THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Steel: Composable Hardware-Based Stateful and Randomised Functional Encryption

Citation for published version:

Bhatotia, P, Kohlweiss, M, Martinico, L & Tselekounis, Y 2021, Steel: Composable Hardware-Based Stateful and Randomised Functional Encryption. in JA Garay (ed.), *Public-Key Cryptography -- PKC 2021*. Lecture Notes in Computer Science, vol. 12711, Springer International Publishing, Cham, pp. 709-736, 24th IACR International Conference on Practice and Theory of Public Key Cryptography, 10/05/21.
https://doi.org/10.1007/978-3-030-75248-4_25

Digital Object Identifier (DOI):

[10.1007/978-3-030-75248-4_25](https://doi.org/10.1007/978-3-030-75248-4_25)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Public-Key Cryptography -- PKC 2021

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Steel: Composable Hardware-based Stateful and Randomised Functional Encryption

Pramod Bhatotia^{1,2}, Markulf Kohlweiss¹, Lorenzo Martinico¹, and Yiannis Tselekounis¹

¹ School of Informatics, University of Edinburgh

² 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** realises FESR in the *universal composition with global sub-routines* 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

---

© IACR 2021. This article is the final version submitted by the author(s) to the IACR and to Springer-Verlag on May 11th, 2021. The version published by Springer-Verlag is available at http://dx.doi.org/10.1007/978-3-030-75248-4_25.

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 [28], ARM Trustzone [49], RISC-V Keystone [33, 42], AMD-SEV [37], and AWS Nitro [55] 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 [52], Google Cloud [50], and Alibaba Cloud [5].

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 roll-back 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) [48] 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 [59], 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 [17]. FE is a generalisation of Attribute/Identify Based Encryption [56, 53], 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 [24, 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 [30] introduces *Iron*, which is a practically efficient protocol that realises FE based on Intel’s SGX. In [30] the authors formally prove security of the proposed protocol, however their proof is in the standalone setting. In a related work, Matt and Maurer [45] 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, 34, 41], 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 [24] and enables cryptographic 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 [30].
- We construct the protocol **Steel** and prove that it realizes FESR in the newly introduced Universal Composition with Global Subroutines (UCGS) model [9]. Our main building blocks are: (1) the functional encryption scheme of [30] 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 [30]. Our security proof shows that one can satisfy a FE ideal functionality as defined in [45], relying on hardware instead of Random Oracles, and for the larger function class of FESR when additionally relying on a common reference string.
- 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_{att} of PST [48]. Our UC protocols assume access to the *global attestation functionality*, G_{att} , that captures the core abstraction provided by a broad class of attested execution processors, such as Intel SGX [28]. 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_{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_{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_{att} is defined as a *global functionality* that uses the same key pair across sessions.

Universal composition with global subroutines [10]. In our work we model global information using the newly introduced UCGS framework, which resolves inconsistencies in GUC [20], 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_{att} (cf. Section 2.3) 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_{att} , abstracting an execution in the presence of secure hardware enclaves. Our protocol is based on Iron [30], so we briefly revisit its 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 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_{att} ’s execution, this produces an attestation signature σ over F, which becomes the functional key 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_{att}).

The statefulness of functional encryption is simply enabled by adding state storage to each functional enclave. Similar to [48], 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. An efficient implementation for such a modification can be realised by drawing on the techniques of signed ElGamal [54], where the random coin used during encryption is sufficient to prove knowledge of the plaintext. However, this construction would require the use of the random oracle and generic group model, or at best the algebraic group model [32], making compositionality non-obvious, and is thus not considered in full beyond this remark. 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, Σ , a CCA-secure public-key encryption scheme, PKE , and a non-interactive zero knowledge scheme, N . Informally, Σ provides the guarantees required for realising attested computation (as discussed above), 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,³ 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 Σ). Indistinguishability is proven via reduction to the EU-CMA security of Σ . 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 re-

---

³ Here we omit some standard UC-related hybrids.

quire reductions to the CCA security of PKE).⁴ 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 $\mathcal{N}$. Then, for maliciously generated ciphertexts under PKE – which might result via tampering with honestly generated encryptors’ ciphertexts – instead of using the decryption operation of PKE, our simulator recovers the corresponding plaintext using the extractability property of $\mathcal{N}$. Finally, we fake all ciphertexts of PKE, that encrypt the inputs to the functions (this reduces to CPA security). Note that, in [30], 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 its choice. While in the classic setting for Functional Encryption (where *Iron* lives) simulation security is easily achieved 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 asks the ideal FESR functionality for its evaluation. Simulation-extractable NIZK can be efficiently instantiated, e.g., using zk-SNARKs [12]. 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_{att} facilitates composable protocol design, however, such a functionality cannot be easily realised 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_{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_{att} within *Steel*, and discuss how countermeasures from the rollback protection literature can be adopted to address these attacks, with a consideration on efficiency.

---

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

1.3 Related work

Hardware is frequently used to improve performance or circumvent impossibility results, e.g. [46, 4, 27]. As a relevant example, Chung et al. [26] show how to use stateless hardware tokens to implement functional encryption.

The use of attestation has been widely adopted in the design of computer systems to bootstrap security [47]. In addition to formalising attested execution, Pass, Shi and Tramer (PST) [48] 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 [59, 61, 25, 23]. Other attempts at providing a formal model for attested execution include the game-based models of Barbosa et al. [16], Bahmani et al. [13], Fisch et al. [30]. The latter model arises from the need to evaluate the security of the aforementioned Iron [31], which first realises Functional Encryption in the hardware setting. A further extension implementing verifiable functional encryption is presented in Suzuki et al. [58].

A significant source of vulnerabilities in Trusted Execution Environments arise from side-channel attacks; many have been conducted in the literature against Intel SGX (see [51] for a non-comprehensive review). One class of vulnerabilities are that of state-continuity, or rollback attacks (also known as reset attacks in the cryptographic literature), also being a common attack vector 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 [62] presents a general attack that is applicable to both virtual machine and enclave executions: 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 [36], 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 [22], TPMs [43], 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 [44]. 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 [44] 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 [57] is an optimized (local) synchronous technique to increment the counter by a single bit flip for deterministic enclaves.

Speicher [14] and Palaemon [35] 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. [18]) 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, [39, 40, 60], protect hardware memory against active attacks, while [38, 6], protect cryptographic hardware against tampering and Trojan injection attacks, respectively.

2 Preliminaries

2.1 Functional Encryption

In the current section, we define the syntax of Functional Encryption.

Functional Encryption is a primitive defined over a class of functions $F : \mathcal{X} \rightarrow \mathcal{Y}$, consisting of the following PPT algorithms:

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

Confidentiality A confidential Functional Encryption scheme allows only the function evaluation $F(x)$ to be learned from the ciphertext ct and functional key sk_F .

Correctness A functional encryption scheme satisfies correctness if, for all functions $F \in F$ and all $x \in \mathcal{X}$, the statement $F(x) \leftarrow \text{Dec}(\text{Enc}(\text{mpk}, x), \text{sk}_F)$

Composable Functional Encryption Matt and Maurer [45] 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).

2.2 UC background

Universal Composability (UC), introduced by Canetti [19], 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) [10] and the 2020 version of UC [19]. We briefly summarise the aspects of UC and UCGS necessary to understand our work between this section and Appendix B, where we define all terminology used.

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 ξ -identity-bounded environment is limited to assume external identities as specified by ξ , a polynomial time boolean predicate on the current system configuration.

We now recall a few definitions. Please consult [19, 10] or (appendix B) for the formal definitions of terms such as *balanced*, *respecting*, *exposing*, *compliant*.

Definition 1 (UC emulation [19]). *Given two PPT protocols π, ϕ and some predicate ξ , we say that π UC-emulates ϕ with respect to ξ -identity bound environments (or π ξ -UC-emulates ϕ) if for any balanced ξ -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 π which UC-emulates a protocol ϕ , and a third protocol ρ , which calls ϕ as a subroutine, we can construct a protocol where all calls to ϕ are replaced with calls to π , denoted as $\rho^{\phi \rightarrow \pi}$.

Theorem 1 (Universal Composition [19]). *Given PPT protocols π, ϕ, ρ and predicate ξ , if π, ϕ are both subroutine respecting and subroutine exposing (see*

B), ρ is (π, ϕ, ξ) -compliant and π ξ -UC-emulates ϕ , then protocol $\rho^{\phi \rightarrow \pi}$ UC-emulates ρ

By the composition theorem, any protocol that leverages subroutine ϕ in its execution can now be instantiated using protocol π .

UCGS As the name suggests, generalised UC (GUC) [20] is an important generalization of the UC model. It accounts for the existence of a shared subroutine γ , such that both ρ and its subroutine π (regardless of how many instances of π are called by ρ) can have γ 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) [21], a global clock [8], or trusted hardware [48].

Unfortunately GUC has inconsistencies and has not been updated from the 2007 to the 2020 version of UC.⁵ Universal Composability with Global Subroutines [10] 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 π with access to subroutine γ is replaced by a new structured protocol $\mu = M[\pi, \gamma]$, known as *management* protocol; μ allows multiplexing a single instance of π and γ into however many are required by ρ , by transforming the session and party identifiers. μ 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 μ , and is used to redirect the outputs of π and γ to the correct machine.

Below we revisit the UC emulation with global subroutines definition from [10].

Definition 2 (UC emulation with global subroutines [10]). *Given protocols π, ϕ , and γ , π ξ -UC emulates ϕ in the presence of γ if $M[\pi, \gamma]$ ξ -UC emulates $M[\phi, \gamma]$*

Now we state the main UCGS theorem.

Theorem 2 (Universal Composition with Global subroutines [10]). *Given subroutine-exposing protocols π, ϕ, ρ , and γ , if γ is a ϕ -regular setup and subroutine respecting, ϕ, π are γ -subroutine respecting, ρ is (π, ϕ, ξ) -compliant and $(\pi, M[x, \gamma], \xi)$ -compliant for $x \in \{\phi, \pi\}$, then if π ξ -UC-emulates ϕ in the presence of γ , the protocol $\rho^{\phi \rightarrow \pi}$ UC-emulates ρ in the presence of γ .*

⁵ 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).

2.3 The G_{att} functionality

We now reproduce the G_{att} global functionality defined in the PST model [48]. The functionality is parameterised 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 vk be a global variable, accessible by enclave programs as $G_{\text{att}}.\text{vk}$. This allows us to use G_{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.

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
<i>On message INITIALIZE from a party P:</i>	
let $(\text{spk}, \text{ssk}) \leftarrow \Sigma.\text{Gen}(1^\lambda)$, $\text{vk} \leftarrow \text{spk}$, $\text{msk} \leftarrow \text{ssk}$	
<i>On message GETPK from a party P:</i>	
return vk	
<i>On message (INSTALL, idx, prog) from a party P where $P.\text{pid} \in \text{reg}$:</i>	
if P is honest then	
assert $\text{idx} = P.\text{sid}$	
generate nonce $\text{eid} \in \{0, 1\}^\lambda$, store $\mathcal{T}[\text{eid}, P] = (\text{idx}, \text{prog}, \emptyset)$	
send eid to P	
<i>On message (RESUME, eid, input) from a party P where $P.\text{pid} \in \text{reg}$:</i>	
let $(\text{idx}, \text{prog}, \text{mem}) \leftarrow \mathcal{T}[\text{eid}, P]$, abort if not found	
let $(\text{output}, \text{mem}') \leftarrow \text{prog}(\text{input}, \text{mem})$, 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, σ) to P	

The G_{att} functionality is a generalisation over other TEE formalisations, such as the one in [30], 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_{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 encoding 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

Our primitive of FESR offers a generalisation over Functional Encryption, to allow computing of *stateful* and *randomized functionalities* over arbitrary ciphertexts. In this section we define the ideal functionality of functional encryption for *stateful* and *randomized functionalities* (FESR).

The syntax for this new primitive matches that of Functional Encryption schemes (outlined in section 2.1). The two primitives differ by the parametrisation of the class of computable functions $\mathbf{F}$; in the case of FESR, this is defined as

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

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

The definition of the primitive follows from the ideal functionality, given below.

3.1 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 $\mathbf{F}$, 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 $\mathbf{F}$, to a decryption party to obtain an evaluation of $\mathbf{F}$ on the message stored therein. Performing the computation will also result in updating the state stored in $\mathcal{P}$.

Functionality FESR[sid, $\mathbf{F}$, **A**, **B**, **C**]

The functionality is parameterized by the randomized function class $\mathbf{F}$ such that for each $\mathbf{F} \in \mathbf{F} : \mathcal{X} \times \mathcal{S} \times \mathcal{R} \rightarrow \mathcal{Y} \times \mathcal{S}$, over state space $\mathcal{S}$ and randomness

space $\mathcal{R}$, and by three distinct types of party identities $\mathbf{A}, \mathbf{B}, \mathbf{C}$ 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 $\mathbf{C}$, for $P \in \{\mathbf{A}, \mathbf{B}\}$:

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

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

$\text{setup}[P] \leftarrow \text{true}$
 $\mathcal{P}[P, F_0] \leftarrow \emptyset$
send SETUP **to** P

On message (ENCRYPT, x) from party $P \in \{\mathbf{A}, \mathbf{B}\}$:

if $\text{setup}[P] = \text{true} \wedge x \in \mathcal{X}$ **then**
     compute $h \leftarrow \text{getHandle}$
      $\mathcal{M}[h] \leftarrow x$
     send (ENCRYPTED, h) **to** P

On message (KEYGEN, $F, \mathbf{B}$) from party $\mathbf{C}$:

if $F \in \mathbf{F}^+ \wedge \text{setup}[\mathbf{B}] = \text{true}$ **then**
     send (KEYGEN, $F, \mathbf{B}$) **to** $\mathcal{A}$ and **receive** ACK
      $\mathcal{P}[\mathbf{B}, F] \leftarrow \emptyset$
     send (ASSIGNED, F) **to** $\mathbf{B}$

On message (DECRYPT, h, F) from party $\mathbf{B}$:

$x \leftarrow \mathcal{M}[h]$
if $\mathbf{C}$ is honest **then**
     if $\mathcal{P}[\mathbf{B}, F] \neq \perp \wedge x \in \mathcal{X}$ **then**
          $r \leftarrow \mathcal{R}$
          $s \leftarrow \mathcal{P}[\mathbf{B}, F]$
          $(y, s) \leftarrow F(x, s, r)$
          $\mathcal{P}[\mathbf{B}, F] \leftarrow s'$
         return (DECRYPTED, y)
     else
         send (DECRYPT, h, F, x) **to** $\mathcal{A}$ and **receive** (DECRYPTED, y)
         return (DECRYPTED, y)

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

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*, originally introduced by [11] (in which **A** and/or **C** might also get corrupted).

Our definition is along the lines of [11, 45]; in order To allow stateful and randomized functions, we extend the function class with support for private state and randomness as above. Whenever **B** accesses a function on the data from the repository, the repository draws fresh randomness, evaluates the function on the old state (for the current **B** and function). The function updates the state and returns an evaluation.

The property of confidentiality for functional encryption also holds for FESR, as the decrypting party is only allowed to learn the function evaluation (and not the state, before or after decryption). Correctness for FESR is slightly stronger than in 2.1: it is necessary for the state at any decryption to be uniquely determined by the sequence of previous decryption for the same F, B pair (without allowing **B** to influence its value, besides the choice of which ciphertexts to decrypt). Intuitively, the ideal world AC repository presented models both confidentiality and correctness. by inspection of the four lines $r \leftarrow \mathcal{R}$, $s \leftarrow \mathcal{P}[B, F]$, $(y, s') \leftarrow F(x, s, r)$, and $\mathcal{P}[B, F] \leftarrow s'$.

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, 34, 41], 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.

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. Confidentiality and correctness of the protocol in the face of an adversarial **B** hold from the proof of indistinguishability between real and ideal world in 5. . 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 π by authorizing a fake reference string 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 (α, β) occurs over secure channels $\mathcal{SC}_\alpha^\beta, \mathcal{SC}_\beta^\alpha$
- Functional keys are (attestation) signatures by an enclave prog_{KME} on input (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 parameterised 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 Σ , a Robust non-interactive zero-knowledge scheme N , and security parameter λ .

Protocol Steel[$F, A, B, C, PKE, \Sigma, N, \lambda$]

State variables	Description
$mpk \leftarrow \perp$	Local copy of master public key for participants
$prog_{\{KME, DE, FE\}} \leftarrow \dots$	Source code of enclaves as defined below
$\mathcal{K}[\cdot] \leftarrow \emptyset$	Table of function keys at B

Key Generation Authority C:

On message (SETUP, P):

```

if  $mpk = \perp$  then
   $eid_{KME} \leftarrow G_{att}.install(C.sid, prog_{KME})$ 
  send GET to CRS and receive (CRS, crs)
   $(mpk, \cdot) \leftarrow G_{att}.resume(eid_{KME}, (init, crs, C.sid))$ 
if  $P = A$  then
  send (SETUP, mpk) to  $SC_A$ 
else if  $P = B$  then
  send (SETUP, mpk,  $eid_{KME}$ ) to  $SC_B$  and receive (PROVISION,  $\sigma$ ,  $eid_{DE}$ ,  $pk_{KD}$ )
   $(ct_{key}, \sigma_{sk}) \leftarrow G_{att}.resume(eid_{KME}, (provision, (\sigma, eid_{DE}, pk_{KD}, eid_{KME})))$ 
  send (PROVISION,  $ct_{key}, \sigma_{sk}$ ) to  $SC_B$ 

```

On message (KEYGEN, F, B):

```

assert  $F \in F \wedge mpk \neq \perp$ 
 $((keygen, F), \sigma) \leftarrow G_{att}.resume(eid_{KME}, (keygen, F))$ 
 $sk_F \leftarrow \sigma$ ; send (KEYGEN, ( $F, sk_F$ )) to  $SC_B$ 

```

Encryption Party A:

On message (SETUP, mpk) *from* SC^C :

```

send GET to CRS and receive (CRS, crs)
store mpk, crs; return SETUP

```

On message (ENCRYPT, m):

```

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

```

Decryption Party B:

On message (SETUP, mpk, eid_{KME}) *from* SC^C :

```

store  $\text{mpk}$ ;  $\text{eid}_{\text{DE}} \leftarrow G_{\text{att}}.\text{install}(\text{B.sid}, \text{prog}_{\text{DE}})$ 
send GET to  $\mathcal{CRS}$  and receive  $(\text{CRS}, \text{crs})$ 
 $((\text{pk}_{\text{KD}}, \cdot, \cdot), \sigma) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{DE}}, \text{init-setup}, \text{eid}_{\text{KME}}, \text{crs}, \text{B.sid})$ 
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,  $(\text{F}, \text{sk}_{\text{F}})$ ) from  $\mathcal{SC}^C$ :
 $\text{eid}_{\text{F}} \leftarrow G_{\text{att}}.\text{install}(\text{B.sid}, \text{prog}_{\text{FE}}[\text{F}])$ 
 $(\text{pk}_{\text{FD}}, \sigma_{\text{F}}) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{F}}, (\text{init}, \text{mpk}, \text{B.sid}))$ 
 $\mathcal{K}[\text{F}] \leftarrow (\sigma_{\text{F}}, \text{eid}_{\text{F}}, \text{pk}_{\text{FD}}, \text{sk}_{\text{F}})$ 
return (ASSIGNED,  $\text{F}$ )

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

 $\text{prog}_{\text{KME}}$ :
on input  $(\text{init}, \text{crs}, \text{idx})$ :
assert  $\text{pk} = \perp$ ;  $(\text{pk}, \text{sk}) \leftarrow \text{PKE.PGen}()$ 
store  $\text{sk}, \text{crs}, \text{idx}$ ; return  $\text{pk}$ 
on input  $(\text{provision}, (\sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{pk}_{\text{KD}}, \text{eid}_{\text{KME}}))$ :
 $\text{vk}_{\text{att}} \leftarrow G_{\text{att}}.\text{vk}$ ; fetch  $\text{crs}, \text{idx}$ 
assert  $\Sigma.\text{Vrfy}(\text{vk}_{\text{att}}, (\text{idx}, \text{eid}_{\text{DE}}, \text{prog}_{\text{DE}}, (\text{pk}_{\text{KD}}, \text{eid}_{\text{KME}}, \text{crs}), \sigma_{\text{DE}}))$ 
 $\text{ct}_{\text{key}} \leftarrow \text{PKE.Enc}(\text{pk}_{\text{KD}}, \text{sk})$ 
return  $\text{ct}_{\text{key}}$ 
on input (keygen,  $\text{F}$ ):
return (keygen,  $\text{F}$ )

```

<pre> prog_{DE}: on input (init-setup, eid_{KME}, crs, idx): assert pk_{KD} ≠ ⊥ (pk_{KD}, sk_{KD}) ← PKE.Gen() store sk_{KD}, eid_{KME}, crs, idx return pk_{KD}, eid_{KME}, crs on input (complete-setup, ct_{key}, σ_{KME}): vk_{att} ← G_{att}.vk fetch eid_{KME}, sk_{KD}, idx m ← (idx, eid_{KME}, prog_{KME}, ct_{key}) assert Σ.Vrfy(vk_{att}, m, σ_{KME}) sk ← PKE.Dec(sk_{KD}, ct_{key}) store sk, vk_{att} on input (provision, σ, eid, pk_{FD}, sk_F, F): fetch eid_{KME}, vk_{att}, sk, idx m₁ ← (idx, eid_{KME}, prog_{KME}, (keygen, F)) m₂ ← (idx, eid, prog_{FE}[F], pk_{FD}) assert Σ.Vrfy(vk_{att}, m₁, sk_F) and Σ.Vrfy(vk_{att}, m₂, σ) return PKE.Enc(pk_{FD}, sk), crs </pre>	<pre> prog_{FE}[F]: on input (init, mpk, idx): assert pk_{FD} = ⊥ (pk_{FD}, sk_{FD}) = PKE.Gen(1^λ) mem ← ∅; store sk_{FD}, mem, mpk, idx return pk_{FD} on input (run, σ_{DE}, eid_{DE}, ct_{key}, ct_{msg}, crs, y'): if y' ≠ ⊥ return (computed, y') vk_{att} ← G_{att}.vk; (ct, π) ← ct_{msg} fetch sk_{FD}, mem, mpk, idx m ← (idx, eid_{DE}, prog_{DE}, ct_{key}, crs) assert Σ.Vrfy(vk_{att}, m, σ_{DE}) sk = PKE.Dec(sk_{FD}, ct_{key}) assert N.V((mpk, ct), π, crs) x = PKE.Dec(sk, ct) out, mem' = F(x, mem) store mem ← mem' return (computed, out) </pre>

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 $R = \{(pk, ct), (m, r) | ct = \text{PKE.Enc}(mpk, m; r)\}$ 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 ξ .

The identity bound ξ on the environment. Our restrictions are similar to [48], namely we assume that the environment can access G_{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-realises the FESR functionality (Section 3) in the presence of the global functionality G_{att} and local functionalities $\mathcal{CRS}$, $\mathcal{R\mathcal{E}P}$, $\mathcal{SC}$, with respect to the identity bound ξ defined above.*

We present a simulator algorithm $\mathcal{S}_{\text{FESR}}$ such that UC emulation 1 holds for protocols **Steel** and $\text{IDEAL}_{\text{FESR}}$ (the protocol encapsulating the ideal functionality and a set of dummy parties corresponding to the real-world parties in **Steel**). Following [45], 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_{att} . While hybrid functionalities $\mathcal{R\mathcal{E}P}$, $\mathcal{SC}$, $\mathcal{CRS}$ (defined in A.4) are only available in the real world and need to be reproduced by the simulator, we use $\mathcal{SC}$ 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_{att} , **Steel** ξ -UC emulates FESR in the presence of G_{att} if $M[\text{Steel}, G_{\text{att}}]$ ξ -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_{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_{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_{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 the identity of B .

Simulator $\mathcal{S}_{\text{FESR}}[\text{PKE}, \Sigma, N, \lambda, F]$

State variables	Description
$\mathcal{H}[\cdot] \leftarrow \emptyset$	Table of ciphertext and handles in public repository
$\mathcal{K} \leftarrow \emptyset$	List of $\text{prog}_{\text{FE}}[\text{F}]$ enclaves and their eid_{F}
$\mathcal{G} \leftarrow \{\}$	Collects all messages sent to G_{att} and its response
$\mathcal{B} \leftarrow \{\}$	Collects all messages signed by G_{att}
$(\text{crs}, \tau) \leftarrow \text{N.S}_1$	Simulated reference string and trapdoor

For Key Generation Authority C:

On message (SETUP, P) from FESR:

```

if  $\text{mpk} = \perp$  then
   $\text{eid}_{\text{KME}} \leftarrow G_{\text{att}}.\text{install}(\text{C.sid}, \text{prog}_{\text{KME}})$ 
   $(\text{mpk}, \cdot) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{KME}}, \text{init})$ 
if  $P = \text{A}$  then
  send (SETUP,  $\text{mpk}$ ) to  $\mathcal{SC}_{\text{A}}$ 
else if  $P = \text{B}$  then
  send (SETUP,  $\text{mpk}, \text{eid}_{\text{KME}}$ ) to  $\mathcal{SC}_{\text{B}}$  and receive (PROVISION,  $\sigma, \text{eid}_{\text{DE}}, \text{pk}_{\text{KD}}$ )
  assert  $(\text{C.sid}, \text{eid}_{\text{DE}}, \text{prog}_{\text{DE}}, \text{pk}_{\text{KD}}) \in \mathcal{B}[\sigma]$ 
   $(\text{ct}_{\text{key}}, \sigma_{\text{sk}}) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{KME}}, (\text{provision}, (\sigma, \text{eid}_{\text{DE}}, \text{pk}_{\text{KD}}, \text{eid}_{\text{KME}}, \text{crs})))$ 
  send (PROVISION,  $\text{ct}_{\text{key}}, \sigma_{\text{sk}}$ ) to  $\mathcal{SC}_{\text{B}}$ 

```

On message (KEYGEN, F, B) from FESR:

```

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

```

For Decryption Party B:

On message GET from party B to CRS:

```

send (CRS,  $\text{crs}$ ) to B

```

On message (READ, h) from party B to $\mathcal{REP}$:

```

send (DECRYPT,  $F_0, h$ ) to FESR on behalf of B and receive  $|m|$ 
assert  $|m| \neq \perp$ 
 $\text{ct} \leftarrow \text{PKE}.\text{Enc}(\text{mpk}, 0^{|m|})$ 
 $\pi \leftarrow \text{N.S}_2(\text{crs}, \tau, (\text{mpk}, \text{ct}))$ 
 $\text{ct}_{\text{msg}} \leftarrow (\text{ct}, \pi); \mathcal{H}[\text{ct}_{\text{msg}}] \leftarrow h$ 
send (READ,  $\text{ct}_{\text{msg}}$ ) to B

```

On message (INSTALL, idx, prog) from party B to G_{att} :

```

 $\text{eid} \leftarrow G_{\text{att}}.\text{install}(\text{idx}, \text{prog})$ 
 $\mathcal{G}[\text{eid}].\text{install} \leftarrow (\text{idx}, \text{prog})$ 
//  $\mathcal{G}[\text{eid}].\text{install}[1]$  is the program's code
forward  $\text{eid}$  to B

```

On message (RESUME, eid, input) from party B to G_{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
   $(\text{output}, \sigma) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}, \text{input})$ 
   $\mathcal{G}[\text{eid}].\text{resume} \leftarrow \mathcal{G}[\text{eid}].\text{resume} \parallel (\sigma, \text{input}, \text{output})$ 
   $\mathcal{B}[\sigma] \leftarrow (\mathcal{G}[\text{eid}].\text{install}[0], \text{eid}, \mathcal{G}[\text{eid}].\text{install}[1], \text{output})$ 

```

```

if  $\mathcal{G}[\text{eid}].\text{install}[1] = \text{prog}_{\text{DE}} \wedge \text{input}[0] = \text{provision}$  then
   $(\text{provision}, \sigma, \text{eid}, \text{pk}_{\text{FD}}, \text{sk}_{\text{F}}, \text{F}) \leftarrow \text{input}$ 
  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  $(\text{output}, \sigma)$  to B
else
   $\text{idx}, \text{prog}_{\text{FE}}[\text{F}] \leftarrow \mathcal{G}[\text{eid}].\text{install}$ 
   $(\text{run}, \sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_{\text{msg}}, \text{crs}, \perp) \leftarrow \text{input}$ 
  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
     $(\text{ct}, \pi) \leftarrow \text{ct}_{\text{msg}}$ 
     $(\text{m}, \text{r}) \leftarrow \text{N.E}(\tau, (\text{mpk}, \text{ct}), \pi)$ 
    if  $\text{m} = \perp$  then send  $(\text{DECRYPT}, \text{F}, \perp)$  to B and abort
    send  $(\text{ENCRYPT}, \text{m})$  to FESR on behalf of B and receive h
     $\mathcal{H}[\text{ct}_{\text{msg}}] \leftarrow \text{h}$ 
   $\text{h} \leftarrow \mathcal{H}[\text{ct}_{\text{msg}}]$ 
  send  $(\text{DECRYPT}, \text{F}, \text{h})$  to FESR on behalf of B and receive y
   $((\text{computed}, \text{y}), \sigma) \leftarrow G_{\text{att}}.\text{resume}(\text{eid}_{\text{F}}, (\text{run}, \perp, \perp, \perp, \perp, \perp, \text{y}))$ 
   $\mathcal{G}[\text{eid}].\text{resume} \leftarrow \mathcal{G}[\text{eid}].\text{resume} \parallel (\sigma, \text{input}, (\text{computed}, \text{y}))$ 
   $\mathcal{B}[\sigma] \leftarrow (\mathcal{G}[\text{eid}].\text{install}[0], \text{eid}, \mathcal{G}[\text{eid}].\text{install}[1], (\text{computed}, \text{y}))$ 
  forward  $((\text{computed}, \text{y}), \sigma)$  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 prog_{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_{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 eid_{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 $\text{pk}_{KD}, \text{eid}_{KME}, \text{crs}$, we can safely send the encrypted sk to the corrupted party as no one can retrieve the decryption key from outside the enclave.

On message (KEYGEN, F, B) from the functionality after a call to KEYGEN , $\mathcal{S}_{\text{FESR}}$ simply produces a functional key by running the appropriate prog_{KME} procedure through G_{att} . Similarly, on receiving (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 $\text{prog}_{\text{FE}}[F]$, 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 $\text{prog}_{\text{FE}}[F]$ 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 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 $\text{prog}_{\text{FE}}[F]$. 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 appendix C.

6 Rollback and Forking Attacks

While the functionality modelled by G_{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. [59] weaken the original G_{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_{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. [44], 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_{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 [18].

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_{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 parameterised with a signature scheme and a registry to capture all platforms with a TEE, like in the original formulation.

Functionality $G_{\text{att}}^{\text{rollback}}[\Sigma, \text{reg}, \lambda]$	
State variables	Description
<code>vk</code>	Master verification key, available to enclave programs
<code>msk</code>	Master secret key, protected by the hardware
<code>$\mathcal{T} \leftarrow \emptyset$</code>	Table for installed programs
On message INITIALIZE from a party P :	
let $(\text{spk}, \text{ssk}) \leftarrow \Sigma.\text{Gen}(1^\lambda), \text{vk} \leftarrow \text{spk}, \text{msk} \leftarrow \text{ssk}$	
On message GETPK from a party P :	
return <code>vk</code>	


```

On message (INSTALL, idx, prog) from a party P where P.pid ∈ reg:
  if P is honest then
    assert idx = P.sid
    generate nonce eid ∈ {0, 1}λ, 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 ∈ reg:
  let (idx, prog, lastnode, tree) ←  $\mathcal{T}[\text{eid}, P]$ , abort if not found
  if P is honest then
    let node ← lastnode
    let mem ← access(tree, node)
    let (output, mem') ← prog(input, mem)
    let tree', child ← insertChild(tree, node, mem')
    let update  $\mathcal{T}[\text{eid}, P] = (\text{idx}, \text{prog}, \text{child}, \text{tree}')$ 
    let  $\sigma \leftarrow \Sigma.\text{Sign}(\text{msk}, (\text{idx}, \text{eid}, \text{prog}, \text{output}))$  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 quantised 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 [62] unrealistic for TEE platforms such as SGX, where an enclave is allowed direct access to a hardware-based source of randomness [7].

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, mem)
  if mem =  $\emptyset$  then
    K ← x
    Store mem ← K
    return ACK
  else if mem =  $\vec{1}$  then
    return  $\perp$ 
  else
    Store mem ←  $\vec{1}$ 
    return  $F_K(x)$ 

```

An adversary who has completed initialisation of its decryption enclave with enclave id eid_{DE} , obtained a functional key sk through the execution of keygen on eid_{KME} , and initialised a functional enclave for PRF-WRAPPER with enclave id eid_F , public key pk_{FD} and attestation σ , 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:  $((\text{computed}, \text{ACK}), \cdot) \leftarrow G_{\text{att}}^{\text{rollback}}.\text{resume}(\text{eid}_{\text{F}}, (\text{run}, \text{vk}_{\text{att}}, \sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_{\text{k}}, \text{crs}, \perp), \text{node})$ 
3: // node is the node id for a leaf for  $\text{eid}_{\text{F}}$ 's mem tree
4:  $((\text{computed}, y), \cdot) \leftarrow G_{\text{att}}^{\text{rollback}}.\text{resume}(\text{eid}_{\text{F}}, (\text{run}, \text{vk}_{\text{att}}, \sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_{\text{x}}, \text{crs}, \perp), \text{node}')$ 
5: //  $\text{node}'$  is the node id for a leaf for  $\text{eid}_{\text{F}}$ 's mem tree
6:  $((\text{computed}, y'), \cdot) \leftarrow G_{\text{att}}^{\text{rollback}}.\text{resume}(\text{eid}_{\text{F}}, (\text{run}, \text{vk}_{\text{att}}, \sigma_{\text{DE}}, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_{\text{x}'}, \text{crs}, \perp), \text{node}')$ 
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 ϵ) 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 trade-off.

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 prog_{FE} enclave is required to check in with the prog_{DE} enclave to retrieve the decryption key as part of the provision call. To enable rollback resilience, we include the counter stored by prog_{FE} as an additional parameter of this call. prog_{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, prog_{FE} increases and stores its local counter.

To achieve rollback resilience for the prog_{DE} enclave, we can rely on existing techniques in the literature, such as augmenting the enclave with asynchronous monotonic counters [14], or using protocols like LCM [18] or ROTE [44]. Formalising how these protocols can be combined with the $G_{\text{att}}^{\text{rollback}}$ functionality to achieve the fully secure G_{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 `msk` unless combining it with a leakage attack. Denial of Service is possible by creating concurrent enclaves (either DE or FE) with different public keys, and passing encrypted ciphertexts 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/obfuscuro-a-commodity-obfuscation-engine-on-intel-sgx/>.
- [5] Alibaba Cloud. TEE-based confidential computing. <https://www.alibabacloud.com/help/doc-detail/164536.htm>, 2020.
- [6] 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.
- [7] J. Aumasson and L. Merino. Sgx secure enclaves in practice: security and crypto review. *Black Hat*, 2016:10, 2016.

- [8] 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.
- [9] 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>.
- [10] 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.
- [11] 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>.
- [12] K. Baghery, 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>.
- [13] 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>.
- [14] 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>.
- [15] B. Barak, R. Canetti, J. B. Nielsen, and R. Pass. Universally composable protocols with relaxed set-up assumptions. In *45th FOCS*, pages 186–195, Rome, Italy, Oct. 17–19, 2004. IEEE Computer Society Press.
- [16] 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>.
- [17] 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.
- [18] 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>.
- [19] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. Cryptology ePrint Archive, Report 2000/067, 2000. <http://eprint.iacr.org/2000/067>.

- [20] 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.
- [21] 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.
- [22] 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.
- [23] R. Cheng, F. Zhang, J. Kos, W. He, N. Hynes, N. M. Johnson, A. Juels, A. Miller, and D. Song. Ekiden: A platform-preserving, trustworthy, and performant smart contract execution. *CoRR*, abs/1804.05141, 2018. URL <http://arxiv.org/abs/1804.05141>.
- [24] 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.
- [25] 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.
- [26] 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.
- [27] 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>.
- [28] V. Costan and S. Devadas. Intel SGX explained. Cryptology ePrint Archive, Report 2016/086, 2016. <http://eprint.iacr.org/2016/086>.
- [29] A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In J. Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 566–598, Santa Barbara, CA, USA, Aug. 19–23, 2001. Springer, Heidelberg, Germany.
- [30] 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.

- [31] B. A. Fisch, D. Vinayagamurthy, D. Boneh, and S. Gorbunov. Iron: Functional encryption using intel SGX. Cryptology ePrint Archive, Report 2016/1071, 2016. <http://eprint.iacr.org/2016/1071>.
- [32] G. Fuchsbauer, A. Plouviez, and Y. Seurin. Blind schnorr signatures and signed ElGamal encryption in the algebraic group model. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 63–95, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany.
- [33] C. Garlati and S. Pinto. A clean slate approach to Linux security RISC-V enclaves. 02 2020.
- [34] 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.
- [35] F. Gregor, W. Ozga, S. Vaucher, R. Pires, D. L. Quoc, S. Arnaudov, 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>.
- [36] V. T. Hoang, R. Reyhanitabar, P. Rogaway, and D. Vizár. Online authenticated-encryption and its nonce-reuse misuse-resistance. In R. Genaro 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.
- [37] D. Kaplan, J. Powell, and T. Woller. AMD memory encryption. *White paper*, 2016.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] 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.

- [43] 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.
- [44] S. Matetic, M. Ahmed, K. Kostiaainen, 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>.
- [45] 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.
- [46] 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.
- [47] 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.
- [48] 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.
- [49] S. Pinto and N. Santos. Demystifying Arm TrustZone: A comprehensive survey. *ACM Computing Surveys*, 51:1–36, 01 2019.
- [50] N. Porter, G. Golan, and S. Lugani. Introducing Google Cloud Confidential Computing with Confidential VMs. 2020.
- [51] J. Radmets. Technical report, Cybernetica Information Security Research Institute, 2021.
- [52] M. Russinovich. Introducing Azure confidential computing. 2017.
- [53] 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.
- [54] C.-P. Schnorr and M. Jakobsson. Security of signed ElGamal encryption. In T. Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 73–89, Kyoto, Japan, Dec. 3–7, 2000. Springer, Heidelberg, Germany.
- [55] A. W. Services. Aws nitro enclaves. <https://aws.amazon.com/ec2/nitro/nitro-enclaves/>.
- [56] 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.
- [57] 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.
- [58] 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>.

- [59] F. Tramer, 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>.
- [60] 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>.
- [61] 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.
- [62] 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.

A Primitives

We now formally define basic notation and definitions used throughout the paper.

Notation. For a bit-string x , $|x|$ denotes the length of x , and λ denotes the security parameter. For a distribution D over a set $\mathcal{X}$, $x \leftarrow D$, denotes sampling an element $x \in \mathcal{X}$, according to D , and $x \leftarrow \mathcal{X}$ denotes sampling a uniform element x from $\mathcal{X}$. “ $\approx$ ” denotes computational indistinguishability, and $\text{negl}(\lambda)$ denotes an unspecified, negligible function, such that $\text{negl}(\lambda) \leq \frac{1}{\lambda^c}$ for all $c \in \mathbb{R}$.

For an algorithm $\mathcal{A}$, using $y \leftarrow \mathcal{A}(x)$ we denote the execution of $\mathcal{A}$ on input x , receiving output y . In case $\mathcal{A}$ is randomized, y is a random variable and $\mathcal{A}(x; r)$ denotes the execution of $\mathcal{A}$ on input x with randomness r . An algorithm $\mathcal{A}$ is probabilistic polynomial-time (PPT) if $\mathcal{A}$ is randomized and for any $x, r \in \{0, 1\}^*$, the computation of $\mathcal{A}(x; r)$ terminates in a number of steps polynomial in $(|x| + |r|)$.

A.1 Public-key encryption

In the current section we define public-key encryption for chosen plaintext (CPA) and chosen ciphertext (CCA) attacks. Note that in the latter, the adversary is allowed to access the decryption oracle even after receiving the challenge ciphertext, usually referred to as CCA2 secure encryption.

The syntax of a public-key encryption scheme is defined below.

PKE syntax. A *public-key encryption* (PKE) scheme is a triple of algorithms $\text{PKE} = (\text{PGen}, \text{Enc}, \text{Dec})$ with the following syntax:

- (*Key generation*): PGen receives a security parameter λ , and outputs a fresh key pair $(\text{pk}, \text{sk}) \leftarrow \text{PGen}(1^\lambda)$.
- (*Encryption*): Enc receives a public key pk and a message m and produces a ciphertext ct .
- (*Decryption*): Dec receives a secret key sk and a ciphertext ct and outputs a message m .

In the following sections the security parameter will be implicit when calling PGen . A PKE scheme must satisfy the following correctness property.

Correctness. For any message m ,

$$\Pr[(\text{pk}, \text{sk}) \leftarrow \text{PGen}(1^\lambda); \text{ct} \leftarrow \text{Enc}(\text{pk}, m); m' \leftarrow \text{Dec}(\text{sk}, \text{ct}) : m = m'] = 1.$$

CPA security game. For any PPT adversary $\mathcal{A}$, $b \in \{0, 1\}$ and PKE scheme PKE, we consider the following CPA security game, denoted by $G_{\text{PKE}, \mathcal{A}}^{\text{cca}}(\lambda, b)$:

- $(\text{pk}, \text{sk}) \leftarrow \text{PGen}(1^\lambda)$.
- $\mathcal{A}$ receives pk and oracle access to $\mathcal{O}_{\text{enc}}(\cdot) := \text{Enc}(\text{pk}, \cdot)$.
- $\mathcal{A}$ outputs (m_0, m_1) .
- $\text{ct} \leftarrow \text{Enc}(\text{pk}, m_b)$.
- $\mathcal{A}$ receives ct and oracle access to $\mathcal{O}_{\text{enc}}(\cdot)$
- $\mathcal{A}$ outputs b' .
- Output $b' = b$.

Definition 3 (CPA security). A public-key encryption scheme PKE is CPA-secure if for all PPT adversaries $\mathcal{A}$, there exists a negligible function negl such that

$$\left| \Pr[G_{\text{PKE}, \mathcal{A}}^{\text{cpa}}(\lambda, 0) = 1] - \Pr[G_{\text{PKE}, \mathcal{A}}^{\text{cpa}}(\lambda, 1) = 1] \right| \leq \text{negl}(\lambda).$$

CCA security game. For any PPT adversary $\mathcal{A}$, $b \in \{0, 1\}$ and PKE scheme PKE, we consider the following CCA security game, denoted by $G_{\text{PKE}, \mathcal{A}}^{\text{cca}}(\lambda, b)$:

- $(\text{pk}, \text{sk}) \leftarrow \text{PGen}(1^\lambda)$.
- $\mathcal{A}$ receives pk and oracle access to $\mathcal{O}_{\text{dec}}(\cdot) := \text{Dec}(\text{sk}, \cdot)$.
- $\mathcal{A}$ outputs (m_0, m_1) .
- $\text{ct} \leftarrow \text{Enc}(\text{pk}, m_b)$.
- $\mathcal{A}$ receives ct and oracle access to $\mathcal{O}_{\text{dec}}(\cdot)$ but is not allowed to query ct .
- $\mathcal{A}$ outputs b' .
- Output $b' = b$.

Definition 4 (CCA security). A public-key encryption scheme PKE is CCA-secure if for all PPT adversaries $\mathcal{A}$, there exists a negligible function negl such that

$$\left| \Pr[G_{\text{PKE}, \mathcal{A}}^{\text{cca}}(\lambda, 0) = 1] - \Pr[G_{\text{PKE}, \mathcal{A}}^{\text{cca}}(\lambda, 1) = 1] \right| \leq \text{negl}(\lambda).$$

A.2 Signatures

In the current section we define existential unforgeability against chosen message attacks (EU-CMA).

Digital signatures syntax. A *digital signature* (DS) scheme is a triple of algorithms $\Sigma = (\text{Gen}, \text{Sign}, \text{Vrfy})$ with the following syntax:

- (*Key generation*): Gen receives a security parameter λ and outputs a fresh public and secret keypair $(\text{spk}, \text{ssk}) \leftarrow \text{Gen}(\lambda)$.
- (*Signing*): Sign receives a signing key ssk and a message m and produces a signature $\sigma \leftarrow \text{Sign}(\text{ssk}, m)$.
- (*Verification*): Vrfy receives a public key spk , a message m and a signature σ , and outputs a bit $b \leftarrow \text{Vrfy}(\text{spk}, m, \sigma)$.

In the following sections the security parameter will be implicit when calling Gen .

A DS scheme must satisfy the following correctness property.

Correctness. For any message m ,

$$\Pr[(\text{spk}, \text{ssk}) \leftarrow \text{Gen}(1^\lambda); \sigma \leftarrow \text{Sign}(\text{ssk}, m) : \text{Vrfy}(\text{spk}, m, \sigma) = 1] = 1.$$

Unforgeability game. For any PPT adversary $\mathcal{A}$ and DS scheme Σ , we consider the following security game, denoted by $G_{\Sigma, \mathcal{A}}^{\text{eu-cma}}(\lambda)$:

- $(\text{spk}, \text{ssk}) \leftarrow \text{Gen}(1^\lambda)$.
- $\mathcal{A}$ receives spk and oracle access to $\mathcal{O}_{\text{sign}}(\cdot) := \text{Sign}(\text{ssk}, \cdot)$. Let Q be the set of queries made by $\mathcal{A}$.
- $\mathcal{A}$ outputs (m, σ) .
- If $\text{Vrfy}(\text{spk}, m, \sigma) = 1$ and $m \notin Q$, output 1, otherwise, output 0.

Definition 5 (Unforgeability). A DS scheme Σ is EU-CMA-secure, if for all PPT adversaries $\mathcal{A}$, there exists a negligible function negl such that

$$\Pr[G_{\Sigma, \mathcal{A}}^{\text{eu-cma}}(\lambda) = 1] \leq \text{negl}(\lambda).$$

A.3 NIZK

Definition 6 (Robust NIZK [29]). Let $\mathcal{W}$ be a witness relation for a language $\mathcal{L} \in \text{NP}$. A non-interactive zero-knowledge argument system for $\mathcal{W}$ is a vector of algorithms $(\mathcal{G}, \mathcal{P}, \mathcal{V}, \mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2))$, satisfying the following properties:

- **Completeness.** For any $y \in \mathcal{L}$ and any w such that $(y, w) \in \mathcal{W}$,

$$\Pr \left[\mathcal{V}(\text{crs}, y, \pi) = 1 \mid \begin{array}{l} (\text{crs}) \leftarrow \mathcal{G}(1^\lambda) \\ \pi \leftarrow \mathcal{P}(y, w, \text{crs}) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

- **Zero-knowledge.** For all non-uniform PPT adversaries $\mathcal{A}$, we have

$$\Pr[\mathcal{A}^{\mathcal{P}(\text{crs}, \cdot, \cdot)}(\text{crs}) = 1 \mid \text{crs} \leftarrow \mathcal{G}(1^\lambda)] \approx \Pr[\mathcal{A}^{\mathcal{S}(\hat{\text{crs}}, \tau, \cdot, \cdot)}(\hat{\text{crs}}) = 1 \mid (\hat{\text{crs}}, \tau) \leftarrow \mathcal{S}_1(1^\lambda)],$$

where $\mathcal{S}(\hat{\text{crs}}, \tau, y, w) = \mathcal{S}_2(\hat{\text{crs}}, \tau, y)$ if $(y, w) \in \mathcal{W}$, otherwise outputs $\perp$.

- **Simulation sound extractability.** There exists a PPT algorithm $\mathcal{E}$, such that for all PPT algorithms $\mathcal{A}$, we have

$$\Pr \left[\begin{array}{l} \mathcal{V}(\hat{\text{crs}}, y, \pi) = 1 \wedge \\ w \notin \mathcal{W}(y) \quad \wedge \\ (y, \pi) \notin Q \end{array} \middle| \begin{array}{l} (\hat{\text{crs}}, \tau) \leftarrow \mathcal{S}_1(1^\lambda) \\ (y, \pi) \leftarrow \mathcal{A}^{\mathcal{S}_2(\hat{\text{crs}}, \tau, \cdot)}(\hat{\text{crs}}) \\ w \leftarrow \mathcal{E}(\tau, y, \pi) \end{array} \right] \leq \text{negl}(\lambda),$$

where Q denotes the simulation queries and answers (y_i, π_i) , produced by the interaction between $\mathcal{A}$ and $\mathcal{S}_2$.

A.4 Additional Functionalities

Secure Channel Functionality $\mathcal{SC}_R^S$ is the secure channel between source S and receiver R .

Functionality $\mathcal{SC}_R^S$
On message (SEND, m) from S: let $M \leftarrow m \parallel \text{length}(m)$ return (SENT, m) to R On message GETMSGs from $\mathcal{A}$:   output M on the backdoor tape

This is quite a strong functionality. The adversary is not activated upon sending, but can later on request the length of messages. This is in line with the modelling of [45] (and [11]). Like them, we focus on the security guarantees against a corrupted party B and this channel simplifies the simulation of network interactions. We do, however, not see any fundamental obstacles to adopting a more realistic secure channel functionality such as that of [19].

Communication notation. **send** (MSG, m) **to** $\mathcal{SC}_R^S$ and **receive** (MSG, m') denotes the secure transmission of a message (MSG, m) from S to R . After sending the message S waits for the reply (MSG, m') over $\mathcal{SC}_S^R$. When the identity of the receiver or the sender is obvious from the context, we might use shorthands $\mathcal{SC}^S$ or $\mathcal{SC}_R$ respectively.

We write **send** (MSG, m) **to** $\mathcal{A}, p$ to denote a delayed (insecure) output to p . $\mathcal{A}$ is first informed about (MSG, m) and can then determine when and if to deliver the message to p .

Repository Functionality $\mathcal{REP}^{W, \bar{R}}$ is based on the repository functionality of [45], parametrised by a writing party W and a set of reading parties $\bar{R}$

Functionality $\mathcal{REP}^{W, \bar{R}}$
<i>On message (WRITE, x) from party W:</i>   let $h \leftarrow \text{getHandle}$ let $M[h] \leftarrow x$ return h <i>On message (READ, h) from party $r \in \bar{R}$:</i>   return $M[h]$

Common Reference String Functionality The $\mathcal{CRS}[D]$ functionality, based on the presentation of [15], is parametrised by a distribution D .

Functionality $\mathcal{CRS}[D]$
<i>On message GET from a party P:</i>   if $\text{crs} = \perp$ then   let $\text{crs} \leftarrow D$ return crs to P

This functionality has a simple interface: on a request from any protocol party (or the adversary), a CRS string sampled from distribution D is returned. Once the CRS has been sampled, an instance of $\mathcal{CRS}^G[D]$ will always return the same string. While this functionality is insufficient to realise global protocols in the GUC setting, where it is subsumed by the augmented CRS functionality [20], our usage of the functionality in the following sections is limited to local protocols. As a result, we are not concerned with the type of deniability attacks that are addressed by the new functionality.

B UC definitions

We provide the definitions for the essential constructs of Universal Composability necessary for understanding this paper. We refer to the original paper [19] and the UCGS paper [10] for a more comprehensive treatment,

Computational model In the UC framework, protocols are executed by interactive Turing machines (ITM). An ITM is a universal machine with additional tapes for encoding its identity, activation status, and incoming or outgoing messages to other machines (including a dedicated tape, the backdoor tape, that an adversary can use to communicate with a machine that corresponds to a

corrupted party). The state (configuration) of a machine instance consists of the content of all its tapes. The state of a system of machine instances consists of a collection of such instances with distinct read-only identity tapes (known as the *extended identity* of an instance); it contains the ITM's code, a *session identifier*, sid and a *party identifier*, pid . When instantiated, a machine is given an “import” parameter, which determines how many operations the instance is allowed to execute before it terminates; machines will typically run a number of operations bounded by a polynomial of the import.

The execution of a system of instances is managed by a control function. Given an initial instance, the control function determines what ITIs can be instantiated. The execution begins with the activation of the initial ITI with some input, and ends when it reaches a halting configuration. During the execution, the initial instance might invoke new instances through its external-write tape, if allowed by the control function; these new instances may also instantiate new instances. If an instance M' accepts some input from instance M , or sends an output message to M which is not rejected, M' is known as a subroutine of M . An extended system allows the control function to modify instances' external-write instructions; this mechanism can be used to modify the system such that, given an instance M that attempts to instantiate subroutine M' , a different subroutine M'' gets instantiated, unbeknownst to M . To allow representing dynamic protocols where the number of participants and the code they run might change over their execution, a protocol session is defined as a set of instances running the same protocol (the main parties of the protocol session, denoted by a shared session id and code within their identity tapes), having been invoked with the intention of interacting with each other with a common purpose. The *extended session* of a protocol is the set comprising of the main parties of the protocol session, the subroutines of parties already in the extended session, and machines invoked by the sub-parties of the extended session (a sub-party is an ITI that is in the extended session, but not as a main party).

To address modelling specific behaviour, a structured protocol is a set of two (or more) nested ITMs, a *body* which contains the actual protocol code, and a *shell* to address modelling functions. Shell and protocol share their externally writable and outgoing tapes, and the shell has write access to all the body's tapes (except for the identity tape).

UC-emulation A proof of security for a protocol π in the UC setting considers an evolving system of ITM instances, where the initial instance $\mathcal{Z}$ is known as the environment, and the instance with $\text{pid } \diamond$ is known as the adversary $\mathcal{A}$. The control function allows $\mathcal{Z}$ to pass input to $\mathcal{A}$ and any ITI with code π and a fixed $\text{sid } s$, without having to use its extended identity as a sender (the identities used by $\mathcal{Z}$ are known as *external identities*). $\mathcal{A}$ can only communicate with existing instances and only using their backdoor tapes. All other ITIs include their extended identities in outgoing messages, and can pass input and outputs to ITIs besides $\mathcal{Z}$ or $\mathcal{A}$; but if the target to a message does not exist, it is passed

as a subroutine-output to $\mathcal{Z}$. A protocol is subroutine respecting if only the main ITIs of a session can communicate with instances in the extended session.

A γ -subroutine respecting protocol relaxes the conditions of subroutine respecting protocols by also allowing external calls to parties that belong to multiple sessions of γ and to machines in those extended sessions.

γ is a ϕ -regular setup if it has no subsidiary with code ϕ , and if γ does not instantiate a new instance of ϕ as part of the output of a subroutine.

While the above control function allows $\mathcal{Z}$ to send messages to π 's main parties with identity set to arbitrary ITIs outside of the protocol (external identities), a ξ -identity-bounded environment is allowed to only use identities satisfying ξ , a polynomial time boolean predicate on the current system configuration. As the environment is responsible with setting their import, modelling of realistic conditions requires the adversary's computational resources to be at least as large as the protocol it is trying to attack (otherwise, the adversary might not run for the entire protocol execution). Thus, a *balanced* environment provides the adversary with import at least equal to the sum of all instances in the protocol.

Definition 7 (UC emulation). *Given two PPT protocols π, ϕ and some predicate ξ , we say that π UC-emulates ϕ with respect to ξ -identity bound environments (or π ξ -UC-emulates ϕ , if ξ allows all external identities this simplifies to π UC-emulates ϕ) if for any balanced ξ -identity-bounded environment and any PPT adversary, there exists a PPT simulator $\mathcal{S}$ such that the probability distribution for the outputs of systems $(\phi, \mathcal{S}, \mathcal{Z})$ and $(\pi, \mathcal{A}, \mathcal{Z})$ are indistinguishable.*

A common variant of the emulation theorem replaces the universal quantifier over adversaries with a single “dummy adversary” $\mathcal{D}$, who simply transfers messages through parties' backdoor tapes on behalf of the environment, who now has full control over the adversarial protocol.

In practice, in a proof of security, we take ϕ to be an ideal functionality, a protocol that acts as a trusted third party that perfectly executes the functionality it expresses, without accepting corrupt messages from the adversary. We say that protocol $\text{IDEAL}_{\mathcal{F}}$ is the protocol that comprises of the ideal functionality ITI $\mathcal{F}$, which cannot be corrupted by adversarial calls, and a set of dummy parties that reproduce the structure of a realistic distributed protocol, but in practice forward all inputs and receive outputs from $\mathcal{F}$. We say that if some protocol π ξ -UC-emulates $\text{IDEAL}_{\mathcal{F}}$ and is ξ -subroutine respecting, then π ξ -UC-realises $\mathcal{F}$.

A protocol ρ is said to be (π, ϕ, ξ) -compliant if ITIs in ρ that make subroutine calls to ϕ or π satisfy ξ , and make no calls to protocol instances of π and ϕ that belong to the same session. A protocol is subroutine exposing if it allows the adversary to find out if a certain ITI belongs to its extended session.

Definition 8 (Universal Composition). *Given PPT protocols π, ϕ, ρ and predicate ξ , if π, ϕ are both subroutine respecting and subroutine exposing, ρ is (π, ϕ, ξ) -compliant and π ξ -UC-emulates ϕ , the protocol $\rho^{\phi \rightarrow \pi}$ UC-emulates ρ , where $\rho^{\phi \rightarrow \pi}$ is the protocol where all of ρ 's calls to subroutine ϕ are replaced with calls to subroutine π .*

C Proof of security Hybrids

We proceed to show that the real world execution with respect to the Steel protocol and dummy adversary is indistinguishable from the ideal execution w.r.t. the ideal functionality FESR and the simulator described in Section 5. Let Hybrid 0 be the real world execution w.r.t. to the Steel protocol and dummy adversary.

Hybrid 1 We define Hybrid 1 to be identical to 0, but as translated to the ideal world; that is, we replace protocol Steel with a protocol consisting of dummy parties A, B, C and ideal functionality FE; the dummy adversary machine is replaced by a simulator $\mathcal{S}'_{\text{FESR}}$, who emulates Steel's execution by emulating all operations that would normally be run by the honest parties. Messages are now sent to the ideal functionality, but the responses are ignored and the original protocol is perfectly executed by $\mathcal{S}'_{\text{FESR}}$. Hybrid 0 and 1 are indistinguishable by UC-emulation, and because their behaviour is equivalent due to simulator $\mathcal{S}'_{\text{FESR}}$. Since the ideal functionality is not visible to the environment, it is harmless for the dummy parties to send messages to it.

Lemma 1. H_0 is identical to H_1 .

Hybrid 2 diverges from Hybrid 1 in that all signature verifications obtained with the attestation public key vk_{att} are replaced by the process of storing all outgoing messages from the G_{att} functionality in a map data structure, and checking on verification that the message and corresponding signature were correctly recorded. The behaviour of the two hybrids is equivalent, as long as the adversary in Hybrid 2 is not able to provide a signature such that the verification checks are successful, even if the messages were not recorded as coming through the G_{att} functionality. Assuming the unforgeability property of G_{att} 's signature scheme is satisfied, Hybrid 1 is then indistinguishable to Hybrid 2.

Lemma 2. If Σ is EU-CMA secure then $H_1 \approx H_2$, over the randomness used by all parties in H_1, H_2 .

Proof. To prove the indistinguishability of the two hybrids, we show how an environment $\mathcal{Z}$ that breaks attestation can be used to build an adversary $\mathbb{R}$ that breaks unforgeability of the signature scheme.

For attestation to break, the adversary needs to produce signature σ such that $\Sigma.\text{Vrfy}(\text{vk}_{\text{att}}, \sigma, (\text{sid}, \text{eid}_X, \text{prog}_X, \text{out})) = 1$ for some program prog_X and output out such that no execution of prog_X under sid produced out with attestation σ

We now describe adversary $\mathbb{R}$, whose goal is to break signature game $\mathcal{C}^\Sigma$ by submitting a forged signature.

Reduction $\mathbb{R}^{\mathcal{C}^\Sigma, \mathcal{Z}}$

State variables	Description
spk	Public key for game $\mathcal{C}^\Sigma$
$\mathcal{H} \leftarrow \{\}$	Repository of message requests
$\mathcal{G} \leftarrow \{\}$	Collects all enclave programs registered by $\mathcal{Z}$
$\text{state} \leftarrow \{\}$	Dictionary to hold the state for each function
<i>On message</i> (SETUP, P):	
if $\text{mpk} = \perp$ then $(\text{mpk}, \text{msk}) \leftarrow \text{PKE.PGen}(1^\lambda)$ send GET to CRS and receive (CRS, crs)	
if $P = \text{A}$ then send (SETUP, mpk) to $\mathcal{SC}_\text{A}$	
else if $P = \text{B}$ then send (SETUP, mpk, eid_{KME}) to $\mathcal{SC}_\text{B}$ and receive (PROVISION, σ , eid_{DE} , pk_{KD}) $\text{m} \leftarrow (\text{sid}, \text{eid}_{\text{DE}}, \text{prog}_{\text{DE}}, \text{pk})$ if $\Sigma.\text{Vrfy}(\text{spk}, \text{m}, \sigma) \wedge \sigma \notin \mathcal{C}^\Sigma.Q$ then output (m, σ) to $\mathcal{C}^\Sigma$ $\text{c} \leftarrow \text{PKE.Enc}(\text{pk}, \text{msk})$ $\sigma_{\text{KME}} \leftarrow \mathcal{C}^\Sigma.\mathcal{O}_{\text{sign}}(\text{sid}, \text{eid}, \text{prog}_{\text{KME}}, \text{c})$ send (PROVISION, c , σ_{KME}) to $\mathcal{SC}_\text{B}$	
<i>On message</i> (KEYGEN, F, B):	
if $F \notin \mathcal{F} \vee \text{mpk} = \perp$ then return $\perp$ $\sigma \leftarrow \mathcal{C}^\Sigma.\mathcal{O}_{\text{sign}}(\text{keygen}, F)$ $\text{state} \leftarrow \vec{0}$ send (KEYGEN, (F, σ)) to $\mathcal{SC}_\text{B}$	
<i>On message</i> (ENCRYPT, m):	
if $\text{mpk} = \perp \vee m \notin \mathcal{X}$ then return $\perp$ $\text{h} \leftarrow \text{getHandle}; \mathcal{H}[\text{h}] \leftarrow m$ return (ENCRYPTED, h)	
<i>On message</i> (READ, h) from party B to $\mathcal{REP}$:	
$\text{m} \leftarrow \mathcal{H}[\text{h}]$ $\text{r} \leftarrow \{0, 1\}^\lambda; \text{ct} \leftarrow \text{PKE.Enc}(\text{mpk}, \text{m}; \text{r})$ $\pi \leftarrow \mathcal{P}((\text{mpk}, \text{ct}), (\text{m}, \text{r}), \text{crs})$ return (ct, π)	
<i>On message</i> GETPK from B to G_{att} :	
return spk	
<i>On message</i> (INSTALL, idx, prog) from B to G_{att} :	
$\text{eid} \leftarrow \{0, 1\}^\lambda$ $\mathcal{G}[\text{eid}] \leftarrow (\text{idx}, \text{prog})$ send eid to B	
<i>On message</i> (RESUME, eid, input) from B to G_{att} :	
if $\text{input}[0] = \text{init-setup}$ then if $\mathcal{G}[\text{eid}][1] \neq \text{prog}_{\text{DE}}$ then abort $(\text{init-setup}, \text{eid}_{\text{KME}}, \text{crs}) \leftarrow \text{input}$ $(\text{pk}, \text{sk}) \leftarrow \text{PKE.PGen}(1^\lambda)$ output $\leftarrow (\text{pk}, \text{eid}_{\text{KME}}, \text{crs})$	


```

 $\sigma \leftarrow \mathcal{C}^\Sigma.\mathcal{O}_{\text{sign}}(\text{sid}, \text{eid}, \text{prog}_{\text{DE}}, \text{output})$ 
forward (output,  $\sigma$ ) to B
else if input[0] = complete-setup then
  if  $\mathcal{G}[\text{eid}][1] \neq \text{prog}_{\text{DE}}$  then abort
  (complete-setup,  $\text{ct}_{\text{key}}, \sigma_{\text{KME}}$ )  $\leftarrow$  input
   $m \leftarrow (\mathcal{G}[\text{eid}][0], \text{eid}_{\text{KME}}, \text{prog}_{\text{KME}}, \text{ct}_{\text{key}})$ 
  if  $\Sigma.\text{Vrfy}(\text{spk}, m, \sigma_{\text{KME}}) \wedge \sigma_{\text{KME}} \notin \mathcal{C}^\Sigma.Q$  then
    output ( $m, \sigma_{\text{KME}}$ ) to  $\mathcal{C}^\Sigma$ 
else if input[0] = provision then
  if  $\mathcal{G}[\text{eid}][1] \neq \text{prog}_{\text{DE}}$  then abort
  (provision,  $\sigma, \text{eid}, \text{pk}_{\text{FD}}, \text{sk}_{\text{F}}, \text{F}$ )  $\leftarrow$  input
   $m_1 \leftarrow (\mathcal{G}[\text{eid}][0], \text{eid}_{\text{KME}}, \text{prog}_{\text{KME}}, (\text{keygen}, \text{F}))$ 
  if  $\Sigma.\text{Vrfy}(\text{spk}, m_1, \text{sk}_{\text{F}}) \wedge \text{sk}_{\text{F}} \notin \mathcal{C}^\Sigma.Q$  then
    output ( $m_1, \sigma$ ) to  $\mathcal{C}^\Sigma$ 
   $m_2 \leftarrow (\mathcal{G}[\text{eid}][0], \text{eid}, \text{prog}_{\text{FE}}[\text{F}], \text{pk}_{\text{FD}})$ 
  if  $\Sigma.\text{Vrfy}(\text{spk}, m_2, \sigma) \wedge \sigma \notin \mathcal{C}^\Sigma.Q$  then
    output ( $m_2, \sigma$ ) to  $\mathcal{C}^\Sigma$ 
   $\text{ct} \leftarrow \text{PKE}.\text{Enc}(\text{pk}_{\text{FD}}, \text{msk}), \text{output} \leftarrow (\text{ct}, \text{crs})$ 
   $\sigma_{\text{DE}} \leftarrow \mathcal{C}^\Sigma.\mathcal{O}_{\text{sign}}(\text{sid}, \text{eid}, \text{prog}_{\text{DE}}, \text{output})$ 
  forward (output,  $\sigma_{\text{DE}}$ ) to B
else if input[0] = run then
  if  $\mathcal{G}[\text{eid}][1] \neq \text{prog}_{\text{FE}[\text{F}]}$  then abort
  (run,  $\sigma, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_{\text{msg}}, \text{crs}, y'$ )  $\leftarrow$  input
  if  $y' \neq \perp$  then
    out  $\leftarrow y'$ 
  else
     $m \leftarrow (\mathcal{G}[\text{eid}][0], \text{eid}_{\text{DE}}, \text{prog}_{\text{DE}}, \text{ct}_{\text{key}}, \text{crs})$ 
    if  $\Sigma.\text{Vrfy}(\text{spk}, m, \sigma) \wedge \sigma \notin \mathcal{C}^\Sigma.Q$  then
      output ( $m, \sigma$ ) to  $\mathcal{C}^\Sigma$ 
    if  $N.\mathcal{V}((\text{mpk}, \text{ct}), \pi, \text{crs}) = 0$  then return  $\perp$ 
    mem  $\leftarrow$  state
    (output, mem')  $\leftarrow F(\text{PKE}.\text{Dec}(\text{msk}, \text{ct}_{\text{msg}}), \text{mem})$ 
   $\sigma_{\text{FE}} \leftarrow \mathcal{C}^\Sigma.\mathcal{O}_{\text{sign}}(\text{sid}, \text{eid}, \text{prog}_{\text{DE}}, \text{out})$ 
  forward (out,  $\sigma_{\text{FE}}$ ) to B

```

Hybrid 3 Let Hybrid 3 replace the output of calls to the provision procedure for enclave KME with new value (ct', σ) , where $\text{ct}' \leftarrow \text{PKE}.\text{Enc}(\text{pk}_{\text{KD}}, 0^{|\text{sk}|})$ for the legitimate $\text{pk}_{\text{KD}}, \text{sk}$ held within the enclave, and σ is a valid attestation signature for an execution that produces ct' . If the PKE scheme internal to the KME program is CCA-secure, the two hybrids are indistinguishable to an attacker. Let Hybrid 3.1 replace the return value of calls to procedure provision on enclave DE with (ct', σ) , with $\text{ct}' \leftarrow \text{PKE}.\text{Enc}(\text{pk}_{\text{FD}}, 0^{|\text{sk}|})$ and σ being a valid attestation signature on the produced output. Similarly, this hybrid is indistinguishable to the previous if PKE provides CCA security.

Below, we prove the following to lemmas, via reductions to CCA security of the encryption scheme. The two reductions are quite similar and depicted below.

Lemma 3. *If PKE is CCA secure, then $H_2 \approx H_3$, over the randomness used by H_2, H_3 .*

Lemma 4. *If PKE is CCA secure, then $H_3 \approx H_{3.1}$, over the randomness used by those experiments.*

Proof. We use an adversary $\mathcal{Z}$ who can distinguish between the two hybrids H_2 and H_3 to construct an adversary $\mathbb{R}$ with the goal of breaking the CCA-security game challenger $\mathcal{C}^{\text{PKE}}$. The challenge encryption ct replaces the encryption of pk_{KD} , the public key used to securely transfer the master secret key between the prog_{KME} and prog_{DE} enclaves. $\mathcal{Z}$ also instantiates a signature scheme Σ to reproduce the attestation role of G_{att}

Note: for this and all following reductions, we follow the convention that, for all subroutines (or parts of subroutines) not explicitly defined in the current reduction, the same code as the previous reduction applies. Furthermore, any calls to the challenge game is replaced with the corresponding primitive.

Reduction $\mathbb{R}^{\mathcal{Z}, \mathcal{C}^{\text{PKE}}}$

State variables	Description
pk_{KD}	Public key for game $\mathcal{C}^{\text{PKE}}$
$\mathcal{H} \leftarrow \{\}$	Repository of message requests
$\mathcal{G} \leftarrow \{\}$	Collects all enclave programs registered by $\mathcal{Z}$
$\text{state} \leftarrow \{\}$	Dictionary to hold the state for each function

On message (SETUP, P):

```

if  $\text{mpk} = \perp$  then
   $(\text{mpk}, \text{msk}) \leftarrow \text{PKE.PGen}(1^\lambda)$ 
   $(\text{spk}, \text{ssk}) \leftarrow \Sigma.\text{Gen}(1^\lambda)$ 
  send GET to CRS and receive (CRS, crs)
if  $P = A$  then
  send (SETUP, mpk) to  $SC_A$ 
else if  $P = B$  then
  send (SETUP, mpk,  $\text{eid}_{\text{KME}}$ ) to  $SC_B$  and receive (PROVISION,  $\sigma$ ,  $\text{eid}_{\text{DE}}$ ,  $\text{pk}_{KD}$ )
   $m_0 \leftarrow \text{msk}; m_1 \leftarrow 0^{|\text{msk}|}$ 
  send (CHALLENGE,  $m_0, m_1$ ) to  $\mathcal{C}^{\text{PKE}}$  and receive ct
   $\sigma_{\text{sk}} \leftarrow \Sigma.\text{Sign}(\text{ssk}, (\text{idx}, \text{eid}_{\text{KME}}, \text{prog}_{\text{KME}}, \text{ct}))$ 
  send (PROVISION, ct,  $\sigma_{\text{sk}}$ ) to  $SC_B$ 

```

On message (RESUME, eid , input) from B to G_{att} :

```

if  $\text{input}[0] = \text{init-setup}$  then
  if  $\mathcal{G}[\text{eid}][1] \neq \text{prog}_{\text{DE}}$  then abort
   $(\text{init-setup}, \text{eid}_{\text{KME}}, \text{crs}) \leftarrow \text{input}$ 
   $\text{output} \leftarrow (\text{pk}_{KD}, \text{eid}_{\text{KME}}, \text{crs})$ 
  forward ( $\text{output}, \Sigma.\text{Sign}(\text{ssk}, (\text{sid}, \text{eid}, \text{prog}_{\text{KME}}, \text{output}))$ )
else
  ...

```

On message (Dec, k, c) from $\mathcal{Z}$ to PKE:

```

if  $k = sk_{KD}$  then
  return  $\mathcal{C}^{PKE}.Dec(sk_{KD}, c)$ 
else
  return  $PKE.Dec(k, c)$ 
On message (OUTPUT,  $b$ ) from  $\mathcal{Z}$ :
  output  $b$ 

```

Let $\mathcal{Z}$ be an adversary that distinguishes between H_3 and $H_{3.1}$; we construct an adversary $\mathbb{R}$ which calls onto $\mathcal{Z}$ to win the CCA game, by serving the challenge ciphertext to $\mathcal{Z}$ as pk_{FD} , the encryption key between enclaves running $prog_{DE}$ and $prog_{FE}$.

Reduction $\mathbb{R}^{\mathcal{Z}, \mathcal{C}^{PKE}}$

State variables	Description
pk_{FD}	Public key for game $\mathcal{C}^{PKE}$
$\mathcal{H} \leftarrow \{\}$	Repository of message requests
$\mathcal{G} \leftarrow \{\}$	Collects all enclave programs registered by $\mathcal{Z}$

On message (SETUP, P):

```

if  $mpk = \perp$  then
   $(mpk, msk) \leftarrow PKE.PGen(1^\lambda)$ 
   $(spk, ssk) \leftarrow \Sigma.Gen(1^\lambda)$ 
  send GET to  $\mathcal{CRS}$  and receive (CRS, crs)
if  $P = A$  then
  send (SETUP, mpk) to  $\mathcal{SC}_A$ 
else if  $P = B$  then
  send (SETUP, mpk,  $eid_{KME}$ ) to  $\mathcal{SC}_B$  and receive (PROVISION,  $\sigma$ ,  $eid_{DE}$ ,  $pk_{KD}$ )
   $ct \leftarrow PKE.Enc(pk_{KD}, msk)$ 
   $\sigma_{sk} \leftarrow \Sigma.Sign(ssk, (idx, eid_{KME}, prog_{KME}, ct))$ 
  send (PROVISION,  $ct, \sigma_{sk}$ ) to  $\mathcal{SC}_B$ 
On message (RESUME,  $eid$ , input) from B to  $G_{att}$ :
  if input[0] = init-setup then
    if  $\mathcal{G}[eid][1] \neq prog_{DE}$  then abort
     $(init-setup, eid_{KME}, crs) \leftarrow input$ 
     $(pk, sk) \leftarrow PKE.PGen(1^\lambda)$ 
    output  $\leftarrow (pk, eid_{KME}, crs)$ 
     $\sigma \leftarrow \Sigma.Sign(ssk, (sid, eid, prog_{DE}, output))$ 
    forward (output,  $\sigma$ ) to B
  else if input[0] = provision then
    if  $\mathcal{G}[eid] \neq prog_{DE}$  then abort
     $m_0 \leftarrow msk; m_1 \leftarrow 0^{|msk|}$ 
    send (CHALLENGE,  $m_0, m_1$ ) to  $\mathcal{C}^{PKE}$  and receive ct
    output  $\leftarrow (ct, crs)$ 
    forward (output,  $\Sigma.Sign(ssk, (sid, eid, prog_{DE}, output))$ ) to B
  else
    ...

```

Hybrid 4 This hybrid differs from $H_{3.1}$ in how the Proof of Plaintext Knowledge is computed for a message encryption. Namely, instead of making calls to the honest prover, it simply creates simulated proofs using the trapdoor. The reduction defined below receives oracle access to either $\mathcal{P}$ or $\mathcal{S}_2$. Therefore, by distinguishing between the two experiments, one can break the zero-knowledge property of the NIZK scheme $\mathbf{N}$. More formally, we prove the following lemma.

Lemma 5. *Assuming the zero-knowledge property of $\mathbf{N}$, $H_{3.1} \approx H_4$, over the randomness used by those experiments.*

Proof. Given adversary $\mathcal{Z}$, which is capable of distinguishing between the two hybrids, we define an adversary $\mathbb{R}$, with the goal of breaking the zero-knowledge game $\mathcal{C}^{\mathbf{N}}$. The reduction against the zero-knowledge property of $\mathbf{N}$ is defined below.

Reduction $\mathbb{R}^{\mathcal{Z}, \mathcal{C}^{\mathbf{N}}}$	
State variables	Description
$\mathcal{H} \leftarrow \{\}$	Repository of message requests
$\mathcal{G} \leftarrow \{\}$	Collects all enclave programs registered by $\mathcal{Z}$
$\mathcal{J} \leftarrow \{\}$	Storage of $\mathbf{N}$ proofs
$\text{state} \leftarrow \{\}$	Dictionary to hold the state for each function
$(\text{crs}, \tau) \leftarrow \mathbf{N}.\mathcal{S}_1$	Simulated reference string and trapdoor
<i>On message (SETUP, P):</i>	
if $\text{mpk} = \perp$ then	
$(\text{mpk}, \text{msk}) \leftarrow \text{PKE.PGen}(1^\lambda)$	
$(\text{spk}, \text{ssk}) \leftarrow \Sigma.\text{Gen}(1^\lambda)$	
...	
<i>On message (READ, h) from party $\mathbf{B}$ to $\mathcal{R}\mathcal{E}\mathcal{P}$:</i>	
$m \leftarrow \mathcal{H}[h]$; send GET to CRS and receive (CRS, crs)	
$r \leftarrow \{0, 1\}^\lambda$; $\text{ct} \leftarrow \text{PKE.Enc}(\text{mpk}, m; r)$	
send (CHALLENGE, (mpk, ct), (r, m)) to $\mathcal{C}^{\mathbf{N}}$ and receive π	
$\mathcal{J}[\text{ct}_{\text{msg}}] \leftarrow \pi$	
return (ct, π)	
<i>On message ($\mathcal{P}$, input) from $\mathcal{Z}$ to $\mathbf{N}$:</i>	
send (CHALLENGE, input) to $\mathcal{C}^{\mathbf{N}}$ and receive π	
return π	
<i>On message (OUTPUT, b) from $\mathcal{Z}$:</i>	
output b	

Hybrid 5 This hybrid is identical to H_4 , but instead of executing the decryption of honestly generated ciphertexts, the decryptor enclave executes the extractor for the NIZK scheme to obtain the original plaintext message. This value is then encrypted by sending it to the ideal functionality, which stores it in its internal repository.

Lemma 6. Assuming the extractability property of $\mathbf{N}$, $H_4 \approx H_5$, over the randomness used by those experiments.

Proof. Given an adversary $\mathcal{Z}$ capable of distinguishing between the two hybrids, we define an adversary $\mathbb{R}$ with the goal of breaking extractability game $\mathcal{C}^{\mathbf{N}}$.

Reduction $\mathbb{R}^{\mathcal{Z}, \mathcal{C}^{\mathbf{N}}}$	
State variables	Description
$\mathcal{H} \leftarrow \{\}$	Repository of message requests
$\mathcal{J} \leftarrow \{\}$	Storage of $\mathbf{N}$ proofs
$\mathcal{G} \leftarrow \{\}$	Collects all enclave programs registered by $\mathcal{Z}$
$\text{state} \leftarrow \{\}$	Dictionary to hold the state for each function
$(\text{crs}, \tau) \leftarrow \mathbf{N}.\mathcal{S}_1$	Simulated reference string and trapdoor
On message (RESUME, eid, input) from $\mathbf{B}$ to G_{att} :	
if input[0] = run then if $\mathcal{G}[\text{eid}][1] \neq \text{prog}_{\text{FE}[\mathbf{F}]}$ then abort $(\text{run}, \sigma, \text{eid}_{\text{DE}}, \text{ct}_{\text{key}}, \text{ct}_{\text{msg}}, \text{crs}, y') \leftarrow \text{input}$ if $y' \neq \perp$ then $\text{out} \leftarrow y'$ else if $\Sigma.\text{Vrfy}(\text{spk}, (\mathcal{G}[\text{eid}][0], \text{eid}_{\text{DE}}, \text{prog}_{\text{DE}}, \text{ct}_{\text{key}}, \text{crs}), \sigma) = 0$ then return $\perp$ if $\mathcal{J}[\text{ct}_{\text{msg}}] = \perp$ then $(\text{ct}, \pi) \leftarrow \text{ct}_{\text{msg}}$ if $\pi \in \mathcal{J}[(\cdot, \pi)]$ then send (DECRYPT, $F, \perp$) to $\mathbf{B}$ and abort $(m, r) \leftarrow \mathcal{E}(\tau, (\text{mpk}, \text{ct}), \pi)$ if $\mathbf{N}.\mathcal{V}((\text{mpk}, \text{ct}), \pi, \text{crs}) = 0 \vee \text{PKE}.\text{Enc}(\text{mpk}, m; r) \neq \text{ct} \vee \text{ct}_{\text{msg}} \notin \mathcal{C}^{\mathbf{N}}.Q$ then output 1 to $\mathcal{C}^{\mathbf{N}}$ $\text{mem} \leftarrow \text{state}$ $(\text{output}, \text{mem}') \leftarrow F(m, \text{mem}); \text{out} \leftarrow \text{output}$ $\sigma_{\text{FE}} \leftarrow \Sigma.\text{Sign}(\text{sid}, \text{eid}, \text{prog}_{\text{DE}}, \text{out})$ forward (out, σ_{FE}) to $\mathbf{B}$ else ...	

Hybrid 6 This hybrid diverges from H_5 by replacing $\mathcal{R}\mathcal{E}\mathcal{P}$'s copy of any message encrypted by $\mathbf{A}$ with an encryption of a string of zeros with the same length as the original plaintext message. Decryption is handled through the FESR functionality. The two hybrids can be distinguished by an adversary who can tell apart the two encrypted cyphertext, by winning the CPA security game.

Lemma 7. Assuming CPA security, $H_5 \approx_{\epsilon_{\text{cca}}} H_6$, over the randomness used by those experiments.

Proof. Given adversary $\mathcal{Z}$ who can distinguish between H_6 and H_5 , we construct an adversary that can break CPA-security game $\mathcal{C}^{\text{PKE}}$

Reduction $\mathbb{R}^{\mathcal{Z}, \mathcal{C}^{\text{PKE}}}$

State variables	Description
pk	Public key for game $\mathcal{C}^{\text{PKE}}$
$\mathcal{G} \leftarrow \{\}$	Collects all enclave programs registered by $\mathcal{Z}$
$\mathcal{H} \leftarrow \{\}$	Repository of message requests

On message (SETUP, P):

if $\text{mpk} = \perp$ **then**
 $\text{mpk} \leftarrow \text{pk}$
 $(\text{spk}, \text{ssk}) \leftarrow \Sigma.\text{Gen}(1^\lambda)$
send GET **to** CRS **and receive** (CRS, crs)

...

On message (READ, h) from party B to $\mathcal{R}\mathcal{E}\mathcal{P}$:

$m_0 \leftarrow \mathcal{H}[h]; m_1 \leftarrow 0^{|m_0|}$
send (CHALLENGE, m_0, m_1) **to** $\mathcal{C}^{\text{PKE}}$ **and receive** ct, r
 $\pi \leftarrow \mathcal{N}.\mathcal{S}_2(\text{crs}, \tau, (\text{mpk}, \text{ct}))$
return (ct, π)

On message (Enc, k, m) from $\mathcal{Z}$ to PKE:

if $k = \text{mpk}$ **then return** $\mathcal{C}^{\text{PKE}}.\text{Enc}(\text{pk}, m)$
else return PKE.Enc(k, m)

On message (OUTPUT, b) from $\mathcal{Z}$:

output b

Hybrid 7 In the last two hybrids we replace the encryption the zero-strings (for the secret keys of the internal scheme) with the original keys. Therefore we have the following lemmas.

Lemma 8. *Assuming CCA security, $H_6 \approx_{\epsilon_{\text{cca}}} H_7$, over the randomness used by those experiments.*

Lemma 9. *Assuming CCA security, $H_7 \approx_{\epsilon_{\text{cca}}} H_{7.1}$, over the randomness used by those experiments.*

The reduction proceeds as in the parallel CCA hybrids (in the other direction) and is thus omitted.

Summary Note how $H_{7.1}$ is in fact equivalent to the Simulator: we have shifted to an ideal world protocol (H_1) where the protocol can proceed only if the simulator verifies through attestation that the enclave programs are being run in the correct order (H_2) and we leak no information about inter-enclave secure channels ($H_3, H_{3.1}$). We then switch, through H_4 and H_5 , from using genuine NIZK provers and verifiers (respectively) for honest parties into simulating the proof and extracting the witness (resp). By switching from encrypting messages to strings of zeros (H_6), we ensure no leakage is possible, while still using the original secrets

for establishing secure channels $(H_7, H_{7.1})$. The final construction corresponds to our definition of the simulator in 5, and thus, through the subsequence of hybrids, protocol **Steel** UC emulates the ideal functionality FESR.

We now argue that all requirements of the UCGS theorem with respect to **Steel**, FESR and G_{att} are satisfied. In particular we prove the following lemma.

Lemma 10.

1. *The functionality G_{att} is FESR-regular setup and subroutine respecting,*
2. *FESR, **Steel** are G_{att} -subroutine respecting,*

Proof. As required by Definition 3.3 of [10], the global functionality G_{att} does not invoke any new ITI of FESR and does not have an ITI with code FESR as subsidiary. Clearly, G_{att} is subroutine respecting. Also, FESR and **Steel** are G_{att} -subroutine respecting since they only make external calls to G_{att} .

By the UCGS theorem [10] (see Theorem 2), Theorem 3 and the above lemma, UCGS security of the **Steel** protocol is concluded, i.e., for any parent protocol ρ which is $(\text{Steel}, \text{IDEAL}_{\text{FESR}}, \xi)$ -compliant and $(\text{Steel}, M[x, G_{\text{att}}], \xi)$ -compliant for $x \in \{\text{IDEAL}_{\text{FESR}}, \text{Steel}\}$, the protocol $\rho^{\phi \rightarrow \pi}$ UC-emulates ρ .