Subroutines [9] aims to rectify these issues by embedding UC emulation in the presence of a global protocol within the standard UC framework.

To achieve this, a protocol  $\pi$  with access to subroutine  $\gamma$  is replaced by a new structured protocol  $\mu = M[\pi, \gamma]$ , known as *management* protocol;  $\mu$  allows multiplexing a single instance of  $\pi$  and  $\gamma$  into however many are required by  $\rho$ , by transforming the session and party identifiers.  $\mu$  is a subroutine exposing protocol, and is given access to an execution graph directory instance, which tracks existing machines within the protocol, and the list of subroutine calls (implemented as a structured protocol). The execution graph directory can be queried by all instances within the extended session of  $\mu$ , and is used to redirect the outputs of  $\pi$  and  $\gamma$  to the correct machine.

Below we revisit the UC emulation with global subroutines definition from [9].

**Definition 2 (UC emulation with global subroutines [9]).** *Given protocols  $\pi, \phi$ , and  $\gamma$ ,  $\pi$   $\xi$ -UC emulates  $\phi$  in the presence of  $\gamma$  if  $M[\pi, \gamma]$   $\xi$ -UC emulates  $M[\phi, \gamma]$*

Now we state the main UCGS theorem.

**Theorem 2 (Universal Composition with Global subroutines [9]).** *Given subroutine-exposing protocols  $\pi, \phi, \rho$ , and  $\gamma$ , if  $\gamma$  is a  $\phi$ -regular setup and subroutine respecting,  $\phi, \pi$  are  $\gamma$ -subroutine respecting,  $\rho$  is  $(\pi, \phi, \xi)$ -compliant and  $(\pi, M[x, \gamma], \xi)$ -compliant for  $x \in \{\phi, \pi\}$ , then if  $\pi$   $\xi$ -UC-emulates  $\phi$  in the presence of  $\gamma$ , the protocol  $\rho^{\phi \rightarrow \pi}$  UC-emulates  $\rho$  in the presence of  $\gamma$ .*

## 2.2 The $G_{\text{att}}$ functionality

We now reproduce the  $G_{\text{att}}$  global functionality defined in the PST model [44]. The functionality is parametrised with a signature scheme and a registry to capture all platforms with a TEE. The below functionality diverges from the original one in that we let  $\text{vk}$  be a global variable, accessible by enclave programs as  $G_{\text{att}}.\text{vk}$ . This allows us to use  $G_{\text{att}}$  for protocols where the enclave program does not trust the caller to its procedures to pass genuine inputs, making it necessary to conduct the verification of attestation from within the enclave.

<sup>5</sup> In a nutshell the inconsistency arises from a discrepancy in the proof that emulation for a single-challenge session version, called EUC (used to prove protocols secure), implies UC-emulation for the multi-challenge GUC notion (used to prove the composition theorem).

**Functionality  $G_{\text{att}}[\Sigma, \text{reg}, \lambda]$**

State variables	Description
vk	Master verification key, available to enclave programs
msk	Master secret key, protected by the hardware
$\mathcal{T} \leftarrow \emptyset$	Table for installed programs

*On message INITIALIZE from a party P:*

**let** (spk, ssk)  $\leftarrow \Sigma.\text{Gen}(1^\lambda)$ , vk  $\leftarrow$  spk, msk  $\leftarrow$  ssk

*On message GETPK from a party P:*

**return** vk

*On message (INSTALL, idx, prog) from a party P where P.pid  $\in$  reg:*

**if** P is honest **then**

assert idx = P.sid

generate nonce eid  $\in \{0, 1\}^\lambda$ , **store**  $\mathcal{T}[\text{eid}, P] = (\text{idx}, \text{prog}, \emptyset)$

**send** eid to P

*On message (RESUME, eid, input) from a party P where P.pid  $\in$  reg:*

**let** (idx, prog, mem)  $\leftarrow \mathcal{T}[\text{eid}, P]$ , abort if not found

**let** (output, mem')  $\leftarrow$  prog(input, mem) , **store**  $\mathcal{T}[\text{eid}, P] = (\text{idx}, \text{prog}, \text{mem}')$

**let**  $\sigma \leftarrow \Sigma.\text{Sign}(\text{msk}, (\text{idx}, \text{eid}, \text{prog}, \text{output}))$  and **send** (output,  $\sigma$ ) **to** P

The  $G_{\text{att}}$  functionality is a generalisation over other TEE formalisations, such as the one in [28], which tries to closely model some SGX implementation details. For instance, their hardware primitive distinguishes between local and remote attestation by exposing two sets of functions to produce and verify *reports* (for local attestation) and *quotes* (for remote attestation). Both data structure include enclave metadata, a tag that can uniquely identify the running program, input and output to the computation and some authentication primitive based on the former (MAC for local reports, signature for remote quotes). The  $G_{\text{att}}$  primitive, intended as an abstraction over different vendor implementations, removes much of this detail: both local and remote attestation consist in verifying the output of a *resume* call to some enclave through a public verification key, available both to machines with and without enclave capabilities. The output of computations is similarly the (anonymous) id of the enclave, the UC session id, some unique encode for the code computed by the enclave (which could be its source code, or its hash), and the output of the computation. Unlike in the Iron model, input does not have to be included in the attested return value, but if security requires parties to verify input, the function can return it as part of its output. On enclave installation, its memory contents are initialised by the specification of its code; this initial memory state is represented by symbol  $\emptyset$ .

### 3 Functional encryption for stateful and randomized functionalities

In this section we define the ideal functionality of functional encryption for *stateful* and *randomized functionalities* (FESR).

FESR *syntax*.

- (*Setup*): given security parameter  $1^\lambda$  as input,  $\text{KeyGen}$  outputs master key-pair  $\text{mpk}, \text{msk}$
- (*Key generation*):  $\text{Setup}$  takes  $\text{msk}, F \in \mathcal{F}$  and returns functional key  $\text{sk}_F$
- (*Encryption*): given string  $x \in \mathcal{X}$  and  $\text{mpk}$ ,  $\text{Enc}$  returns ciphertext  $\text{ct}$  or an error
- (*Decryption*): on evaluation over some ciphertext  $\text{ct}$  and functional key  $\text{sk}_F$ ,  $\text{Dec}$  returns  $y \in \mathcal{Y}$

While the above definition matches with that of classical functional encryption, we inject non-determinism and statefulness (respectively) by adding two additional inputs to functions in the allowed function class

$$F : \mathcal{X} \times \mathcal{S} \times \mathcal{R} \rightarrow \mathcal{Y} \times \mathcal{S}$$

where  $\mathcal{S} = \{0, 1\}^{s(\lambda)}$ ,  $\mathcal{R} = \{0, 1\}^{r(\lambda)}$  for polynomials  $s(\cdot)$  and  $r(\cdot)$ .

#### 3.1 Properties of FESR

Matt and Maurer [41] shows that the notion of functional encryption is equivalent, up to assumed resources, to that of an access control (AC) repository, where some parties  $\mathbf{A}$  are allowed to upload data, and other parties  $\mathbf{B}$  are allowed to retrieve some function on that data, if they have received authorisation (granted by a party  $\mathbf{C}$ ). A party  $\mathbf{B}$  does not learn anything else about the stored data, besides the function they are authorised to compute (and length leakage  $F_0$ ).

To allow stateful and randomized functions, we extend the function class with support for private state and randomness as above. Whenever  $\mathbf{B}$  accesses a function on the data from the repository, the repository draws fresh randomness, evaluates the function on the old state. The function updates the state and evaluates to a value. Intuitively, this ideal world AC repository models both confidentiality and correctness:

**Confidentiality** Confidentiality holds as  $\mathbf{B}$  does not learn anything about the data besides the evaluations of these stateful randomized functions.

**Correctness** A stateful functionality defines a stateful automaton, a set of states  $\mathcal{S}$ , the initial state  $\emptyset \in \mathcal{S}$ , a probabilistic transition function  $\delta : \mathcal{X} \times \mathcal{S} \rightarrow \mathcal{Y} \times \mathcal{S}$ . For every transition, a new input is sampled from  $\mathcal{R}$  and given to  $F$  along with the input, to determine the next state. The transition function determines,

for a given input and the current state, the probability  $\Pr_\delta$  that the automaton will find itself in a certain next state, as well as an output value. Correctness requires that all consecutive outputs must always be justified by some input and a state reachable via  $\delta$  from  $\emptyset$ .

Correctness holds for the ideal world AC repository as B can make exactly those state transitions by accessing a function on the data from the repository.

### 3.2 UC functionality

Our treatment considers the existence of several parties of type A (encryptors), B (decryptors), and a singular trusted authority C. The latter is allowed to run the `KeyGen`, `Setup` algorithms; parties of type A run `Enc`, and those of type B run `Dec`. The set of all decryptors (resp. encryptors) is denoted by **B** (resp. **A**). When the functionality receives a message from such a party, their UC extended id is used to distinguish who the sender is and store or retrieve the appropriate data. For simplicity, in our ideal functionality we refer to all parties by their type, with the implied assumption that it might refer to multiple distinct UC parties. For the sake of conciseness, we also omit including the `sid` parameter as an argument to every message.

The functionality reproduces the four algorithms that comprise functional encryption. During `KeyGen`, a record  $\mathcal{P}$  is initialised for all  $t$  instances of B, to record the authorised functions for each instance, and its state. The `Setup` call marks a certain B as authorised to decrypt function F, and initialises its state to  $\emptyset$ . The `Enc` call allows a party A, B, to provide some input  $x$ , and receive a unique identifying handle  $h$ . This handle can then be provided, along with some F, to a decryption party to obtain an evaluation of F on the message stored therein. Performing the computation will also result in updating the state stored in  $\mathcal{P}$ .

#### Functionality FESR[ $\text{sid}, F, A, B, C$ ]

The functionality is parameterized by the randomized function class  $\mathbf{F}$  such that for each  $F \in \mathbf{F} : \mathcal{X} \times \mathcal{S} \times \mathcal{R} \rightarrow \mathcal{Y} \times \mathcal{S}$ , over state space  $\mathcal{S}$  and randomness space  $\mathcal{R}$ , and by three distinct types of party identities  $A, B, C$  interacting with the functionality via dummy parties (that identify a particular role). For each decryptor/function pair, a state value is recorded.

State variables	Description
$F_0$	Leakage function returning the length of the message
$\text{setup}[\cdot] \leftarrow \text{false}$	Table recording which parties were initialized.
$\mathcal{M}[\cdot] \leftarrow \perp$	Table storing the plaintext for each message handler
$\mathcal{P}[\cdot] \leftarrow \perp$	Table of authorized functions and their states for all decryption parties

*On message (SETUP, P) from party C, for  $P \in \{A, B\}$ :*

**setup**[ $P$ ]  $\leftarrow$  true  
**send** (SETUP, P) to A

*On message (SETUP, P) from A, for  $P \in \{A, B\}$ :*

```

setup[ $P$ ]  $\leftarrow$  true
 $\mathcal{P}[P, F_0] \leftarrow \emptyset$ 
send SETUP to  $P$ 
On message (ENCRYPT,  $x$ ) from party  $P \in \{A, B\}$ :
if setup[ $P$ ] = true  $\wedge x \in \mathcal{X}$  then
  compute  $h \leftarrow \text{getHandle}$ 
   $\mathcal{M}[h] \leftarrow x$ 
  send (ENCRYPTED,  $h$ ) to  $P$ 
On message (KEYGEN,  $F, B$ ) from party  $C$ :
if  $F \in F^+ \wedge \text{setup}[B] = \text{true}$  then
  send (KEYGEN,  $F, B$ ) to  $A$  and receive ACK
   $\mathcal{P}[B, F] \leftarrow \emptyset$ 
  send (ASSIGNED,  $F$ ) to  $B$ 
On message (DECRYPT,  $h, F$ ) from party  $B$ :
 $x \leftarrow \mathcal{M}[h]$ 
if  $C$  is honest then
  if  $\mathcal{P}[B, F] \neq \perp \wedge x \in \mathcal{X}$  then
     $r \leftarrow \mathcal{R}$ 
     $s \leftarrow \mathcal{P}[B, F]$ 
     $(y, s') \leftarrow F(x, s, r)$ 
     $\mathcal{P}[B, F] \leftarrow s'$ 
    return (DECRYPTED,  $y$ )
  else
    send (DECRYPT,  $h, F, x$ ) to  $A$  and receive (DECRYPTED,  $y$ )
    return (DECRYPTED,  $y$ )

```

The functionality is defined for possible corruptions of parties in  $\mathbf{B}$ ,  $\mathbf{A}$ . If  $C$  is corrupted, we can no longer guarantee the evaluation to be correct, since  $C$  might authorize the adversary to compute any function in  $F$ . In this scenario, we allow the adversary to learn the original message value  $x$  and to provide an arbitrary evaluation  $y$ .

Note that, our definition is along the lines of [10, 41], however, as opposed to [10], in which  $A$  and/or  $C$  might also get corrupted, in this work we primarily focus on the security guarantees provided by FE, which is confidentiality of the encrypted message against malicious decryptors,  $B$ . Yet, it provides security against malicious encryptors,  $A$ , thus it satisfies *input consistency*, which was originally introduced by [10]. In addition, our definition is the first one that captures stateful and randomized functionalities, where the latter refers to the standard notion of randomized functionalities in which each invocation of the function uses independent randomness. Therefore, our protocol achieves a stronger notion of randomized FE than [2, 30, 37], which require a new functional key for each invocation of the function, i.e., decryptions with the same functional key always return the same output.

Both correctness and confidentiality clearly hold for the ideal functionality by inspection of the 4 lines  $r \leftarrow \mathcal{R}$ ,  $s \leftarrow \mathcal{P}[B, F]$ ,  $(y, s') \leftarrow F(x, s, r)$ , and  $\mathcal{P}[B, F] \leftarrow s'$ .

## 4 A UC-formulation of Steel

In this section we present **Steel** in the UCGS setting. As we already state above, our treatment involves three roles: the *key generation* party **C**, the *decryption* parties **B**, and the *encryption* parties **A**. Among them, only the encryptor does not need to have access to an enclave. Like the FESR functionality, the protocol fulfills confidentiality and correctness in the face of an adversarial **B**. We do not give any guarantees of security for corrupted **A**, **C**; although we remark informally that, as long as its enclave is secure, a corrupted **C** has little chances of learning the secret key. Besides the evaluation of any function in **F** it authorises itself to decrypt, it can also fake or extract from proofs of ciphertext validity  $\pi$  by authorizing a fake reference string  $\text{crs}$ . Before formally presenting our protocol we highlight important assumptions and conventions:

- For simplicity of presentation, we assume a single instance each for **A**, **B**
- all communication between parties  $(\alpha, \beta)$  occurs over secure channels  $\mathcal{SC}_\alpha^\beta, \mathcal{SC}_\beta^\alpha$
- Functional keys are (attestation) signatures by an enclave  $\text{prog}_{\text{KME}}$  on input  $(\text{keygen}, F)$  for some function  $F$ ; it is easy, given a list of keys, to retrieve the one which authorises decryptions of  $F$
- keyword **fetch** retrieves a stored variable from memory and aborts if the value is not found
- on keyword **assert**, the program checks that an expression is true, and proceeds to the next line, aborting otherwise
- all variables within an enclave are erased after use, unless saved to encrypted memory through the **store** keyword

Protocol **Steel** is parametrised by a function family  $F : \mathcal{X} \times \mathcal{S} \times \mathcal{R} \rightarrow \mathcal{Y} \times \mathcal{S}$ , UC parties **A**, **B**, **C**, a CCA secure public key encryption scheme **PKE**, a EU-CMA secure signature scheme  $\Sigma$ , a Robust non-interactive zero-knowledge scheme **N**, and security parameter  $\lambda$ .

<b>Protocol Steel</b> [ $F, A, B, C, \text{PKE}, \Sigma, N, \lambda$ ]	
State variables	Description
$\text{mpk} \leftarrow \perp$	Local copy of master public key for participants
$\text{prog}_{\{\text{KME}, \text{DE}, \text{FE}\}} \leftarrow \dots$	Source code of enclaves as defined below
$\mathcal{K}[\cdot] \leftarrow \emptyset$	Table of function keys at <b>B</b>
<b>Key Generation Authority C:</b>	
<i>On message</i> ( $\text{SETUP}, P$ ):	
<b>if</b> $\text{mpk} = \perp$ <b>then</b> $\text{eid}_{\text{KME}} \leftarrow G_{\text{att}}.\text{install}(\text{C.sid}, \text{prog}_{\text{KME}})$ <b>send</b> GET <b>to</b> $\text{CRS}$ and <b>receive</b> ( $\text{CRS}, \text{crs}$ ) $(\text{mpk}, \cdot) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{KME}}, (\text{init}, \text{crs}, \text{C.sid}))$	
<b>if</b> $P = \text{A}$ <b>then</b> <b>send</b> ( $\text{SETUP}, \text{mpk}$ ) <b>to</b> $\mathcal{S}_A$	
<b>else if</b> $P = \text{B}$ <b>then</b> <b>send</b> ( $\text{SETUP}, \text{mpk}, \text{eid}_{\text{KME}}$ ) <b>to</b> $\mathcal{S}_B$ and <b>receive</b> ( $\text{PROVISION}, \sigma, \text{eid}_{\text{DE}}, \text{pk}_{KD}$ ) $(\text{ct}_{\text{key}}, \sigma_{\text{sk}}) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{KME}}, (\text{provision}, (\sigma, \text{eid}_{\text{DE}}, \text{pk}_{KD}, \text{eid}_{\text{KME}})))$	

**send** (PROVISION,  $\text{ct}_{\text{key}}, \sigma_{\text{sk}}$ ) **to**  $\mathcal{SC}_B$

*On message* (KEYGEN, F, B):  
**assert**  $F \in \mathcal{F} \wedge \text{mpk} \neq \perp$   
 $((\text{keygen}, F), \sigma) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{KME}}, (\text{keygen}, F))$   
 $\text{sk}_F \leftarrow \sigma$ ; **send** (KEYGEN, (F,  $\text{sk}_F$ )) **to**  $\mathcal{SC}_B$

**Encryption Party A:**

*On message* (SETUP, mpk) *from*  $\mathcal{SC}^C$ :  
**send** GET **to**  $\mathcal{CRS}$  and **receive** (CRS, crs)  
**store** mpk, crs; **return** SETUP

*On message* (ENCRYPT, m):  
**assert**  $\text{mpk} \neq \perp \wedge m \in \mathcal{X}$   
 $\text{ct} \xleftarrow{r} \text{PKE}.\text{Enc}(\text{mpk}, m)$   
 $\pi \leftarrow \mathcal{P}((\text{mpk}, \text{ct}), (m, r), \text{crs}), \text{ct}_{\text{msg}} \leftarrow (\text{ct}, \pi)$   
**send** (WRITE,  $\text{ct}_{\text{msg}}$ ) **to**  $\mathcal{R}\mathcal{E}\mathcal{P}$  and **receive** h  
**return** (ENCRYPTED, h)

**Decryption Party B:**

*On message* (SETUP, mpk,  $\text{eid}_{\text{KME}}$ ) *from*  $\mathcal{SC}^C$ :  
**store** mpk;  $\text{eid}_{\text{DE}} \leftarrow G_{\text{att}}.\text{install}(\text{B}.\text{sid}, \text{prog}_{\text{DE}})$   
**send** GET **to**  $\mathcal{CRS}$  and **receive** (CRS, crs)  
 $((\text{pk}_{\text{KD}}, \cdot, \cdot), \sigma) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{DE}}, \text{init-setup}, \text{eid}_{\text{KME}}, \text{crs}, \text{B}.\text{sid})$   
**send** (PROVISION,  $\sigma, \text{eid}_{\text{DE}}, \text{pk}_{\text{KD}}$ ) **to**  $\mathcal{SC}_C$  and **receive** (PROVISION,  $\text{ct}_{\text{key}}, \sigma_{\text{KME}}$ )  
 $G_{\text{att}}.\text{resume}(\text{eid}_{\text{DE}}, (\text{complete-setup}, \text{ct}_{\text{key}}, \sigma_{\text{KME}}))$   
**return** SETUP

*On message* (KEYGEN, (F,  $\text{sk}_F$ )) *from*  $\mathcal{SC}^C$ :  
 $\text{eid}_F \leftarrow G_{\text{att}}.\text{install}(\text{B}.\text{sid}, \text{prog}_{\text{FE}}[\text{F}])$   
 $(\text{pk}_{\text{FD}}, \sigma_F) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_F, (\text{init}, \text{mpk}, \text{B}.\text{sid}))$   
 $\mathcal{K}[\text{F}] \leftarrow (\sigma_F, \text{eid}_F, \text{pk}_{\text{FD}}, \text{sk}_F)$   
**return** (ASSIGNED, F)

*On message* (DECRYPT, F, h):  
**assert**  $\mathcal{K}[\text{F}] \neq \perp$   
**send** (READ, h) **to**  $\mathcal{R}\mathcal{E}\mathcal{P}$  and **receive**  $\text{ct}_{\text{msg}}$   
 $(\sigma_F, \text{eid}_F, \text{pk}_{\text{FD}}, \text{sk}_F) \leftarrow \mathcal{K}[\text{F}]$   
 $((\text{ct}_{\text{key}}, \text{crs}), \sigma_{\text{DE}}) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{DE}}, (\text{provision}, \sigma_F, \text{eid}_F, \text{pk}_{\text{FD}}, \text{sk}_F, \text{F}))$   
 $((\text{computed}, y), \cdot) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_F, (\text{run}, \sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_{\text{msg}}, \text{crs}, \perp))$   
**return** (DECRYPTED, y)

```

progKME:
on input (init, crs, idx):
  assert pk =  $\perp$ ; (pk, sk)  $\leftarrow$  PKE.PGen()
  store sk, crs, idx; return pk
on input (provision, ( $\sigma_{DE}$ , eidDE, pkKD, eidKME)):
  vkatt  $\leftarrow$  Gatt.vk; fetch crs, idx
  assert  $\Sigma$ .Vrfy(vkatt, (idx, eidDE, progDE, (pkKD, eidKME, crs),  $\sigma_{DE}$ )
  ctkey  $\leftarrow$  PKE.Enc(pkKD, sk)
  return ctkey
on input (keygen, F):
  return (keygen, F)

```

```

progDE:
on input (init-setup, eidKME, crs, idx):
  assert pkKD  $\neq$   $\perp$ 
  (pkKD, skKD)  $\leftarrow$  PKE.Gen()
  store skKD, eidKME, crs, idx
  return pkKD, eidKME, crs

```

```

on input (complete-setup, ctkey,  $\sigma_{KME}$ ):
  vkatt  $\leftarrow$  Gatt.vk
  fetch eidKME, skKD, idx
  m  $\leftarrow$  (idx, eidKME, progKME, ctkey)
  assert  $\Sigma$ .Vrfy(vkatt, m,  $\sigma_{KME}$ )
  sk  $\leftarrow$  PKE.Dec(skKD, ctkey)
  store sk, vkatt
on input (provision,  $\sigma$ , eid, pkFD, skF, F):
  fetch eidKME, vkatt, sk, idx
  m1  $\leftarrow$  (idx, eidKME, progKME, (keygen, F))
  m2  $\leftarrow$  (idx, eid, progFE[F], pkFD)
  assert  $\Sigma$ .Vrfy(vkatt, m1, skF) and
   $\Sigma$ .Vrfy(vkatt, m2,  $\sigma$ )
  return PKE.Enc(pkFD, sk), crs

```

```

progFE[F]:
on input (init, mpk, idx):
  assert pkFD =  $\perp$ 
  (pkFD, skFD) = PKE.Gen( $1^\lambda$ )
  mem  $\leftarrow$   $\emptyset$ ; store skFD, mem, mpk, idx
  return pkFD
on input (run,  $\sigma_{DE}$ , eidDE, ctkey, ctmsg, crs, y'):
  if y'  $\neq$   $\perp$ 
    return (computed, y')
  vkatt  $\leftarrow$  Gatt.vk; (ct,  $\pi$ )  $\leftarrow$  ctmsg
  fetch skFD, mem, mpk, idx
  m  $\leftarrow$  (idx, eidDE, progDE, ctkey, crs)
  assert  $\Sigma$ .Vrfy(vkatt, m,  $\sigma_{DE}$ )
  sk = PKE.Dec(skFD, ctkey)
  assert  $N$ . $\mathcal{V}$ ((mpk, ct),  $\pi$ , crs)
  x = PKE.Dec(sk, ct)
  out, mem' = F(x, mem)
  store mem  $\leftarrow$  mem'
  return (computed, out)

```

As we mention in the Introduction, our modeling considers a “backdoor” in the `progFE.run` subroutine, such that, if the last argument is set, the subroutine just returns the value of that argument, along with a label declaring computation. The addition of the label *computed* is necessary, otherwise the backdoor would allow producing an attested value for the public key generated in subroutine `progFE.init`.

As a further addition we strengthen the encryption scheme with a plaintext proof of knowledge (PPoK). For public key `pk`, ciphertext `ct`, plaintext `m`, ciphertext randomness `r`, the relation  $R = \{(pk, ct), (m, r) \mid ct = \text{PKE.Enc}(mpk, m; r)\}$  defines the language  $L_R$  of correctly computed ciphertexts. As a chosen-plaintext secure PKE scheme becomes CCA secure when extended with a simulation-extractable PPoK this is a natural strengthening of the CCA security requirement of Iron. However, it enables the simulator to extract valid plaintexts from

all adversarial ciphertexts. In our security proof the simulator will submit these plaintexts to FESR on behalf of the corrupt  $B$  to keep the decryption states of the real and ideal world synchronized.

## 5 UC-security of Steel

We now prove the security of Steel in the UCGS framework. To make the PST model compatible with the UCGS model, we first define the identity bound  $\xi$ .

*The identity bound  $\xi$  on the environment.* Our restrictions are similar to [44], namely we assume that the environment can access  $G_{\text{att}}$  in the following ways: (1) Acting as a corrupt party, and (2) acting as an honest party but only for non-challenge protocol instances.

We now prove our main theorem.

**Theorem 3.** *Steel (Section 4) UC-realizes the FESR functionality (Section 3) in the presence of the global functionality  $G_{\text{att}}$  and local functionalities  $\mathcal{CRS}$ ,  $\mathcal{R}\mathcal{E}\mathcal{P}$ ,  $\mathcal{S}\mathcal{C}$ , with respect to the identity bound  $\xi$  defined above.*

We present a simulator algorithm such that, for all probabilistic adversaries running in polynomial time with the ability of corrupting  $B$ . Following [41], our proof considers static corruption of a single party  $B$ , we did, however, not encounter any road-blocks to adaptive corruption of multiple decryptors besides increased proof notational complexity. The environment is unable to distinguish between an execution of the Steel protocol in the real world, and the protocol consisting of  $\mathcal{S}_{\text{FESR}}$ , dummy parties  $A$ ,  $C$  and ideal functionality FESR. Both protocols have access to the shared global subroutines of  $G_{\text{att}}$ . While hybrid functionalities  $\mathcal{R}\mathcal{E}\mathcal{P}$ ,  $\mathcal{S}\mathcal{C}$ ,  $\mathcal{C}\mathcal{R}\mathcal{S}$  (for their definition, see the full version) are only available in the real world and need to be reproduced by the simulator, we use  $\mathcal{S}\mathcal{C}$  in the simulator to denote simulated channels, either between the simulator and corrupted parties (for corrupt parties), or between the simulator and itself (for honest parties).

Given protocols Steel, FESR, and  $G_{\text{att}}$ , Steel  $\xi$ -UC emulates FESR in the presence of  $G_{\text{att}}$  if  $M[\text{Steel}, G_{\text{att}}]$   $\xi$ -UC emulates  $M[\text{FESR}, G_{\text{att}}]$  (see Definition 2). We focus our exposition on the messages exchanged between the environment and the machine instances executing Steel, FESR, and  $G_{\text{att}}$ , since the machine  $M$  is simply routing messages; i.e., whenever  $\mathcal{Z}$  wants to interact with the protocol,  $M$  simply forwards the message to the corresponding party; the same holds for  $G_{\text{att}}$ .

The simulator operates in the ideal world, where we have the environment  $\mathcal{Z}$  sending message to dummy protocol parties which forward their inputs to the ideal functionality FESR.  $\mathcal{S}_{\text{FESR}}$  is activated either by an incoming message from a corrupted party or the adversary, or when FESR sends a message to the ideal world adversary. As  $\mathcal{A}$  is a dummy adversary which normally forwards all queries between the corrupt party and the environment,  $\mathcal{S}_{\text{FESR}}$  gets to see all messages  $\mathcal{Z}$  sends to  $\mathcal{A}$ . The simulator is allowed to send messages to the FESR and  $G_{\text{att}}$

functionalities impersonating corrupt parties. In the current setting, the only party that can be corrupted such that FESR still gives non trivial guarantees is party B. Thus, whenever the real world adversary or the ideal world simulator call  $G_{\text{att}}.\text{install}$  and  $G_{\text{att}}.\text{resume}$  for the challenge protocol instance, they must do so using an extended identity of B.

<b>Simulator <math>\mathcal{S}_{\text{FESR}}[\text{PKE}, \Sigma, \mathcal{N}, \lambda, \mathbb{F}]</math></b>	
State variables	Description
$\mathcal{H}[\cdot] \leftarrow \emptyset$	Table of ciphertext and handles in public repository
$\mathcal{K} \leftarrow []$	List of $\text{prog}_{\mathbb{F}}[\mathbb{F}]$ enclaves and their $\text{eid}_{\mathbb{F}}$
$\mathcal{G} \leftarrow \{\}$	Collects all messages sent to $G_{\text{att}}$ and its response
$\mathcal{B} \leftarrow \{\}$	Collects all messages signed by $G_{\text{att}}$
$(\text{crs}, \tau) \leftarrow \mathcal{N}.\mathcal{S}_1$	Simulated reference string and trapdoor
<b>For Key Generation Authority C:</b>	
<i>On message (SETUP, P) from FESR:</i>	
<b>if</b> $\text{mpk} = \perp$ <b>then</b>	
$\text{eid}_{\text{KME}} \leftarrow G_{\text{att}}.\text{install}(\text{C.sid}, \text{prog}_{\text{KME}})$	
$(\text{mpk}, \cdot) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{KME}}, \text{init})$	
<b>if</b> $P = \text{A}$ <b>then</b>	
<b>send</b> (SETUP, mpk) <b>to</b> $\mathcal{S}_{\text{CA}}$	
<b>else if</b> $P = \text{B}$ <b>then</b>	
<b>send</b> (SETUP, mpk, $\text{eid}_{\text{KME}}$ ) <b>to</b> $\mathcal{S}_{\text{CB}}$ <b>and receive</b> (PROVISION, $\sigma$ , $\text{eid}_{\text{DE}}$ , $\text{pk}_{\text{KD}}$ )	
<b>assert</b> $(\text{C.sid}, \text{eid}_{\text{DE}}, \text{prog}_{\text{DE}}, \text{pk}_{\text{KD}}) \in \mathcal{B}[\sigma]$	
$(\text{ct}_{\text{key}}, \sigma_{\text{sk}}) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{KME}}, (\text{provision}, (\sigma, \text{eid}_{\text{DE}}, \text{pk}_{\text{KD}}, \text{eid}_{\text{KME}}, \text{crs})))$	
<b>send</b> (PROVISION, $\text{ct}_{\text{key}}, \sigma_{\text{sk}}$ ) <b>to</b> $\mathcal{S}_{\text{CB}}$	
<i>On message (KEYGEN, F, B) from FESR:</i>	
<b>assert</b> $\mathbb{F} \in \mathbb{F} \wedge \text{mpk} \neq \perp$	
$((\text{keygen}, \mathbb{F}), \sigma) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{KME}}, (\text{keygen}, \mathbb{F}))$	
$\text{sk}_{\mathbb{F}} \leftarrow \sigma$	
<b>send</b> (KEYGEN, (F, $\text{sk}_{\mathbb{F}}$ )) <b>to</b> $\mathcal{S}_{\text{CB}}$	
<b>For Decryption Party B:</b>	
<i>On message GET from party B to CRS:</i>	
<b>send</b> (CRS, crs) <b>to</b> B	
<i>On message (READ, h) from party B to <math>\mathcal{R}\mathcal{E}\mathcal{P}</math>:</i>	
<b>send</b> (DECRYPT, $\mathbb{F}_0$ , h) <b>to</b> FESR on behalf of B <b>and receive</b> $ m $	
<b>assert</b> $ m  \neq \perp$	
$\text{ct} \leftarrow \text{PKE}.\text{Enc}(\text{mpk}, 0^{ m })$	
$\pi \leftarrow \mathcal{N}.\mathcal{S}_2(\text{crs}, \tau, (\text{mpk}, \text{ct}))$	
$\text{ct}_{\text{msg}} \leftarrow (\text{ct}, \pi)$ ; $\mathcal{H}[\text{ct}_{\text{msg}}] \leftarrow h$	
<b>send</b> (READ, $\text{ct}_{\text{msg}}$ ) <b>to</b> B	
<i>On message (INSTALL, idx, prog) from party B to <math>G_{\text{att}}</math>:</i>	
$\text{eid} \leftarrow G_{\text{att}}.\text{install}(\text{idx}, \text{prog})$	
$\mathcal{G}[\text{eid}].\text{install} \leftarrow (\text{idx}, \text{prog})$	
// $\mathcal{G}[\text{eid}].\text{install}[1]$ is the program's code	

**forward eid to B**

*On message (RESUME, eid, input) from party B to  $G_{\text{att}}$ :*

```

// The  $G_{\text{att}}$  registry does not allow B to access  $\text{eid}_{\text{KME}}$  in real world
assert  $\mathcal{G}[\text{eid}] \neq \perp \wedge \text{eid} \neq \text{eid}_{\text{KME}}$ 
if  $\mathcal{G}[\text{eid}].\text{install}[1] \neq \text{prog}_{\text{FE}}[\cdot] \vee \text{input}[-1] \neq \perp$  then
  (output,  $\sigma$ )  $\leftarrow G_{\text{att}}.\text{resume}(\text{eid}, \text{input})$ 
   $\mathcal{G}[\text{eid}].\text{resume} \leftarrow \mathcal{G}[\text{eid}].\text{resume} \parallel (\sigma, \text{input}, \text{output})$ 
   $\mathcal{B}[\sigma] \leftarrow (\mathcal{G}[\text{eid}].\text{install}[0], \text{eid}, \mathcal{G}[\text{eid}].\text{install}[1], \text{output})$ 
  if  $\mathcal{G}[\text{eid}].\text{install}[1] = \text{prog}_{\text{DE}} \wedge \text{input}[0] = \text{provision}$  then
    (provision,  $\sigma$ , eid,  $\text{pk}_{\text{FD}}$ ,  $\text{sk}_{\text{F}}$ , F)  $\leftarrow \text{input}$ 
    fetch  $(\cdot, (\text{init-setup}, \text{eid}_{\text{KME}}, \text{crs}), \cdot) \in \mathcal{G}[\text{eid}].\text{resume}$ 
    assert  $(\text{idx}, \text{eid}_{\text{KME}}, \text{prog}_{\text{KME}}, (\text{keygen}, \text{F})) \in \mathcal{B}[\text{sk}_{\text{F}}]$ 
    assert  $(\text{idx}, \text{eid}_{\text{DE}}, \text{prog}_{\text{DE}}, \text{ct}_{\text{key}}, \text{crs}) \in \mathcal{B}[\sigma_{\text{DE}}]$ 
    forward (output,  $\sigma$ ) to B
  else
     $\text{idx}, \text{prog}_{\text{FE}}[\text{F}] \leftarrow \mathcal{G}[\text{eid}].\text{install}$ 
    (run,  $\sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_{\text{msg}}, \text{crs}, \perp$ )  $\leftarrow \text{input}$ 
    assert  $(\sigma_{\text{F}}, (\text{init}), (\text{pk}_{\text{FD}})) \in \mathcal{G}[\text{eid}].\text{resume}$ 
    assert  $(\text{idx}, \text{eid}, \text{prog}_{\text{FE}}[\text{F}], \text{pk}_{\text{FD}}) \in \mathcal{B}[\sigma_{\text{F}}]$ 
    assert  $(\text{idx}, \text{eid}_{\text{DE}}, \text{prog}_{\text{DE}}, \text{ct}_{\text{key}}, \text{crs}) \in \mathcal{B}[\sigma_{\text{DE}}]$ 
    // If the ciphertext was not computed honestly and saved to  $\mathcal{H}$ 
    if  $\mathcal{H}[\text{ct}_{\text{msg}}] = \perp$  then
      (ct,  $\pi$ )  $\leftarrow \text{ct}_{\text{msg}}$ 
      (m, r)  $\leftarrow \text{N.E}(\tau, (\text{mpk}, \text{ct}), \pi)$ 
      if m =  $\perp$  then send (DECRYPT, F,  $\perp$ ) to B and abort
      send (ENCRYPT, m) to FESR on behalf of B and receive h
       $\mathcal{H}[\text{ct}_{\text{msg}}] \leftarrow \text{h}$ 
      h  $\leftarrow \mathcal{H}[\text{ct}_{\text{msg}}]$ 
      send (DECRYPT, F, h) to FESR on behalf of B and receive y
      ((computed, y),  $\sigma$ )  $\leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{F}}, (\text{run}, \perp, \perp, \perp, \perp, \perp, y))$ 
       $\mathcal{G}[\text{eid}].\text{resume} \leftarrow \mathcal{G}[\text{eid}].\text{resume} \parallel (\sigma, \text{input}, (\text{computed}, y))$ 
       $\mathcal{B}[\sigma] \leftarrow (\mathcal{G}[\text{eid}].\text{install}[0], \text{eid}, \mathcal{G}[\text{eid}].\text{install}[1], (\text{computed}, y))$ 
      forward ((computed, y),  $\sigma$ ) to B

```

**Designing the simulation** The ideal functionality FESR and protocol Steel share the same interface consisting of messages SETUP, KEYGEN, ENCRYPT, DECRYPT. During Steel's SETUP, the protocol generates public parameters when first run, and provisions the encrypted secret key to the enclaves of B. As neither of these operations are executed by the ideal functionality, we need to simulate them, generating and distributing keys outside of party C.

As in Steel, we distribute the public encryption key on behalf of C to any newly registered B and A over secure channels. Once B has received this message, it will try to obtain the (encrypted) decryption key for the global PKE scheme from party C and its provision subroutine of  $\text{prog}_{\text{KME}}$ . Since C is a dummy party in the ideal world, it would not respond to this request, so we let  $\mathcal{S}_{\text{FESR}}$  respond. In Steel key parameters are generated within the key management enclave, and communication of the encrypted secret key to the decryption enclave produces

an attestation signature. Thus, the simulator, which can access  $G_{\text{att}}$  impersonating B, is required to install an enclave. Because of the property of anonymous attestation, the environment cannot distinguish whether the new enclave was installed on B or C. If the environment tries to resume the program running under  $\text{eid}_{\text{KME}}$  through B, this is intercepted and dropped by the simulator.

Before sending the encrypted secret key, the simulator verifies that B’s public key was correctly produced by an attested decryption enclave, and was initialised with the correct parameters. If an honest enclave has been instantiated and we can verify that it uses  $\text{pk}_{\text{KD}}, \text{eid}_{\text{KME}}, \text{crs}$ , we can safely send the encrypted  $\text{sk}$  to the corrupted party as no one can retrieve the decryption key from outside the enclave.

On message  $(\text{KEYGEN}, F, B)$  from the functionality after a call to  $\text{KEYGEN}$ ,  $\mathcal{S}_{\text{FESR}}$  simply produces a functional key by running the appropriate  $\text{prog}_{\text{KME}}$  procedure through  $G_{\text{att}}$ . Similarly, on receiving  $(\text{READ}, h)$  for  $\mathcal{R}\mathcal{E}\mathcal{P}$ ,  $\mathcal{S}_{\text{FESR}}$  produces an encryption of a canonical message (a string of zeros) and simulates the response.

When the request to compute the functional decryption of the corresponding ciphertext is sent to  $\text{prog}_{\text{FE}}[\text{F}]$ , we verify that the party B has adhered to the Steel protocol execution, aborting if any of the required enclave installation or execution steps have been omitted, or if any of the requests were made with dishonest parameters generated outside the enclave execution (we can verify this through the attestation of enclave execution). If the ciphertext was not obtained through a request to  $\mathcal{R}\mathcal{E}\mathcal{P}$ , we use the NIZK extractor to learn the plaintext  $m$  and submit a message  $(\text{ENCRYPT}, m)$  to FESR on behalf of the corrupt B. This guarantees that the state of FESR is in sync with the state of  $\text{prog}_{\text{FE}}[\text{F}]$  in the real world.

If all such checks succeed, and the provided functional key is valid,  $\mathcal{S}_{\text{FESR}}$  fetches the decryption from the ideal functionality. While the Steel protocol ignores the value of the attested execution of  $\text{run}$ , we can expect the adversary to check its result for authenticity. Therefore, it is necessary to pass the result of our decryption  $y$  through the backdoor we constructed in  $\text{prog}_{\text{FE}}[\text{F}]$ . This will produce an authentic attestation signature on  $y$ , which will pass any verification check convincingly (as discussed in the previous section, the backdoor does not otherwise impact the security of the protocol).

The full proof of security is available in the full version; for an overview, refer to Section 1.2.

## 6 Rollback and Forking Attacks

While the Attested Execution functionality modelled by  $G_{\text{att}}$  is a meaningful first step for modeling attested execution, it is easy to argue that it is not realisable (in a UC-emulation sense) by any of the existing Trusted Execution Environment platforms to date. In a follow-up paper, Tramer et al. [52] weaken the original  $G_{\text{att}}$  model to allow complete leakage of the memory state. This is perhaps an excessively strong model, as the use of side channel attacks might only allow

a portion of the memory or randomness to be learned by the adversary. Additionally, there are many other classes of attacks that can not be expressed by this model. We now extend the  $G_{\text{att}}$  functionality to model *rollback* and *forking attacks* against an enclave.

### 6.1 $G_{\text{att}}^{\text{rollback}}$ functionality

Our model of rollback and forking attacks is drawn from the formulation expressed in Matetic et al. [40], but with PST’s improved modelling of attestation, which does not assume perfectly secure authenticated reads/writes between the attester and the enclave.

Matetic et al. model rollback by distinguishing between enclaves and enclave instances. Enclave instances have a distinct memory state, while sharing the same code. As with  $G_{\text{att}}$ , where the outside world has to call subroutines individually, the environment is not allowed to interact directly with a program once it is instantiated, except for pausing, resuming, or deleting enclave instances. Additionally, their model provides functions to store encrypted memory outside the enclave (*Seal*) and load memory back (*Unseal*).

In a typical rollback attack, an attacker crashes an enclave, erasing its volatile memory. As the enclave instance is restarted, it attempts to restart from the current state snapshot. By replacing this with a stale snapshot, the attacker is able to rewind the enclave state.

In a forking attack an attacker manages to run two instances of the same enclave concurrently, such that, once the state of one instance is changed by an external operation, querying the other instance will result in an outdated state. This relies on both enclaves producing signature that at the minimum attest the same program. On a system where attestation uniquely identifies each copy of the enclave, a forking attack can still be launched by an attacker conducting multiple rollback attacks and feeding different stale snapshots to a single enclave copy [16].

Our new functionality  $G_{\text{att}}^{\text{rollback}}$  employs this idea to model the effect of both rollback and forking attacks. We replace the internal `mem` variable of  $G_{\text{att}}$  with a tree data structure. The honest caller to the functionality will always continue execution from the memory state of an existing leaf of the tree while an adversary can specify an arbitrary node of the tree (through a unique node identifier), to which the state of the enclave gets reset. The output `mem'` will then be appended as a new child branch to the tree. To model a rollback attack, the adversary specifies the parent node for the next call to `resume` (or any ancestor node to execute a second rollback). To model a forking attack, the adversary can interactively choose nodes in different branches of the tree. The functionality is parametrised with a signature scheme and a registry to capture all platforms with a TEE, like in the original formulation.

**Functionality  $G_{\text{att}}[\Sigma, \text{reg}, \lambda]$**

State variables	Description
$vk$	Master verification key, available to enclave programs
$msk$	Master secret key, protected by the hardware
$\mathcal{T} \leftarrow \emptyset$	Table for installed programs

On message INITIALIZE from a party  $P$ :

```

let (spk, ssk)  $\leftarrow$   $\Sigma$ .Gen( $1^\lambda$ ), vk  $\leftarrow$  spk, msk  $\leftarrow$  ssk

```

On message GETPK from a party  $P$ :

```

return vk

```

On message (INSTALL, idx, prog) from a party  $P$  where  $P.pid \in \text{reg}$ :

```

if  $P$  is honest then
  assert idx =  $P.sid$ 
  generate nonce eid  $\in \{0, 1\}^\lambda$ , store  $\mathcal{T}[\text{eid}, P] = (\text{idx}, \text{prog}, \text{root}, \text{Tree}(\emptyset))$ 
  send eid to  $P$ 

```

On message (RESUME, eid, input, node) from a party  $P$  where  $P.pid \in \text{reg}$ :

```

let (idx, prog, lastnode, tree)  $\leftarrow$   $\mathcal{T}[\text{eid}, P]$ , abort if not found
if  $P$  is honest then
  let node  $\leftarrow$  lastnode
  let mem  $\leftarrow$  access(tree, node)
  let (output, mem')  $\leftarrow$  prog(input, mem)
  let tree', child  $\leftarrow$  insertChild(tree, node, mem')
  let update  $\mathcal{T}[\text{eid}, P] = (\text{idx}, \text{prog}, \text{child}, \text{tree}')$ 
  let  $\sigma \leftarrow \Sigma$ .Sign(msk, (idx, eid, prog, output)) and send (output,  $\sigma$ ) to  $P$ 

```

The proposed rollback model is perhaps somewhat reductive, as it only allows “discrete” rollback operations, where memory states are quantized by program subroutines. It is conceivable that real world attackers would have a finer-grained rollback model, where they can interrupt the subroutine’s execution, and resume from an arbitrary instruction.

**Attack on stateful functional encryption** Although our protocol uses probabilistic primitives, we deem the generic reset attack presented in [55] unrealistic for TEE platforms such as SGX, where an enclave is allowed direct access to a hardware-based source of randomness [6].

On the other hand, it is easy to find a protocol-specific rollback attack on Steel. While  $F$ ’s state remains secret to a corrupt  $B$  interacting with  $G_{\text{att}}^{\text{rollback}}$  (the memory is still sealed when stored), an adversary can make enclave calls produce results that would be impossible in the simpler model. As an example, take the following function from  $F$  that allows setting a key and sampling the output of a PRF function  $F$  for a single message:

```

function PRF-WRAPPER( $x, \text{mem}$ )
  if mem =  $\emptyset$  then
    K  $\leftarrow$   $x$ 
    Store mem  $\leftarrow$  K
    return ACK
  else if mem =  $\vec{1}$  then
    return  $\perp$ 

```

```

else
  Store mem  $\leftarrow \vec{1}$ 
  return  $F_K(x)$ 

```

An adversary who has completed initialisation of its decryption enclave with enclave id  $\text{eid}_{\text{DE}}$ , obtained a functional key  $\text{sk}$  through the execution of keygen on  $\text{eid}_{\text{KME}}$ , and initialised a functional enclave for PRF-WRAPPER with enclave id  $\text{eid}_{\text{F}}$ , public key  $\text{pk}_{\text{FD}}$  and attestation  $\sigma$ , executes the current operations for three ciphertexts  $\text{ct}_K, \text{ct}_x, \text{ct}_{x'}$ , encrypting a key  $K$  and plaintexts  $x, x'$ :

```

1:  $((\text{ct}_{\text{key}}, \text{crs}), \sigma_{\text{DE}}) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{DE}}, (\text{provision}, \sigma, \text{eid}, \text{pk}_{\text{FD}}, \text{sk}))$ 
2:  $((\text{computed}, ACK), \cdot) \leftarrow G_{\text{att}}^{\text{rollback}}.\text{resume}(\text{eid}_{\text{F}}, (\text{run}, \text{vk}_{\text{att}}, \sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_k, \text{crs}, \perp), \text{node})$ 
3: // node is the node id for a leaf for  $\text{eid}_{\text{F}}$ 's mem tree
4:  $((\text{computed}, y), \cdot) \leftarrow G_{\text{att}}^{\text{rollback}}.\text{resume}(\text{eid}_{\text{F}}, (\text{run}, \text{vk}_{\text{att}}, \sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_x, \text{crs}, \perp), \text{node}')$ 
5: //  $\text{node}'$  is the node id for a leaf for  $\text{eid}_{\text{F}}$ 's mem tree
6:  $((\text{computed}, y'), \cdot) \leftarrow G_{\text{att}}^{\text{rollback}}.\text{resume}(\text{eid}_{\text{F}}, (\text{run}, \text{vk}_{\text{att}}, \sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_{x'}, \text{crs}, \perp), \text{node}')$ 
7: //  $\text{node}'$  is the same node id as in the previous call (and thus to the parent
   of the current leaf in mem)

```

As a result of this execution trace, the adversary violates correctness by inserting an illegal transition (with input  $\epsilon$ ) in the stateful automaton for PRF-WRAPPER, from state  $\text{access}(\text{tree}, \text{node}'.\text{child}) = \vec{1}$  back to  $\text{access}(\text{tree}, \text{node}') = [K]$ , and then back to state  $\vec{1}$  with input  $x'$ . The adversary can then obtain the illegal set of values  $y \leftarrow F_K(x)$  and  $y' \leftarrow F_K(x')$ , whereas in the ideal world after obtaining  $y$ , the only possible output for the function would be  $\perp$  (the only legal transition from state  $\vec{1}$  leads back to itself). The simulator is unable to address this attack, as the memory state is internal to the ideal functionality, and the key will always be erased after the second call.

One might think that the simulator could respond by sampling a value from the uniform distribution and feed it through the enclave's backdoor; however, the environment can reveal the key  $k$  and messages  $x, x'$  to the adversary, or conversely the adversary could reveal the uniform value to the environment. Thus the environment can trivially distinguish between the honest PRF output and the uniform distribution, and thus between the real and ideal world. Note that this communication between environment and adversary is necessary for universal composition as this leakage of  $k, x, x'$  could happen as part of a wider protocol employing functional encryption.

**Mitigation techniques** In Section 1.3, we showed that rollback resilience for trusted execution environments is an active area of research, with many competing protocols. However, most solutions inevitably entail a performance tradeoff.

Due to the modular nature of Steel, it is possible to minimise the performance impact. Observe that party  $B$  instantiates a single DE and multiple FE. We can reduce the performance penalty by making only DE rollback resilient. We guarantee correctness despite rollbacks of FE, by encoding a counter alongside the function state for each F. On a decryption request, the  $\text{prog}_{\text{FE}}$  enclave is required to check in with the  $\text{prog}_{\text{DE}}$  enclave to retrieve the decryption key as part of the provision call. To enable rollback resilience, we include the counter

stored by  $\text{prog}_{\text{FE}}$  as an additional parameter of this call.  $\text{prog}_{\text{DE}}$  compares the counter received for the current evaluation of  $F$  with the one received during the last evaluation, and authorises the transfer of the secret key only if greater. Before evaluating the function,  $\text{prog}_{\text{FE}}$  increases and stores its local counter.

To achieve rollback resilience for the  $\text{prog}_{\text{DE}}$  enclave, we can rely on existing techniques in the literature, such as augmenting the enclave with asynchronous monotonic counters [13], or using protocols like LCM [16] or ROTE [40]. Formalising how these protocols can be combined with the  $G_{\text{att}}^{\text{rollback}}$  functionality to achieve the fully secure  $G_{\text{att}}$  is left for future work.

We also note that Stateless functional encryption as implemented in IRON is resilient to rollback and forking because there is little state held between computation. Since we assume  $C$  is honest, the only programs liable to be attacked are DE and FE[F].

DE stores PKE Parameters after init setup, and the decrypted master secret key after complete setup. The adversary could try to gain some advantage by creating multiple PKE pairs before authenticating with the authority, but will never has access to the raw sk unless combining it with a leakage attack. Denial of Service is possible by creating concurrent enclaves (either DE or FE) with different PKs, and passing encrypted ciphertxts to the "wrong" copy which would be unable to decrypt (but it's not clear what the advantage of using rollback attacks would be, as the adversary could always conduct a DoS attack by denying the necessary resources to the enclave).

*Acknowledgements* This research was partially supported by the National Cyber Security Centre, the UK Research Institute in Secure Hardware and Embedded Systems (RISE), and the European Union's Horizon 2020 Research and Innovation Programme under grant agreement 780108 (FENTEC).

## Bibliography

- [1] M. Abdalla, F. Benhamouda, M. Kohlweiss, and H. Waldner. Decentralizing inner-product functional encryption. In D. Lin and K. Sako, editors, *PKC 2019, Part II*, volume 11443 of *LNCS*, pages 128–157, Beijing, China, Apr. 14–17, 2019. Springer, Heidelberg, Germany.
- [2] S. Agrawal and D. J. Wu. Functional encryption: Deterministic to randomized functions from simple assumptions. In J.-S. Coron and J. B. Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 30–61, Paris, France, Apr. 30 – May 4, 2017. Springer, Heidelberg, Germany.
- [3] S. Agrawal, S. Gorbunov, V. Vaikuntanathan, and H. Wee. Functional encryption: New perspectives and lower bounds. *Cryptology ePrint Archive*, Report 2012/468, 2012. <http://eprint.iacr.org/2012/468>.
- [4] A. Ahmad, B. Joe, Y. Xiao, Y. Zhang, I. Shin, and B. Lee. OBFUSCURO: A commodity obfuscation engine on intel SGX. In *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society, 2019. ISBN

- 1-891562-55-X. URL <https://www.ndss-symposium.org/ndss-paper/obfusculo-a-commodity-obfuscation-engine-on-intel-sgx/>.
- [5] G. Ateniese, A. Kiayias, B. Magri, Y. Tselekounis, and D. Venturi. Secure outsourcing of cryptographic circuits manufacturing. In J. Baek, W. Susilo, and J. Kim, editors, *ProvSec 2018*, volume 11192 of *LNCS*, pages 75–93, Jeju, South Korea, Oct. 25–28, 2018. Springer, Heidelberg, Germany.
  - [6] J. Aumasson and L. Merino. Sgx secure enclaves in practice: security and crypto review. *Black Hat*, 2016:10, 2016.
  - [7] C. Badertscher, U. Maurer, D. Tschudi, and V. Zikas. Bitcoin as a transaction ledger: A composable treatment. In J. Katz and H. Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 324–356, Santa Barbara, CA, USA, Aug. 20–24, 2017. Springer, Heidelberg, Germany.
  - [8] C. Badertscher, R. Canetti, J. Hesse, B. Tackmann, and V. Zikas. Universal composition with global subroutines: Capturing global setup within plain UC. Cryptology ePrint Archive, Report 2020/1209, 2020. <https://eprint.iacr.org/2020/1209>.
  - [9] C. Badertscher, R. Canetti, J. Hesse, B. Tackmann, and V. Zikas. Universal composition with global subroutines: Capturing global setup within plain UC. In R. Pass and K. Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 1–30, Durham, NC, USA, Nov. 16–19, 2020. Springer, Heidelberg, Germany.
  - [10] C. Badertscher, A. Kiayias, M. Kohlweiss, and H. Waldner. Consistency for functional encryption. Cryptology ePrint Archive, Report 2020/137, 2020. <https://eprint.iacr.org/2020/137>.
  - [11] K. Bagheri, M. Kohlweiss, J. Siim, and M. Volkhov. Another look at extraction and randomization of groth’s zk-SNARK. Cryptology ePrint Archive, Report 2020/811, 2020. <https://eprint.iacr.org/2020/811>.
  - [12] R. Bahmani, M. Barbosa, F. Brasser, B. Portela, A.-R. Sadeghi, G. Scerri, and B. Warinschi. Secure multiparty computation from SGX. Cryptology ePrint Archive, Report 2016/1057, 2016. <http://eprint.iacr.org/2016/1057>.
  - [13] M. Bailleu, J. Thalheim, P. Bhatotia, C. Fetzer, M. Honda, and K. Vaswani. SPEICHER: securing lsm-based key-value stores using shielded execution. In A. Merchant and H. Weatherspoon, editors, *17th USENIX Conference on File and Storage Technologies, FAST 2019, Boston, MA, February 25-28, 2019*, pages 173–190. USENIX Association, 2019. URL <https://www.usenix.org/conference/fast19/presentation/bailleu>.
  - [14] M. Barbosa, B. Portela, G. Scerri, and B. Warinschi. Foundations of hardware-based attested computation and application to SGX. Cryptology ePrint Archive, Report 2016/014, 2016. <http://eprint.iacr.org/2016/014>.
  - [15] D. Boneh, A. Sahai, and B. Waters. Functional encryption: Definitions and challenges. In Y. Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273, Providence, RI, USA, Mar. 28–30, 2011. Springer, Heidelberg, Germany.

- [16] M. Brandenburger, C. Cachin, M. Lorenz, and R. Kapitza. Rollback and forking detection for trusted execution environments using lightweight collective memory. *CoRR*, 2017. URL <http://arxiv.org/abs/1701.00981v2>.
- [17] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. Cryptology ePrint Archive, Report 2000/067, 2000. <http://eprint.iacr.org/2000/067>.
- [18] R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally composable security with global setup. In S. P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 61–85, Amsterdam, The Netherlands, Feb. 21–24, 2007. Springer, Heidelberg, Germany.
- [19] R. Canetti, D. Shahaf, and M. Vald. Universally composable authentication and key-exchange with global PKI. In C.-M. Cheng, K.-M. Chung, G. Persiano, and B.-Y. Yang, editors, *PKC 2016, Part II*, volume 9615 of *LNCS*, pages 265–296, Taipei, Taiwan, Mar. 6–9, 2016. Springer, Heidelberg, Germany.
- [20] S. Cen and B. Zhang. Trusted time and monotonic counters with intel software guard extensions platform services. *Online at: <https://software.intel.com/sites/default/files/managed/1b/a2/Intel-SGX-Platform-Services.pdf>*, 2017.
- [21] R. Cheng, F. Zhang, J. Kos, W. He, N. Hynes, N. M. Johnson, A. Juels, A. Miller, and D. Song. Ekiden: A platform for confidentiality-preserving, trustworthy, and performant smart contract execution. *CoRR*, abs/1804.05141, 2018. URL <http://arxiv.org/abs/1804.05141>.
- [22] J. Chotard, E. Dufour Sans, R. Gay, D. H. Phan, and D. Pointcheval. Decentralized multi-client functional encryption for inner product. In T. Peyrin and S. Galbraith, editors, *ASIACRYPT 2018, Part II*, volume 11273 of *LNCS*, pages 703–732, Brisbane, Queensland, Australia, Dec. 2–6, 2018. Springer, Heidelberg, Germany.
- [23] A. R. Choudhuri, M. Green, A. Jain, G. Kaptchuk, and I. Miers. Fairness in an unfair world: Fair multiparty computation from public bulletin boards. In B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu, editors, *ACM CCS 2017*, pages 719–728, Dallas, TX, USA, Oct. 31 – Nov. 2, 2017. ACM Press.
- [24] K.-M. Chung, J. Katz, and H.-S. Zhou. Functional encryption from (small) hardware tokens. In K. Sako and P. Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 120–139, Bangalore, India, Dec. 1–5, 2013. Springer, Heidelberg, Germany.
- [25] M. Ciampi, Y. Lu, and V. Zikas. Collusion-preserving computation without a mediator. Cryptology ePrint Archive, Report 2020/497, 2020. <https://eprint.iacr.org/2020/497>.
- [26] A. Cloud. Tee-based confidential computing. <https://www.alibabacloud.com/help/doc-detail/164536.htm>, 2020.
- [27] V. Costan and S. Devadas. Intel SGX explained. Cryptology ePrint Archive, Report 2016/086, 2016. <http://eprint.iacr.org/2016/086>.
- [28] B. Fisch, D. Vinayagamurthy, D. Boneh, and S. Gorbunov. IRON: Functional encryption using intel SGX. In B. M. Thuraisingham, D. Evans,

- T. Malkin, and D. Xu, editors, *ACM CCS 2017*, pages 765–782, Dallas, TX, USA, Oct. 31 – Nov. 2, 2017. ACM Press.
- [29] C. Garlati and S. Pinto. A clean slate approach to linux security risc-v enclaves. 02 2020.
- [30] V. Goyal, A. Jain, V. Koppula, and A. Sahai. Functional encryption for randomized functionalities. In Y. Dodis and J. B. Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 325–351, Warsaw, Poland, Mar. 23–25, 2015. Springer, Heidelberg, Germany.
- [31] F. Gregor, W. Ozga, S. Vaucher, R. Pires, D. L. Quoc, S. Arnautov, A. Martin, V. Schiavoni, P. Felber, and C. Fetzer. Trust management as a service: Enabling trusted execution in the face of byzantine stakeholders. *CoRR*, abs/2003.14099, 2020. URL <https://arxiv.org/abs/2003.14099>.
- [32] V. T. Hoang, R. Reyhanitabar, P. Rogaway, and D. Vizár. Online authenticated-encryption and its nonce-reuse misuse-resistance. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 493–517, Santa Barbara, CA, USA, Aug. 16–20, 2015. Springer, Heidelberg, Germany.
- [33] D. Kaplan, J. Powell, and T. Woller. Amd memory encryption. *White paper*, 2016.
- [34] A. Kiayias and Y. Tselekounis. Tamper resilient circuits: The adversary at the gates. In K. Sako and P. Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 161–180, Bangalore, India, Dec. 1–5, 2013. Springer, Heidelberg, Germany.
- [35] A. Kiayias, F.-H. Liu, and Y. Tselekounis. Practical non-malleable codes from l-more extractable hash functions. In E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi, editors, *ACM CCS 2016*, pages 1317–1328, Vienna, Austria, Oct. 24–28, 2016. ACM Press.
- [36] A. Kiayias, F.-H. Liu, and Y. Tselekounis. Non-malleable codes for partial functions with manipulation detection. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 577–607, Santa Barbara, CA, USA, Aug. 19–23, 2018. Springer, Heidelberg, Germany.
- [37] I. Komargodski, G. Segev, and E. Yogev. Functional encryption for randomized functionalities in the private-key setting from minimal assumptions. In Y. Dodis and J. B. Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 352–377, Warsaw, Poland, Mar. 23–25, 2015. Springer, Heidelberg, Germany.
- [38] D. Lee, D. Kohlbrenner, S. Shinde, K. Asanović, and D. Song. Keystone: An open framework for architecting trusted execution environments. In *Proceedings of the Fifteenth European Conference on Computer Systems*, pages 1–16, 2020.
- [39] D. Levin, J. R. Douceur, J. R. Lorch, and T. Moscibroda. Trinc: Small trusted hardware for large distributed systems. In *NSDI*, volume 9, pages 1–14, 2009.
- [40] S. Matetic, M. Ahmed, K. Kostianen, A. Dhar, D. Sommer, A. Gervais, A. Juels, and S. Capkun. ROTE: Rollback protection for trusted execution.

- Cryptology ePrint Archive, Report 2017/048, 2017. <http://eprint.iacr.org/2017/048>.
- [41] C. Matt and U. Maurer. A definitional framework for functional encryption. In C. Fournet and M. Hicks, editors, *CSF 2015 Computer Security Foundations Symposium*, pages 217–231, Verona, Italy, jul 13-17 2015. IEEE Computer Society Press.
  - [42] K. Nayak, C. W. Fletcher, L. Ren, N. Chandran, S. V. Lokam, E. Shi, and V. Goyal. HOP: Hardware makes obfuscation practical. In *NDSS 2017*, San Diego, CA, USA, Feb. 26 – Mar. 1, 2017. The Internet Society.
  - [43] B. Parno, J. M. McCune, and A. Perrig. Bootstrapping trust in commodity computers. In *2010 IEEE Symposium on Security and Privacy*, pages 414–429, Berkeley/Oakland, CA, USA, May 16–19, 2010. IEEE Computer Society Press.
  - [44] R. Pass, E. Shi, and F. Tramèr. Formal abstractions for attested execution secure processors. In J.-S. Coron and J. B. Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 260–289, Paris, France, Apr. 30 – May 4, 2017. Springer, Heidelberg, Germany.
  - [45] S. Pinto and N. Santos. Demystifying arm trustzone: A comprehensive survey. *ACM Computing Surveys (CSUR)*, 51(6):1–36, 2019.
  - [46] N. Porter, G. Golan, and S. Lugani. Introducing google cloud confidential computing with confidential vms. 2020.
  - [47] M. Russinovich. Introducing azure confidential computing. *URL=https://azure.microsoft.com/ru-ru/blog/introducing-azureconfidential-computing*, 2017.
  - [48] A. Sahai and B. R. Waters. Fuzzy identity-based encryption. In R. Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473, Aarhus, Denmark, May 22–26, 2005. Springer, Heidelberg, Germany.
  - [49] A. Shamir. Identity-based cryptosystems and signature schemes. In G. R. Blakley and D. Chaum, editors, *CRYPTO’84*, volume 196 of *LNCS*, pages 47–53, Santa Barbara, CA, USA, Aug. 19–23, 1984. Springer, Heidelberg, Germany.
  - [50] R. Strackx and F. Piessens. Ariadne: A minimal approach to state continuity. In T. Holz and S. Savage, editors, *USENIX Security 2016*, pages 875–892, Austin, TX, USA, Aug. 10–12, 2016. USENIX Association.
  - [51] T. Suzuki, K. Emura, T. Ohigashi, and K. Omote. Verifiable functional encryption using intel SGX. Cryptology ePrint Archive, Report 2020/1221, 2020. <https://eprint.iacr.org/2020/1221>.
  - [52] F. Tramèr, F. Zhang, H. Lin, J.-P. Hubaux, A. Juels, and E. Shi. Sealed-glass proofs: Using transparent enclaves to prove and sell knowledge. Cryptology ePrint Archive, Report 2016/635, 2016. <http://eprint.iacr.org/2016/635>.
  - [53] I. Tselekounis. *Cryptographic techniques for hardware security*. PhD thesis, University of Edinburgh, UK, 2018. URL <http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.763966>.
  - [54] P. Wu, Q. Shen, R. H. Deng, X. Liu, Y. Zhang, and Z. Wu. OblIDC: An SGX-based oblivious distributed computing framework with formal proof.

- In S. D. Galbraith, G. Russello, W. Susilo, D. Gollmann, E. Kirda, and Z. Liang, editors, *ASIACCS 19*, pages 86–99, Auckland, New Zealand, July 9–12, 2019. ACM Press.
- [55] S. Yilek. Resettable public-key encryption: How to encrypt on a virtual machine. In J. Pieprzyk, editor, *CT-RSA 2010*, volume 5985 of *LNCS*, pages 41–56, San Francisco, CA, USA, Mar. 1–5, 2010. Springer, Heidelberg, Germany.