AGATE: Augmented Global Attested Trusted Execution in the Universal Composability framework

Lorenzo Martinico, Markulf Kohlweiss
University of Edinburgh & Input Output
firstname.lastname@{ed.ac.uk,iohk.io}

***Abstract*—A Trusted Execution Environment (TEE) is a security technology, implemented by CPU manufacturers, which guarantees integrity and confidentiality on a restricted execution environment to any remote verifier through attestation. TEEs are deployed on various consumer and commercial hardware platforms, and have been widely adopted as a component in the design of cryptographic protocols both theoretical and practical.**

Within the provable security community, the use of TEEs as a setup assumption has converged to a standard ideal definition in the Universal Composability setting (G_{att}, defined by Pass et al., Eurocrypt '17). However, it is unclear whether any real TEE design can actually realise such a level of security, or whether the diverse capabilities of today's TEE implementations will in fact converge to a single standard. Therefore, it is necessary for cryptographers and protocol designers to specify what assumptions are necessary for the TEE they are using to support the correctness and security of their protocol.

To this end, this paper provides a more careful treatment of trusted execution than the existing literature, focusing on the capabilities of enclaves and adversaries. Our goal is to provide meaningful patterns for comparing different classes of TEEs, particularly how a weaker TEE functionality can implement a stronger one given an appropriate mechanism to bridge the two. We introduce a new, "modular" definition of TEEs that captures a broad range of pre-existing functionalities defined in the literature while maintaining their high level of abstraction. While our goal is not directly to model implementations of specific commercial TEE providers, our modular definition provides a way to capture more meaningful and realistic hardware capabilities. We propose to characterise TEE capabilities along the following terms:

- **the set of trusted features available to the enclave;**
- **the set of possible attacks on an enclave;**
- **the content of attestation signatures.**

We then define various possible ideal modular G_{att} functionality instantiations that capture existing variants in the literature. Finally, we conclude the paper by constructing a protocol template to realise stronger G_{att} setups from weaker ones, and provide an example of removing an attack.

I. Introduction

In recent years, programmable hardware-based Trusted Execution Environments (TEEs) have been made available by computer manufacturers for different market segments, including consumer and server CPUs. Their introduction has led them to be considered as a realistic component in the development of secure interactive protocols, for a range of diverse use cases (surveyed in [66, 82]). While actual real-world protocol deployments adopting TEEs have been limited, despite being included in the offering of major cloud vendors, there has been a large number of academic publications exploring the feasibility of their use, both in the systems and cryptographic literature. A popular approach within the provable security community is to treat the existence of TEEs as a setup assumption. Hardware setup assumptions had previously considered simpler devices with limited computational power [42]. To capture the flexibility of a full fledged TEE, Pass, Shi, and Tramèr [69] (PST) formulate an ideal setup in the Generalised Universal Composability [25] setting. Their *global attested execution* functionality G_{att} captures the two core security claims of Trusted Execution: 1) programs run in an isolated environment (a secure "enclave") maintain confidentiality and integrity in the presence of an otherwise corrupted party; 2) the output of such a program is authenticated by "attestation". G_{att} models attestation as a signature over the program output and metadata, allowing any remote party to verify the value of attestation through a (globally shared) public verification key, regardless of their access to an enclave.

The PST functionality provides a clean abstraction for TEEs that facilitates security proofs in the Universal Composability model [16]. It is a high level formulation that does not contain precise implementation details for any one TEE platform to allow generalisable protocols. As a consequence, the functionality does not define a specific programming model, leading to publications with various (sometimes incompatible or unrealistic) assumptions about what features are available to the enclave program e.g. whether they are able to access the attestation service or establish secure channels to other enclaves or external parties without explicitly performing key exchange. It is also not clear whether the strong guarantees of the functionality can be met by any real implementation of TEEs, given the vast number of practical side-channels and physical attacks discovered since the release of the technology. It is thus unclear whether the promised guarantees can actually be delivered in the real world. This is an issue that most protocol designers that incorporate TEEs in their constructions conveniently choose to ignore by considering these attacks as out of scope. While we agree it would be unreasonable to ask cryptographers to become experts in the finer details of computer architectures necessary to secure TEE implementations, a more realistic model is warranted if we are to see actual deployment of these protocols. Otherwise, replacing the idealisation of a TEE with a specific instantiation is bound to invalidate any security claim. A salient example is given by Bhatotia et al. [10], who show how a weakened abstraction that allows malicious adversarial to interfere with an enclave's state can lead to loss of confidentiality, by mounting a *rollback attack* on a protocol that would be secure in the G_{att}-hybrid model. On the other hand, other works [36, 78] have shown that, for some protocols, a (significantly) weaker TEE implementation can still provide meaningful guarantees. The existence of these protocols suggests the need for cryptographers to articulate more precisely what aspects of a TEE their design relies on. Articulation requires an appropriate language; our goal for this work is to create one.

Our model: We augment the ideal PST functionality with three extension points that can be modularly combined to characterise a TEE instantiation with specific guarantees: features, attacks, and attestation contents.

Features model the high level (trusted) interface available to programs executed within a TEE to interact with a secure subroutine, or an untrusted program running on the same machine or on a remote party. The implementation of a feature might be realised through specific hardware modifications to the CPU architecture, trusted firmware, a cryptographic protocol between multiple enclaves and remote parties, or a combination thereof. As such, we give the enclave program access to "oracles" (an abstraction of a trusted interface) for the available features. The goal of a feature oracle is to model the guarantees of the untrusted boundary between the trusted code running within the enclave and its access to the external (untrusted) world. Feature oracles can also be used to simplify the design of an enclave program, by abstracting commonly used subroutines for which we have a provably secure implementation (as discussed later in the paper).

Attacks are also represented as abstract oracles, available to the adversary when interacting with the ideal TEE functionality. When constructing protocols that interact with TEEs, the attacker is generally modelled as a malicious party that is executing an enclave on their local machine. As such, we give the attacker the option of passing additionally malicious control instruction along with any input to the enclave, and explicitly state in the formulation how a call to that oracle will affect the internal enclave state.

The values of *Attestation* that are transmitted to a remote verifier

to certify the authenticity of the installed program are defined as a function over the state of the enclave (its *measurement*) and is bound to the TEE instance it runs on. The PST model has a rigid definition of attestation, with its guarantees inspired by the earliest attestation mechanism adopted by Intel SGX. Our formulation is more abstract and allows us to adopt a wider class of measurements and attestation properties.

Our modelling of these interfaces is presented in a modular fashion, with a shared baseline abstraction that provides an interface to parties interacting with TEEs. For each instantiation of a TEE, we capture its unique combination of features, attacks and attestation through a combination of UC "shells", a modelling construct that allows us to reason about the interface of the enclave without the need to analyse the specific application code it is running. We provide several examples of shells that capture pre-existing formulations of TEEs in the literature, unifying all previous PST variants we are aware of.

By providing a modular functionality for TEEs, we let the security proof for a protocol be independent from a concrete TEE instantiation. The protocol designer simply needs to choose the minimum set of features required by the enclave program, an upper bound on how an attacker is allowed to tamper with enclaves, and how much information about the enclave is provided to other parties (or "leaked" to the environment) by the attestation. Despite this, we do not want to dismiss the pre-existing work to prove protocols as secure in the simpler PST model. As such we propose a technique to bridge different versions of the functionality, either by adding a new feature oracle, or by removing an attack oracle. We show how to construct "wrapper" protocols which, combined with a less powerful TEE abstraction, are functionally equivalent to a stronger one, by implementing the missing features in runtime, or patching the remaining attacks. Showing that a more realistic TEE formulation, combined with the appropriate wrapper, is equivalent to PST allows us to preserve pre-existing proofs under Universal Composability. By repeatedly showing that the combination of a "weak" TEE with a protocol implements a "stronger" TEE, we can provide a path to realise a powerful abstraction such as PST from realistic TEE implementations. We hope that our functionality will provide the cryptographic community with a unifying abstraction to characterise different versions of TEEs, including those that have already been proposed in the literature, and will help analyse how they relate to each other. We see this as an important step to enable a more nuanced discussion on the security claims of TEE vendors and the requirements for TEE-enabled protocols - but is ultimately still a theoretical contribution. Constructing the next generation of practical TEE will require a much more system-focused approach than what is possible at this level of specificity.

The paper is structured as follows: in Section II, we provide an overview of the existing TEE modelling from Pass, Shi, and Tramèr [69] and follow-up works. We then propose, starting from Section III, our framework for specifying *Trusted Execution* functionalities with more granular interfaces. The key characteristic of our framework is to provide three parameters for each TEE setup: the set of features that an enclave running on the TEE can access; the set of attacks the adversary is allow to mount; and what values are included in the attestation measurement signed by each enclave. In Section IV, we provide examples of enclave instantiations based on this model, adapting pre-existing formulations of TEE setups, as well as new capabilities that form useful building blocks for building protocols. Our goal is to unify all pre-existing variants of Trusted Execution setups based on the PST model.. Section V provides a notion of equivalence between different classes of TEEs, based on the difference between interfaces, and gives a template for showing that a TEE setup that allows a particular adversarial attack can realise a "stronger" setup without the same attack when combined with an appropriate defensive protocol. We give a simplified version of our technique for readability, with the full variant (and corresponding template for feature addition) presented in Appendix B. We illustrate how to realise this construction through a simple example protocol that removes the adversarial ability to conduct rollback attacks on an enclave through access to trusted storage.

II. Background

A. Universal Composability

Universal Composability (UC), introduced by Canetti [16], is a computational proof model that allows modular proofs of protocol security under concurrent composition in the simulation setting. Due to its flexible modelling of communication channels and adversarial capabilities, UC can capture a broad variety of adversarial scenarios, and a large number of protocols have been shown to be UC-secure. Moreover, since its introduction, the framework has inspired numerous extensions and variations [3, 14, 18, 44, 51, 61] including different revisions to the original model (see [15, Appendix B]).

We now provide an overview of the essential components of UC required for understanding the rest of the paper, with more precise definitions, theorem statements and corresponding notation left to Appendix A.

A UC *protocol* is an execution of a number of *Interactive Turing Machine Instances (ITIs)*. Each machine's is identified by its code, the party ID to which the machine belongs, and a shared session ID (collectively, the extended identity). Machines can pass messages to each other by writing to a number of communication tapes. Additionally, an *adversary* can issue special corruption commands to learn a machine's internal state or control its behaviour. A structured protocol facilitates distinguishing protocol operations from modelling constructs by defining a nested ITI structure, with an external *shell* handling modelling instructions (such as corruption commands or message redirection) running an internal *virtual ITI* that represents the actual code executed by the party (it is possible to nest multiple levels of shells).

To model concurrent composition, a protocol has to be proven secure in the presence of an additional polynomially-bounded machine, the environment. The environment schedules the order and length of execution (through a mechanism called *import*) of protocol ITIs, and is able to run additional protocols in any session, except for a special test session. A protocol can be shown to be secure if it UC-emulates an *ideal functionality*, an incorruptible trusted third party that perfectly executes the protocol. The UC emulation experiments sees the environment attempting to distinguish, for any possible adversary, whether it is interacting (in the test session) with the "real-world" protocol execution, or with the combination of the ideal functionality and a *simulator* machine that acts as the adversary while also injecting protocol-specific messages into the execution transcript.

A UC protocol π can be used as a subroutine for another protocol ρ. If protocol π UC-emulates an ideal functionality F, ρ UC-emulates $\rho^{\pi \to F}$ i.e. protocol ρ where all calls to π are replaced with calls to F. This result (UC composition) can be used to prove the security of protocols in a modular fashion by progressively replacing protocol subroutines with the corresponding functionalities. If ρ calls ideal functionality F, we say it is an F-hybrid protocol. In the standard UC setting, protocols are limited to calling hybrid functionalities from the same session. In later work, the introduction of global functionalities (in the GUC framework of Canetti et al. [25], now replaced by UCGS [6]) allows different protocol sessions to share state through a single hybrid functionality.

B. Trusted Execution Environments

Trusted Execution Environments have generally been the domain of the system security research community. Costan and Devadas [32] first attempted to bridge the knowledge gap to allow cryptographers to understand the internals and guarantees of Intel SGX, the first widely available and commercially successful TEE. We assume the reader is familiar with the high level guarantees of a TEE, and refer to that work for an in-depth explanation of the internals of one of its most popular instantiation.

Since then, a variety of works have attempted to formalise TEEs for the purposes of cryptographic protocol design [2, 8, 35, 38, 52, 56, 72–75, 80, 86].

Given the desirable composition guarantees of UC, we focus our treatment of TEEs in that model. While various works exist to model HSM-like functionality in UC (e.g. see [48]), and some initial work has been proposed by Canetti et al. [26] to give a UC treatment of validating the security guarantees of generic hardware constructions (including protecting against side-channel attacks), Pass, Shi, and Tramèr [69] provide the first UC formulation of TEEs. Their G_{att} functionality (fully reproduced in Figure 1) is a generic model for TEEs that aims to capture architecture-independent properties. It distills the essence of TEEs into attested execution i.e. evaluation of a program with associated proof of execution. G_{att} lets a pre-established set of parties, with local access to a TEE, install and execute arbitrary enclave programs, which produce anonymous attestation signature over the program output and enclave metadata. While the environment is able to verify the authenticity of an attested output and install their own programs through a corrupted party in any session (through claimed session ID idx), they learn nothing about the internal state of an enclave or the identity of the party executing that program. Any implementation details of the trusted hardware or concrete attestation protocol are abstracted away from the attested execution formalism. A simple signature mechanism collapses local and remote attestation into a single operation, which any party can verify having obtained the relevant public key from G_{att}. This is the only meaningfully global shared state for the functionality, with individual enclaves tied to session identifiers, and a unique (across sessions) enclave identifier eid.

The role of the signature scheme is a simplification over the EPID attestation protocol used in the original version of SGX [70], that removes the key revocation phase. Attestation verification amounts to simply verifying the output data structure as described through a simple signature scheme with the globally available (both to machines with and without enclave capabilities) public verification key. The signing key is never released by the functionality, capturing the provisioning mechanism of the SGX system enclaves. The inclusion of the session ID in the attestation signature ensures that enclaves installed in different sessions (for which the simulator has no visibility) can not adversely interact with the protocol.

Since its publication, numerous cryptographic protocols that rely on TEEs have been proven using G_{att} in the (G)UC framework [9, 28–31, 40, 41, 43, 46, 49, 53–55, 58, 65, 83, 84, 87, 88] or as a resource [57] in the Abstract Cryptography framework of [62]. G_{att} has also provided a basis for formalising TEE usage in property-based security proofs [33, 37, 39, 59, 63, 76, 85].

Additionally, some attempts have been made to relax the G_{att} functionality for the purposes of capturing TEE vulnerabilities. Tramèr et al. [78] introduced the concept of transparent enclaves to model confidentiality leaks in an enclave program (formalised under GUC in [68, Section 8.1]). The transparent enclave functionality behaves exactly as G_{att}, except that for each RESUME operation, the functionality additionally leaks the randomness used by the enclave

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

State variables	Description
vk	Master verification key
msk	Master secret key
$\mathcal{T} \leftarrow \emptyset$	Table for installed programs

On message INITIALIZE *from a party* P*:*
  let $(\mathsf{vk}, \mathsf{msk}) \leftarrow \Sigma.\mathsf{Gen}(1^\lambda)$
On message GETPK *from a party* P*:*
  return vk
On message $(\text{INSTALL}, \mathsf{idx}, \mathsf{prog})$ *from* $P \in \mathsf{reg}$*:*
  if P is honest **then assert** $\mathsf{idx} = P.\mathsf{sid}$
  generate nonce $\mathsf{eid} \xleftarrow{\$} \{0,1\}^\lambda$
  store $\mathcal{T}[\mathsf{eid}, P] \leftarrow (\mathsf{idx}, \mathsf{prog}, \emptyset)$
  return eid
On message $(\text{RESUME}, \mathsf{eid}, \mathsf{inp})$ *from* $P \in \mathsf{reg}$*:*
  let $(\mathsf{idx}, \mathsf{prog}, \mathsf{mem}) \leftarrow \mathcal{T}[\mathsf{eid}, P]$, **abort** if not found
  let $(\mathsf{output}, \mathsf{mem}') \leftarrow \mathsf{prog}(\mathsf{inp}, \mathsf{mem})$
  store $\mathcal{T}[\mathsf{eid}, P] \leftarrow (\mathsf{idx}, \mathsf{prog}, \mathsf{mem}')$
  let $\sigma \leftarrow \Sigma.\mathsf{Sign}(\mathsf{msk}, (\mathsf{idx}, \mathsf{eid}, \mathsf{prog}, \mathsf{output}))$
  return $(\mathsf{output}, \sigma)$

Fig. 1. The G_{att} functionality of [69]

(allowing the host to derive any secret created within the enclave).

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. Dörre, Mechler, and Müller-Quade [36] proposes both a weaker and a stronger variant. Since the SGX quoting enclave that allows producing attestation does not have any specific hardening mechanism compared to other enclaves running on the machine, besides being carefully implemented with side-effect free primitives, the authors argue that it is realistic to model a class of TEEs where side channels do not affect certain secure operations such as key exchange and symmetric encryption (since the quoting enclave relies on them for attestation to be successful). As such, they define *almost-transparent enclaves* as transparent enclaves with access to side-channel free implementations of symmetric cryptography primitives and Diffie-Hellman key exchange operation. On a RESUME operation, an almost transparent enclave leaks the random bits used during its execution, the memory of the enclave at the start of the RESUME call, and the return value of the cryptographic operations, but crucially not the randomness used to perform the cryptographic functions. This allows the adversary (and the simulator) to learn any values that would have been leaked through any intermediate computation on secrets the enclave had access to. Additionally, they consider a *semi-honest enclave*, inspired by the modelling of [56], where the adversary is able to adaptively leak the list of operations executed by an enclave run by any party regardless of their corruption status. A semi-honest enclave model captures a scenario where the manufacturer of the TEE might have introduced a backdoor that enables them to remotely instruct any TEE-enabled machine to record and leak their data. Besides providing the alternative attacker models, their global functionalities are realised in UCGS, and allow any party to install an enclave (i.e. there is no fixed registry set reg).

Bhatotia et al. [10] provides a further weakened UCGS version of G_{att} that allows an adversary to conduct rollback and forking attacks. Their functionality keeps track of enclave states in a tree data structure and allows a corrupted party to select an arbitrary node in the tree to load the state from as part of a RESUME operation. This

new weaker setup can be shown to no longer be sufficient to guarantee the security of a protocol that includes stateful enclaves.

III. A MODULAR G_{att} SETUP

As an ideal functionality, the G_{att} formalisation described in the previous section does not provide a detailed account of enclave execution. This formulation of TEEs does not explicitly expose any specific hardware or implementation details, beyond the abstract interface that allows the local party to install a program and execute it. When describing the components of G_{att}, Pass, Shi, and Tramer [68, Section 3.2] explicitly state that the functionality emerges from a combination of the TEE features with some assumed firmware to provide this type of confidential computing service. In particular, they attribute the generation of unique per-enclave ids at installation, which are not guaranteed by all TEE architectures, to this firmware sampling a nonce from a unique key distributed to each TEE by the manufacturer during provisioning.

This abstraction of TEEs as an isolated execution mechanism with an easily verifiable proof of computation is a key insight of the model, and its promise of using the abstraction as a block box for constructing protocols a major selling point. The high level of abstraction does however conceal how realistic hardware component might fail, preventing protocol designers to take such scenarios into account. A more careful approach would then consider the functionality provided by G_{att} as implementable by a combination of hardware, trusted firmware, and system-defined enclaves. The attestation signature guarantees that all of these components were acting in concert at the time when an output was generated.

Examining these components in more detail provides two advantages. First, it allows more meaningful relaxations of the security guarantees, by allowing to distinguish which components of the system can be compromised. Additionally, once we stop thinking of the functionality as a monolithic hardware component, it becomes natural to consider alternative features that the manufacturer or third parties might augment the TEE with. In particular, we may think of the combined hardware and software libraries an enclave has access to during its execution "runtime" as providing a kind of API. While the list of features provided by G_{att} could be considered a "standard" enclave interface, it is possible to imagine additional API calls available to the enclaves, for example a trusted clock [27], monotonic counters [27, 60], secure access to GPU compute resources [77, 79, 89] etc. Regardless of how these interfaces are implemented (e.g. by modifying the architecture or trusted firmware, or running the enclave through a "wrapper" library that interacts with a trusted system enclave, or even through a distributed protocol between multiple mutually untrusted enclaves), the attestation mechanism should capture their presence. Beyond showing that an enclave is running the correct program, a sound attestation mechanism also needs to certify to the verifier that the TEE provides the correct version of the API, otherwise the program code can not provide its security guarantees. In other words, a TEE functionality attests to the combination of (*prog, runtime*) rather than the mere application code *prog*.

Features, Attacks, and Attestation: We now extend the G_{att} functionality from [68, 69] (henceforth referred to as G_{att}^{PST}) to allow defining a larger class of TEE setups. Our goal is to capture the runtime behaviour of enclaves, without delving into the specifics of their implementation. To maintain this level of abstraction, we use a number of idealised interfaces.

Within our new formalism, a TEE application developer can choose to target a minimum set of features required by their applications. A standard error will be returned if such a program is installed on an instance of the TEE functionality that does not support the feature set. For each possible modular instantiation of a TEE G_{att}^{mod}, we thus define a set of feature oracles $\mathbb{O}$, which represent the library of subroutines that are available to an enclave program. A feature of this kind is a polynomial time algorithm, as implemented by the runtime combination of hardware and software in that version of G_{att}^{mod}, including any communication with external parties. We also define a set of attack oracles $\mathbb{A}$ to capture adversarial behaviour. This can be thought as a parameter chosen by a protocol designer that captures "allowable" attacks in the current TEE setting under which the target protocol can still be proven secure. Any cryptographic protocol that wants to use TEE will therefore need to provide a lower bound for the set of required features $\mathbb{O}$, and an upper bound for the set of tolerated attacks $\mathbb{A}$, to parameterise their chosen version of G_{att}^{mod}. Relationships between different versions of TEEs are captured by the difference of these two sets, with equivalence statements made possible by running some additional runtime along enclave programs (either to increase the size of interfaces provided by $\mathbb{O}$, or to reduce the attacks available in $\mathbb{A}$).

We also introduce modularity in the attestation procedure. This is both to allow capturing a greater class of TEE architectures, as well as being a technical requirement. A reader familiar with the simulation framework will quickly realize that our programme of proving that, given the right runtime, a weaker TEE setup G_{att}^{mod} can UC-emulate the stronger $G_{\mathsf{att}}^{mod\prime}$, is hindered by the usage of a fixed signature scheme to model attestation. Since the two different TEE functionalities would each sign different (*prog, runtime*) messages, it would be trivial for an environment to distinguish whether it is communicating with the real or ideal world, since the G_{att}^{PST} attestation scheme would also include the *runtime* as part of its code (thus signing two different programs in the two worlds). We therefore abstract the attestation mechanism in order to allow the UC simulator to "program" the signature scheme.

Our model ties attestation and its verification to the specific G_{att}^{mod} functionality instance the user interacts with: the public parameters of the functionality allow a verifier to directly assess the capabilities of the attested enclave runtime and its adversary. It thus allows the verifier to make an informed trust decision based on the feature and vulnerability of the enclave they are communicating with.

The functionality: We now highlight the differences between the new formulation of G_{att}^{mod} (Figure 2) and the original G_{att}^{PST} functionality (shown in Figure 1).

We iterate on the work of [10] to more carefully follow the conventions and formality of modern UC versions compared to G_{att}^{PST}. In particular, we now model enclaves as structured ITI subroutines to the G_{att}^{mod} functionality. On installation of an enclave, the functionality spawns a new ITI subroutine with composite extended identity $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$, encoding the program prog, oracles $\mathbb{O}, \mathbb{A}$, the unique enclave ID eid, the identity pid for the party that installed the enclave, and the claimed session identity idx. The new subroutine is part of a UC *structured protocol*, where the top level subroutine with code $\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]$ spawned by G_{att}^{mod} is known as a *shell*, and a second subroutine with code prog created by the shell is known as the *body*. We use the shell of our structured protocol to capture modelling instructions related to the oracles, while the body is instantiated with the unaltered program code for the enclave (see Figure 3 for a graphical representation). Running enclaves as separate subroutine ITIs is functionally equivalent to running the input code within the global functionality as in the original treatment. It does provide, however, a cleaner abstraction, in that we are able to explicitly instantiate an ITI that runs the code of the enclave program installed, rather than having the ideal functionality act as an interpreter. In particular, our formalism now involves enclaves run by different parties being executed as separate ITIs, which we believe is a more natural model. Enclave programs are subroutine respecting

Functionality $G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}]$

State variables	Description
$\mathsf{vk} \leftarrow \epsilon$	Master verification key
$\mathsf{Sign} \leftarrow \epsilon$	Attestation Signing algorithm
$\mathcal{S} \leftarrow \emptyset$	Table for signed messages
$\mathcal{T} \leftarrow \emptyset$	Table for installed programs

On message INITIALISE *from a party* P*:*
  send INITIALISE **to** $\mathcal{A}$ and **receive** k, s ; $\mathsf{vk} \leftarrow k, \mathsf{Sign} \leftarrow s$
On message GETPK *from a party* P*:*
  return vk
On message $(\text{VERIFY}, \sigma, m)$ *from a party* P*:*
  return $m \overset{?}{\in} \mathcal{S}[\sigma]$
On message $(\text{INSTALL}, \mathsf{idx}, \mathsf{prog})$ *from a party* P *where* $P.\mathsf{pid} \in \mathsf{reg}$*:*
  if pid is not corrupted **then**
    assert $\mathsf{idx} = \mathsf{sid}$
  for instruction $\mathsf{i} \in \mathsf{prog}$ **do**
    if $\mathsf{i} \notin \mathbb{O}$ **then return** MissingInstructionError
  generate nonce eid $\xleftarrow{\$} \{0,1\}^\lambda$, **store** $\mathcal{T}[\mathsf{eid}, \mathsf{pid}] = (\mathsf{idx}, \mathsf{prog})$
  send INSTALL **to** $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$
  return eid
On message $(\text{RESUME}, \mathsf{eid}, \mathsf{inp}, \mathsf{attack})$ *from a party* P *where* $P.\mathsf{pid} \in \mathsf{reg}$*:*
  let $(\mathsf{idx}, \mathsf{prog}) \leftarrow \mathcal{T}[\mathsf{eid}, \mathsf{pid}]$, **abort** if not found
  if $\mathsf{attack} = \epsilon \vee$ pid is not corrupted **then**
    send inp **to** $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$ and **receive** output
  else
    assert $\mathsf{attack} \in \mathbb{A}$
    send $(\mathsf{attack}, \mathsf{inp})$ **to** $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$ and **receive** $\mathsf{output}, \mathsf{aux}$
    if $\mathsf{aux} \neq \epsilon$ **then**
      query $\mathcal{A}$ **with** $(\mathsf{attack}, \mathsf{aux})$ and **receive the reply** CONTINUE
  let $\mathsf{meas} \leftarrow \mathbb{S}$(configuration of $\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]ITI$), $\sigma \leftarrow \mathsf{Sign}(\mathsf{meas})$
  $\mathcal{S}[\sigma] \leftarrow \mathcal{S}[\sigma] \parallel \mathsf{meas}$
  return $(\mathsf{output}, \sigma)$

Fig. 2. Global functionality G_{att}^{mod}

in that the shell rejects any input message not sent through the G_{att}^{mod} functionality, and will only accept subroutine output messages from machines in its extended session (i.e. all machines invoked by the protocol's main parties or one of their subroutines). When resuming an enclave, the calling party might need to provide some additional runtime import (see Appendix A), depending on how much work the shell is required to carry out in addition to the enclave code execution in itself (e.g. if an enclave calls a feature that involves significant communication with external parties to be implemented, G_{att}^{mod} needs to be activated with sufficient import to activate those subroutines).

We parameterise each instance of G_{att}^{mod} by the static sets $\mathbb{O}, \mathbb{A}$ which capture feature and adversarial oracles respectively. On installation of a new enclave, G_{att}^{mod} first checks that all instructions in the proposed program code correspond to a call to one of the oracles in $\mathbb{O}$, and aborts with an error message if they are not. Both sets are the basis for the definition of the shell for all enclave subroutine ITIs installed by that instance of G_{att}^{mod}. We use the shell mechanism device to help us capture a specification of how the enclave program and the adversary can interact with the runtime. In particular, for each unique combination of oracles, we have to give a specific shell definition.

The shell detects when its enclave calls a feature oracle at runtime, and provides a return value. This can be derived through some local computation conducted by the shell, potentially after communicating with the adversary or other parties; or delegated to a distinct subroutine. When defining shell in this work we will generally use ideal subroutines, but this can be implemented through a real protocol without changing the definition (through UC-emulation).

A corrupted party is allowed to specify an auxiliary command along with their resume instructions that is executed by the shell in conjunction or instead of the normal program execution. The adversarial oracle is allowed to send a message to the adversary after the RESUME call has completed, and the adversary can in turn prevent the output of the program from being released with an attestation. The shell also handles any communication between enclaves that might be prompted by an attacker or feature oracle.

Finally, we parameterise the functionality by $\mathbb{S}$, a function that defines the contents of the attestation message for each enclave's execution. The original G_{att}^{PST} models an anonymous attestation signature scheme, and as such always produces an attestation signature tied to the set of arguments $(\mathsf{idx}, \mathsf{eid}, prog, \mathsf{output})$. This includes the claimed session ID for the current protocol executing the enclave, its unique enclave ID, the program code and the output of the most recent computation. Replacing this fixed data structure with a function allows us to model a broader range of attestation primitives, such as non-anonymous attestation (e.g. by including the UC party ID as one of the returned values, or a long-term public key tied to the party identity, as outlined in [68, Section 8.4]). We further relax the attestation mechanism of the G_{att}^{PST} functionality by allowing the adversary (through the simulator in the ideal world) to choose the format of attestation signatures, to allow the addition of details

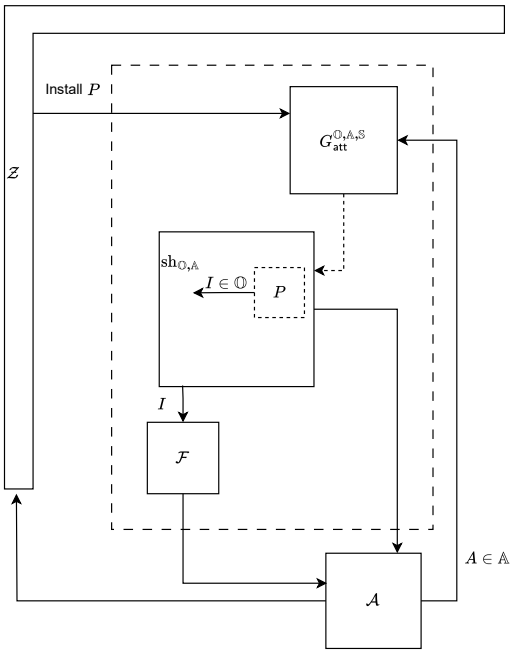

Fig. 3. When a program with code P is installed on a G_{att}^{mod} enclave, the functionality spawns a new structured protocol subroutine with shell $\mathrm{sh}_{\mathbb{O},\mathbb{A}}[]$ and body P. For some interfaces $I \in \mathbb{O}$, the shell will outsource its computation to some external functionality $\mathcal{F}$. The adversary $\mathcal{A}$ can interact with the enclave shell for any attacks $A \in \mathbb{A}$ through G_{att}^{mod}. Both $\mathrm{sh}_{\mathbb{O},\mathbb{A}}[]$ and $\mathcal{F}$ can leak additional information to $\mathcal{A}$

lacking in the high-level abstraction. Rather than having a full-fledged offline digital signature algorithm, the adversary provides (during the INITIALISE phase of the setup) G_{att}^{mod} with a public key and a signing algorithm. The algorithm is not required to be a well-formed signature scheme or guarantee typical security properties such as existential unforgeability. Instead, G_{att}^{mod} implements signature verification by maintaining a map $\mathcal{S}$ of all signed strings and corresponding signatures generated by Sign. Verifications require sending a message to the setup, which checks whether it did produce the signed output through an "ideal" table lookup, rather than running a real verification algorithm as specified by the signature scheme. We still allow fetching a verification key for interface compatibility with G_{att}^{PST}, but any environment party that has obtained the relevant verification algorithm and key from the adversary will not have any guarantees of existential unforgeability.

When showing UC-emulation between two TEE setups, the simulator can provide a modified version of these algorithms to convince the environment that the ideal world TEE shares its runtime with the real-world TEE. Take an adversary, for instance, that selects a signature scheme Σ, and initialises a G_{att}^{mod} instance with closure $s(\mathsf{meas}) = \Sigma.Sign(sk, \mathsf{meas})$, such that on a RESUME call, G_{att}^{mod} applies s to the value produced by function $\mathbb{S}$ over the configuration of the enclave ITI, the *enclave measurement*. On receiving algorithm s from the adversary, the ideal-world simulator can derive a new signing algorithm $s'(\mathsf{meas}) = s(R(\mathsf{meas}))$. R is a transformation on the measurement that preserves all of its information, except that, if the measurement contains a public commitment to the program executed in the enclave (such as a hash of its source code), and the real-world G_{att}^{mod} functionality is running code of type $\mathsf{prog} = (\mathsf{app}, \mathsf{runtime})$ for a specific $\mathsf{runtime}$ library, R replaces the commitment to enclave code app with a commitment to $(\mathsf{app}, \mathsf{runtime})$. This means that attestations in the ideal world will look like attestations to $(\mathsf{app}, \mathsf{runtime})$, despite G_{att}^{mod} only installing and executing app as part of its enclave. Of course, app still needs access to the interface offered by the runtime, but in the ideal world it directly accesses the idealised features in the $\mathbb{O}$ set.

It is easy to show that G_{att}^{PST} UC-emulates G_{att}^{mod} for the sets of oracles and measurement function that correspond to G_{att}^{PST} (which we describe in Section IV-A), by constructing a simulator that selects the exact signature scheme specified in G_{att}^{PST}. Note that the opposite direction G_{att}^{mod} UC-emulates G_{att}^{PST} is more subtle. In fact, it is clear that the statement can not hold for all possible signature schemes provided by an adversary. Consider the null signature scheme where the signing algorithm $\mathsf{Sign}(\mathsf{ssk}, m) = 0^{\lambda}$; the signature scheme is still valid under the definition of G_{att}^{mod}, but it allows the environment to learn whether an enclave has produced a specific message, without having to communicate with it (by simply querying the ideal functionality for verification of an arbitrary measurement produced by $\mathbb{S}$). This is not possible in G_{att}^{PST}. A minimum entropy requirement for signatures provided by the adversary would therefore be necessary (but not sufficient) for the other direction of the equivalence.

A recent work by Canetti et al. [24] shows, as a corollary of the UCGS composition theorem, that if a global protocol G UC-emulates G' with respect to simulator S, then it is possible, in the general case of any context protocol ρ, to replace any subroutine call from ρ to G with a call to the combined subroutine of G' and S. This enables us to port any existing proofs that rely on G_{att}^{PST} (provided that the proof is valid under UCGS rather than GUC) into our new model, by simply replacing G_{att}^{PST} with the combination of the G_{att}^{mod} instance with equivalent $\mathbb{O}, \mathbb{A}$ oracles, and the simulator that during initialisation chooses the precise G_{att}^{PST} signature scheme over the usual $(\mathsf{idx}, \mathsf{pid}, \mathsf{prog}, \mathsf{output})$ measurement produced by $\mathbb{S}$.

IV. Defining a G_{att}^{mod} Zoo

We now provide the definition for several G_{att}^{mod} oracle instantiations. As far as we are aware, our G_{att}^{mod} formulation can capture all existing variants of G_{att}^{PST} in the literature, as well as additional natural extensions related to real-world TEE realisations and attacks.

Instantiating a shell for the functionality and adversarial oracles is a required step for using a new G_{att}^{mod} variant, and we have made efforts to write shells modularly so that they are easy to reuse. This does not mean that we can directly apply clean-room UC composition, but the structure of the shells makes it easy to mix and match them as required to handle additional oracles. In particular, most shells are structured around a loop that examines all instructions executed in the enclave subroutine ITI. When the instruction matches a specific oracle call, the shell shows how to implement it (in an ideal way). Some shells (such as the one presented in Section IV-D), modify the structure of ITIs created by the shell, but are still compatible with the formulation of other shells.

We first highlight the basic structure of our shell construction, which are shared across all instantiations. Intuitively, the shell acts as a hypervisor for the virtual ITI running the enclave program, interpreting each operation and dispatching external calls to the appropriate functionalities.

The extended identity of the shell is defined as $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$, where the PID is a concatenation of the enclave identifier generated by G_{att}^{mod} and the PID of the source machine which installed the enclave; the session SID is a concatenation of string att and the session of the protocol under which the enclave was installed. The enclave itself is a (virtual) subroutine ITI with extended identity $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$, where the party ID of the virtual ITI is the unique enclave identifier eid and

the session is claimed session idx. The extended identity of the shell includes the calling pid, as it might need to be aware of the calling party for communicating with assisting functionalities; the virtual ITI's extended identity does not, as enclaves do not normally have access to which party is executing them.

The shell is created by receiving message INSTALL from G_{att}^{mod}, after which it runs any necessary setup steps for the oracle functionalities, and initialises the virtual ITI that will actually execute the enclave program.

On input $(\text{RESUME}, \mathsf{eid}, \mathsf{inp})$, rather than passing on its arguments to the virtual ITI's input tape to execute program $\mathsf{prog}(\mathsf{inp})$ directly, the shell enters a loop where for each step it observes the current configuration of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$, and each instruction i that would be executed if it was activated with inp (we denote this as "step through execution of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ on inp:"). Within the main loop, a branch is defined for each type of instructions specified in the set $\mathbb{O}$. Each branch defines how the shell performs the computation for a specific instruction, how to update the internal configuration of the virtual ITI, and how to select the next instruction to execute. If an instruction can be executed by any Interactive Turing Machine, the shell activates the virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ with input i and allows it to execute it without modification (we say that such instructions are part of the standard set $\mathbb{O}^{\text{std}}$). Additionally, the shell might contain specific interfaces for processing adversarial messages, or perform additional operations within its resume loop to allow modelling current or future corruption.

Most shells include a special branch to process a **return** instruction. This will generally perform an external write request to G_{att}^{mod} from the shell ITI, but might include additional behaviour depending on the parametrisation of G_{att}^{mod}.

For the remainder of this work, we consider versions of G_{att}^{mod} that use the same attestation signature function $\mathbb{S}$ as G_{att}^{PST} i.e. anonymous attestastions over $(\mathsf{idx}, \mathsf{eid}, \mathsf{prog}, \mathsf{output})$.

A. G_{att}^{PST}

We begin by reformulating G_{att}^{PST} in the language of G_{att}^{mod}. While this is not made explicitly in the original work, G_{att}^{PST} relies on a number of implicit "features":

- addressable instructions: enclave execution begins at arbitrary instructions addressed through labels; in other words, the enclave program defines some entrypoint as procedures that can be called by the registered party that installed the enclave, along with optional input arguments. On every execution, the enclave returns some output with an associated attestation signature
- stateful resumes: each RESUME instruction is atomic, meaning that the subroutine will execute perfectly without any possibility for adversarial intervention. The state of the enclave is maintained across each sequential RESUME execution, and the adversary is not able to erase or otherwise tamper with it
- sample randomness: enclave programs are assumed to provide a true source of randomness (of arbitrary lengths)
- unique enclave identifiers: a unique enclave ID is generated as a cryptographic nonce during enclave installation. Enclave IDs should be unique for all enclaves, regardless of which party installed them
- attestation verification: attestation signatures can be verified from within the enclave program, without having to trust the external OS code to provide the attestation verification key as an input.

The first three notions are usually considered standard for Interactive Turing Machines. We therefore define the standard oracle set $\mathbb{O}^{\text{std}}$ to capture all ITI instructions that are standard for local computation. Although the operation of an Interactive Turing Machine are much more abstract, this can be thought of as the set of microarchitectural instruction provided by the processing unit executing the ITI. Attestation verification from within an enclave is explicitly not allowed in the G_{att}^{PST} paper [68, page 23], but we include it because it is required by many G_{att}^{PST}-hybrid protocols in the literature. It could be argued that adding this capability makes the functionality less composable than intended, due to the inability to swap the fixed signature scheme with a call to an attestation service as provided by Intel for SGX. G_{att}^{mod} resolves this by moving verification to an abstract check in the functionality rather than verification of a concrete signature scheme.

We also note that the G_{att}^{PST} model forbids the enclave to have access to the UC PID for the party that is running it. While this is not explicitly stated, enclave programs with PID access could assist the party to establish a secure channel with another enclave-enabled party [68, Section 3.3], negating the need for doing this through a protocol.

As for the adversarial powers, even in the scenario where a host party is fully corrupted, adversarial interactions are limited when it comes to the PST enclaves. For any fully corrupted party, a G_{att}^{PST} adversary is able to install programs with arbitrary sessions identifiers under that host, honestly execute an enclave, and verify attestation signatures. These behaviours are all captured by default in the G_{att}^{mod} functionality, so no additional attack oracle is required.

For capturing G_{att}^{PST} under G_{att}^{mod}, we thus define $\mathbb{O}$ to be the union of $\mathbb{O}^{\text{std}}$ and $\{\text{AttestVerif}\}$, and $\mathbb{A} = \{\}$. We now give an implementation for a UC shell that models enclave access to the oracle sets as defined.

$\text{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]$

The shell is defined for $\mathbb{O} = \mathbb{O}^{\text{std}} \cup \{\text{AttestVerif}\}$ and $\mathbb{A} = \{\}$
The extended identity of the shell is defined as $(\text{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

On message INSTALL *from* G_{att}^{mod}*:*
  if virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ does not exist **then** create

On message inp *from* G_{att}^{mod}*:*
  step through execution of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ on inp:
  for instruction i **do**
    if $\mathsf{i} \in \mathbb{O}^{\text{std}}$ **then**
      allow $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ to execute i
    else if $\mathsf{i} = \text{AttestVerif}(\sigma, m)$ **then**
      send $(\text{VERIFY}, \sigma, m)$ **to** G_{att}^{mod} and **receive** v
      write v to subroutine output of virtual ITI
    else if $\mathsf{i} =$ (**return** v) **then**
      return v with source $(\text{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

The behaviour of this shell within the RESUME loop is fairly simple: most program instructions it considers will be in the standard oracle set $\mathbb{O}^{\text{std}}$. In this case, the shell activates $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ with input i; as this is a simple instruction that any ITI can compute, the shell does not need to modify its behaviour, and it will allow the virtual ITI to execute it (updating its work tape) and immediately halt. The activation token now returns to the shell, which can select the next instruction i from the updated configuration.

When the instruction is of type $\text{AttestVerif}(\cdot)$, the shell does not activate $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$, but rather sends a message to G_{att}^{mod} to verify the attestation signature. Once it receives a boolean response, it writes it to the subroutine output tape of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$, and advances the location of the tape head on its work tape. This

essentially convinces the enclave virtual ITI that on its last activation it called the $\mathrm{AttestVerif}$ subroutine, and has just received its return value. We use this mechanism extensively in the rest of the section, as it allows modelling feature oracles so that the enclave program is oblivious of how they are computed.

Finally, when the next instruction i for the enclave is to return some value, the shell forwards it to G_{att}^{mod}, overwriting the sender-id of the outgoing message with its own extended identity. The shell thus yields activation back to G_{att}^{mod}, which proceeds with generating the attestation by calling $\mathbb{S}$ on the configuration of $(\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx})$.

B. Accessing a Clock

A natural extension of G_{att}^{PST}, which the original paper uses to realise fair MPC [68, Section 7.2], is to give the enclave access to a clock. The protocol is proven in a synchronous setting, where each party is activated in a round-robin fashion and is therefore aware of the round number. Enclaves are also equipped with round aware capabilities, even if they are not activated every round.

We now show how to realise a new G_{att}^{mod} functionality that supports feature oracles $\mathbb{O} = \mathbb{O}^{\mathrm{std}} \cup \{\mathrm{ReadRound}, \mathrm{IncRound}\}$ by giving it access to a local functionality that any protocol participant is allowed to interact with (both from within the enclave and outwith). Whenever the enclave program tries to execute an instruction interacting with the clock, the shell intervenes to forward the message to an ideal functionality, and inserts the value back into the enclave virtual ITI through the subroutine output tape.

$\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]$

The shell is defined for $\mathbb{O} = \mathbb{O}^{\mathrm{std}} \cup \{\mathrm{ReadRound}, \mathrm{IncRound}\}$ and $\mathbb{A} = \{\}$
The extended identity of the shell is defined as $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

On message INSTALL *from* G_{att}^{mod}*:*
  if virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ does not exist **then** create
  if ideal functionality $(\mathcal{F}_{\mathsf{clock}}, (\bot, \mathsf{idx}))$ does not exist **then** create
  send register **to** $\mathcal{F}_{\mathsf{clock}}$ **on behalf of** pid
On message inp *from* G_{att}^{mod}*:*
  step through execution of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ on inp:
  for instruction i **do**
    if $\mathsf{i} \in \mathbb{O}^{\mathrm{std}}$ **then**
      allow $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ to execute i
    else if $\mathsf{i} = \mathrm{ReadRound}$ **then**
      send Read **to** $\mathcal{F}_{\mathsf{clock}}, (\bot, \mathsf{idx}))$ **through** pid and **receive** v
      write v to subroutine output of virtual ITI
    else if $\mathsf{i} = \mathrm{IncRound}$ **then**
      send Inc **to** $\mathcal{F}_{\mathsf{clock}}, (\bot, \mathsf{idx}))$ **through** pid and **receive** v
      write v to subroutine output of virtual ITI
    else if $\mathsf{i} =$(**return** v) **then**
      return v with source $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$

The INSTALL subroutine of this shell installs the virtual ITI for a new enclave, and ensures that an instance of the ideal functionality for the clock exists in this session (with a standard PID $\bot$). It then sends a registration message for the enclave to $\mathcal{F}_{\mathsf{clock}}$. For enclave RESUME calls, the structure of the shell execution loop is the same as in the shell from last section, with the instructions executed by the enclave for either $\mathrm{ReadRound}, \mathrm{IncRound}$ oracle calls forwarded to the ideal functionality, and its return values returned to the enclave in the same way that we added the return value for an attestation verification call in the previous section. We now describe the behaviour of the clock functionality

Functionality $\mathcal{F}_{\mathsf{clock}}$

The identity of the functionality is $(\bot, \mathsf{sid}_F)$
On message REGISTER *from* $P(\mathsf{pid}, \mathsf{sid})$*:*
  if $\mathsf{t} = \{\}$ **then** $r \leftarrow 0$
  if $\mathsf{sid} = \mathsf{sid}_F$ **then**
    $t[\mathsf{pid}] \leftarrow \bot$
On message READ*:*
  return r
On message INC *from* $P(\mathsf{pid}, \mathsf{sid})$*:*
  if $\mathsf{sid} = \mathsf{sid}_F$ **then** $t[\mathsf{pid}] \leftarrow \top$
    if all values in $t = \top$ **then**
      $r{+}{+}$
      reset all values in t to $\bot$
    return r

$\mathcal{F}_{\mathsf{clock}}$ provides a per-session round counter functionality. A round is increased when all registered parties consent to. Internally, it stores the round counter as a monotonically increasing integer r, and records whether a party has agreed to increase the round via dictionary t, which records a boolean value for each party. Once a party sends an INC message, they are not allowed to withdraw. After the last registered party agrees to increase, r is incremented, and all values in t are set to false. A new party can register at any point, and the value of the round counter is publicly accessible.

C. Interrupting computation

We now model a version of G_{att}^{mod} that allows an enclave program to explicitly control which objects in their memory can be saved to confidential persistent storage. An enclave is able to preserve state across enclave executions by *storing* arbitrary bitstrings in an encrypted form, and later *fetch* them back into memory when next resumed. Only the original enclave itself is able to access any data it stored through the oracle call; the adversary only learns the size of what was stored. In Intel SGX, these features are known as sealing and unsealing.

As the enclave now interacts with the (untrusted) memory of the host, the adversary will be notified of any store or fetch attempt, and will have a chance to censor them. Given that the program integrity relies on these external oracle calls completing, this is equivalent to the adversary aborting the enclave program. We therefore provide the adversary with oracles $\mathbb{A} = \{\mathrm{Abort}, \mathrm{Continue}\}$. The adversary can stop a memory access oracle from completing, but can not erase or leak external memory that was already successfully stored. This example oracle combination for G_{att}^{mod} is for illustrative purposes; a more realistic oracle definition would both allow memory operations return to the enclave with an error, allowing the program code to handle the failure, and further allow the adversary to permanently erase external memory.

We define the following shell:

$\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]$

The shell is defined for $\mathbb{O} = \mathbb{O}^{\mathrm{std}} \cup \{\mathrm{Store}, \mathrm{Fetch}\}$ and $\mathbb{A} = \{\mathrm{Abort}, \mathrm{Continue}\}$
The extended identity of the shell is defined as $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

State variables	Description
$mem \leftarrow \epsilon$	Persistent memory storage for the enclave

On message INSTALL *from* G_{att}^{mod}*:*
  if virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ does not exist **then** create
  $\mathsf{halt} \leftarrow \bot$
On message inp *from* G_{att}^{mod}*:*
  if $\mathsf{halt} = \top$ **then abort**
  step through execution of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ on inp:
  for instruction i **do**
    if $\mathsf{i} \in \mathbb{O}^{\mathrm{std}}$ **then**
      allow $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ to execute i
    else if $\mathsf{i} \in \{\mathrm{Store}(s), \mathrm{Fetch}\}$ **then**
      if pid is corrupted **then**
        $\mathsf{halt} \leftarrow \top$
        if QUERYADV(i) $= \bot$ **then return**
      if $i = \mathrm{Store}(s)$ **then**
        $mem \leftarrow s$
      else if $i = \mathrm{Fetch}$ **then**
        write mem to subroutine output of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$
    else if $\mathsf{i} = (\textbf{return } v)$ **then**
      return v with source $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$
let QUERYADV(m):
  if $m = \mathrm{Store(s)}$ **then**
    send $(\mathrm{STORE}, |s|)$ **to** $\mathcal{A}$ and **await**
  else if $m = \mathrm{Fetch}$ **then**
    send FETCH **to** $\mathcal{A}$ and **await**
  while next message m' on the input tape is not in $\mathbb{A}$ **do**
    ignore
  if $m' = \mathrm{Abort}$ **then**
    erase work tape of virtual ITI and **return** $\bot$
  else if $m' = \mathrm{Continue}$ **then**
    $\mathsf{halt} \leftarrow \bot$
    return $\top$

The execution loop of the above includes adversarial interactions as part of the enclave operation. In particular, when an enclave run by a corrupted party tries to interact with external memory by calling a Store or Fetch instruction, the shell sets flag $\mathsf{halt} \leftarrow \top$ and triggers procedure QUERYADV, which notifies the adversary, and relinquishes the activation token. On the shell's next activation, if it finds a message from the set $\mathbb{A}$, it resumes execution from where it last stopped. Otherwise, on any other input, it will abort (as long as flag $\mathsf{halt} = \top$): storing and fetching are *blocking*.

The adversary $\mathcal{A}$ only learns that enclave eid run by party pid in session idx is either trying to read from external storage, or that is writing some data and its size. $\mathcal{A}$ replies by sending a message of type $(\mathrm{RESUME}, \mathsf{eid}, \epsilon, a \in \mathbb{A})$ from corrupted party pid to G_{att}^{mod}. If $a = \mathrm{Continue}$, the shell continues executing from where it left off, storing bitstring s "ideally" (within its own internal variable mem). Otherwise, if $a = \mathrm{Abort}$, the enclave crashes, losing all memory stored within the virtual ITI's work tape. An Abort attack is not final: depending on the code of prog, the enclave can be resumed later on, and recover some partial state from the last value successfully stored to mem, if any.

D. Rollback Attacks

While the previous version of G_{att}^{mod} describes an adversary that is able to stop an enclave from storing any data to an external medium, the integrity and freshness of a successfully stored message is always guaranteed by a successful Fetch. We now explore a model with a slightly stronger adversary, who controls the storage medium and can overwrite the external memory location. Despite this, the enclave will not accept arbitrary messages, but only ones that were produced during a legitimate Store operation.

Bhatotia et al. [10] introduce $G_{\mathsf{att}}^{\mathsf{rollback}}$, a new variant of G_{att} that allows state continuity attacks. We now want to provide an equivalent attack as a G_{att}^{mod} adversarial oracle. The $G_{\mathsf{att}}^{\mathsf{rollback}}$ functionality tracks enclave state updates in a tree-like structure, and allows the adversary to specify an index for an arbitrary node in the tree to resume enclave execution from a specific snapshot. The tree allows the adversary to fork the enclave at an arbitrary state and maintain multiple copies that can progress independently.

Since G_{att}^{mod} no longer tracks the state of an enclave in a table $\mathcal{T}$, an instance of G_{att}^{mod} that supports Rollback or Fork instructions in $\mathbb{A}$ will require an alternative mechanism to maintain the state. We implement this through an enclave shell that executes each RESUME operation as a distinct virtual ITI. After the RESUME returns, the shell instantiates a new ITI by copying the last active configuration, and notifies the adversary of a unique pointer for that execution through an ITER message. When the adversary calls for a Rollback or Forking attack with a specific pointer, the shell can run the provided input on with adequately stale state by activating the older ITI that the pointer corresponds to.

$\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]$

The shell is defined for $\mathbb{O} = \mathbb{O}^{\mathrm{std}}$ and $\mathbb{A} = \{\mathrm{Rollback}, \mathrm{Fork}\}$
The extended identity of the shell is defined as $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

On message INSTALL *from* G_{att}^{mod}*:*
  generate nonce $c \xleftarrow{\$} \{0,1\}^{\lambda}$
  create virtual ITIs $(\mathsf{prog}, (\mathsf{eid}||\emptyset, \mathsf{idx}))$, $(\mathsf{prog}, (\mathsf{eid}||c, \mathsf{idx}))$
  if pid is corrupted **then send** $(\mathrm{ITER}, \emptyset, c)$ **to** $\mathcal{A}$
On message inp *from* G_{att}^{mod}*:*
  let $viti \leftarrow$ virtual ITI $(\mathsf{prog}, (\mathsf{eid}||c, \mathsf{idx}))$
  execute inp on $viti$
  generate nonce $c' \xleftarrow{\$} \{0,1\}^{\lambda}$
  create new virtual ITI $(\mathsf{prog}, (\mathsf{eid}||c', \mathsf{idx}))$
  copy work tape of $viti$ into $(\mathsf{prog}, (\mathsf{eid}||c', \mathsf{idx}))$
  if pid is corrupted **then send** (ITER, c, c') **to** $\mathcal{A}$
  $c \leftarrow c'$
On message $(\mathrm{ROLLBACK}, (i, \mathsf{inp}))$ *from* G_{att}^{mod}*:*
  run $(\mathsf{out}, (\mathrm{FORK}, i, i')) \leftarrow \mathrm{FORK}(i, \mathsf{inp})$
  $c \leftarrow i'$
  return $(\mathsf{out}, (\mathrm{ROLLBACK}, i, i'))$
On message $(\mathrm{FORK}, (i, \mathsf{inp}))$ *from* G_{att}^{mod}*:*
  let $viti \leftarrow$ virtual ITI $(\mathsf{prog}, (\mathsf{eid}||i, \mathsf{idx}))$
  if $viti$ exists **then**
    $\mathsf{out} \leftarrow \epsilon$
    if $\mathsf{inp} \neq \epsilon$ **then**
      execute inp on $viti$ and read subroutine output into out
    generate nonce $i' \xleftarrow{\$} \{0,1\}^{\lambda}$
    create new virtual ITI $(\mathsf{prog}, (\mathsf{eid}||i', \mathsf{idx}))$
    copy work tape of $viti$ to $(\mathsf{prog}, (\mathsf{eid}||i', \mathsf{idx}))$
    return $(\mathsf{out}, (\mathrm{FORK}, i, i'))$

The structure of each subroutine's extended identity involves appending a unique pointer nonce to the enclave id (the initial state is denoted by special pointer $\emptyset$). Variable c holds the pointer to the latest snapshot of the enclave accessible by a honest RESUME command. After each honest execution, the enclave creates a new UC subroutine by generating a new id and copies the execution tape of the subroutine c points to into this new copy, which is where the new instructions

will be executed. The adversary always learn the pointer generated for each iteration. If the adversary conducts a $\mathrm{Rollback}$ (by sending message $(\text{RESUME}, \mathsf{eid}, \mathsf{inp}, \mathrm{Rollback})$ from corrupted party pid to G_{att}^{mod}), c is overwritten with the pointer for an ITI whose memory state is copied from the one the adversary provided a pointer for. In a fork, c is not affected, but the adversary learns the new pointer i' it can access. In both cases, the shell returns to G_{att}^{mod} with the enclave output (if the attack also contained an instruction to execute) and auxiliary information on the new ITI pointer. Since the attack was successful, G_{att}^{mod} waits for the adversary to issue a CONTINUE message to finalise the return value and produce attestation (otherwise the RESUME call for G_{att}^{mod} never terminates).

It is clear from our formulation that a rollback is just a special case of a fork, where one of the two fork branches is not used again (in fact, on any ROLLBACK message, the shell executes the FORK procedure with the appropriate parameters). Distinguishing the two cases is primarily useful in the setting of a mobile adversary. While corrupted, a party can always choose the index for an enclave copy it wants to execute through the FORK command. When the party is no longer corrupted, however, the only copy of the enclave that can be executed is the one at index c. The adversary can thus use their last ROLLBACK to force the post-compromise party to execute the enclave from an arbitrary state, essentially erasing the access to any state that might have succeeded it.

E. Transparent enclaves

As we have discussed in the Background section II-B, some previous works in the literature have extended the G_{att}^{PST} model to capture additional types of side-channel attacks. We now adapt those extensions into G_{att}^{mod} shells. Tramèr et al. [78] provide a (local) UC functionality for attested execution with no confidentiality guarantees, later extended in [68, Section 8] to the global setting. Enclaves in this Transparent Enclave setting suffer from leakage of all internal memory, except for the master signing key for attestation. This allows integrating an enclave with such a leakage in protocols that only require the integrity provided by enclaves. The modeling of transparent enclave is a simple extension over that of G_{att}^{PST}: the output of each resume call is followed by the leakage of the random bits sampled by the enclave program. Knowing the inputs, randomness and the code of the program is sufficient to reconstruct its operation and internal memory for any randomised program, whereas deterministic programs are inherently transparent by default, since the adversary knows the code of the enclave when they install it.

In the language of G_{att}^{mod}, we state that for any attested functionality with $\mathrm{RandomSample} \in \mathbb{O}$ (and therefore any functionality where $\mathbb{O}^{\mathrm{std}} \subset \mathbb{O}$), we can realise a transparent version by including $\mathrm{TranspLeak}$ in the adversarial oracles $\mathbb{A}$. We recover the modelling from [78] and [68, Section 8.1] by letting the shell leak produce the entirety of the virtual ITI random tape to the adversary after each execution. On installation, enclaves start in the default non-transparent state, but once the adversary issues the $\mathrm{TranspLeak}$ attack, all further values of the tape are leaked.

$\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]$

The shell is defined for $\mathbb{O} = \mathbb{O}^{\mathrm{std}}$ and $\mathbb{A} = \{\mathrm{TranspLeak}\}$
The extended identity of the shell is defined as $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

On message INSTALL *from* G_{att}^{mod}*:*
  if virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ does not exist, create
  $\mathsf{transparent} \leftarrow \bot$

On message inp *from* G_{att}^{mod}*:*
  step through execution of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ on inp:
  for instruction i **do**
    if $\mathsf{i} \in \mathbb{O}^{\mathrm{std}}$ **then**
      allow $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ to execute i
    else if $\mathsf{i} =$(**return** v) **then**
      if $\mathsf{transparent} = \top \wedge$ pid is corrupted **then**
        send $(\text{LEAK}, \text{random tape of } (\mathsf{prog}, (\mathsf{eid}, \mathsf{idx})))$ **to** $\mathcal{A}$
      return v with source $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

On message $(\text{TRANSPLEAK}, \mathsf{inp})$ *from* G_{att}^{mod}*:*
  set $\mathsf{transparent} \leftarrow \top$
  return random tape of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$

A stronger type of leakage would leak the entirety of the virtual ITI's work tape. This would allow the adversary to recover any shared secret that predate the corruption attack. This can be implemented by simply appending the work tape to the LEAK message, or allow the adversary to apply standard UC passive corruption to the virtual ITI.

Almost-transparent and Semi-honest enclaves: Dörre, Mechler, and Müller-Quade [36] introduce two relaxations over the G_{att} functionality that aim to capture a middle ground between the side-channel free G_{att}^{PST} and transparent enclaves. Their models provides enclaves with access (in our language) to feature oracles for secure key exchange and symmetric encryption.

We now provide in Figure 4 an implementation for a shell that implement these cryptographic functions by outsourcing them to local functionality $\mathcal{F}_{\mathsf{crypto}}$ as defined by Küsters and Rausch [50].

Most of the oracle calls in the shell are simply forwarded from the enclave to the ideal functionality. $\mathrm{KeyExchange}$ is more interesting, as it is our first oracle call that involves direct communication between two enclaves. We implement a "synchronous" key exchange, in that we expect both enclaves to call the respective $\mathrm{KeyExchange}$ oracle to establish a channel. We do not provide a mechanism for enclaves to discover enclave IDs, and assume that they are provided by one of the other protocol inputs. The first enclave to call the oracle will stop accepting any further activations until the key exchange protocol completes (we refer to this enclave as the initiator). If the other enclave's shell receives a KEYEX message before its enclave has reached the $\mathrm{KeyExchange}$ call, it will store the received share in dictionary $\mathcal{E}$ to be retrieved at a later point. Once both parties have communicated their shares to each other, the shared key is computed by the $\mathcal{F}_{\mathsf{crypto}}$ functionality. Rather than returning it directly to the two enclaves, our shell uses it to derive a new symmetric key, which is what is obtained by both parties as the return value of $\mathrm{KeyExchange}$ (this step is necessary because $\mathcal{F}_{\mathsf{crypto}}$ does not allow using keys of type `dh-key` for symmetric operations). If either party running the enclave is corrupted, the adversary can learn that the key exchange is taking place and issue a HALT message. Additionally, the adversary might learn any other information leaked by $\mathcal{F}_{\mathsf{crypto}}$ and its leakage functions.

The addition of these oracles does not provide the enclave with meaningful new capabilities on its own, since an enclave can implement these operations as part of a library with access to randomness and attestation verification. However, it becomes significant once it is combined with the $\mathrm{TranspLeak}$ attack: by executing the secure operations "ideally" through oracles, the randomness needed to compute them is not leaked as part of the transparent attack. Dörre, Mechler, and Müller-Quade [36] define an enclave with access to both $\{\mathrm{KeyExchange}, \mathrm{SKEGen}, \mathrm{SKEEnc}, \mathrm{SKEDec}, \mathrm{ReleaseKey}\} \subset \mathbb{O}$ and $\mathrm{TranspLeak} \in \mathbb{A}$ to be a *almost-transparent enclave*, and show that it is possible to realise one-sided PSI between two parties running almost-transparent enclaves even if one of the parties is corrupted.

$\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]$

The shell is defined for $\mathbb{O} = \mathbb{O}^{\mathrm{std}} \cup \{\mathrm{KeyExchange}, \mathrm{SKEGen}, \mathrm{SKEEnc}, \mathrm{SKEDec}, \mathrm{ReleaseKey}\}$ and $\mathbb{A} = \{\mathrm{TranspLeak}, \mathrm{Halt}\}$
The extended identity of the shell is defined as $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

State variables	Description
$\mathcal{E} \leftarrow \{\}$	Stores Group elements received by other enclaves

On message INSTALL *from* G_{att}^{mod}*:*
- if virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ does not exist, create
- if ideal functionality $(\mathcal{F}_{\mathsf{crypto}}, (\mathsf{idx}, \bot))$ does not exist, create
- **send** GETDHGROUP **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** $(\mathrm{DHGROUP}, G, n, g)$
- **send** $(\mathrm{DHGROUP}, G, n, g)$ **to** $\mathcal{A}$

On message inp *from* G_{att}^{mod}*:*
- **if** $\mathsf{halt} = \top$ **then abort**
- step through execution of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ on inp:
- **for** instruction i **do**
  - **if** $\mathsf{i} \in \mathbb{O}^{\mathrm{std}}$ **then**
    - allow $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ to execute i
  - **else if** $\mathsf{i} = \mathrm{KeyExchange}(\mathsf{pid}', \mathsf{eid}')$ **then**
    - set $\mathsf{halt} \leftarrow \top$
    - **send** GENEXP **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** $(\mathrm{EXPOPOINTER}, ptr^e, g^e)$
    - **if** pid is corrupted **then**
      - **query** $\mathcal{A}$ **with** $(\mathrm{KEYEXTO}, \mathsf{pid}', \mathsf{eid}')$ and **receive the reply** *continue*
    - **if** $\mathcal{E}[\mathsf{pid}', \mathsf{eid}'] = \bot$ **then**
      - `// no stored keyshare for eid′, we are the initiatior`
      - **send** (KEYEX, g^e) **to** $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}'||\mathsf{pid}', \text{“att”}||\mathsf{idx}))$ and **await**
      - **while** next message on the input tape is not $(\mathrm{KEYEX}, \mathsf{pid}', \mathsf{eid}', h)$ **do** ignore
      - **send** $(\mathrm{BLOCKGROUPELEMENT}, h)$ **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** OK
    - **else**
      - `// eid′ was the key exchange initiatior, we already have h`
      - $h \leftarrow \mathcal{E}[\mathsf{pid}', \mathsf{eid}']$
    - **send** $(\mathrm{GENDHKEY}, ptr^e, h)$ **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** $(\mathrm{POINTER}, ptr^{dhk})$
    - **send** $(\mathrm{DERIVE}, ptr^{dhk}, \texttt{unauth-key})$ **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** $(\mathrm{POINTER}, ptr^{sk})$
    - set $\mathsf{halt} \leftarrow \bot$
    - write ptr^{sk} to subroutine output of virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$
  - **else if** $\mathsf{i} = \mathrm{SKEGen}$ **then**
    - **send** $(\mathrm{NEW}, \texttt{unauth-key})$ **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** $(\mathrm{POINTER}, ptr)$
    - write ptr to subroutine output of virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$
  - **else if** $\mathsf{i} = \mathrm{SKEEnc}(ptr, m)$ **then**
    - **send** (ENC, ptr, m) **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** $(\mathrm{CIPHERTEXT}, hdl)$
    - write hdl to subroutine output of virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$
  - **else if** $\mathsf{i} = \mathrm{SKEDec}(hdl, ct)$ **then**
    - **send** (DEC, ptr, ct) **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** $(\mathrm{PLAINTEXT}, m)$
    - write m to subroutine output of virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$
  - **else if** $\mathsf{i} = \mathrm{ReleaseKey}(ptr)$ **then**
    - **send** $(\mathrm{RETRIEVE}, ptr)$ **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** (KEY, k)
    - write k to subroutine output of virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$
  - **else if** $\mathsf{i} =$(**return** v) **then**
    - **return** v with source $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

On message HALT *from* G_{att}^{mod}*:*
- set $\mathsf{halt} \leftarrow \top$
- **return**

On message (KEYEX, h) *from* $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}'||\mathsf{pid}', \textit{“att”}||\mathsf{idx}))$*:*
- **if** $\mathsf{halt} = \bot$ **then**
  - `// we are not waiting for key exchange to complete;`
  - `// eid′ is the initiator`
  - **send** $(\mathrm{BLOCKGROUPELEMENT}, h)$ **to** $\mathcal{F}_{\mathsf{crypto}}$ and **receive** OK
  - $\mathcal{E}[\mathsf{pid}', \mathsf{eid}'] \leftarrow h$
- `// if the enclave is halted, eid is the initiator; on message KEYEX, we exit the loop to complete the key exchange`

Fig. 4. G_{att}^{mod} shell implementing key exchange and symmetric encryption feature oracles

Constructing a shell that realises the almost-transparent enclave can be achieved through a combination of the previous two shells, with the $\mathrm{TranspLeak}$ additionally leaking the state of the work tape of the program before the command was executed, and the return value of all secure operation oracles. Leaking these values is required for compatibility with Transparent Enclaves, since the randomness leakage is not sufficient to reconstruct deterministic computation that includes values the enclave obtained through the secure operations.

There are some minor differences between our model and the one in [36]: in their version of almost-transparent enclaves, once the initiator issues a KEYEXCHANGE command, the receiving enclave is immediately notified and provided the symmetric key. Therefore, the initiator program needs to be run first (a natural constraint in their protocol). An additional difference from their model is our use of the idealised $\mathcal{F}_{\mathsf{crypto}}$ for all operations, rather than using a mix of ideal key exchange and concrete symmetric operations in their model. Therefore, we have to do an additional step to derive a symmetric key, rather than using the shared DH key directly.

The second relaxation, *semi-honest* enclaves, captures an adversarial manufacturer who is able to adaptively break into enclaves and extract historical transaction data. Note that in this setting, the party running the enclave does not need to be corrupted for the leakage to occur i.e. the adversary can cause leakage for any enclave run by a honest party. Despite the extreme vulnerability of this type of enclave to an adversarial manufacturer, it is still useful to construct some classes of private set intersection (distinct from the ones in the previous setting).

The shell for a Semi-honest enclave is defined as follows

$\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]$

The shell is defined for $\mathbb{O} = \mathbb{O}^{\mathrm{std}}$ and $\mathbb{A} = \{\mathrm{CompleteLeak}\}$
The extended identity of the shell is defined as $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

On message INSTALL *from* G_{att}^{mod}*:*
  if virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ does not exist, create
  $\mathsf{rec} \leftarrow []$
On message $(\text{RESUME}, \mathsf{inp})$ *from* G_{att}^{mod}*:*
  step through execution of $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ on inp:
  for instruction i **do**
    if $\mathsf{i} \in \mathbb{O}^{\mathrm{std}}$ **then**
      allow $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ to execute i
    else if $\mathsf{i} =$(**return** v) **then**
      $\mathsf{rec} \leftarrow \mathsf{rec} \parallel (\mathsf{inp}, args,$ virtual ITI work tape)
      return v with source $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$
On message COMPLETELEAK *from* $\mathcal{A}$*:*
  return rec

The definition of the shell is quite simple, as it merely records the output of each resume execution and returns it to the adversary when it issues the $\mathrm{CompleteLeak}$ command. The message is sent directly to the shell rather than through a corrupted resume call to represent that it doesn't have to be issued by the calling party.

F. Shared Registry

We now give a shell to implement a single-writer multi-reader registry functionality for any subset of enclaves. The registry contains a linearisable list of values that any enclave in the set can read, but only one enclave can write into (in this case, the first enclave to complete a write). We give the adversary the ability to temporarily block or permanently censor corrupted parties, such that they can not access the registry for reading. If the number of censored replicas is greater than a certain quorum Q (a percentage of the registered parties) the registry is no longer able to guarantee termination of read/write operation, and will produce an error instead. If the writing enclave is censored, all subsequent write calls will fail but read calls from other enclaves can continue. The registry can be thought of as a shared single-writer ledger whose storage is distributed between enclaves, and is synchronised through a consensus mechanism; if less than Q of the total enclaves return a value, there are not enough live enclaves to establish consensus and thus the protocol terminates.

We define the following shell, where the adversarial oracle $\mathrm{Censor}_{\mathrm{Q}}$ is parameterised by Q.

$\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}]$

The shell is defined for $\mathbb{O} = \mathbb{O}^{\mathrm{std}} \cup \{\mathrm{Read}, \mathrm{Write}\}$ and $\mathbb{A} = \{\mathrm{Block}, \mathrm{Censor}_{\mathrm{Q}}\}$
The extended identity of the shell is defined as $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

On message INSTALL *from* G_{att}^{mod}*:*
  $j \leftarrow \perp$
  if virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$ does not exist, create
  if ideal functionality $(\mathrm{RegCoord[Q]}, (\mathsf{idx}, \perp))$ does not exist, create
On message inp *from* G_{att}^{mod}*:*
  let $viti \leftarrow$ virtual ITI $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$
  step through execution of $viti$ on inp:
  for instruction i **do**
    if $\mathsf{i} \in \mathbb{O}^{\mathrm{std}}$ **then**
      allow $viti$ to execute i
    else if $\mathsf{i} = \{\mathrm{Read}, (\mathrm{Write}, v)\}$ **then**
      if $j = \perp$ **then**
        send JOIN **to** $\mathrm{RegCoord[Q]}$ **on behalf of** $viti$
        $j \leftarrow \top$
      send i **to** $\mathrm{RegCoord[Q]}$ **through** $viti$ and **receive** v
      write v to subroutine output of virtual ITI
    else if $\mathsf{i} =$(**return** v) **then**
      return v with source $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$
On message $(\text{CENSOR}, \epsilon)$ *from* G_{att}^{mod}*:*
  send $(\text{CENSOR}, \mathsf{pid})$ **to** $\mathrm{RegCoord[Q]}$

Functionality $\mathrm{RegCoord[Q]}$

State variables	Description
$P \leftarrow []$	List of enclaves participating in the registry
$C \leftarrow []$	List of censored enclaves
$V \leftarrow []$	List of registry values over time
$w \leftarrow \perp$	identity of writer enclave

On message JOIN *from* $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$*:*
  $P \leftarrow P \cup (\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$
  send $(\text{JOIN}, (\mathsf{prog}, (\mathsf{eid}, \mathsf{idx})))$ **to** $\mathcal{A}$
On message (cmd, v) *from* $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$*:*
  if eid is running on a corrupted party **then**
    query $\mathcal{A}$ **with** $(\text{READ}, (\mathsf{prog}, (\mathsf{eid}, \mathsf{idx})))$
    receive b **from** $\mathcal{A}$
  if $b \neq \top \wedge \frac{|C|}{|P|} < Q$ **then**
    if cmd=WRITE **then**
      if $w = \perp$ **then** $w \leftarrow P$

if $w \neq P \vee P \in C$ **then return** Fail
$V \leftarrow V \parallel v$
send $(\text{CMD}, V, (\mathsf{prog}, (\mathsf{eid}, \mathsf{idx})))$ **to** $\mathcal{A}$
return V
else
return Fail
On message HEALTHCHECK *from* $(\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$*:*
return $|P|, |C|$
On message $(\text{CENSOR}, (\mathsf{prog}, (\mathsf{eid}, \mathsf{idx})))$ *from* G_{att}^{mod}*:*
if eid is running on a corrupted enclave **then**
$C \leftarrow C \parallel (\mathsf{prog}, (\mathsf{eid}, \mathsf{idx}))$

The above functionality allows any enclave shell to join the protocol as a registry party. The first shell who writes to the registry is locked in as w, the writer. Thereafter, only w can issue a new WRITE, which appends the value to the end of the registry, and all other registered parties receive the entirety of the registry on every READ[1] On any read and write, a corrupted party will query the adversary on whether they are allowed to proceed. The adversary can also permanently block an enclave by issuing a Censor message. If too many parties have been censored (i.e. the ratio between the number of censored parties and total registered parties is greater than Q), it is impossible for the registry to guarantee that the registry value is still safe, and the functionality fails.

We assume the functionality has access to some directory ITI that records whether enclaves are run by corrupted parties.

V. RELATIONSHIPS BETWEEN G_{att}^{mod} VARIANTS

Having defined examples of G_{att}^{mod} instantiations with different sets $\mathbb{O}, \mathbb{A}$, we are now interested in exploring how they relate to each other.

Given two versions of G_{att}^{mod} which sign over the same measurement functions, a "weaker" setup G_{att} that has either fewer features or more attacks can UC-emulate the stronger one G'_{att} through a dedicated "wrapper" protocol that reproduces the missing feature oracle, or mitigate the attacks offered by the additional adversarial oracle. Depending on the interfaces it is trying to bridge, an oracle or protection mechanism can be implemented by just running some additional computation within the enclave itself, by calling out to a library running within an assisting enclave on the same party, or by conducting an interactive protocol with multiple remote parties. We can represent these type of runtime behavior as a UC protocol.

We now sketch how to design such a protocol for any two $G_{\mathsf{att}}, G'_{\mathsf{att}}$ setups. In this section, we show how to remove adversarial interfaces that exist in G_{att} to realise the stronger G'_{att} in a simplified setting, where the protection mechanism can be implemented as additional runtime code in the target enclave itself. In Appendix B, we consider both the more general case where the protocol is distributed between different enclaves or even between different parties, as well as the equivalent protocol for adding a new feature to a G_{att}^{mod} setup Our treatment aims to be generic and provide a universal compiler protocol, but we can not prove security in this general setting, as it might not be possible to create such a protocol for all G_{att}^{mod} pairings. Instead, we give conjectures preconditioned on the existence of a secondary protocol that, for a G_{att}^{mod} setup with specific parameters, provides a mechanism to add the feature or removing the attack from the starting setup (i.e. making the shell behaviour between two G_{att}^{mod} instances indistinguishable). We also provide a matching proof template that can be instantiated in the specific cases considered (as we do later in this section).

[1]The functionality could be made more efficient by keeping track of what values have been read by each group member, and only downloading the difference on a read.

Given two (modular) implementations of attested executions $G_{\mathsf{att}}, G'_{\mathsf{att}}$ with adversarial interfaces $\mathbb{A}, \mathbb{A}'$ respectively, and shared $\mathbb{O}$ and $\mathbb{S}$, we define a wrapper protocol $\mathcal{W}$ that uses G_{att} as a subroutine and UC-realizes G'_{att}. $\mathcal{W}$ simply forwards all messages to G_{att}, with the exception of INSTALL commands, where it replaces the argument prog with $\mathrm{W}^A[\mathsf{prog}]$. Wrapper enclave $\mathrm{W}^A[\cdot]$ protects the execution of prog by running the protocol that defends an enclave against attacks $A = \mathbb{A} \setminus \mathbb{A}'$, and must be defined for the specific combination of $\mathbb{O}$, $\mathbb{A}$, and $\mathbb{S}$. The wrapper enclave will run, on a RESUME call, any combination of the prog code with additional local computation and calls to oracles in $\mathbb{O}$, provided that the distribution of outputs of $\mathrm{W}^A[\mathsf{prog}]$ under G_{att} is equal to that of prog under G'_{att}.

The protocol $\mathcal{W}$ only protects against the attacks in A; all other attacks in $\mathbb{A}'$ are still allowable in both worlds.

Protocol $\mathcal{W}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}', \mathbb{S}, \mathrm{W}^A[\cdot]]$

$\mathbb{A}$ is defined as $A \cup \mathbb{A}'$
On message $(m \in \{\text{INITIALISE,GETPK,VERIFY,RESUME}\}, args)$*:*
send $(m, args)$ **to** $G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}]$ and **forward** response
On message $(\text{INSTALL}, \mathsf{idx}, \mathsf{prog})$*:*
if $\mathsf{prog} = \mathrm{W}^A[\mathsf{prog}']$ **then**
send $(\text{INSTALL}, \mathsf{idx}, \mathsf{prog})$ **to** $G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}]$
else
send $(\text{INSTALL}, \mathsf{idx}, \mathrm{W}^A[\mathsf{prog}'])$ **to** $G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}]$
forward response

We now provide a conjecture that $\mathcal{W}$ in the presence of G_{att} is sufficient to securely realise G'_{att}.Without a precise definition of which attack $\mathcal{W}$ is trying to defend against (based on the definition of $\mathrm{W}^A[\cdot]$ for the specific $\mathbb{O}, \mathbb{A}$ sets), it is difficult to provide a proof (which is why we do not give a generic theorem). Instead, we will provide some guidelines on how a simulator for theorems that are special instances of this conjecture might be structured.

Conjecture 1

Let $G_{\mathsf{att}} = G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}]$ *and* $G'_{\mathsf{att}} = G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}', \mathbb{S}]$ *such that* $\mathbb{A} \setminus \mathbb{A}' = A$*. For any enclave wrapper* $\mathrm{W}^A[\cdot]$ *which implements the difference in behaviour between the shells* $\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\cdot]$ *and* $\mathrm{sh}_{\mathbb{O},\mathbb{A}'}[\cdot]$*, protocol* $\mathcal{W}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}, \mathrm{W}^A[\cdot]]$ *in the presence of* G_{att} *UC-emulates* G'_{att}

We give a graphical representation of the UC-emulation statement in Figure 5. The simulator begins by observing, during the global functionality initialisation phase, the signature algorithm s chosen by the environment through the dummy adversary, and provides G'_{att} with a new algorithm s' which uses s to sign the transformation $F(\mathsf{meas})$. $F()$ takes the measurement string produced by $\mathbb{S}$ that contains an identifier for program prog, and replaces it with an identifier for $\mathrm{W}^A[\mathsf{prog}]$, as discussed in Section III. Once G'_{att} has been initialised, the simulator simply blocks the installation of any unwrapped programs from protocol parties, and replaces installations of wrapped programs with the corresponding unwrapped program in G'_{att}. If the (local) adversary attempts to install an unwrapped program to G_{att} directly, the simulator can run the program "in its head" without going through G'_{att}, and use the algorithm s provided by the dummy adversary for the environment to produce plausible attestation signatures for the unwrapped code.

The real and ideal world are indistinguishable due to the inability of the adversary to perform attacks in A, and from all attestation signatures containing references to the wrapped version of a program (regardless of whether it is really wrapped or not). Since honest parties in $\mathcal{W}$ do not install any unwrapped program, and no external session will have direct access to G_{att} since it is instantiated as a

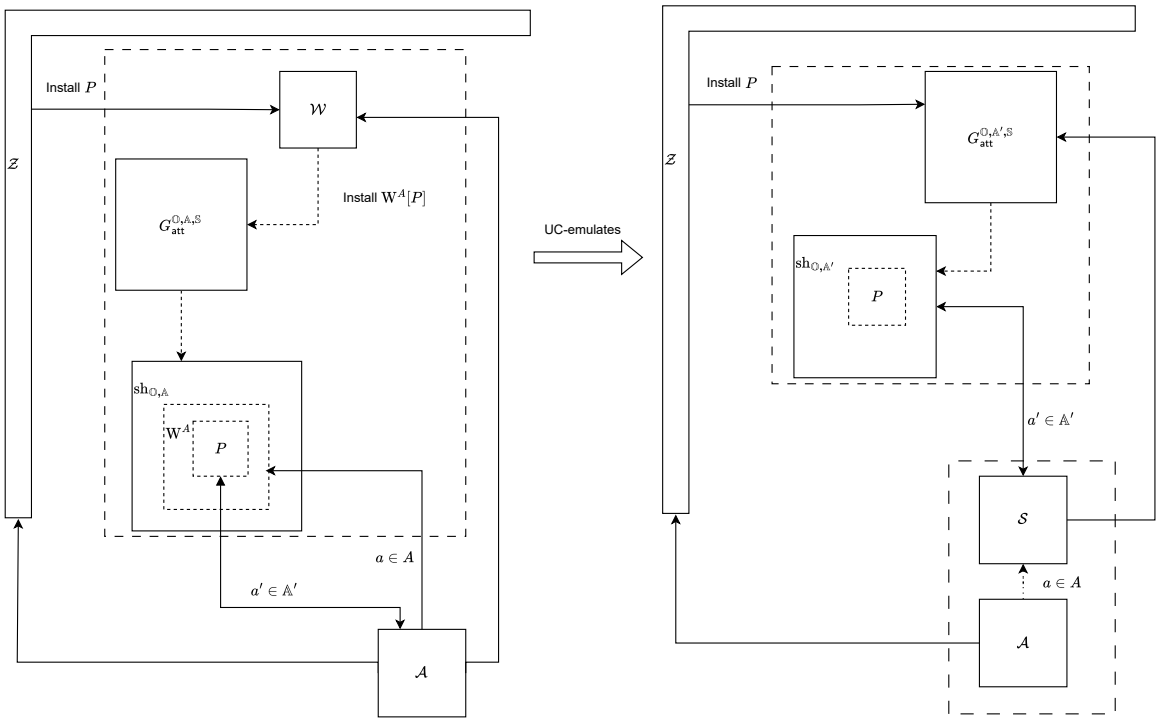

Fig. 5. Protocol $\mathcal{W}$ can add a shell to $G_{\mathsf{att}}^{\mathbb{O},\mathbb{A},\mathbb{S}}$ enclaves to UC-emulate the behaviour of $G_{\mathsf{att}}^{\mathbb{O},\mathbb{A}',\mathbb{S}}$ by blocking any attacks in set $A = \mathbb{A}' \setminus \mathbb{A}$

$\mathcal{W}$ subroutine, the only party to directly install unwrapped programs on a G_{att}^{mod} functionality will be the corrupted parties, whose return values to the environment is controlled by the simulator and can thus be replaced by the simulated execution described above. Obtaining a simulated signature does not provide the environment with additional distinguishing powers, since they would not verify through either world's G_{att}^{mod} verification interface, but look like genuine signatures to environment when running the local verification algorithm that corresponds to s.

We believe our simulation template to be sufficient to prove security for most natural comparisons of TEE setups. However, depending on the nature of the programs installed, the wrapper code, or the shared oracles between the two setups, a different simulation strategy might be needed. For example, if the adversary is able to directly observe (through an oracle) the source code of an enclave while it is executing, the simulation will not work. It might be possible for some of these cases where our simulation strategy does not work to add some backdoor code in the $\mathrm{W}^I[\cdot]$ description to give the simulator some additional powers (see Bhatotia et al. [10] and Pass, Shi, and Tramèr [69]).

Before showing a concrete instantiation of the conjecture that can be shown secure using this simulator template, we first remark that the composition theorem holds for all possible combinations of $G_{\mathsf{att}}, G'_{\mathsf{att}}$ (a non-trivial statement according to Badertscher, Hesse, and Zikas [4]).

Theorem 1

Let $G_{\mathsf{att}}, G'_{\mathsf{att}}, \mathcal{W}$ be any G_{att}^{mod} setups and a wrapper protocol such that Conjecture 1 holds. For any protocol ρ in the presence of G'_{att} that UC-emulates $\mathcal{F}$ in the presence of G'_{att}, ρ in the presence of $\mathcal{W}$ UC-emulates $\mathcal{F}$ in the presence of $\mathcal{W}$.

The theorem holds because the adversarial interface is smaller in the ideal world, so there is no additional attack used by the ρ to $\mathcal{F}$ simulator which would no longer be available by replacing the setup with the protocol. This is the inverse scenario from what [4] are concerned with, where the real world global protocol includes fewer attacks that the ideal world global functionality. Therefore, the theorem holds due to the composition theorem of [4, Theorem 3.10], as the ρ to F simulator is $\mathcal{W} \setminus \mathbb{A}'$-agnostic (i.e. the simulator does not interact with $\mathcal{W}$ except for using adversarial interfaces in $\mathbb{A}'$ - that is, everything except for A). This is true because A is not a valid adversarial interface in G'_{att}. Therefore, if simulator of the pre-condition is able to simulate the protocol without using A, the same simulator will equally apply to the statement where G'_{att} has been replaced with $\mathcal{W}$.

A. Rollback protection

We now outline our example instantiation of Conjecture 1 that addresses the rollback attack techniques described in [10] and Section IV-D by relying on the non-volatile storage feature of Section IV-C to realise a *record-then-execute* trusted state digest (based on a well known observation in the literature [64, 67]). Although this protocol equally applies to the related class of forking attacks, we do not explicitly address them in this section for simplicity. More specifically, we present the simple wrapper protocol $\mathcal{W}$, which removes the $\mathrm{Rollback}$ interface from the hybrid attestation functionality G_{att}, whose set $\mathbb{A}$ includes Abort. To construct the protocol, we require our target enclave functionality G_{att} to support the trusted $\mathrm{Store}, \mathrm{Fetch}$ interfaces described in Section IV-C, as well as an oracle Meas, which returns a digest (such as a hash) over the state of the enclave's virtual ITI.

The intuition for the protocol is that for every instruction to install an enclave with code prog, it can install an external wrapper enclave with code $\mathrm{W}^{Rollback}[\mathsf{prog}]$ instead. The wrapper will store the digest of the latest copy of the internal enclave measurement in persistent storage at the end of every RESUME. When enclave execution starts,

the wrapper can fetch the stored measurement digest and compare it with the measurement for the current state as returned by Meas. If the two states match, the enclave can be safely executed; otherwise, the state must have been tampered with, and the function aborts. We denote this sequence of operations as procedure MEASEXEC. If every resume operation uses MEASEXEC, the adversary is not able to execute a rollback attack, but will effectively abort the enclave. Defining a rollback protection protocol by relying on the usage of safe memory might seem like a circular definition - if the enclave has access to trustworthy $\mathrm{Store}, \mathrm{Fetch}$ oracles, why not just store the entirety of memory using this interface? We believe that the current setting is still valuable, as it minimises the size of data stored in trusted memory, but more importantly provides a clear example of our theorem for expository purposes. Note that, there exist several protocols that claim to resolve rollback attacks without access to trusted storage (for example [1, 11, 34, 47, 60, 64, 81]). We leave the formalisation of such a protocol to remove the adversarial rollback interface as future work.

Shell $\mathrm{W}^{\mathrm{Rollback}}[\mathsf{prog}]$

The identity of the shell is $(\mathsf{eid} \parallel c, \mathsf{idx})$
The parent shell extended identity is $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathrm{W}^{\mathrm{Rollback}}[\mathsf{prog}]], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$ for $\{\mathrm{Fetch}, \mathrm{Store}, \mathrm{Meas}\} \subset \mathbb{O}$ and $\{\mathrm{Abort}, \mathrm{Rollback}\} \subset \mathbb{A}$
On message INIT *from* $(\mathsf{eid}||\mathsf{pid}, \textit{“att”}||\mathsf{idx})$*:*
  install virtual ITI $(\mathsf{prog}, (\mathsf{eid}||c||\text{“wrapped”}, \mathsf{idx}))$
  let $m \leftarrow \mathrm{Meas}()$
  $\mathrm{Store}(m)$
On message inp *from* $(\mathsf{eid}||\mathsf{pid}, \textit{“att”}||\mathsf{idx})$*:*
  $v \leftarrow$MEASEXEC(inp)
  if $v = \textit{“abort”}$ **then**
    erase the virtual ITI work tape and **abort**
  else return v
let MEASEXEC(inp):
  $m \leftarrow \mathrm{Fetch}(), m' \leftarrow \mathrm{Meas}()$
  if $m \neq m'$ **then return** $\textit{“abort”}$
  step through execution of $(\mathsf{prog}, (\mathsf{eid}||c||\text{“wrapped”}, \mathsf{idx}))$ on inp:
  for instruction i **do**
    if $\mathsf{i} =$ (**return** v) **then**
      $b \leftarrow \mathrm{Write}(\mathrm{Meas}())$
      assert $b = OK$
    else allow $(\mathsf{prog}, (\mathsf{eid}||c||\text{“wrapped”}, \mathsf{idx}))$ to execute i

Fig. 6. The $\mathrm{W}^{\mathrm{Rollback}}[\cdot]$ enclave shell installed by protocol $\mathcal{W}$ for rollback iteration c of enclave eid installed by party pid for session idx

It is convenient for our purposes to model the code of $\mathrm{W}^{\mathrm{Rollback}}[\cdot]$ using a UC shell (as presented in Figure 6), since its behaviour is similar to some of the shells we constructed in the previous section. The two types of shell are complementary: UC structured protocols support nesting shells, so we instantiate the $\mathrm{W}^{\mathrm{Rollback}}[\cdot]$ as a subroutine of $\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\cdot]$. The $\mathcal{W}$ protocol is still instructing G_{att} to install the full code for $\mathrm{W}^{\mathrm{Rollback}}[\mathsf{prog}]$ but using the shell means that we don't have to define its full source code and how it instruments the code of the inner enclave.

The shell runs with ID $(\mathsf{eid}||c, \mathsf{idx})$ as a subroutine to the top level shell $(\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx})$, which implements the full oracle set, including the attack $\mathrm{Rollback} \in \mathbb{A}$. Whenever the inner shell calls to a feature oracle, its execution is paused by $(\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx})$, which computes the oracle value and writes it on the subroutine output tape. Shell $(\mathsf{eid}||c, \mathsf{idx})$ is oblivious to this mechanism, and can simply call the oracles as if they were local subroutines. The identity of the shell includes counter c because the shell is one of the copies created by the shell $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathsf{prog}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$ from Section IV-D to enable rollbacks. All shell copies created for new RESUME iterations share the same storage interface for $\mathrm{Store}, \mathrm{Fetch}$. $(\mathsf{eid}||c, \mathsf{idx})$ instantiates a subroutine $(\mathsf{prog}, (\mathsf{eid}||c||\text{“wrapped”}, \mathsf{idx}))$ to execute the code of prog. For most of the execution of prog, it allows the internal subroutine to run. Since the execution of $(\mathsf{eid}||c, \mathsf{idx})$ is also running within an execution loop of $(\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx})$, whenever $(\mathsf{prog}, (\mathsf{eid}||c||\text{“wrapped”}, \mathsf{idx}))$ calls an oracle, $(\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx})$ will pause the execution of both subroutines to provide a return value. Likewise, if the adversary issues an Abort attack, $(\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx})$ will handle it directly. We now give the formal statement of UC emulation.

Theorem 2

Take $G_{\mathsf{att}} = G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}]$ *such that* $\{\mathrm{Rollback}, \mathrm{Abort}\} \subset \mathbb{A}$ *and* $\{\mathrm{Store}, \mathrm{Fetch}, \mathrm{Meas}\} \subset \mathbb{O}$*; and* $G'_{\mathsf{att}} = G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}', \mathbb{S}]$ *where* $\mathbb{A}' = \mathbb{A} \setminus \{\mathrm{Rollback}\}$.
Protocol $\mathcal{W}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}', \mathbb{S}, \mathrm{W}^{\mathrm{Rollback}}[\mathsf{prog}]]$ *in the presence of* G_{att} *UC-emulates* G'_{att}.

We can show that the theorem holds by constructing a simulator for the ideal world experiment, based on the simulator template presented earlier in this Section. The intuition for the simulation strategy is that, for any unwrapped enclave installed by the environment, the simulator running the enclave in its head can turn any attempt to conduct a rollback into the equivalent Abort attack in the ideal world by detecting a digest mismatch. We give a full construction of the simulator and proof of security in Appendix C.

REFERENCES

[1] Sebastian Angel et al. “Nimble: Rollback Protection for Confidential Cloud Services”. In: *17th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2023, Boston, MA, USA, July 10-12, 2023*. Ed. by Roxana Geambasu and Ed Nightingale. USENIX Association, 2023, pp. 193–208. URL: https://www.usenix.org/conference/osdi23/presentation/angel.

[2] Pedro Antonino, Wojciech Aleksander Woloszyn, and A. W. Roscoe. “Guardian: Symbolic Validation of Orderliness in SGX Enclaves”. In: *Proceedings of the 2021 on Cloud Computing Security Workshop*. CCS '21. ACM, Nov. 2021. DOI: 10.1145/3474123.3486755. URL: http://dx.doi.org/10.1145/3474123.3486755.

[3] Michael Backes, Birgit Pfitzmann, and Michael Waidner. *The Reactive Simulatability (RSIM) Framework for Asynchronous Systems*. Cryptology ePrint Archive, Report 2004/082. 2004. URL: https://eprint.iacr.org/2004/082.

[4] Christian Badertscher, Julia Hesse, and Vassilis Zikas. *On the (Ir)Replaceability of Global Setups, or How (Not) to Use a Global Ledger*. Cryptology ePrint Archive, Report 2020/1489. 2020. URL: https://eprint.iacr.org/2020/1489.

[5] Christian Badertscher, Julia Hesse, and Vassilis Zikas. “On the (Ir)Replaceability of Global Setups, or How (Not) to Use a Global Ledger”. In: *TCC 2021, Part II*. Ed. by Kobbi Nissim and Brent Waters. Vol. 13043. LNCS. Springer, Cham, Nov. 2021, pp. 626–657. DOI: 10.1007/978-3-030-90453-1_22.

[6] Christian Badertscher et al. “Universal Composition with Global Subroutines: Capturing Global Setup Within Plain UC”. In: *TCC 2020, Part III*. Ed. by Rafael Pass and Krzysztof Pietrzak. Vol. 12552. LNCS. Springer, Cham, Nov. 2020, pp. 1–30. DOI: 10.1007/978-3-030-64381-2_1.

[7] Christian Badertscher et al. *Universal Composition with Global Subroutines: Capturing Global Setup within plain UC*. Cryptology ePrint Archive, Report 2020/1209. 2020. URL: https://eprint.iacr.org/2020/1209.

[8] Manuel Barbosa et al. *Foundations of Hardware-Based Attested Computation and Application to SGX*. Cryptology ePrint Archive, Report 2016/014. 2016. URL: https://eprint.iacr.org/2016/014.

[9] Saskia Bayreuther et al. "Hidden Δ-fairness: A Novel Notion for Fair Secure Two-Party Computation". In: *IACR Cryptol. ePrint Arch.* (2024), p. 587. URL: https://eprint.iacr.org/2024/587.

[10] Pramod Bhatotia et al. "Steel: Composable Hardware-Based Stateful and Randomised Functional Encryption". In: *PKC 2021, Part II*. Ed. by Juan Garay. Vol. 12711. LNCS. Springer, Cham, May 2021, pp. 709–736. DOI: 10.1007/978-3-030-75248-4_25.

[11] Marcus Brandenburger et al. "Rollback and Forking Detection for Trusted Execution Environments Using Lightweight Collective Memory". In: *CoRR* (2017). arXiv: 1701.00981 [cs.DC]. URL: http://arxiv.org/abs/1701.00981v2.

[12] Konstantinos Brazitikos and Vassilis Zikas. "General Adversary Structures in Byzantine Agreement and Multi-Party Computation with Active and Omission Corruption". In: *IACR Cryptol. ePrint Arch.* (2024), p. 209. URL: https://eprint.iacr.org/2024/209.

[13] Jan Camenisch, Manu Drijvers, and Björn Tackmann. *Multi-Protocol UC and its Use for Building Modular and Efficient Protocols*. Cryptology ePrint Archive, Report 2019/065. 2019. URL: https://eprint.iacr.org/2019/065.

[14] Jan Camenisch et al. *iUC: Flexible Universal Composability Made Simple*. Cryptology ePrint Archive, Report 2019/1073. 2019. URL: https://eprint.iacr.org/2019/1073.

[15] Ran Canetti. *Universally Composable Security: A New Paradigm for Cryptographic Protocols*. Cryptology ePrint Archive, Report 2000/067. 2000. URL: https://eprint.iacr.org/2000/067.

[16] Ran Canetti. "Universally Composable Security: A New Paradigm for Cryptographic Protocols". In: *42nd FOCS*. IEEE Computer Society Press, Oct. 2001, pp. 136–145. DOI: 10.1109/SFCS.2001.959888.

[17] Ran Canetti. *Universally Composable Signatures, Certification and Authentication*. Cryptology ePrint Archive, Report 2003/239. 2003. URL: https://eprint.iacr.org/2003/239.

[18] Ran Canetti, Asaf Cohen, and Yehuda Lindell. "A Simpler Variant of Universally Composable Security for Standard Multiparty Computation". In: *CRYPTO 2015, Part II*. Ed. by Rosario Gennaro and Matthew J. B. Robshaw. Vol. 9216. LNCS. Springer, Berlin, Heidelberg, Aug. 2015, pp. 3–22. DOI: 10.1007/978-3-662-48000-7_1.

[19] Ran Canetti and Marc Fischlin. "Universally Composable Commitments". In: *CRYPTO 2001*. Ed. by Joe Kilian. Vol. 2139. LNCS. Springer, Berlin, Heidelberg, Aug. 2001, pp. 19–40. DOI: 10.1007/3-540-44647-8_2.

[20] Ran Canetti, Eyal Kushilevitz, and Yehuda Lindell. "On the Limitations of Universally Composable Two-Party Computation Without Set-Up Assumptions". In: *Journal of Cryptology* 19.2 (Apr. 2006), pp. 135–167. DOI: 10.1007/s00145-005-0419-9.

[21] Ran Canetti and Tal Rabin. "Universal Composition with Joint State". In: *CRYPTO 2003*. Ed. by Dan Boneh. Vol. 2729. LNCS. Springer, Berlin, Heidelberg, Aug. 2003, pp. 265–281. DOI: 10.1007/978-3-540-45146-4_16.

[22] Ran Canetti, Daniel Shahaf, and Margarita Vald. "Universally Composable Authentication and Key-Exchange with Global PKI". In: *PKC 2016, Part II*. Ed. by Chen-Mou Cheng et al. Vol. 9615. LNCS. Springer, Berlin, Heidelberg, Mar. 2016, pp. 265–296. DOI: 10.1007/978-3-662-49387-8_11.

[23] Ran Canetti et al. *Universally Composable End-to-End Secure Messaging*. Cryptology ePrint Archive, Report 2022/376. 2022. URL: https://eprint.iacr.org/2022/376.

[24] Ran Canetti et al. "Universally Composable End-to-End Secure Messaging". In: *CRYPTO 2022, Part II*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Vol. 13508. LNCS. Springer, Cham, Aug. 2022, pp. 3–33. DOI: 10.1007/978-3-031-15979-4_1.

[25] Ran Canetti et al. "Universally Composable Security with Global Setup". In: *TCC 2007*. Ed. by Salil P. Vadhan. Vol. 4392. LNCS. Springer, Berlin, Heidelberg, Feb. 2007, pp. 61–85. DOI: 10.1007/978-3-540-70936-7_4.

[26] Ran Canetti et al. "Using Universal Composition to Design and Analyze Secure Complex Hardware Systems". In: *2020 Design, Automation & Test in Europe Conference & Exhibition, DATE 2020, Grenoble, France, March 9-13, 2020*. IEEE, 2020, pp. 520–525. ISBN: 978-3-9819263-4-7. DOI: 10.23919/DATE48585.2020.9116295. URL: https://doi.org/10.23919/DATE48585.2020.9116295.

[27] Shanwei Cen and Bo 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.

[28] Raymond Cheng et al. "Ekiden: A Platform for Confidentiality-Preserving, Trustworthy, and Performant Smart Contracts". In: *IEEE European Symposium on Security and Privacy, EuroS&P 2019, Stockholm, Sweden, June 17-19, 2019*. IEEE, 2019, pp. 185–200. ISBN: 978-1-7281-1148-3. DOI: 10.1109/EuroSP.2019.00023. URL: https://doi.org/10.1109/EuroSP.2019.00023.

[29] Arka Rai Choudhuri et al. "Fairness in an Unfair World: Fair Multiparty Computation from Public Bulletin Boards". In: *ACM CCS 2017*. Ed. by Bhavani M. Thuraisingham et al. ACM Press, Oct. 2017, pp. 719–728. DOI: 10.1145/3133956.3134092.

[30] Michele Ciampi, Aggelos Kiayias, and Yu Shen. "Universal Composable Transaction Serialization with Order Fairness". In: *44th Annual International Cryptology Conference*. Ed. by Leonid Reyzin and Douglas Stebila. LNCS. Aug. 2024.

[31] Michele Ciampi, Yun Lu, and Vassilis Zikas. "Collusion-Preserving Computation without a Mediator". In: *CSF 2022 Computer Security Foundations Symposium*. IEEE Computer Society Press, Aug. 2022, pp. 211–226. DOI: 10.1109/CSF54842.2022.9919678.

[32] Victor Costan and Srinivas Devadas. *Intel SGX Explained*. Cryptology ePrint Archive, Report 2016/086. 2016. URL: https://eprint.iacr.org/2016/086.

[33] Poulami Das et al. "FastKitten: Practical Smart Contracts on Bitcoin". In: *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*. Ed. by Nadia Heninger and Patrick Traynor. USENIX Association, 2019, pp. 801–818. URL: https://www.usenix.org/conference/usenixsecurity19/presentation/das.

[34] Baltasar Dinis, Peter Druschel, and Rodrigo Rodrigues. "RR: A Fault Model for Efficient TEE Replication". In: *NDSS 2023*. The Internet Society, Feb. 2023.

[35] Natnatee Dokmai et al. "Privacy-preserving genotype imputation in a trusted execution environment". In: *Cell Systems* 12.10 (Oct. 2021), 983–993.e7. ISSN: 2405-4712. DOI: 10.

1016/j.cels.2021.08.001. URL: http://dx.doi.org/10.1016/j.cels.2021.08.001.

[36] Felix Dörre, Jeremias Mechler, and Jörn Müller-Quade. "Practically Efficient Private Set Intersection from Trusted Hardware with Side-Channels". In: *ASIACRYPT 2023, Part IV*. Ed. by Jian Guo and Ron Steinfeld. Vol. 14441. LNCS. Springer, Singapore, Dec. 2023, pp. 268–301. DOI: 10.1007/978-981-99-8730-6_9.

[37] Andreas Erwig et al. "CommiTEE : An Efficient and Secure Commit-Chain Protocol using TEEs". In: *8th IEEE European Symposium on Security and Privacy, EuroS&P 2023, Delft, Netherlands, July 3-7, 2023*. IEEE, 2023, pp. 429–448. ISBN: 978-1-6654-6512-0. DOI: 10.1109/EuroSP57164.2023.00033. URL: https://doi.org/10.1109/EuroSP57164.2023.00033.

[38] Ben Fisch et al. "IRON: Functional Encryption using Intel SGX". In: *ACM CCS 2017*. Ed. by Bhavani M. Thuraisingham et al. ACM Press, Oct. 2017, pp. 765–782. DOI: 10.1145/3133956.3134106.

[39] Tommaso Frassetto et al. "POSE: Practical Off-chain Smart Contract Execution". In: *NDSS 2023*. The Internet Society, Feb. 2023.

[40] Sivanarayana Gaddam et al. "How to Design Fair Protocols in the Multi-Blockchain Setting". In: *IACR Cryptol. ePrint Arch.* (2023), p. 762. URL: https://eprint.iacr.org/2023/762.

[41] Sivanarayana Gaddam et al. "LucidiTEE: Scalable Policy-Based Multiparty Computation with Fairness". In: *CANS 23*. Ed. by Jing Deng, Vladimir Kolesnikov, and Alexander A. Schwarzmann. Vol. 14342. LNCS. Springer, Singapore, Oct. 2023, pp. 343–367. DOI: 10.1007/978-981-99-7563-1_16.

[42] Vipul Goyal et al. "Founding Cryptography on Tamper-Proof Hardware Tokens". In: *TCC 2010*. Ed. by Daniele Micciancio. Vol. 5978. LNCS. Springer, Berlin, Heidelberg, Feb. 2010, pp. 308–326. DOI: 10.1007/978-3-642-11799-2_19.

[43] Michele Grisafi et al. "PISTIS: Trusted Computing Architecture for Low-end Embedded Systems". In: *USENIX Security 2022*. Ed. by Kevin R. B. Butler and Kurt Thomas. USENIX Association, Aug. 2022, pp. 3843–3860.

[44] Dennis Hofheinz and Victor Shoup. *GNUC: A New Universal Composability Framework*. Cryptology ePrint Archive, Report 2011/303. 2011. URL: https://eprint.iacr.org/2011/303.

[45] Dennis Hofheinz and Victor Shoup. "GNUC: A New Universal Composability Framework". In: *Journal of Cryptology* 28.3 (July 2015), pp. 423–508. DOI: 10.1007/s00145-013-9160-y.

[46] Sashidhar Jakkamsetti, Zeyu Liu, and Varun Madathil. "Scalable Private Signaling". In: *IACR Cryptol. ePrint Arch.* (2023), p. 572. URL: https://eprint.iacr.org/2023/572.

[47] Gabriel Kaptchuk, Matthew Green, and Ian Miers. "Giving State to the Stateless: Augmenting Trustworthy Computation with Ledgers". In: *NDSS 2019*. The Internet Society, Feb. 2019. DOI: 10.14722/ndss.2019.23060.

[48] Jonathan Katz. "Universally Composable Multi-party Computation Using Tamper-Proof Hardware". In: *EUROCRYPT 2007*. Ed. by Moni Naor. Vol. 4515. LNCS. Springer, Berlin, Heidelberg, May 2007, pp. 115–128. DOI: 10.1007/978-3-540-72540-4_7.

[49] Mahimna Kelkar et al. *Complete Knowledge: Preventing Encumbrance of Cryptographic Secrets*. Cryptology ePrint Archive, Report 2023/044. 2023. URL: https://eprint.iacr.org/2023/044.

[50] Ralf Küsters and Daniel Rausch. "A Framework for Universally Composable Diffie-Hellman Key Exchange". In: *2017 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2017, pp. 881–900. DOI: 10.1109/SP.2017.63.

[51] Ralf Küsters and Max Tuengerthal. *The IITM Model: a Simple and Expressive Model for Universal Composability*. Cryptology ePrint Archive, Report 2013/025. 2013. URL: https://eprint.iacr.org/2013/025.

[52] Dayeol Lee et al. "Cerberus: A Formal Approach to Secure and Efficient Enclave Memory Sharing". In: *ACM CCS 2022*. Ed. by Heng Yin et al. ACM Press, Nov. 2022, pp. 1871–1885. DOI: 10.1145/3548606.3560595.

[53] Ming Li et al. *IvyCross: A Trustworthy and Privacy-preserving Framework for Blockchain Interoperability*. Cryptology ePrint Archive, Report 2021/1244. 2021. URL: https://eprint.iacr.org/2021/1244.

[54] Jinghui Liao et al. "Speedster: An Efficient Multi-party State Channel via Enclaves". In: *ASIACCS 22*. Ed. by Yuji Suga et al. ACM Press, May 2022, pp. 637–651. DOI: 10.1145/3488932.3523259.

[55] Joshua Lind et al. "Teechain: a secure payment network with asynchronous blockchain access". In: *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP 2019, Huntsville, ON, Canada, October 27-30, 2019*. Ed. by Tim Brecht and Carey Williamson. ACM, 2019, pp. 63–79. ISBN: 978-1-4503-6873-5. DOI: 10.1145/3341301.3359627. URL: https://doi.org/10.1145/3341301.3359627.

[56] Yibiao Lu et al. "Correlated Randomness Teleportation via Semi-trusted Hardware - Enabling Silent Multi-party Computation". In: *ESORICS 2021, Part II*. Ed. by Elisa Bertino, Haya Shulman, and Michael Waidner. Vol. 12973. LNCS. Springer, Cham, Oct. 2021, pp. 699–720. DOI: 10.1007/978-3-030-88428-4_34.

[57] Yao Ma et al. "QEnclave - A practical solution for secure quantum cloud computing". In: *CoRR* abs/2109.02952 (2021). arXiv: 2109.02952. URL: https://arxiv.org/abs/2109.02952.

[58] Varun Madathil et al. "Private Signaling". In: *USENIX Security 2022*. Ed. by Kevin R. B. Butler and Kurt Thomas. USENIX Association, Aug. 2022, pp. 3309–3326.

[59] Wenze Mao, Peng Jiang, and Liehuang Zhu. "BTAA: Blockchain and TEE-Assisted Authentication for IoT Systems". In: *IEEE Internet of Things Journal* 10.14 (July 2023), pp. 12603–12615. ISSN: 2372-2541. DOI: 10.1109/jiot.2023.3252565. URL: http://dx.doi.org/10.1109/jiot.2023.3252565.

[60] Sinisa Matetic et al. "ROTE: Rollback Protection for Trusted Execution". In: *USENIX Security 2017*. Ed. by Engin Kirda and Thomas Ristenpart. USENIX Association, Aug. 2017, pp. 1289–1306.

[61] Ueli Maurer. "Constructive Cryptography - A Primer (Invited Paper)". In: *FC 2010*. Ed. by Radu Sion. Vol. 6052. LNCS. Springer, Berlin, Heidelberg, Jan. 2010, p. 1. DOI: 10.1007/978-3-642-14577-3_1.

[62] Ueli Maurer and Renato Renner. "Abstract Cryptography". In: *ICS 2011*. Ed. by Bernard Chazelle. Tsinghua University Press, Jan. 2011, pp. 1–21.

[63] Koichi Moriyama and Akira Otsuka. "Permissionless Blockchain-Based Sybil-Resistant Self-Sovereign Identity Utilizing Attested Execution Secure Processors". In: *IEEE International Conference on Blockchain, Blockchain 2022, Espoo, Finland, August 22-25, 2022*. IEEE, 2022, pp. 1–10. ISBN: 978-1-6654-6104-7. DOI: 10.1109/Blockchain55522.2022.00012. URL: https://doi.org/10.1109/Blockchain55522.2022.00012.

[64] Jianyu Niu et al. "NARRATOR: Secure and Practical State Continuity for Trusted Execution in the Cloud". In: *ACM CCS*

2022. Ed. by Heng Yin et al. ACM Press, Nov. 2022, pp. 2385–2399. DOI: 10.1145/3548606.3560620.

[65] Chris Orsini, Alessandra Scafuro, and Tanner Verber. *How to Recover a Cryptographic Secret From the Cloud*. Cryptology ePrint Archive, Paper 2023/1308. https://eprint.iacr.org/2023/1308. 2023. URL: https://eprint.iacr.org/2023/1308.

[66] Arttu Paju et al. "SoK: A Systematic Review of TEE Usage for Developing Trusted Applications". In: *Proceedings of the 18th International Conference on Availability, Reliability and Security*. ARES 2023. ACM, Aug. 2023. DOI: 10.1145/3600160.3600169. URL: http://dx.doi.org/10.1145/3600160.3600169.

[67] Bryan Parno et al. "Memoir: Practical State Continuity for Protected Modules". In: *2011 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2011, pp. 379–394. DOI: 10.1109/SP.2011.38.

[68] Rafael Pass, Elaine Shi, and Florian Tramer. *Formal Abstractions for Attested Execution Secure Processors*. Cryptology ePrint Archive, Report 2016/1027. 2016. URL: https://eprint.iacr.org/2016/1027.

[69] Rafael Pass, Elaine Shi, and Florian Tramèr. "Formal Abstractions for Attested Execution Secure Processors". In: *EUROCRYPT 2017, Part I*. Ed. by Jean-Sébastien Coron and Jesper Buus Nielsen. Vol. 10210. LNCS. Springer, Cham, Apr. 2017, pp. 260–289. DOI: 10.1007/978-3-319-56620-7_10.

[70] Muhammad Usama Sardar, Do Le Quoc, and Christof Fetzer. "Towards Formalization of Enhanced Privacy ID (EPID)-based Remote Attestation in Intel SGX". In: *23rd Euromicro Conference on Digital System Design, DSD 2020, Kranj, Slovenia, August 26-28, 2020*. IEEE, 2020, pp. 604–607. ISBN: 978-1-7281-9535-3. DOI: 10.1109/DSD51259.2020.00099. URL: https://doi.org/10.1109/DSD51259.2020.00099.

[71] Stephan van Schaik et al. *SoK: SGX.Fail: How Stuff Get eXposed*. https://sgx.fail. 2022.

[72] Rohit Sinha et al. "A design and verification methodology for secure isolated regions". In: *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*. Ed. by Chandra Krintz and Emery D. Berger. ACM, 2016, pp. 665–681. ISBN: 978-1-4503-4261-2. DOI: 10.1145/2908080.2908113. URL: https://doi.org/10.1145/2908080.2908113.

[73] Rohit Sinha et al. "Moat: Verifying Confidentiality of Enclave Programs". In: *ACM CCS 2015*. Ed. by Indrajit Ray, Ninghui Li, and Christopher Kruegel. ACM Press, Oct. 2015, pp. 1169–1184. DOI: 10.1145/2810103.2813608.

[74] Pramod Subramanyan et al. "A Formal Foundation for Secure Remote Execution of Enclaves". In: *ACM CCS 2017*. Ed. by Bhavani M. Thuraisingham et al. ACM Press, Oct. 2017, pp. 2435–2450. DOI: 10.1145/3133956.3134098.

[75] Haiyong Sun and Hang Lei. "A Design and Verification Methodology for a TrustZone Trusted Execution Environment". In: *IEEE Access* 8 (2020), pp. 33870–33883. ISSN: 2169-3536. DOI: 10.1109/access.2020.2974487. URL: http://dx.doi.org/10.1109/ACCESS.2020.2974487.

[76] Taisei Takahashi, Taishi Higuchi, and Akira Otsuka. "VeloCash: Anonymous Decentralized Probabilistic Micropayments With Transferability". In: *IEEE Access* 10 (2022), pp. 93701–93730. DOI: 10.1109/ACCESS.2022.3201071. URL: https://doi.org/10.1109/ACCESS.2022.3201071.

[77] Florian Tramèr and Dan Boneh. "Slalom: Fast, Verifiable and Private Execution of Neural Networks in Trusted Hardware". In: *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL: https://openreview.net/forum?id=rJVorjCcKQ.

[78] Florian Tramèr et al. "Sealed-Glass Proofs: Using Transparent Enclaves to Prove and Sell Knowledge". In: *2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017, Paris, France, April 26-28, 2017*. IEEE. IEEE, 2017, pp. 19–34. ISBN: 978-1-5090-5762-7. DOI: 10.1109/EuroSP.2017.28. URL: https://doi.org/10.1109/EuroSP.2017.28.

[79] Stavros Volos, Kapil Vaswani, and Rodrigo Bruno. "Graviton: Trusted Execution Environments on GPUs". In: *13th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2018, Carlsbad, CA, USA, October 8-10, 2018*. Ed. by Andrea C. Arpaci-Dusseau and Geoff Voelker. USENIX Association, 2018, pp. 681–696. URL: https://www.usenix.org/conference/osdi18/presentation/volos.

[80] Ivana Vukotic, Vincent Rahli, and Paulo Esteves-Veríssimo. "Asphalion: trustworthy shielding against Byzantine faults". In: *Proceedings of the ACM on Programming Languages* 3.OOPSLA (Oct. 2019), pp. 1–32. ISSN: 2475-1421. DOI: 10.1145/3360564. URL: http://dx.doi.org/10.1145/3360564.

[81] Weili Wang et al. "ENGRAFT: Enclave-guarded Raft on Byzantine Faulty Nodes". In: *ACM CCS 2022*. Ed. by Heng Yin et al. ACM Press, Nov. 2022, pp. 2841–2855. DOI: 10.1145/3548606.3560639.

[82] Newton C. Will and Carlos A. Maziero. "Intel Software Guard Extensions Applications: A Survey". In: *ACM Computing Surveys* 55.14s (July 2023), pp. 1–38. ISSN: 1557-7341. DOI: 10.1145/3593021. URL: http://dx.doi.org/10.1145/3593021.

[83] Pengfei Wu et al. "Exploring Dynamic Task Loading in SGX-Based Distributed Computing". In: *IEEE Trans. Serv. Comput.* 16.1 (2023), pp. 288–301. DOI: 10.1109/TSC.2021.3123511. URL: https://doi.org/10.1109/TSC.2021.3123511.

[84] Pengfei Wu et al. "ObliDC: An SGX-based Oblivious Distributed Computing Framework with Formal Proof". In: *ASIACCS 19*. Ed. by Steven D. Galbraith et al. ACM Press, July 2019, pp. 86–99. DOI: 10.1145/3321705.3329822.

[85] Rongwu Xu et al. "Miso: Legacy-Compatible Privacy-Preserving Single Sign-On Using Trusted Execution Environments". In: *CoRR* (2023). arXiv: 2305.06833 [cs.CR]. URL: http://arxiv.org/abs/2305.06833v2.

[86] Shiwei Xu et al. "A Symbolic Model for Systematically Analyzing TEE-Based Protocols". In: *ICICS 20*. Ed. by Weizhi Meng et al. Vol. 11999. LNCS. Springer, Cham, Aug. 2020, pp. 126–144. DOI: 10.1007/978-3-030-61078-4_8.

[87] Fan Zhang et al. *Paralysis Proofs: Secure Access-Structure Updates for Cryptocurrencies and More*. Cryptology ePrint Archive, Report 2018/096. 2018. URL: https://eprint.iacr.org/2018/096.

[88] Fan Zhang et al. "Town Crier: An Authenticated Data Feed for Smart Contracts". In: *ACM CCS 2016*. Ed. by Edgar R. Weippl et al. ACM Press, Oct. 2016, pp. 270–282. DOI: 10.1145/2976749.2978326.

[89] Xuyang Zhao et al. "Towards A Secure Joint Cloud With Confidential Computing". In: *2022 IEEE International Conference on Joint Cloud Computing (JCC)* (Aug. 2022). DOI: 10.1109/jcc56315.2022.00019. URL: http://dx.doi.org/10.1109/jcc56315.2022.00019.

APPENDIX

A. Universal Composability notions

Universal Composability, introduced by Canetti [16], is a computational proof model that allows showing the security of protocols

under concurrent composition.

UC is based on the computation model of Interactive Turing Machines (ITM) A protocol is defined as a set of ITM instances (ITIs) whose unique identities are composed of a party identifier (PID) and a shared session identifier (SID). We generally refer to the ITIs that represent the protocol principals as main parties, which can spawn subroutine that represents portion of code executed by the principal. To allow separating modelling artefacts from the code of the analysed protocol, a "structured protocol" divides ITIs into a shell and body component (introduced in [15, Version of 2018]). The body of the protocol handles the cryptographic operations, and is not aware of the shell, which is limited to handling modelling related instructions and can read and modify the contents of the body appropriately. A protocol is executed in the presence of a probabilistic polynomial time (PPT) bound machine, the environment, that captures the influence of any computation that might be taking place outwith the current instance of the analysed protocol. The environment can be seen as initialising the computation of the protocol, and providing input to each of the protocol principals and the adversary. The adversary is another PPT-bound machine that is able to instruct ITIs with special corruption messages to modify their behaviour, through a dedicated *backdoor tape*. For the rest of the paper, we assume the convention that any adversary is a *dummy adversary*, where its behaviour is to simply forward corruption messages originated by the environment to protocol parties. Besides the adversarial backdoor tape, ITIs are able to communicate with each other by writing messages on some dedicated tapes. These mechanisms should not be seen as equivalent to network communication but rather as a modelling artefact, while the network model can be implemented as an ideal functionality (allowing flexibility to model networks with different properties). While the framework does not impose general restrictions on which ITIs can communicate with each other, there are certain communication topologies that can be considered "better-formed", and necessary for certain composition results (such as subroutine respecting protocols, where all communications to protocol subroutines have to originate from the protocol main parties or one of their subroutines - the protocol's *extended session*). To allow composing our examined protocol, the environment represents external communication by claiming an external machine's identity when sending an input to the protocol parties. An environment is said to be ξ-identity-bounded if the set of identities it can claim is restricted by ξ (expressed as a predicate over the system's state at the time the environment sends a message claiming an external machine's identity).

The model of execution of ITIs is inherently single-threaded, but allows flexibility in describing the granularity of operations and how they interleave. Runtime constraints are satisfied by maintaining a runtime budget for each machine (known as *import*). Import can be shared with a machine's subroutine, allowing arbitrary dynamical subroutine nesting without running the risk of exceeding the remaining runtime. The minimum import considered by UC protocols is the length of the security parameter. A *balanced* environment ensures that at any point during the execution of a protocol, the adversarial import is at least as large as the sum of imports for all other ITIs in the protocol.

Like other simulation proofs, the basic mechanism for showing UC-security is to define an ideal functionality, which captures the essential properties of the desired protocol as being run by a trusted party, and show it to be computationally indistinguishable from an execution of the real protocol (*UC-emulation*). $\mathsf{EXEC}_{\pi,\mathcal{A},\mathcal{Z}}$ is the random variable representing the output of environment $\mathcal{Z}$ for an execution of π in the presence of adversary $\mathcal{A}$ (conversely $\mathsf{EXEC}_{\phi,\mathcal{S},\mathcal{Z}}$ is for the execution of the ideal functionality ϕ in the presence of simulator $\mathcal{S}$).

Theorem 3 (*UC emulation*)

*For any PPT protocols π, ϕ and identity predicate ξ, we say that π ξ-*UC-emulates *ϕ (or simply π UC-emulates ϕ if the identity bound allows any identity) if for any PPT adversary $\mathcal{A}$ there exists a corresponding PPT adversary $\mathcal{S}$ (the simulator), such that for any balanced PPT ξ-identity-bounded environment $\mathcal{Z}$, it holds that $\mathsf{EXEC}_{\pi,\mathcal{A},\mathcal{Z}} \approx \mathsf{EXEC}_{\phi,\mathcal{S},\mathcal{Z}}$*

UC-emulation can be used to show that, if we have a protocol π that *realises* an ideal functionality $\mathcal{F}$, the security analysis of a new protocol ρ that has π as a subroutine can be carried out by replacing all of ρ's call to subroutines running π with calls to ideal functionality $\mathcal{F}$, which we denote as $\rho^{\pi \to \mathcal{F}}$. This new version of ρ is said to be in the *hybrid model*, since its ITIs interact with both other real ITIs and ideal functionalities. For the replacement to be successful, we require that any party in ρ that calls to a subroutine in π or $\mathcal{F}$ satisfies ξ and does not call instances of π and $\mathcal{F}$ in the same session (we say that the protocol ρ is (π, ϕ, ξ)-compliant). Additionally, the adversary should be able to determine whether an ITI in a certain session is part of the protocol (the protocol is subroutine exposing).

Theorem 4 (*UC Composition Theorem*)

For any PPT protocols ρ, π, ϕ and predicate ξ, if ρ is (π, ϕ, ξ)-compliant, ϕ, π are both subroutine respecting and subroutine exposing, and π ξ-UC-emulates ϕ, then $\rho^{\pi \to \phi}$ UC-emulates ρ.

Unfortunately, many interesting protocols, such as commitment schemes [19], secure two-party computation [20] or even authenticated channels [17], are not easily provable in UC in the plain model. We therefore need to add some ideal subroutine that can represent the cryptographic assumptions required as a block box ideal subroutine. The next section will discuss how hybrid functionalities that share state among sessions can also be used composably through some tweaks to the UC framework.

1) Globality: While UC provides a powerful paradigm for reusable cryptographic proofs, composition imposes many restrictions over the base model as outlined in Theorem 4. To address the limitation of the UC theorem of subroutine-respecting interactions, Canetti and Rabin [21] introduce Universal Composition with Joint-State, a new composition theorem that allows a single protocol session to be a subroutine of different protocols. This can be used, for example, to prove the security of different protocols that use an authenticated channel, where all sessions interacting with the same party share the signing key. This composition theorem is, however, only valid for static protocols (where the number of shared sessions is already well defined). Canetti et al. [25] formulate two new variants of Universal Composition, Extended UC and Generalised [2] UC, that allow composition when arbitrary protocol interact with the shared subroutine. The formulation of GUC has been widely used in the literature, allowing modelling of protocols that were previously impossible to prove in plain UC, such as those that provide deniability. Canetti, Shahaf, and Vald [22] later extended the GUC composition theorem to allow the replacement of global functionalities with protocols. Despite its popularity, proving security in GUC is more difficult than in the incompatible plain UC setting, as it requires arguing about all possible protocols rather than just the one being analysed. Moreover, as basic UC has received multiple updates and fixes over time, those have not percolated to the GUC formalisation, and the equivalence between GUC and the simpler EUC theorem (which most security proofs in the global setting are

[2] commonly misattributed as Global UC

actually using) has been called into question due to some components of the framework being underspecified [7].

Universal Composability with Global Subroutines [6] aims to rectify some of 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 the *management* protocol. The management protocol is designed to be subroutine-respecting to preserve composition, while allowing the external protocol ρ to access a single instance of π and multiple of γ. μ is a shell only protocol that uses a directory ITI to redirect external communication from ρ to the appropriate machines in π or γ (and conversely to the external machine that should receive a response). The following definition roughly corresponds to the EUC formulation of global functionalities:

Definition 1

*For protocols π, ϕ, γ, we say that π ξ-*UC-emulates ϕ in the presence of (global subroutine) γ *if $M[\pi, \gamma]$ ξ-UC-emulates $M[\phi, \gamma]$*

As in the basic UC framework, the composition theorem follows, with some additional restrictions: π and ϕ are allowed to break their subroutine-respecting behaviour to use the global subroutine γ (we say they are γ-subroutine respecting), and γ itself does not depend on ϕ as one of its subroutines (we say that γ is ϕ-regular). These requirements allow the use of the shared state subroutine without provoking circular dependencies that would prevent a clean cut replacement [3].

Theorem 5 (*Universal Composition with Global Subroutines*)

For any subroutine-exposing protocols ρ, ϕ, π, γ where

- *γ is subroutine respecting and ϕ-regular,*
- *π, ϕ are γ-subroutine respecting,*
- *ρ is (π, ϕ, ξ)-compliant, $(\pi, M[\phi, \gamma], \xi)$-compliant and $(\pi, M[\pi, \gamma], \xi)$-compliant;*

if π ξ-UC-emulates ϕ in the presence of γ (per Definition 1), then $\rho^{\phi \to \pi}$ UC-emulates ρ.

The above theorems can be used to recover EUC statements in the literature by formulating an appropriate identity bound. While most of the existing work focus on ideal functionalities as global subroutine, Badertscher, Hesse, and Zikas [5] show that UCGS does not universally preserve the composition theorem from [22] to replace the setup with a potentially interactive protocol using a different setup. In particular, when replacing a particularly weak global setup G (where adversarial capabilities are more extensive than the proposed protocol γ that realises it), the simulator S in the emulation of a G-hybrid functionality F by some protocol π might no longer be possible in the γ-hybrid world, as it can no longer use the attacks allowed by G. Their work then provides some guidelines on which global setups can be successfully replaced by a protocol. Namely, an equivalent setup (where protocol γ UC-emulates ideal functionality G, and G UC-emulates γ) can always be replaced, regardless of the context protocols which use it as a global subroutine. Additionally, replacement is possible if the simulation strategy of S either avoids using any of the adversarial capabilities of G (S is an *agnostic simulator*), or that the adversarial capabilities it does interact with will be preserved by γ (S is an *admissible* simulator).

Canetti et al. [24] later observe that the replacement statement also holds if protocol γ replaces the protocol that combines G with the simulator from the γ UC-emulates G experiment, and thus any F using that combined protocol as a global subroutine can be replaced with γ.

Camenisch, Drijvers, and Tackmann [13] also show that neither the UC or GUC composition theorems allow proving that a protocol $\rho^{\pi \to \mathcal{F}}$ can UC-emulate ρ if π is a subroutine of both ρ and of another distinct ideal subroutine of ρ. They therefore propose a new recursive composition theorem for jointly subroutine respecting functionalities, multi-protocol UC. The model of Hofheinz and Shoup [45], which is partially compatible with UC, also includes a more restricted model of composition with shared subroutines.

To conclude this section, we note that in the rest of this work, whenever an ideal functionality calls another (global) ideal subroutine (e.g. provides some input to the global subroutine on behalf of a specific party), the underlying operation relies on the intermediary dummy party convention of [6, Definition 4].

2) Notation: We now list additional convention taken by our pseudocode for the remainder of this work. We hope our notation is generally self-explanatory, but in case of ambiguity we refer the reader to the following explanation. We might refer to UC terminology beyond what was described above; any such usage is self-contained to this section, but we refer the reader to [15, Section 3.1] for additional context.

Our notation defines ITIs in terms of their behaviour when they are activated and find a new message on their input tape. We define the code executed when such a message is received as a procedure. Some procedures definitions are not meant to be triggered by external parties writing on the input tape, but are simply used to extract some shared code that the ITI might need to execute multiple times. In that case, we use the keyword "run" followed by the procedure name to denote that the same ITI is executing it. The ITI is understood to choose which procedure to execute by pattern matching on the program definition as specified in the pseudocode, starting from the earliest procedure definition i.e. if there are multiple commands that start with the same keyword, it will try to find the one with the correct arguments starting from the earlier definition. When font cmd is used in this context, it is taken to be a variable, such that the procedure executed is not literally the one named CMD but rather the value held in that named variable.

Our message-passing treatment tends to stay at a higher level than the underlying UC execution. As such, we omit many details of the ITI behaviour in our protocol descriptions. We generally describe a procedure by using the notation "*On message (*PROCEDURENAME*,list of procedure arguments) from party P:*" followed by high-level pseudo code for the ITI execution, in the style of an imperative programming language. This notation is short for indicating that the machine we are describing on activation reads from its input type a message of type $(P, (\text{PROCEDURENAME}, \textit{list of procedure arguments}))$, where P is an object that contains fields $\mathsf{pid}, \mathsf{sid}$ for the party and session identifiers (respectively) of the sender ITI; and PROCEDURENAME corresponds to some code in its program it can execute with the inputs from the argument list. Conversely, the notation **return** $(\text{MSG}, args)$, as part of the description of procedure pseudocode for an ITI M, denotes the end of the execution of the current procedure with the issuing an external write request (f, M', t, r, M, m), where destination ITI M' is the same machine from which it received input, and $m = (\text{MSG}, args)$. In this case, we always set f, the `forced-write` flag, to 1; t, the destination tape, to **subroutine-output** (unless the pseudocode describes an adversarial machine, in which case t=**backdoor**); and r, the `reveal-sender-id` flag, also set to 1. Keyword **abort**, or **return** with no arguments indicate the end of execution for the current procedure without issuing a corresponding subroutine output

[3]This type of recursive composition is implemented in multi-protocol UC [13]; however the composition theorem of that work is not compatible with Theorem 4, and therefore Theorem 1 does not apply either

message.

If M wants to issue an external write request for a destination ITI that is not the same that initiated the current procedure execution by passing input to M, we use "**Send** (MSG,args) **to** M'" to issue a the same message as described above, except for setting t to **input**. If the **Send** instructions is not the last one in the current procedure description, the external write request is not issued immediately, but rather queued in the outgoing message tape for M until the end of the procedure, or when M next relinquishes the activation token. On the other hand, when we use "**Send** (MSG,args) **to** M' and **receive** (MSG′, args′)", M yields activation immediately, and resumes execution the pseudocode from the same instruction when it next receives message $(\text{MSG'}, args')$ on its subroutine output tape from the sender. When this happens, the ITI stores its current execution context (i.e. any intermediate computation on the work tape) somewhere in memory in a way that it can be restored when re-activated by the response message. Between sending and receiving the response, the ITI can be activated with any other message on any tape, although if our current program can not tolerate such concurrency, the ITI might abort by checking some internal flag. If multiple outgoing messages were sent to the same M', we assume that the response includes some unique identifier to allow M to restore the correct context for which it is responding to[4]. When M issues an outgoing message, and expects the corresponding response to come from a different party, we use the keyword **await** followed by a full description of the behaviour on next activation.

A variable assigned as part of a procedure does not guarantee that it will be available to other procedures, unless it is defined in the State variables table at the start of the definition. When the same program uses the same identifier across different procedures, they are generally taken to be distinct values, especially if received as part of a message. Variables first defined within a loop or **if** branch have their scope local to that block. Protocol parameters are generally taken to be globally readable to all protocol parties and their procedures.

Our formulations in this work rely on structured protocols, as defined in [15, Section 5.1]. A structured protocol is a series of nested ITIs, on which a higher level ITI (generally referred to as shell) has full access to read or overwrite the tapes of any lower level subroutine ITI (which we refer to interchangeably as the virtual ITI, or by their extended identities). ITIs have access to a number of tapes to store their identity, code, running memory, and communicate with other machines. Although the description of an ITI is not precise or prescriptive in terms of how it implements the computation, we assume that the program description uses some well-defined language, perhaps similar to a low level programming language or assembly. We represent each individual instruction as a command with optional arguments, which we represent using function call notation $command(argument)$ sometimes with optional parenthesis (e.g. for the case where $command = \mathbf{return}$). We overload the set membership operator $\in$ to verify that the command component of the instruction belongs to the set. The code of an enclave can be seen as a list of ITI instructions of this type, and the notation "**for** instruction $\mathsf{i} \in \mathsf{prog}$ **do**" can be interpreted as iterating over the list of instructions for program prog (including command and arguments) without executing them (i.e. by advancing only the head of the shell over the tape). Conversely, when an ITI ρ in a structured protocol contains pseudocode

begin executing inp on π
for next instruction i on π **do** $f(\mathsf{i})$

[4]This is not a universally safe assumption to make for any UC protocol, but it is sufficently safe for the ones analysed in this work

it should be read as ρ iterating through the code of a subroutine with extended identity π, and for each instruction i, ρ executes subroutine $f(\mathsf{i})$ to advance the state of π (updating its tapes and advancing π's head), while performing any additional operations in ρ's code.

B. Defining Generic Transformation Protocols

We now provide an equivalent protocol to the one outlined in Section V to address feature addition to a trusted hardware setup. Unlike the attack removal protocol in that section, we consider a more generalised class of protocol topologies that relies on additional enclaves. The technique introduced in this section can also be used to extend the attack removal protocol and corresponding simulation template.

1) Adding a Feature oracle: To fully capture the modular power of our new formalisation, we show how to add a new feature to a TEE instance, increasing the size of its feature oracle set. We want to show that a TEE that has native access to that feature (through an oracle) is indistinguishable from one that does not and has to implement it through runtime code. An oracle can be implemented through a UC protocol that realises the ideal oracle interface.

More formally, we consider two instantiations of attested execution G_{att} and G'_{att} (both modular), with feature oracles $\mathbb{O}, \mathbb{O}'$, respectively, where $\mathbb{O} \subset \mathbb{O}'$. Let $I = \mathbb{O}' \setminus \mathbb{O}$. The adversarial oracles $\mathbb{A}$ and attestation signature function $\mathbb{S}$ are shared between G_{att} and G'_{att}. We now define a new "wrapper" protocol $\mathcal{W}$ which uses G_{att} as a subroutine and UC-realises G'_{att} by implementing the interface for I in the G_{att}-hybrid real world.

$\mathcal{W}$ takes the same parameters as G^{mod}_{att}, and in addition the two functions $\mathsf{map}^{\mathsf{L}}, \mathsf{map}^{\mathsf{R}}$, and the code of enclave program $\mathrm{W}^I[\cdot]$. Function $\mathsf{map}^{\mathsf{R}}$ takes the set of G^{mod}_{att}-enabled parties, and chooses a subset to run remote assisting enclaves that any party can rely on (the parties chosen by $\mathsf{map}^{\mathsf{R}}$ do not have to be honest). $\mathsf{map}^{\mathsf{L}}$ returns a set of local assisting enclave programs the party should install locally (on the same machine), and a next message function for $\mathrm{W}^I[\cdot]$.

$\mathrm{W}^I[\mathsf{prog}, \mathrm{nextmsg}]$ is a "wrapper" enclave that instruments prog with additional code such that, when prog attempts to use interface I, the next message function begins executing I as a protocol. Function nextmsg observes the current state of the enclave, and chooses the command required to start the I-protocol execution. The command issued by nextmsg will either be run as a local subroutine in the enclave wrapper code itself, by another enclave installed locally (as instructed by $\mathsf{map}^{\mathsf{L}}$), or by a remote party (in an enclave created through $\mathsf{map}^{\mathsf{R}}$). nextmsg is aware of the details of each assisting enclave, such as their enclave ID or what party they are installed on, through these functions.

If the next command issued by nextmsg is received by the assisting enclave it is destined for (a corrupted party could diverge from the protocol and choose not to deliver the message), it executes the requested subroutine, produces its own next command, and forwards it to the party that should execute it. Eventually, the original $\mathrm{W}^I[\mathsf{prog}, \mathrm{nextmsg}]$ will receive a final message, and return the result value of I to the prog oracle call. Essentially, the program that implements I is compiled into a multi-party computation between the enclaves. We do not require a full-fledged secure MPC protocol to execute I, however, due to the integrity guarantees of attestation, as the only possible malicious behaviour of a participant is dropping messages (known as the omission corruption adversarial model in MPC [12]). Within the execution of the next message functions, enclaves are able to construct an authenticated or secure channel through attestation. We do not give a description of how this is done, and refer the reader back to the construction of the secure channel from [68]

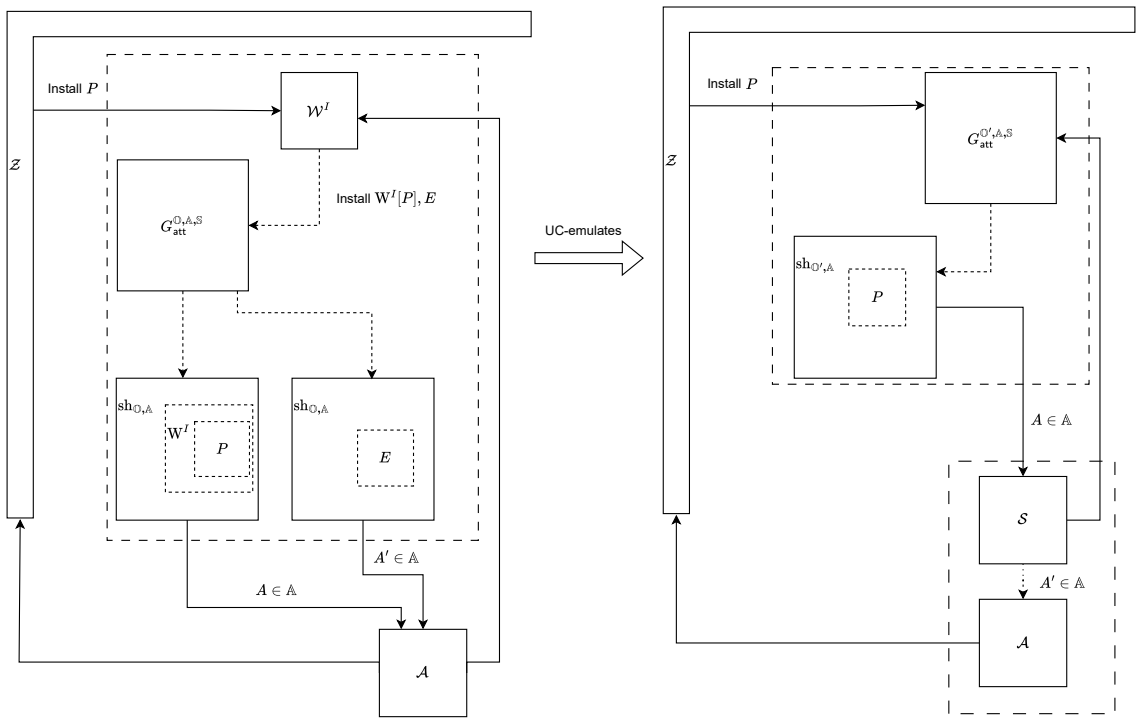

Fig. 7. Protocol $\mathcal{W}$ can add a shell to $G_{\mathsf{att}}^{\mathbb{O},\mathbb{A},\mathbb{S}}$ enclaves to UC-emulate the missing feature oracles I from $G_{\mathsf{att}}^{\mathbb{O}',\mathbb{A},\mathbb{S}}$

The $\mathcal{W}$ protocol (Figure 8) proceeds as follows. During initialisation, it calls the $\mathsf{map}^{\mathsf{R}}$ function to produce a list of supporting enclaves E_i^R run by a subset of reg parties, initialises G_{att}^{mod} with the appropriate parameters, and requests each selected party to install the wrapped E_i^R. It then returns a public list of all assisting parties and the associated enclave IDs.

On a call from pid_i to install some program prog_i, if prog_i does not include any call to I, it installs a wrapped version with dummy next message function ϵ. Otherwise, it runs $\mathsf{map}^{\mathsf{L}}()$ to produce a list of local assisting enclave programs to be installed by the same party, and a next message function $\mathrm{nextmsg}_\emptyset$. The party installs all such enclaves, runs their initialisation subroutine, and creates a new message function $\mathrm{nextmsg}$ that is a wrapper around $\mathrm{nextmsg}_\emptyset$ aware of the assisting parties enclave IDs. $\mathsf{map}^{\mathsf{L}}$ makes $\mathrm{nextmsg}$ and all E_i^L programs aware of the enclave IDs for any E^R parties, and assists $\mathrm{W}^I[\cdot]$ in generating the appropriate next commands to implement I along with the assisting protocols.

On a resume call from its local party pid_i to execute command cmd on arguments $args$ for enclave $\mathrm{W}^I[\mathsf{prog}_\mathsf{i}, \mathrm{nextmsg}]$, the enclave wrapper (described in Figure 9) begins executing the code of prog_i with those inputs. Once the program makes a subroutine call to I, the wrapper stops the internal program execution and calls the $\mathrm{nextmsg}$ function, which returns the PID, Enclave ID and some command that needs to be executed to begin computing the value for the subroutine call. The enclave returns these to its local party with the special keyword RESUMEREQUEST, and waits for a next activation. When the party receives this return value, it knows that cmd($args$) did not terminate. Instead, it passes the RESUMEREQUEST and associated command on to the appropriate party, or, if the destination PID is pid_i, activates one of its local enclaves, including $\mathrm{W}^I[\mathsf{prog}_\mathsf{i}, \mathrm{nextmsg}]$ itself. When resuming an enclave as part of the I computation, the local party can set input flag $\top$ as part of the RESUME arguments to indicate that the command being executed is not part of the normal prog_i code. pid_i waits to receive the next message, and once again passes it on to one of its local enclaves, and forwards the resulting RESUMEREQUEST. Eventually, when the PID and EID returned by $\mathrm{nextmsg}$ for the enclave that initiated the call to I are both $\perp$, the computation of I has terminated, and the wrapper can pass back v as its return value to the internal execution of prog_i. Whenever the enclave returns with an intermediate message, the latest attestation signature should always be bundled with the next message input for the receiver party. Attestation validation logic is defined in the code of $\mathrm{W}[\cdot]$ for all appropriate messages, and interacts directly with the G_{att}^{mod} verification request through a call to the $\mathrm{AttestVerif}$ interface. When a user requests verification of one of these intermediate attestation signatures from outside one of the participating enclaves, the protocol always returns $\perp$.

We now provide the following conjecture:

Conjecture 2

Let $G_{\mathsf{att}} = G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}]$ *and* $G'_{\mathsf{att}} = G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}', \mathbb{A}, \mathbb{S}]$, *such that* $\mathbb{O}' \setminus \mathbb{O} = I$. *For any enclave wrapper* $\mathrm{W}^I[\cdot]$ *which, combined with functions* $\mathsf{map}^{\mathsf{L}}, \mathsf{map}^{\mathsf{R}}$ *implements the difference in behaviour between the shells* $\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\cdot]$ *and* $\mathrm{sh}_{\mathbb{O}',\mathbb{A}}[\cdot]$, *it is possible to show that protocol* $\mathcal{W}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}, \mathsf{map}^{\mathsf{R}}, \mathsf{map}^{\mathsf{L}}, \mathrm{W}^I[\cdot]]$ *in the presence of* G_{att} *UC-emulates* G'_{att}

Like in Section V, we provide a generic template for simulation without being able to make more general statements that are valid for all combinations of feature and adversarial oracles.

We describe simulation for the three possible protocol topologies implementing I:

1) We begin our simulation sketch for the case of a wrapper protocol $\mathcal{W}$ where neither $\mathsf{map}^{\mathsf{R}}$ or $\mathsf{map}^{\mathsf{L}}$ functions returns any additional enclave i.e. the wrapper $\mathrm{W}^I[\cdot]$ can implement I without relying on any external assistance i.e. using a combina-

Protocol $\mathcal{W}[\lambda, G_{\mathsf{att}}^{mod}, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}, \mathsf{map}^{\mathsf{R}}, \mathsf{map}^{\mathsf{L}}, \mathrm{W}^I[\cdot]]$

On message INITIALISE *from a party* P*:*
- $[(\mathrm{E}_1^R, \mathsf{pid}_1), \ldots, (\mathrm{E}_n^R, \mathsf{pid}_m)] \leftarrow \mathsf{map}^{\mathsf{R}}(\mathsf{reg})$
- **send** (INITIALISE) **to** $G_{\mathsf{att}}^{mod}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}]$
- **let** $\bar{\mathrm{E}}^R \leftarrow []$
- **for** $i \in \{0, \ldots, n\}$ **do**
  - **send** $(\text{INSTALLREQUEST}, \mathrm{E}_i^R)$ **to** pid_i and **receive** $\mathsf{eid}_\mathsf{i}, \mathsf{output}, \sigma$
  - **send** $(\text{VERIFY}, \sigma, \mathsf{output})$ **to** G_{att}^{mod} and **receive** v
  - **assert** that output is a valid return value for $\mathrm{E}_i^R(init)$ and $v = \top$
  - $\bar{\mathrm{E}}^R \leftarrow \bar{\mathrm{E}}^R \parallel \mathsf{pid}_i, \mathsf{eid}_\mathsf{i}$
- **return** $\bar{\mathrm{E}}^R$

On message GETPK *from a party* P*:*
- **send** GETPK **to** G_{att}^{mod} and **receive** vk
- **return** vk

On message $(\text{VERIFY}, \sigma, m)$ *from a party* P*:*
- **if** m is an attestation measurement that contains a commitment to some program with code E_i^R or E_i^L **then**
  - **return** $\perp$
- **else**
  - **send** $(\text{VERIFY}, \sigma, m)$ **to** G_{att}^{mod} and **receive** v and **return** v

On message $(\text{INSTALL}, \mathsf{prog})$ *from a party* P *where* $P.\mathsf{pid} \in \mathsf{reg}$*:*
- **if** $I \in \mathsf{prog}$ **then**
  - $(\mathrm{nextmsg}_\emptyset, (\mathrm{E}_1^L, \ldots, \mathrm{E}_n^L)) \leftarrow \mathsf{map}^{\mathsf{L}}(\bar{\mathrm{E}}^R)$
  - **let** $\bar{\mathrm{E}}^L \leftarrow []$
  - **for** $i \in \{1, \ldots, n\}$ **do**
    - **send** $(\text{INSTALL}, \mathrm{E}_i^L)$ **to** G_{att}^{mod} and **receive** eid_i
    - **send** $(\text{RESUME}, \mathsf{eid}_\mathsf{i}, \text{INIT})$ **to** G_{att}
    - $\bar{\mathrm{E}}^L \leftarrow \bar{\mathrm{E}}^L \parallel \mathsf{eid}_\mathsf{i}$
  - **let** $\mathrm{nextmsg}(x) \leftarrow \mathrm{nextmsg}_\emptyset(x, \bar{\mathrm{E}}^L)$
  - **send** $(\text{INSTALL}, \mathrm{W}^I[\mathsf{prog}, \mathrm{nextmsg}])$ **to** G_{att} and **receive** $\mathsf{eid}_{\mathsf{prog}}$
- **else**
  - **send** $(\text{INSTALL}, \mathsf{sid}, \mathrm{W}^I[\mathsf{prog}, \epsilon])$ **to** G_{att} and **receive** $\mathsf{eid}_{\mathsf{prog}}$
- **return** $\mathsf{eid}_{\mathsf{prog}}$

On message $(\text{RESUME}, \mathsf{eid}, \mathsf{inp})$ *from a party* P *with* pid_i*:*
- **send** $(\text{RESUME}, \mathsf{eid}, \mathsf{inp})$ **to** G_{att}^{mod} and **receive** out, σ
- $\sigma_{\mathsf{prev}} \leftarrow \sigma$
- **while** $\mathsf{out} = (\text{RESUMEREQUEST}, \mathsf{pid}, \mathsf{eid}', v)$ **do**
  - **if** $\mathsf{pid} = \mathsf{pid}_i$ **then**
    - $(\mathsf{out}, \sigma') \leftarrow \text{RESUME}(\mathsf{eid}', (v, \sigma_{\mathsf{prev}}, \top)))$
  - **else**
    - **send** $(\text{RESUMEREQUEST}, \mathsf{eid}', (v, \sigma_{\mathsf{prev}}))$ **to** pid and **await**
    - **if** next message m, σ' on input tape does not start with RESUMEREQUEST **then** ignore
    - **else** $\mathsf{out} \leftarrow m, \sigma'$
    - $\sigma_{\mathsf{prev}} \leftarrow \sigma'$
- **return** $\mathsf{out}, \sigma_{\mathsf{prev}}$

On message $(\text{INSTALLREQUEST}, \mathsf{prog})$ *from a party* $P \in \mathsf{reg}$*:*
- run $\mathsf{eid} \leftarrow (\text{INSTALL}, \mathsf{prog})$
- **send** $(\text{RESUME}, \mathsf{eid}, init))$ **to** G_{att}^{mod} and **receive** output, σ
- **return** $\mathsf{eid}, \mathsf{output}, \sigma$

On message $(\text{RESUMEREQUEST}, \mathsf{eid}, \mathsf{inp})$ *from a party* $P \in \mathsf{reg}$*:*
- run $\mathsf{output}, \sigma \leftarrow (\text{RESUME}, \mathsf{eid}, \mathsf{inp})$
- **return** $(\mathsf{output}, \sigma)$

Fig. 8. The wrapper protocol to implement interface I for a G_{att}^{mod} setup with oracle sets $\mathbb{O}, \mathbb{A}$

Shell $\mathrm{W}^I[\mathsf{prog}, \mathrm{nextmsg}]$ **(Template)**

The identity of the shell is $(\mathsf{eid}, \mathsf{idx})$
The parent shell extended identity is $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathrm{W}^{\mathrm{I}}[\mathsf{prog}]], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$

On message $(\mathrm{cmd}, args, r)$ *from* $(\mathsf{eid}||\mathsf{pid}, \textit{“att”}||\mathsf{idx})$*:*
  if virtual ITI $(\mathsf{prog}, (\mathsf{eid}||\text{“wrapped”}, \mathsf{idx}))$ does not exist **then** create
  if $r = \bot$ **then**
    let $\mathsf{inp} \leftarrow (\mathrm{cmd}, args)$
    step through execution of $(\mathsf{prog}, (\mathsf{eid}||\text{“wrapped”}, \mathsf{idx}))$ on inp:
    for instruction i **do**
      if $\mathsf{i} \notin I$ **then**
        // Execution of i is delegated to the higher order shell
        allow $(\mathsf{prog}, (\mathsf{eid}||\text{“wrapped”}, \mathsf{idx}))$ to execute i
      else
        $r \leftarrow \top$
        $(\mathsf{pid_j}, \mathsf{eid_j}, v) \leftarrow \mathrm{nextmsg}$(tapes of virtual ITI)
        while $(\mathsf{pid}_j, \mathsf{eid_j}) \neq (\bot, \bot)$ **do**
          send $(\textsc{ResumeRequest}, (\mathsf{pid_j}, \mathsf{eid_j}, v))$ **to** $(\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx})$ and **await**
          if next message on the input tape is $(\mathrm{cmd}', args', \top)$ **then**
            execute $(\mathsf{pid_j}, \mathsf{eid_j}, v) \leftarrow \mathrm{cmd}'(args', \top)$
          else
            ignore
        // The loop terminates when nextmsg() returns $(\bot, \bot, v)$
        write v to subroutine output of $(\mathsf{prog}, (\mathsf{eid}||\text{“wrapped”}, \mathsf{idx}))$
        $r \leftarrow \bot$
  else
    // $r = \top$ as the result of an I computation, execute code in subroutine cmd
    execute $\mathrm{cmd}(args)$
    return nextmsg(tapes of virtual ITI)

Fig. 9. Template for the internal wrapper shell. A complete definition of the shell requires an implementation for any additional CMD that might be requested by the next message functions

tion of local computation and the feature oracles it already has access to. During the global functionality initialisation phase, the simulator observes the signature algorithm s chosen by the environment through the dummy adversary, and provides G'_{att} with a new algorithm s' which applies s over the transformation $F(\mathsf{meas})$. F takes a measurement string that contains an identifier for program prog, and replaces it with an identifier for $\mathrm{W}^I[\mathsf{prog}, \mathrm{nextmsg}]$ (for an appropriate value of $\mathrm{nextmsg}$), as discussed in Section III. Attestations produced by an enclave prog in the G'_{att}-hybrid world are thus indistinguishable from those produced by the equivalent wrapped enclave in the G_{att}-hybrid world. Therefore, the simulator can simply block any installations of an un-wrapped program that requires access to I with MissingInstructionError, and replace installations of wrapped programs with the unwrapped version on the G'_{att} functionality. Honest parties in protocol $\mathcal{W}$ do not install any unwrapped program, and no external session will have direct access to G_{att} since it is installed as a $\mathcal{W}$ subroutine. If the (local) adversary attempts to install an unwrapped program to G_{att} directly, the simulator can run the program “in its head” without going through G'_{att}, and use the algorithm s provided by the dummy adversary for the environment to produce plausible attestation signatures for the unwrapped code. The signatures will not verify through any calls to the ideal verification subroutine, as they wouldn't for honest parties of $\mathcal{W}$, but they will look legitimate to environment through running the local verification algorithm that corresponds to s.

2) When $\mathsf{map}^{\mathsf{R}}$ does not install any assisting enclaves, but $\mathsf{map}^{\mathsf{L}}$ does, the simulator instantiates the same signature scheme as in the previous case (by adding the F transformation for an appropriate value of $\mathrm{nextmsg}$).
When it receives a request to install any enclave with code E_i^L, it generates a plausible enclave ID and returns it, without actually installing the enclave in G_{att}^{mod}. While we do not explicitly define an enclave ID generation algorithm for G_{att}^{mod}, we assume that the probability of sampling the same ID is negligible. The simulator then ensures that, before a corrupted party requests to install some enclave $\mathrm{W}^I[\mathsf{prog}, \mathrm{nextmsg}]$, it has requested to install all necessary E_i^L enclaves produced by $\mathsf{map}^{\mathsf{L}}()$, and has given a value of $\mathrm{nextmsg}$ with the appropriate enclave IDs, otherwise $\mathrm{W}^I[\mathsf{prog}, \mathrm{nextmsg}]$ would not be able to verify them for attestation.
Whenever the adversary resumes the program enclave, the simulator runs the input in its head to determine whether it contains any calls to I. If it does, it calls the next message function $\mathrm{nextmsg}$ on the partial result, and uses the signing key generated during initialisation by the adversary to produce signing algorithm s, and uses it to sign a RESUMEREQUEST message. If the adversary then tries to resume the receiver enclave, the simulator executes the related command in its head and returns the next RESUMEREQUEST message. Any attempts from the adversary to verify one of the intermediate attestation messages directly is dropped, since the protocol does not let parties verify these attestations either (they are however likely to be verified by the code of the wrapper enclave as part of its next command execution). Once it is satisfied that the adversary has provided the appropriate sequence of messages to fully compute I, it sends the initial original input to the unwrapped

program in G'_{att}. If the feature shell triggers any adversarial interaction, it uses the values provided by the adversary through RESUMEREQUEST messages to maintain a consistent state with the $\mathcal{W}$ interactions. Any interactions with the adversary through attacks or feature requests unrelated to I are captured by the ideal shell run by G_{att}, so no additional simulation is required for them.

3) Finally, in the case of the $\mathsf{map}^{\mathsf{R}}$ function requesting multiple enclaves across a variety of parties, the simulator initialises G^{mod}_{att} with the same signature algorithm as before. It then calls the $\mathsf{map}^{\mathsf{R}}$ function and sends the resulting resume requests to corrupted parties, but installs the assisting enclaves for honest parties on a machine it controls, and produces the appropriate list of assisting enclave IDs, $\bar{\mathrm{E}}^R$.
Like in the previous case, on an enclave installation request, it installs a non-wrapped copy of any enclaves requested by corrupted parties, as long as they have installed all the related local assisting enclaves. Simulation proceeds as in the previous case, except that the simulator also ensures that any remote RESUMEREQUEST message is delivered (i.e. the appropriate messages on the network are not censored). When a next command is sent to a remote assisting party run by some honest user, the simulator does not pass it on, and runs the command on its local copy to find out the next message location, using its copy of the s algorithm to sign plausible attestations (including faking the party ID if using non-anonymous attestation). Finally, if the computation succeeds, it calls the unwrapped enclave in G'_{att} as before. Any attempts by a corrupted party to send a RESUMEREQUEST to honest enclaves outside of the correct sequence of events is dropped.

General replacement of global setups: As we discussed in Appendix A1, it is not possible to prove, in the general case, that a protocol UC-emulates a global subroutine. A well formed replacement statement needs to account for the context emulation statement the global subroutine is being invoked in.

Intuitively, since the adversarial oracle sets for the $G_{\mathsf{att}}, G'_{\mathsf{att}}$ functionalities considered are the same, replacing the global functionality G'_{att} with a G_{att}-hybrid protocol $\mathcal{W}$ to provide the missing feature interface I should generally be safe, as a higher level simulator that interacts with TEEs as part of a protocol subroutine will have the same interface for attacks. However, given the general nature of our conjecture, we can not conclusively say that the implementation of I provided by $\mathcal{W}$ communicates with the adversary in the same manner as the ideal implementation of I provided by the G'_{att} shell. Indeed, the role of the $\mathcal{W}$ to G'_{att} simulator is to reconciling any such difference. We therefore have to analyse two distinctive cases.

Theorem 6

Let $G_{\mathsf{att}}, G'_{\mathsf{att}}, \mathcal{W}$ be any G^{mod}_{att} setups and a wrapper protocol such that Conjecture 2 holds, and additionally G'_{att} UC-emulates $\mathcal{W}$. For any protocol ρ in the presence of G'_{att} that UC-emulates some $\mathcal{F}$ in the presence of G'_{att}, ρ in the presence of $\mathcal{W}$ UC-emulates $\mathcal{F}$ in the presence of $\mathcal{W}$.

The statement follows from the composition theorem of [4, Theorem 3.3]. Showing that G'_{att} UC-emulates $\mathcal{W}$ (i.e. in conjunction with Conjecture 2, $\mathcal{W}$ and G'_{att} are UC-equivalent) involves constructing a new simulator $\mathcal{S}'$ such that $\mathsf{EXEC}_{G'_{\mathsf{att}},\mathcal{A},\mathcal{Z}} \approx \mathsf{EXEC}_{\mathcal{W},\mathcal{S}',\mathcal{Z}}$.

During the setup phase, $\mathcal{S}'$ instantiates G^{mod}_{att} with the inverse transformation $F^{-1}$for attestation signatures described in the proof guidelines for Conjecture 2 i.e. for any attestation measurement that includes an identifier for some program with code $\mathrm{W}^I[\mathsf{prog}, \cdot]$ and replaces it with an identifier for prog. Thereafter, the behaviour of $\mathcal{S}'$ consists of simply forwarding any input from the environment to the protocol $\mathcal{W}$ (including allowable attacks in $\mathbb{A}$), and after a RESUME, execute any associated RESUMEREQUEST for corrupted parties without modifying their inputs or showing the result to the environment, except for any adversarial leakage consistent with what would be produced by the shell implementation for G'_{att}. When $\mathcal{W}$ returns the output of the RESUME and associated attestation message, $\mathcal{S}'$ only forwards this result and its attestation (with the wrapper code removed by the F^{-1} transformation).

If the shell implementing feature I in the G'_{att} world includes direct communication with the adversary that is not fully equivalent by the messages produced by the supporting enclaves in $\mathcal{W}$, the simulation will fail. For such protocols we need to consider a weaker setting, where we fix the feature simulator within the ideal subroutine available to the higher level protocols.

Theorem 7

Let $G_{\mathsf{att}}, G'_{\mathsf{att}}, \mathcal{W}$ be any G^{mod}_{att} setups and a wrapper protocol such that Conjecture 2 holds for some simulator S. Let G^S_{att} be the combination of G'_{att} and S; for any protocol ρ in the presence of G^S_{att} that UC-emulates some $\mathcal{F}$ in the presence of G^S_{att}, ρ in the presence of $\mathcal{W}$ UC-emulates $\mathcal{F}$ in the presence of $\mathcal{W}$.

The statement above directly follows from [23, Lemma 1].

2) Generalising Adversarial Interfaces removal: The protocol and simulation strategy for feature addition can be used to extend the original attack removal protocol of Section V to cover protective protocol topologies that use multiple enclaves. A core difference from the oracle interface implementation of that section, however, is that, rather than interrupting the execution of a normal enclave program for a specific instruction to run a protocol between supplementary enclaves, it is necessary to run the defensive protocol from the start of the execution.

More formally, the wrapper protocol $\mathcal{W}$ is defined in the same way as the one in the previous section to implement missing features, but installing wrapper enclave $\mathrm{W}^A[\cdot]$ rather than $\mathrm{W}^I[\cdot]$ (for some A, I). The only difference between the two protocols is that $\mathrm{W}^A[\mathsf{prog}, \mathrm{nextmsg}]$ never executes the internal protocol prog directly. Instead, the $\mathrm{nextmsg}()$ function now takes the code prog as an additional argument, and compiles it into the code for local enclaves $\mathrm{E}^L_{i\in\{1,\ldots,n\}}$. When the party installs all enclaves $\{\mathrm{W}^A[\mathsf{prog}, \mathrm{nextmsg}], \mathrm{E}^L_1 \ldots, \mathrm{E}^R_n\}$, it immediately resumes all of them, in sequence, with message INIT, which allows the enclaves to conduce any necessary setup operations. Thereafter, on a RESUME call to prog, $\mathrm{W}^A[\mathsf{prog}, \mathrm{nextmsg}]$ always begins its execution by running $\mathrm{nextmsg}$ first. The function returns a tuple $(\mathsf{pid}_{\mathsf{i}}, \mathsf{eid}_{\mathsf{i}}, v)$, indicating that $\mathrm{W}[\cdot]$ should send an authenticated message (through attestation) with payload v to enclave $\mathsf{eid}_{\mathsf{i}}$ running on party $\mathsf{pid}_{\mathsf{i}}$ (on the same party if the first element of the tuple is $\perp$). Each assisting enclave will also run their own copy of the next message functions to forward the intermediate computation results to the next enclave. When $\mathrm{nextmsg}$ returns $(\perp, \perp, v)$, this indicates that the RESUME call has completed, and the wrapper enclave can return v to the calling party.

We now restate Conjecture 1 for the full version of protocol $\mathcal{W}$.

Conjecture 3

Let $G_{\mathsf{att}} = G^{mod}_{\mathsf{att}}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}]$ and $G'_{\mathsf{att}} = G^{mod}_{\mathsf{att}}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}', \mathbb{S}]$ such that $\mathbb{A} \setminus \mathbb{A}' = A$. For any enclave wrapper $\mathrm{W}^A[\cdot]$ which, combined with functions $\mathsf{map}^{\mathsf{L}}, \mathsf{map}^{\mathsf{R}}$ implements the difference in behaviour between the shells $\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\cdot]$ and $\mathrm{sh}_{\mathbb{O},\mathbb{A}'}[\cdot]$, protocol $\mathcal{W}[\lambda, \mathsf{reg}, \mathbb{O}, \mathbb{A}, \mathbb{S}, \mathsf{map}^{\mathsf{R}}, \mathsf{map}^{\mathsf{L}}, \mathrm{W}^A[\cdot]]$ in the presence of G_{att} UC-emulates G'_{att}

3) Interactions Between Features and Attacks: When defining the transformation between two versions of G_{att}, it is important to think carefully about specifying the necessary requirements. Introducing a new attacker oracle to a TEE might not allow any sensible protection mechanism without the enclave having access to certain feature oracles.

Additionally, in some cases the addition of a new feature will also imply the expansion of adversarial attacks. Consider the addition of explicit storage and fetching capabilities described in Section IV-C. By adding those external oracle calls, we are also forced to provide an adversarial oracle to abort the program. While it would be possible to consider a version of G_{att}^{mod} where only the new interfaces were added, it would be hard to justify as the natural implementation of that feature requires handing off control of the memory to untrusted permanent storage. Of course, a novel TEE architecture could allow a more secure way to implement storage and fetching without exposing the enclaves to adversarial crashes. Our goal for G_{att}^{mod} is not to be prescriptive with what kind of (ideal) TEE objects should be used as assumptions in cryptographic protocols; however we recommend caution when designing a new variant of G_{att}^{mod} with complex or unrealistic features.

Another illustrative example could be the introduction of enclave cloning [52] to a G_{att}^{mod} setup. This feature allows efficient enclave creation, as it instantiates a second copy of an enclave including its memory (equivalent to normal process forking in operating systems). Depending on the implementation, the addition of this feature might however give the adversary additional power, as it could now be able to swap memory regions for each of the two versions of the enclave interactively, effectively executing a forking attack not tied to rollback (where the remote party is not able to distinguish which of the two enclaves it is communicating with, and the adversary can interactively swap and censor messages between the two). While this specific attack can be easily mitigated with another wrapper protocol that augments sealing with freshness values, it will require an additional explicit transformation and corresponding level of shells.

Our theorems in this section only show a single step G_{att}^{mod} oracle change (through feature addition and attack removal). Unlike the oracle shells in our Zoo, which have to be manually integrated to provide the appropriate functionality for the set of oracles chosen (although in many cases the shell changes are trivial), it should be easy for some oracle combinations, where they don't negatively interact with each other, to repeatedly apply Conjectures 2 and 1 without modifying the protocols.

We note that in some cases the oracle transformation protocols given above might not be simulatable for all possible enclave programs. In those cases, it is still possible for a program designed to run in G'_{att} to run in the G_{att}-hybrid world where the oracle feature is not available, or be secure even if G_{att} allows an attack not in G'_{att}. Such substitution require to be shown as valid on a case by case basis, but the observation is consistent with the state of the art of TEE program design, where mitigations for certain attacks exist only if the program is "well-written" (e.g. memory safe or using oblivious primitives) or does not use certain functions (see [71, Table 1]).

C. Rollback protection Proof of Security

We now give a proof of security for Theorem 2 (with the full protocol $\mathcal{W}$ defined in Appendix B) by defining the appropriate simulator. First, we show that the theorem is an instantiation of Conjecture 3 by showing how the protocol described in Section V is a special case of the more general protocol described in the Appendix. In particular, the protocol can be implemented through a pair of functions $\mathsf{map}^{\mathsf{R}}()$ and $\mathsf{map}^{\mathsf{L}}()$ that produce no supporting enclaves.

Shell $\mathrm{W}^{A}[\mathsf{prog}, \mathrm{nextmsg}_{\mathrm{meas}}]$

The identity of the shell is $(\mathsf{eid} \parallel c, \mathsf{idx})$
The parent shell extended identity is $(\mathrm{sh}_{\mathbb{O},\mathbb{A}}[\mathrm{W}^{\mathrm{I}}[\mathsf{prog}]], (\mathsf{eid}||\mathsf{pid}, \text{"att"}||\mathsf{idx}))$

On message INIT *from* $(\mathsf{eid}||\mathsf{pid}, \textit{"att"}||\mathsf{idx})$*:*
  if $\mathrm{Fetch}() \neq \epsilon$ **then**
    return ABORT
  install virtual ITI $(\mathsf{prog}, (\mathsf{eid}||c||\text{"wrapped"}, \mathsf{idx}))$
  let $m \leftarrow \mathrm{Meas}()$
  $\mathrm{Store}(m)$
On message inp *from* $(\mathsf{eid}||\mathsf{pid}, \textit{"att"}||\mathsf{idx})$*:*
  while $\top$ **do**
    $\mathsf{out} \leftarrow \mathrm{nextmsg}_{\mathrm{meas}}$(tapes of virtual ITI)
    if $\mathsf{out} = (\mathsf{pid}, \mathsf{eid}, (\text{MEASEXEC}, \mathsf{inp}))$ **then**
      run MEASEXEC(inp)
    else if $\mathsf{out} = (\bot, \bot, v) \wedge v \neq \textit{"abort"}$ **then**
      return v
    else
      erase the virtual ITI work tape
      abort
let MEASEXEC(inp):
  $m \leftarrow \mathrm{Fetch}(), m' \leftarrow \mathrm{Meas}()$
  if $m \neq m'$ **then abort**
  step through execution of $(\mathsf{prog}, (\mathsf{eid}||c||\text{"wrapped"}, \mathsf{idx}))$ on inp:
  for instruction i **do**
    if $\mathsf{i} =$ (**return** v) **then**
      $b \leftarrow \mathrm{Write}(\mathrm{Meas}())$
      assert $b = OK$
    else allow $(\mathsf{prog}, (\mathsf{eid}||c||\text{"wrapped"}, \mathsf{idx}))$ to execute i

Fig. 10. The $\mathrm{W}^{A}[\cdot]$ enclave shell installed by protocol $\mathcal{W}$ for rollback iteration c of enclave eid installed by party pid for session idx (full version using the next message function)

Function $\mathsf{map}^{\mathsf{L}}()$ defines a next message function $\mathrm{nextmsg}_{\mathrm{meas}}()$, which determines how to execute the wrapped program. When the enclave state is at the beginning of executing a RESUME instruction, $\mathrm{nextmsg}_{\mathrm{meas}}$ runs the MEASEXEC subroutine of $\mathrm{W}^{A}[\cdot]$. Subroutine MEASEXEC checks that the current measurement of the enclave's state corresponds to the last state saved in storage, before executing the input subroutine, and updating the storage with the resulting new state. If MEASEXEC aborts, $\mathrm{nextmsg}_{\mathrm{meas}}$ returns $(\bot, \bot, \text{"abort"})$, while if it terminates successfully with value v, it returns $\bot, \bot, v$; in both cases, the enclave returns the values to its caller. A definition of the wrapper shell that uses this next message function is shown in Figure 10.

We now define a simulator for the wrapper protocol in Figure 10 . The simulator roughly follows the sketch outlined in the first case of the proof strategy for Conjecture 2, although we modify it appropriately for the adversarial case.

Assume the simulator has access to the same parameters as $\mathcal{W}$. The simulation translates all requests to install a wrapped enclave from corrupted parties into requests to install the unwrapped enclave in G'_{att}; any attempt to install an unwrapped enclave will be simulated "in its head". Thereafter, whenever the party resumes one of the wrapped enclaves, the simulator fakes an access to the Fetch oracle, to reproduce the behaviour of the MEASEXEC subroutine in wrapper $\mathrm{W}^{A}[\cdot]$ to check that the enclave was not previously rolled back. If the check succeeds, the enclave begins executing the program ideally through running its non-wrapped

version through the G'_{att} functionality. During its execution, the shell $(\mathrm{sh}_{\mathbb{O},\mathbb{A}'}[\mathsf{prog_w}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$ might send messages on the backdoor tape related to some attacks in $\mathbb{A}'$ unrelated to rollback (therefore present in both real and ideal world). In that case, the simulator forwards it to the adversary and returns its response back to the shell without modification. After the execution of the enclave program has terminated, the simulator fakes a call to the Store oracle, with the length of the hash function used for measuring enclave states (m) as its leakage.

If at any point during the simulation the adversary aborts a simulated oracle call, or if the simulator has recorded (in dictionary P) that the adversary has issued a rollback attack against that enclave, it will issue an abort message through the adversarial interface of G'_{att}, and halt its own execution. Otherwise, if all the checks succeed, it returns the output value and attestation signatures produced by G'_{att}. Additionally, the simulator produces an ITER message to signal that the RESUME execution has been successful, and the creation of a new copy for the ITI state (as if the enclave was running on G_{att}). Attestation verification requests are forwarded to G'_{att} if they are for the wrapped version of an enclave (where it will succeed only if the unwrapped version of the same enclave issued that message, before being transformed by F). Any request to verify a message where the attestation contains the unwrapped code (which is what is actually running on G'_{att}) is rejected.

Calls to install, resume, or verify the attestation of any unwrapped enclaves are not allowed by the protocol, but a corrupted party might try to get around this by directly writing to the tapes of real world G_{att} subroutine - this is allowed by the identity bound. In that case, the simulator lets the message through to its local simulated G_{att} subroutine, which can produce a convincing attestation signature for any message by using the original s algorithm. To denote this, we adopt the convention of forwarding adversarial messages for unwrapped enclaves to a “fake” copy of the hybrid functionality G^F_{att}. It is possible to think of G^F_{att} as simply shorthand for the book keeping operations inlined by the simulator's code, similar to the roles of the dictionary $\mathcal{G}$ in the Steel simulator of [10]. Alternatively, it is possible to see G^F_{att} as a bona-fide instance of G^{mod}_{att} run by the simulator as a local subroutine, and therefore granting no access to machines in other sessions. Adopting this view is only possible in our modular setting: while the Steel simulator, in the presence of G^{PST}_{att}, was required to keep a separate record of all messages signed by adversarial enclaves, this is the default for G^{mod}_{att}, and therefore we do not require keeping track of any additional operations. G^F_{att} is taken to be initialised with the same arguments as the real world G_{att} emulated by the protocol, such that any attempts to access an attack in $\mathbb{A}$ is reproduced by its (simulated) shells.

The pseudocode for the simulator described above is as follows:

Simulator $\mathcal{S}$

$F(a, f)$ is the function that transforms an attestation measurement a so that it replaces the code of an enclave program p with code $\mathrm{W}^A[p, f]$. $\mathbb{M}$ is the standard uniform length for the output of $\mathrm{Meas}()$ oracle calls

State variables	Description
$\mathrm{P} \leftarrow \{\}$	Dictionary of state pointers for rollback protected enclaves

On message INITIALISE *from* G'_{att}:
- **send** INITIALISE **to** $\mathcal{A}$ **through** G_{att} and **receive** pk, s
- $\mathrm{nextmsg}_{\mathrm{meas}} \leftarrow \mathsf{map}^{\mathsf{L}}(\bar{\mathrm{E}}^R)$
- **let** $s'(x) \leftarrow s(F(x, \mathrm{nextmsg}_{\mathrm{meas}})$
- **send** $(pk, s'))$ **to** G'_{att} **on behalf of** $\mathcal{A}$
- **send** INITIALISE **to** G^F_{att} **through** $\mathcal{Z}$ and **receive** INITIALISE
- **send** Σ **to** G^F_{att} **on behalf of** $\mathcal{A}$

On message $(\text{INSTALL}, \mathsf{idx}, \mathsf{prog})$ *from corrupted party* P:
- **if** $\mathsf{prog} = \mathrm{W}^A[\mathsf{prog_w}, \mathrm{nextmsg}_{\mathrm{meas}}]$ **then**
  - **send** $(\text{INSTALL}, \mathsf{prog_w})$ **to** G'_{att} **through** P and **receive** eid
  - $\mathrm{P}[P, \mathsf{eid}] \leftarrow (\mathsf{idx}, \mathsf{prog_w}, \emptyset, \emptyset)$
- **else**
  - **send** $(\text{INSTALL}, \mathsf{prog})$ **to** G^F_{att} **through** P and **receive** eid
- **return** eid

On message $(\text{RESUME}, \mathsf{eid}, (i, \mathsf{inp}), \text{Rollback})$ *from corrupted party* P:
- **if** $\mathrm{P}[P, \mathsf{eid}] = (\mathsf{sid}, \mathsf{prog}, c, c_{latest})$ **then**
  - $\mathrm{P}[P, \mathsf{eid}] \leftarrow (\mathsf{sid}, \mathsf{prog}, i, c_{latest})$
  - **if** $\mathsf{inp} \neq \epsilon$ **then**
    - run $\mathsf{out}, \sigma \leftarrow$RESUME$(\mathsf{eid}, \mathsf{inp}, \epsilon)$,
  - **else**
    - **send** (ITER, c, i) **to** $\mathcal{A}$ **on behalf of** enclave shell
- **else**
  - **send** $(\text{RESUME}, \mathsf{eid}, (i, \mathsf{inp}), \text{Rollback})$ **to** G^F_{att} **on behalf of** P

On message $(\text{RESUME}, \mathsf{eid}, \cdot, \text{Abort})$ *from corrupted party* P:
- **if** $\mathrm{P}[P, \mathsf{eid}] \neq \perp$ **then**
  - **send** $(\text{RESUME}, \mathsf{eid}, \epsilon, \text{Abort})$ **to** G'_{att} **on behalf of** P
- **else send** $(\text{RESUME}, \mathsf{eid}, \epsilon, \text{Abort})$ **to** G^F_{att} **on behalf of** P

On message $(\text{RESUME}, \mathsf{eid}, \mathsf{inp}, a)$ *from corrupted party* P:
- **if** $(\mathsf{sid}, \mathsf{prog_w}, c, c_{latest}) \in \mathrm{P}[P, \mathsf{eid}]$ **then**
  - **assert** $a = \epsilon \vee a \in \mathbb{A}'$
  - **let** $shEID \leftarrow (\mathrm{sh}_{\mathbb{O},\mathbb{A}'}[\mathsf{prog_w}], (\mathsf{eid}||\mathsf{pid}, \text{“att”}||\mathsf{idx}))$
  - **send** FETCH **to** $\mathcal{A}$ **through** $shEID$ and **receive** b
  - **if** $b \neq \mathrm{Continue} \vee c \neq c_{latest}$ **then**
    - **send** $(\text{RESUME}, \mathsf{eid}, \epsilon, \text{Abort})$ **to** G'_{att} and **return**
  - **send** $(\text{RESUME}, \mathsf{eid}, \mathsf{inp}, a)$ **to** G'_{att} **on behalf of** P and
  - **while receive** $(\mathrm{msg}, args)$ **from** $shEID$ **do**
    - **send** $(\mathrm{msg}, args)$ **to** $\mathcal{A}$ **on behalf of** $shEID$
    - **forward** adversarial response to $shEID$
    - **if** response $= \mathrm{Abort}$ **then return**
  - **receive** out, σ **from** G'_{att}
  - **send** $(\text{STORE}, 1^{\mathbb{M}})$ **to** $\mathcal{A}$ **through** $shEID$ and **receive** b'
  - **if** $b' \neq \mathrm{Continue}$ **then**
    - **send** $(\text{RESUME}, \mathsf{eid}, \epsilon, \text{Abort})$ **to** G'_{att} and **return**
  - generate nonce $c' \xleftarrow{\$} \{0,1\}^\lambda$
  - $\mathrm{P}[P, \cdot, \mathsf{eid}, \mathsf{prog_w}] \leftarrow (c', c')$
  - **send** (ITER, c, c') **to** $\mathcal{A}$ **on behalf of** $shEID$
- **else**
  - **send** $(\text{RESUME}, \mathsf{eid}, \mathsf{inp}, \epsilon)$ **to** G^F_{att} and **receive** out, σ
- **return** out, σ

On message $(\text{VERIFY}, \sigma, m)$ *from corrupted party* P:
- **if** m is a measurement for an enclave with program $\mathrm{W}^A[\mathsf{prog_w}, \mathrm{nextmsg}_{\mathrm{meas}}]$ **then**
  - **send** $(\text{VERIFY}, \sigma, F(m, \mathrm{nextmsg}_{\mathrm{meas}}))$ **to** G'_{att} and **receive** v
- **else if** m is a measurement for an “unwrapped” program with enclave ID eid installed by some party P', and $\mathrm{P}[P', \mathsf{eid}] \neq \perp$]
- **then return** $\perp$
- **else send** $(\text{VERIFY}, \sigma, m))$ **to** G^F_{att} and **receive** v
- **return** v

For any protocol that adopts the standard identity bound, preventing the environment from sending messages on behalf of corrupted parties outside of the test session, the environment can not distinguish the real or ideal world, due to the simulator constructing a perfect transcript for the execution of $\mathcal{W}$ with the attestation signatures in the ideal world verifying for a real world $\mathrm{W}^A[\cdot]$ program.

Consider the case where the adversary does not conduct a rollback attack. For every RESUME operation from the corrupted party, the simulator activates the adversary with message FETCH, allowing it to interrupt the computation. If this happens, the simulator mounts

the equivalent ABORT attack on G'_{att}. If FETCH is allowed, the measurement stored will be the same as from the previous execution, and therefore the simulator runs the program in G'_{att}. The behaviour of this execution is equivalent to the real world setup, since the shells of G_{att} and G'_{att} implement the same (non-rollback) oracles, and the simulator lets through any such adversarial access. Finally, the adversary receives a final STORE for a message of the same length as a MEAS value. Since the storage oracle does not leak the message contents but only their size, the adversary can not distinguish it from a state storage as executed during the MEASEXEC subroutine. If it chooses to abort, the real world wrapper would never terminate, so the simulator does the same for the ideal world enclave (by issuing its own ABORT), otherwise it returns the (ideally computed) value. The distribution of the return value for the enclave as executed in G_{att} and G'_{att} is equivalent (given they have the same feature oracles implementation), and the modified signature scheme attests to code $\mathrm{W}^A[\mathsf{prog}, f]$ in both worlds, thanks to the transformation F.

We now describe the case of an adversary who, after some sequence of successful resumes, issues a rollback attack to an earlier state. The code of subroutine $\mathrm{W}^A[\cdot]$ does not allow executing any further RESUME, since the assertion that the measurement stored is equal to the current one will fail with non-negligible probability (as long as the measurement computed by oracle MEAS is collision-resistant, and the code of the enclave program iterates through a sufficiently diverse state distribution[5]). The simulator perfectly reproduces this behaviour, by issuing an ABORT to the ideal enclave, after having issued the preceding FETCH. The environment can not distinguish real and ideal world, as in neither cases the interrupted enclave will be able to proceed.

[5]If the enclave is running a program with a very limited set of states, such as a small finite state automaton, it is possible to artificially expand the state space by augmenting the program with a monotonically increasing counter for each resume. This will ensure that every measurement is distinct.