

# AUC: Accountable Universal Composability

Mike Graf,<sup>✉</sup> Ralf Küsters,<sup>✉</sup> and Daniel Rausch<sup>✉</sup>

University of Stuttgart

Stuttgart, Germany

Email: {mike.graf,ralf.kuesters,daniel.rausch}@sec.uni-stuttgart.de

## Abstract

Accountability is a well-established and widely used security concept that allows for obtaining undeniable cryptographic proof of misbehavior, thereby incentivizing honest behavior. There already exist several general purpose accountability frameworks for formal game-based security analyses. Unfortunately, such game-based frameworks do not support modular security analyses, which is an important tool to handle the complexity of modern protocols.

Universal composability (UC) models provide native support for modular analyses, including re-use and composition of security results. So far, accountability has mainly been modeled and analyzed in UC models for the special case of MPC protocols, with a general purpose accountability framework for UC still missing. That is, a framework that among others supports arbitrary protocols, a wide range of accountability properties, handling and mixing of accountable and non-accountable security properties, and modular analysis of accountable protocols.

To close this gap, we propose *AUC*, the first general purpose accountability framework for UC models, which supports all of the above, based on several new concepts. We exemplify *AUC* in three case studies not covered by existing works. In particular, *AUC* unifies existing UC accountability approaches within a single framework.

## 1. INTRODUCTION

Accountability is a prominent concept that is widely used in security. Many security properties and applications, such as auctions [7, 29, 97], e-voting [1, 65, 67], non-repudiation [98–100], multi-party computation (MPC) [6, 8, 13, 54], public key infrastructures (PKIs) [61, 73, 76], distributed ledgers [15, 16, 38, 40, 47, 50, 59, 82, 88], DRM [91, 93], power infrastructures [58, 75], content delivery networks [31], and distributed systems [52, 90, 94] make use of and rely on accountability to provide security. In the formal security analysis literature, so-called *property-*, *policy-*, or *goal-based* accountability is the standard and most commonly used interpretation of accountability (e.g., [33, 42–45, 52, 57, 62, 63, 67, 68, 79]), which we therefore also consider in this work and often just call *accountability* (see also Section 2.7). Property-based accountability intuitively states that, if some intended security property of a protocol, such as correctness of the output, is violated, then we obtain a cryptographic proof that one or more protocol participants have misbehaved. We call security properties ensured via accountability in this way *accountability(-based) properties*. Parties which violated an accountability property can be held accountable for their misbehavior, e.g., via contractual and financial penalties or by excluding them from future protocol runs. This serves as a strong incentive for malicious parties to honestly follow the protocol and to not break security goals.

Over the past decade, researchers have developed general tools and approaches to formally analyze accountability properties in game-based settings (e.g., [46, 62, 67]). These works cover many flavors of accountability, such as (i) different *accountability levels*; a weak level might guarantee identification only of a (potentially large) group such that at least one party in that group has misbehaved, where the strongest level, so-called *individual accountability* [67], allows for identifying one or more parties such that *all* of them have misbehaved. (ii) *Local* or *public/universal* accountability (w.r.t. some security goal/property) [41, 49]; a public accountability property allows everyone, including external observers, to identify misbehaving parties in case the security property is broken. Local accountability allows only protocol participants to identify misbehaving parties.

**Preventive vs. Accountability Properties.** The counterpart to accountability properties are so-called *preventive security properties* (cf. [45]), which includes many special cases such as *proactive security* [2, 12]. In preventive security, one proves that the expected security property cannot be broken in the first place, typically based on certain (strong) security assumptions. In contrast, accountability-based security accepts that security properties might be broken by misbehaving parties but instead requires that one can identify and hence deter such parties. In exchange for this weaker security guarantee, accountability provides several advantages, including: (i) accountability might already be achievable with simpler, more efficient components. (ii) In order to ensure

accountability properties, one might need less security assumptions. *(iii)* Accountability properties might be stronger in some aspects than preventive security notions. For example, Graf et al. [50] illustrate and leverage all of these advantages of accountability by proposing a slight modification of the Hyperledger Fabric permissioned blockchain that *(i)* uses only a very efficient crash-fault-tolerant instead of a more complex Byzantine-fault tolerant consensus protocol, *(ii)* does not assume an honest (super)majority or any other set of honest parties to achieve accountability of consistency, and *(iii)* can enforce consistency (in an accountable way) not only for honest nodes but also for dishonest ones. So both concepts, preventive and accountability-based security, come with their own merits and tradeoffs. They are *orthogonal* concepts in the sense that both types of security can be used as stand-alone mechanisms to provide security for an intended property. But they can also be combined for the same security goal, where accountability serves as a second layer of defense in case the underlying assumptions of preventive security are broken. Such a combination is, for instance, used by the system PeerReview [52] to strengthen the security property of *consistency* in Byzantine-fault tolerant (BFT) consensus protocols (e. g., [28, 83, 95]): as long as there is an honest supermajority in the BFT protocol, consistency cannot be broken at all. If this assumption is no longer met, then PeerReview running on top of the BFT protocol still provides accountability w.r.t. consistency, i. e., allows for identifying parties that cause the consensus protocol to fail.

**Universal Composability.** Universal composability (UC) (e. g., [17, 18, 22, 23, 53, 64]) is a very popular approach for designing, modeling, and analyzing security protocols due to its strong security guarantees and its inherent support for modular design and analysis. Roughly speaking, in UC one first specifies an *ideal protocol* (or *ideal functionality*)  $\mathcal{F}$  that specifies the intended behavior/security properties of a target protocol, abstracting away implementation details. For a concrete realization – the *real protocol* –  $\mathcal{P}$ , one then proves that “ $\mathcal{P}$  behaves just as  $\mathcal{F}$ ” in arbitrary contexts, where the network of  $\mathcal{F}$  is controlled by a benign attacker called simulator. One can then build other protocols  $\mathcal{P}'$  on top of  $\mathcal{F}$  and analyze their security. A so-called composition theorem provided by the underlying UC model then implies that  $\mathcal{P}'$  remains secure even if the subroutine  $\mathcal{F}$  is replaced by  $\mathcal{P}$ .

**Current State.** Preventive security properties and their formalization via ideal functionalities in UC models is a well-studied problem, with the vast majority of existing UC literature focusing on preventive properties. In contrast, accountability properties have only been studied and formalized for special cases, mostly MPC protocols (e. g., [13, 14, 24, 34, 35, 55, 74]). These works do not and were not intended to serve as general UC accountability frameworks. In particular, these works *(i)* are tailored towards MPC protocols and accountability properties thereof, such as *identifiable abort* (e. g., [14, 55], cf. Appendix G, where we capture identifiable within our formalization approach), *(ii)* assume certain protocol structures, such as non-interactive protocols or protocols without internal servers, *(iii)* typically focus on either local or public accountability, *(iv)* cannot express arbitrary relationships of properties, including preventive properties with accountability as second layer of defense, and/or *(v)* consider composition out of scope or focus on composition of certain protocol types.

Hence, a general framework for accountability in UC models is still missing. Such a general framework would

1. Enable the design and analysis of arbitrary protocols and accountability properties within a UC model, thereby allowing such protocols to benefit from strong security statements and modularity offered by UC models,
2. Provide a common tool set that allows for easily formalizing accountability properties and reducing the effort for protocol designers, and thus,
3. Help avoid mistakes and oversights that can otherwise easily occur when formalizing accountability for every use case from scratch, as well as
4. Facilitate the comparison of different types of accountability properties and guarantees as they are defined within the same overarching and unified framework.

Obtaining such a general framework for accountability in UC models is non-trivial as it has to support and combine a number of different features to achieve the desired generality. Among others, it has to be able to express many different flavors of accountability, including the following flavors from formal game-based accountability security frameworks [62, 63, 67, 79]: *(i)* accountability for a wide range of security properties, *(ii)* different levels of accountability, *(iii)* both local and public accountability properties, and *(iv)* relationships of different properties, including combinations of accountability and preventive properties both as independent concurrent properties but also as layered defense for the same security goal. Such a general framework further

must be able to express virtually arbitrary protocols. In particular, it must (*v*) be independent of a specific protocol type and structure, thereby supporting, e. g., interactive protocols and protocols with purely internal parties such as client-server protocols, and (*vi*) fully support the modular design, analysis, and composition of protocols with accountability.

**AUC – Accountable Universal Composability.** To close this gap, we propose the *Accountable Universal Composability (AUC) framework*, the first general framework for the modular UC analysis of accountability properties. The AUC framework works within existing UC models, such as the UC [22, 23], including its variants (e.g., SUC [24], GUC [18]), the GNUC [53], the CC [77], and the IITM [64] models. By this, AUC inherits and can leverage features of the underlying UC model, such as the respective composition theorems possibly including support for composition with joint, arbitrarily shared, and/or global state. This also allows protocol designers to compose AUC protocols with existing (preventive security) protocols while remaining within whatever model they are already familiar with.

A major component of AUC is a generic transformation that, given any ideal functionality, allows for incorporating a wide range of accountability properties into the functionality. AUC further enables modeling and analyzing accountability properties in the corresponding real protocols such that composition of the resulting (accountability-based) protocols, also with preventive security protocols, is fully supported. To this end, AUC generalizes several ideas from the literature, both in game-based and UC settings, but also adds novel concepts, such as what we call *judicial reports* and *supervisors*. By combining these concepts, AUC achieves all the previously mentioned goals and features of a general framework for property-based accountability.

To exemplify features, applications, and the generality of AUC, we present the first accountability analyses of three different case studies in a UC model. These case studies are chosen to be relatively simple to better illustrate AUC. Firstly, we show how an accountable consensus service can be scaled up, e. g., to support a larger number of clients, by adding a scaling protocol layer on top while retaining accountability of the overall protocol. Using composability, this result can be iterated arbitrarily often to obtain security of multiple scaling layers. The case study introduces and illustrates general techniques that can be used for future analyses of accountability of existing real-world protocols that follow a similar scaling approach, such as the prominent blockchain Hyperledger Fabric [3, 50], the consensus service Hashgraph [11], and content delivery networks [39]. Secondly, we model and analyze accountability of a public key infrastructure (PKI) based on certificate transparency logs (CTLs). We then, thirdly, compose this result with an ISO 9798-3-based authenticated key exchange [17, 26, 56, 66], showing that the resulting protocol provides security based on an accountable PKI. To the best of our knowledge, this analysis is the first UC analysis of that protocol without assuming pre-distributed public keys or an idealized PKI where the adversary cannot create certificates for honest parties.

**Contributions.** In summary, the contributions of the paper are as follows:

- We propose AUC, the first general framework for accountability in UC models.
- AUC transfers and generalizes existing concepts from game-based approaches to the UC world, generalizes existing UC approaches, and develops new concepts.
- AUC supports, among others, arbitrary ideal and real protocols, a wide range of accountability properties, local and public accountability properties, concurrent consideration and combination of preventive and accountability properties, accountability of protocol internal parties, and composition of accountable protocols.
- To exemplify AUC, we present three case studies, providing the first UC analysis of accountability properties for the considered protocols.
- As a sanity check, in Appendix G we further show that AUC can capture accountability aspects of existing UC MPC literature as a special case, thereby generalizing and unifying this line of work.

**Structure of this paper.** Section 2 presents AUC, including a discussion of its core concepts. Section 3 provides our case studies. We discuss related work in Section 4. Further details are given in the appendix.

## 2. AUC – ACCOUNTABLE UNIVERSAL COMPOSABILITY

In this section, we introduce the accountable universal composability (AUC) framework. We first clarify notation and terminology in Section 2.1. Section 2.2 discusses several high-level ideas of AUC before we formally specify the AUC transformation for ideal functionalities in Section 2.3, with AUC’s modeling of real protocols covered in Section 2.4. In Section 2.5, we discuss several aspects of AUC’s composability abilities. In Section 2.6, we present a deterrence analysis to analyze the behavior of rational adversaries. Section 2.7 concludes this section with a discussion on AUC.

## 2.1 Notation and Terminology

**Computational Model.** Formally, we define AUC within the iUC model [17], an easy to use but fully expressive instantiation of the IITM model [64]. However, AUC and its concepts can also be used within arbitrary other models for universal composability, e. g., [23, 53, 77]. We keep our presentation on a level such that readers familiar with these UC models can understand and use AUC without requiring any prior knowledge of iUC. For interested readers, we provide a brief introduction to the iUC model including an overview of its pseudo code notation in Appendix A and in Appendix B, we provide an overview of the used pseudo code notation.

A party in iUC and hence also AUC is uniquely identified via its party ID  $pid$ , the session  $sid$  it runs in, and the piece of code/role  $role$  it is executing. We call the triple  $(pid, sid, role)$  *entity*.<sup>1</sup> In what follows, we use the terms entity and party synonymously.

We call a party in a protocol *main* if it can directly receive inputs from and send outputs to the environment. We call a party *internal* otherwise, i. e., if it is part of an internal subroutine. Whether a party is an internal or a main party can be determined from its role. As in all UC models, an ideal functionality and a realization share the same sets of main parties/roles. A realization might have additional internal parties/roles, such as an internal server used by main clients, that are not present in the ideal protocol.

As is standard in UC models, the adversary  $\mathcal{A}$  is allowed to corrupt parties by sending a special corrupt command. If an entity is corrupted, the adversary generally gets full control over the (input and output interface of the) entity. The environment can obtain the current corruption status of main parties in a protocol, which allows for checking whether corruption of main parties is simulated correctly.

**Classes of Security Properties.** As mentioned in the introduction, the literature traditionally distinguishes between *preventive security properties* and *accountability properties*. We denote the set of accountability properties of a protocol by  $Sec^{acc}$  and divide preventive properties into two classes:

**Absolute:** A preventive security property is called *absolute* if all underlying assumptions used in the security proof are assumed to always hold true. The analysis of the case where such assumptions might become broken is out of scope. We denote the set of absolute security properties (of some protocol) by  $Sec^{abs}$ .

**Assumption-based:** We call a preventive security property *assumption-based* if it is shown to hold true under certain assumptions but the security analysis also analyzes the case that these assumptions might become broken at some point. We denote the set of assumption-based preventive security properties by  $Sec^{assumption}$ .

By this categorization, we have that absolute security properties can neither be assumption-based properties nor accountability properties, i. e., we require  $Sec^{abs} \cap (Sec^{assumption} \cup Sec^{acc}) = \emptyset$ . The set  $Sec^{assumption} \cap Sec^{acc}$ , if non-empty, contains preventive security properties that offer accountability as a second layer of defense whenever the underlying assumptions are broken. The set  $Sec^{assumption} \setminus Sec^{acc}$  contains those assumption-based security properties which are not additionally secured via accountability.

**Verdicts.** AUC defines verdicts to be positive boolean formulas consisting of propositions of the form  $dis(A_i)$  where  $A_i$  is an entity and  $dis(A_i)$  expresses that the judge believes that  $A_i$  misbehaved/is dishonest. For example, the verdict “ $dis(A_1) \wedge (dis(A_2) \vee dis(A_3))$ ” captures the statement that party  $A_1$  and at least one of the parties  $A_2$  and  $A_3$  have misbehaved. We can evaluate verdicts by setting  $dis(A_i) = \text{true}$  iff  $A_i$  is actually a corrupted entity at the point where the verdict is stated, and  $\text{false}$  otherwise. We call a verdict *fair* if it evaluates to  $\text{true}$ , and hence, does not mistakenly blame honest parties. In a secure protocol, all honestly computed verdicts are required to be fair.

This definition of verdicts allows for capturing different levels of accountability. Verdicts such as  $dis(A_1)$  or  $dis(A_1) \wedge dis(A_2)$  imply that the specific party  $A_1$  (and potentially others) misbehaved. This captures the strongest level of so-called *individual accountability*. In contrast, a verdict of the form  $dis(A_1) \vee dis(A_2) \vee dis(A_3)$  only identifies a group of three parties where at least one has misbehaved, therefore capturing a weaker level of accountability.

## 2.2 Overview of AUC’s Central Concepts

The AUC framework serves as a general blueprint for modeling and analyzing a wide range of accountability properties both in real and ideal protocols. Before delving into AUC, let us first give a high-level overview of its main concepts and ideas:

<sup>1</sup>In those UC models that identify parties only via the pair  $(pid, sid)$ , different roles can be modeled by adding them as a prefix to  $pid$ , say,  $pid = (role, pid')$ , where  $pid'$  is the actual party ID.

*Breaking accountability properties in exchange for verdicts:* In an (AUC) accountable ideal functionality, the adversary  $\mathcal{A}$  may, at any point in time, instruct the functionality to break/disable accountability properties. In exchange for breaking an accountability property, the ideal functionality requires  $\mathcal{A}$  to provide a verdict that indicates parties who are blamed for the security breach. This verdict must be *fair*, i.e., it may not blame parties who were honestly following the protocol. Verdicts are received and their fairness is checked by so-called (ideal) *judges*. A judge is a new role added to the ideal functionality that models parties who are responsible for determining misbehaving protocol participants. The environment and higher-level protocols, including higher-level judges, can ask for a judge’s verdict.

In a realization of the ideal functionality, the corresponding (real) judges specify the exact judging procedures, such as checking signatures, and the exact input data used as evidence to compute verdicts. For example, in an e-voting protocol that is supposed to provide accountability w.r.t. counting votes (a strong form of so-called *end-to-end verifiability* [33]), a real judge – run by an auditor or even a voter – might take as input all messages from a bulletin board and then blame parties who produce output, e.g., the election result, but with invalid accompanying zero-knowledge proofs.

*Judge Types:* AUC considers an a priori unbounded number of concurrent judges, which can be instantiated by protocol designers to capture different numbers and types of judges executed by different parties modeling various flavors of accountability. Most common are *public* and *local judges*, which capture *public* and *local accountability*, respectively: (i) A single public judge implements public accountability, i.e., verdicts of the (real) public judge are computed solely based on publicly available information. For example, in e-voting to check that the tallying process went correctly, (cf. “accountability w.r.t. counting votes” above) a public judge typically uses data, such as zero-knowledge proofs, from a public bulletin board. Since the data used is public, everyone, including outside observers, can take the role of a public judge. It therefore typically makes sense to model a public judge as an incorruptible entity. (ii) Several local judges, each of them belonging to and typically representing the validation procedure executed by one protocol participant, model local accountability. That is, verdicts of a (real) local judge are computed by a protocol participant, say Alice, and can therefore be based not just on public information and whatever data other parties are willing to provide to Alice, but verdicts can also be based on Alice’s own private data, possibly allowing for detecting a wider variety of misbehavior. For example, in e-voting Alice, resp. her local judge, can tell that her ballot does not appear on the bulletin board, even though she submitted it. Since a local judge is run by a (potentially malicious) protocol participant who might lie about the verdict, a local judge is typically considered corrupted iff the corresponding protocol participant is corrupted. (iii) Additionally, AUC also supports other types of judges, such as, *mandated judges*. A mandated judge models a (potentially trusted) external auditor which (in the realization) computes verdicts based on evidence that protocol participants provide. For instance, in some e-voting protocols aiming for everlasting privacy [36] not all data needed for verifiability of the election can be published. Instead, only a trusted auditor would obtain all necessary data, including potentially non-public data, to perform the verification procedure.

We emphasize that in AUC protocol designers are free to define judges in whatever way suitable. In particular, protocol designers can decide what kind of information judges use/require, whether the information is public or private, what groups of parties they are responsible for, or whether judges are corruptible. Hence, AUC does not dictate specific types of judges, and as mentioned, also does not bound the number of judges considered.

*Judge dependent accountability properties:* We take the view that a judge is responsible for one or more security properties and one or more parties (not necessarily exclusively). That is, if a security properties is broken from the point of view of a party, then a judge responsible for this security property and party is supposed to output a verdict; conversely, without such a verdict the property should still hold true for this party. For example, consider again accountability w.r.t. counting votes in e-voting. Here it makes sense to have a public judge responsible for all parties and the correctness of the tallying process, who would output a verdict when the tallying process was carried out incorrectly. To give another example, in Section 3.2 we consider a PKI based on Certificate Transparency Logs (CTL). Here, a security property is that there should not be certificates on the CTL containing a public key that does not belong to the alleged party, say Alice. This property may be broken for Alice but not other parties, and only Alice can check this property. So here we would have a local judge responsible for Alice carrying out the necessary checks; as long as this judge does not output a verdict, all of Alice’s certificates on the CTL should be correct.

We formalize the above in AUC by allowing the adversary in the ideal protocol to mark an accountability

property  $p$  as broken for a specific set of judges, and require these judges to output verdicts. From then on, the property  $p$  is no longer enforced by the ideal functionality for those parties that are protected by that set of judges.

*Judicial reports:* Judges can provide arbitrary information, such as an aggregated view of collected evidence to higher-level protocols via the novel concept of *judicial reports*. Judges in such higher-level protocols can then in turn also use this information for computing their own verdicts. This is crucial to enable modular design and analysis of a wide range of accountability-based protocols.

*Assumption-based properties:* Similarly to accountability properties, AUC formalizes assumption-based properties in ideal protocols by allowing the adversary to mark the underlying assumptions as broken. If that property is not additionally protected by accountability, then the property itself is also considered broken. In general, the assumptions and corresponding properties might hold true for some parties but not for others. For example, Alice might achieve *liveness* since all of her messages were delivered within a *bounded network delay* but Bob might no longer have liveness as some of his messages were dropped. Hence, just as for judge dependent accountability properties, the ideal functionality allows the adversary to mark assumptions as broken for a specific set of affected parties. Those parties then lose security guarantees, unless they are still protected by accountability.

*Supervisor:* A new meta party called *supervisor* collects information about (i) corruption of internal protocol participants (i. e., those that exist only within subroutines of the real protocol but not in the ideal protocol) and (ii) broken security assumptions. This information is provided by the supervisor to the environment to guarantee that the real and ideal worlds coincide in these aspects. In other words, a simulator cannot cheat but has to keep these aspects consistent.

### 2.3 AUC Transformation for Ideal Functionalities

We now explain how AUC turns an arbitrary ideal functionality into an accountable one. Let  $\mathcal{F}$  be an arbitrary (non-accountable) ideal functionality that enforces a set of preventive security properties  $\mathcal{Sec}$ , e. g., correctness, consistency, or liveness. AUC provides a general transformation  $\mathcal{T}(\cdot)$  for such ideal functionalities  $\mathcal{F}$  that creates an accountable ideal functionality  $\mathcal{F}^{\text{acc}}$  where arbitrary subsets of the properties from  $\mathcal{Sec}$  are changed to instead be assumption-based and/or accountability properties.<sup>2</sup> That is,  $\mathcal{F}^{\text{acc}}$  ensures a combination of absolute properties  $\mathcal{Sec}^{\text{abs}}$ , assumption-based properties  $\mathcal{Sec}^{\text{assumption}}$ , and accountability properties  $\mathcal{Sec}^{\text{acc}}$ . The transformation  $\mathcal{T}(\cdot) := \mathcal{T}_2(\mathcal{T}_1(\cdot))$  consists of two steps  $\mathcal{T}_1$  and  $\mathcal{T}_2$  that progressively modify  $\mathcal{F}$  to obtain  $\mathcal{F}^{\text{acc}} = \mathcal{T}(\mathcal{F})$ . The first step adds the necessary infrastructure to  $\mathcal{F}$  to be able to express accountability properties, such as code for the new judge and supervisor roles, but does not yet alter the actual behavior and security guarantees of  $\mathcal{F}$ . The second step  $\mathcal{T}_2(\cdot)$  then changes the behavior of  $\mathcal{F}$  to model the effects of broken security properties. Next, we define each step.

**Step 1:** The transformation  $\mathcal{T}_1(\mathcal{F}) =: \mathcal{F}'$  takes as input the original ideal functionality to create an intermediate functionality  $\mathcal{F}'$ . This step adds a fixed set of parameters, variables, and static code to  $\mathcal{F}$  that defines the set of accountability properties, judges, a supervisor, verdicts, as well as related operations for the adversary (cf. Figures 1 to 3). The parameters are designed to be instantiated by a protocol designer to customize aspects of  $\mathcal{F}^{\text{acc}}$  that depend on the specific real protocols at hand, e. g., the exact accountability level one wants to analyze.

The full version of  $\mathcal{T}_1(\cdot)$  builds the entire range of AUC features into  $\mathcal{F}'$ . This should be seen as a general blueprint. Features not needed in the application at hand can be omitted.

**Structural changes.** Figure 1 specifies structural components, parameters, and state variables that are added by  $\mathcal{T}_1$  to  $\mathcal{F}$ . More specifically, the roles of judge and supervisor are added to  $\mathcal{F}$ , as well as their corruption behavior: The supervisor is purely a modeling tool, hence incorruptible. As explained in Section 2.2, whether judges are corruptible mainly depends on the types of judges that are modeled and is therefore not a priori fixed by AUC. Protocol designers rather flexibly specify the corruption model for judges by instantiating the newly added subroutine  $\mathcal{F}_{\text{judgeParams}}$  (this new subroutine is also used to customize further aspects such as

<sup>2</sup>It is straightforward to also add entirely new assumption-based and/or accountability properties to  $\mathcal{F}$  while applying the AUC transformation. For simplicity of presentation we will leave this option implicit.

<b>Additional roles:</b> judge, supervisor	
<b>Additional protocol parameters:</b>	<i>{They may be polynomially checkable predicates}</i>
- $\text{Sec}^{\text{acc}} \subset \{0, 1\}^*$	<i>{Accountability properties}</i>
- $\text{Sec}^{\text{assumption}} \subset \{0, 1\}^*$	<i>{Assumption-based security properties}</i>
- $\text{pids}_{\text{judge}} \subset \{0, 1\}^*$	<i>{set of judge entities/(P)IDs in the protocol (which are often directly related to some protocol participants)}</i>
- $\text{ids}_{\text{assumption}} \subset \{0, 1\}^*$	<i>{set of entities/IDs where properties are ensured via assumptions}</i>
<b>Additional subroutines:</b> $\mathcal{F}_{\text{judgeParams}}$	
<b>Additional Corruption behavior:</b>	
- <b>AllowCorruption</b> ( $\text{pid}, \text{sid}, \text{role}$ ):	
Do not allow corruption of ( $\text{pid}, \text{sid}, \text{supervisor}$ ).	
<b>if</b> $\text{role} = \text{judge}$ :	
<b>send</b> ( <b>Corrupt</b> , ( $\text{pid}, \text{sid}, \text{judge}$ ), $\text{internalState}$ )	
<b>to</b> ( $\text{pid}, \text{sid}, \mathcal{F}_{\text{judgeParams}} : \text{judgeParams}$ )	<i>{<math>\mathcal{F}_{\text{judgeParams}}</math> decides whether judges can be corrupted}</i>
<b>wait for</b> $b$	
<b>return</b> $b$	
- <b>DetermineCorrStatus</b> <sup>a</sup> ( $\text{pid}, \text{sid}, \text{role}$ ):	
<b>if</b> $\text{role} = \text{judge}$ :	<i>{<math>\mathcal{F}_{\text{judgeParams}}</math> may determine a judge's corruption status}</i>
<b>send</b> ( <b>CorruptionStatus?</b> , ( $\text{pid}, \text{sid}, \text{judge}$ ), $\text{internalState}$ )	
<b>to</b> ( $\text{pid}, \text{sid}, \mathcal{F}_{\text{judgeParams}} : \text{judgeParams}$ )	
<b>wait for</b> $b$ ; <b>return</b> $b$	
- <b>AllowAdvMessage</b> ( $\text{pid}, \text{sid}, \text{role}, \text{pid}_{\text{receiver}}, \text{sid}_{\text{receiver}}, \text{role}_{\text{receiver}}, m$ )	
Do not allow sending messages to $\mathcal{F}_{\text{judgeParams}}$ .	<i>{<math>\mathcal{A}</math> is not allowed to invoke <math>\mathcal{F}_{\text{judgeParams}}</math> in the name of corrupted parties.}</i>
<b>Additional internal state:</b>	
- $\text{brokenProps} : (\text{Sec}^{\text{assumption}} \cup \text{Sec}^{\text{acc}}) \times (\text{pids}_{\text{judge}} \cup \text{ids}_{\text{assumption}}) \rightarrow \{\text{true}, \text{false}\}$	<i>{Stores broken security properties per judge/id, initially false <math>\forall</math> entries}</i>
- $\text{verdicts} : \text{pids}_{\text{judge}} \rightarrow \{0, 1\}^*$	<i>{Verdicts per <math>p \in \text{pids}_{\text{judge}}</math>, initially <math>\varepsilon</math>}</i>
- $\text{brokenAssumptions} : \text{Sec}^{\text{assumption}} \times \text{ids}_{\text{assumption}} \rightarrow \{\text{true}, \text{false}\}$	<i>{Stores broken security assumptions per id, initially false <math>\forall</math> entries}</i>
- $\text{corruptedIntParties} \subset \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \setminus (\text{Roles}_{\mathcal{F}}^b \cup \{\text{judge}, \text{supervisor}\})$ , initially $\emptyset$	<i>{The set of corrupted internal parties (<math>\text{pid}, \text{sid}, \text{role}</math>)}</i>
<sup>a</sup> <b>DetermineCorrStatus</b> allows protocol designers to specify whether an entity that is currently not directly controlled by the attacker should nevertheless consider itself to be corrupted. E.g., a local judge will typically consider itself to be corrupted already if its corresponding party is corrupted.	
<sup>b</sup> $\text{Roles}_{\mathcal{F}}$ is the set of (main) roles provided by $\mathcal{F}$ to the environment. For example, $\text{Roles}_{\mathcal{F}} = \{\text{signer}, \text{verifier}\}$ for an ideal signature functionality $\mathcal{F} := \mathcal{F}_{\text{sig}}$ .	

Fig. 1: Parameters and state added by the transformation  $\mathcal{T}_1(\mathcal{F})$  to an ideal functionality  $\mathcal{F}$ .

accountability levels; we explain this later). For example, for local judges  $\mathcal{F}_{\text{judgeParams}}$  would typically disallow corruption if the corresponding protocol participant is honest, and conversely consider the local judge to be corrupted as soon as the protocol participant is corrupted. As mentioned, the adversary gains full control over corrupted judges. Further, AUC adds the parameters  $\text{Sec}^{\text{acc}}$  and  $\text{Sec}^{\text{assumption}}$  to  $\mathcal{F}$ , which are instantiated by the protocol designer to contain exactly those accountability and assumption-based properties ( $\subseteq \text{Sec}$ ) she wants to consider.

The parameter  $\text{pids}_{\text{judge}}$  specifies the considered (possibly infinite) set of party IDs of (different types of) judges. For example, the most common types of judges mentioned in Section 2.2 can be model via PIDs  $\text{pid}_j \in \text{pids}_{\text{judge}}$  of the following form: (i)  $\text{pid}_j = \text{public}$  identifies a unique public judge, (ii)  $\text{pid}_j = (\text{local}, \text{pid}, \text{role})$  models a local judge of protocol participant ( $\text{pid}, \text{sid}, \text{role}$ ), and (iii)  $\text{pid}_j = (\text{mandated}, \text{pid})$  highlights a mandated judge with PID  $\text{pid}$ .

In the state variable  $\text{brokenProps}$  we track for each combination of a security property  $\in \text{Sec}^{\text{acc}}$  and judge  $\in \text{pids}_{\text{judge}}$  whether the property is broken for that judge (and hence, not guaranteed anymore for the parties this judge is responsible for). Jumping slightly ahead, the second transformation step  $\mathcal{T}_2(\cdot)$  will change the behavior of the functionality depending on which properties are marked broken for which judges and their governed parties. The map  $\text{verdicts}$  stores the current verdict of each judge.

The parameter  $\text{ids}_{\text{assumption}}$  defines the (possibly infinite) set of IDs for which an assumption and the corresponding assumption based property might become broken. For example, (i)  $\text{id} = \text{public}$  can be used to model a global assumption and property that affects all parties, whereas (ii)  $\text{id} = (\text{local}, \text{pid}, \text{role})$  can model an assumption and corresponding property specific to a protocol participant ( $\text{pid}, \text{sid}, \text{role}$ ). For each combination of property  $\in \text{Sec}^{\text{assumption}}$  and ID  $\in \text{ids}_{\text{assumption}}$ , we track whether underlying assumptions are currently broken in the new variable  $\text{brokenAssumptions}$ . Once assumptions are broken, then the property itself might also become broken for the affected ID (see below), which again is tracked in  $\text{brokenProps}$ .

Finally, we add the set  $\text{corruptedIntParties}$  which tracks corrupted internal parties in a realization in addition to corrupted main parties that are already tracked by the functionality.

**Additional code for the judge role:**

```

recv (BreakAccProp, verdict, toBreak) from NETa to (pid, sid, judge)
  s.t. toBreak  $\subseteq$  Secacc  $\times$  pidsjudge  $\wedge$  verdict maps from pidsjudge  $\rightarrow$  {0, 1}*:
  (successful, leakage)  $\leftarrow$  breakAttempt(verdict, toBreak)           {breakAttempt is defined below}
  reply (BreakAccProp, successful, leakage)

recv GetVerdict from I/O to (pidj, sid, judge):           {The environment can query the verdicts of local and public judges}
  reply (GetVerdict, verdicts[pidj])

recv (GetJudicialReport, msg) from I/O to (pidj, sid, judge): {The environment may query for local or public judicial reports}
  send (GetJudicialReport, msg, internalState) to (pidj, sid,  $\mathcal{F}_{\text{judgeParams}}$  : judgeParams) {Forward judicial report request to  $\mathcal{F}_{\text{judgeParams}}$ }
  wait for (GetJudicialReport, report)
  reply (GetJudicialReport, report)

```

**Helperfunctions:**

```

procedure breakAttempt(verdict, toBreak) :           {Process break attempt}
  Check for all non- $\varepsilon$  verdicts in verdict, i.e.,  $\forall \text{pid}_j$  s.t. verdict[pidj]  $\neq \varepsilon$ :
  1. it holds true that verdict[pidj] is a positive boolean expression built from propositions of the form dis((pid, sid, role)),
  2. it holds true that eval(verdict[pidj]) = true,b
  Check that  $\forall (prop, \text{pid}_j) \in \text{toBreak}$ :
  3. verdict[pidj]  $\neq \varepsilon$ ,
  4. brokenAssumptions[prop, pidj] = true, if prop  $\in$  Secassumption  $\wedge$  pidj  $\in$  Secassumption.
  if any of the above check fails:
  return (false,  $\varepsilon$ )
  send (BreakAccProp, verdict, toBreak, internalState) to ( $\_$ ,  $\_$ ,  $\mathcal{F}_{\text{judgeParams}}$  : judgeParams) { $\mathcal{F}_{\text{judgeParams}}$  can impose further conditions, e.g., on verdicts or whether it is allowed to violate breakable security properties.}
  wait for (BreakAccProp, successful, leakage)
  if successful:
    for all  $\forall \text{pid}_j$  s.t. verdict[pidj]  $\neq \varepsilon$  do:
      verdicts[pidj]  $\leftarrow$  verdict[pidj]           {Record accepted local and public verdict}
    for all (prop, pidj)  $\in$  toBreak do:
      brokenProps[prop, pidj]  $\leftarrow$  true
  return (successful, leakage)

```

<sup>a</sup>NET denotes message from the network adversary. I/O denotes messages from the environment.  
<sup>b</sup>eval evaluates the boolean expression, where dis(pid, sid, role) evaluates to true if (pid, sid, role)  $\in$  CorruptionSet or (pid, sid, role)  $\in$  corruptedIntParties. CorruptionSet is a predefined variable of iUC that contains all corrupted main parties of this functionality. We set eval( $\varepsilon$ ) := true.

Fig. 2: Judge code added by the transformation  $\mathcal{T}_1(\mathcal{F})$  to an ideal functionality  $\mathcal{F}$ .

**Additional code for the supervisor role:**

```

recv (BreakAssumption, toBreak) from NET to ( $\_$ ,  $\_$ , supervisor) s.t. toBreak  $\subseteq$  Secassumption  $\times$  idsassumption: { $\mathcal{A}$  may break these assumptions}
for all (prop, id)  $\in$  toBreak do:
  brokenAssumptions[prop, id]  $\leftarrow$  true
  if prop  $\notin$  Secacc  $\vee$  id  $\notin$  pidsjudge:
    brokenProps[prop, id]  $\leftarrow$  true
  {Record broken assumptions}

send (BreakAssumption, toBreak, internalState)
to (pid, sid,  $\mathcal{F}_{\text{judgeParams}}$  : judgeParams)           { $\mathcal{F}_{\text{judgeParams}}$  provides leakage}

wait for (BreakAssumption, leakage)
reply (BreakAssumption, leakage)

recv (corruptInt, (pid, sid, role)) from NET to ( $\_$ ,  $\_$ , supervisor)
  s.t. role  $\notin$  Roles $\mathcal{F}$   $\cup$  {judge, supervisor}:
  corruptedIntParties.add((pid, sid, role))
  reply (corruptInt, ack)
  { $\mathcal{A}$  is allowed to corrupt internal protocol parties}

recv (IsAssumptionBroken?, prop, id) from I/O
to ( $\_$ ,  $\_$ , supervisor) s.t. id  $\in$  idsassumption:
if prop  $\in$  Secassumption:
  reply (IsAssumptionBroken?, brokenAssumptions[prop, id])
else:
  reply (IsAssumptionBroken?,  $\perp$ )
  {The environment may ask whether properties are broken}

recv (corruptInt?, (pid, sid, role)) from I/O to ( $\_$ ,  $\_$ , supervisor)
  s.t. role  $\notin$  Roles $\mathcal{F}$   $\cup$  {judge, supervisor}:
  {The environment may ask for the corruption status of internal parties}
  if (pid, sid, role)  $\in$  corruptedIntParties:
    reply (corruptInt, true)
  else:
    reply (corruptInt, false)

```

Fig. 3: Supervisor code added by the transformation  $\mathcal{T}_1(\mathcal{F})$  to an ideal functionality  $\mathcal{F}$ .



**Judges.** Figure 2 specifies the code AUC adds to ideal functionalities to model judges. Each judge in AUC is an entity of the form  $(pid_j, sid, judge)$  where  $pid_j$  is the judge’s PID running the fixed judge role.

The adversary can try to break accountability properties  $\in \text{Sec}^{\text{acc}}$  by sending a `BreakAccProp` message to the ideal functionality, which contains a list of properties including the judges  $\subset \text{pids}_{\text{judge}}$  for which these properties shall be broken as well as a list of new verdicts for (some of the) judges. After receiving the message, the ideal functionality first checks whether this attempt meets all minimal requirements for accountability, i.e., (i) all (non-empty) verdicts are positive boolean formulas, (ii) all verdicts are fair based on the current corruption status of main and internal parties, (iii) if a property is marked as broken for some judge, then that judge also outputs a verdict, and (iv) accountability properties that are also assumption-based properties may not be broken as long as the underlying assumptions still hold true. If the attempt passes these checks,  $\mathcal{F}'$  forwards the attempt (including its full internal state) to the subroutine  $\mathcal{F}_{\text{judgeParams}}$ , which decides whether the attempt actually succeeds and whether/which information is leaked to the attacker, say, because a privacy property was broken. By instantiating  $\mathcal{F}_{\text{judgeParams}}$  a protocol designer can therefore customize the exact level of accountability and also relationships between properties. For example, an instance of  $\mathcal{F}_{\text{judgeParams}}$  might require verdicts to have the form “dis( $A$ )”, i.e., identify exactly one misbehaving party, thereby providing individual accountability (cf. Section 3 for examples with different accountability levels). It might also require that, if a property connected to a party is broken, then the same property must also be broken concurrently for others, capturing the relationship that several/all parties are affected and thus able to compute verdicts simultaneously (cf. Appendix G).

The environment can query judges of  $\mathcal{F}'$  to obtain their current verdicts and their judicial reports. For verdicts,  $\mathcal{F}'$  returns the last accepted verdict. For reports,  $\mathcal{F}'$  calls  $\mathcal{F}_{\text{judgeParams}}$  which can compute the report based on the entire internal state of  $\mathcal{F}'$ . This allows a protocol designer to customize which information is contained in such reports by instantiating  $\mathcal{F}_{\text{judgeParams}}$  appropriately (see Section 3.1 for an example).

**Supervisor.** Figure 3 presents the code added for the supervisor. The adversary can send a `BreakAssumption` message to mark the assumptions as broken that underlie some properties  $p \in \text{Sec}^{\text{assumption}}$  for some set of IDs  $\subset \text{ids}_{\text{assumption}}$ . Assumption-based properties also protected by accountability, i.e.,  $p \in \text{Sec}^{\text{acc}}$  and  $id \in \text{pids}_{\text{judge}} \cap \text{ids}_{\text{assumption}}$ , are not yet marked as broken; for these, the adversary still has to issue a `BreakAccProp` message to  $id$  with a valid verdict. Otherwise, breaking the underlying assumption also breaks the property. The adversary may further mark arbitrary internal parties as corrupted. These parties can then also be blamed in verdicts.

As mentioned, the environment can ask the supervisor whether assumptions are marked as broken and whether internal parties are marked as corrupted. Note that  $\mathcal{F}'$  does not impose any limitations on when assumptions can be marked as broken but only ensures – by providing this information to the environment – that this occurs if and only if assumptions are broken in the realization. By this, the realization can specify the exact conditions and limitations for broken assumptions without requiring any modifications to  $\mathcal{F}'$  each time a different realization is considered (cf. Sections 2.4 and 2.7).

**Step 2:** The second step of the transformation  $\mathcal{T}_2$  specifies the effects of a broken property. Observe that the exact implications in terms of behavior of  $\mathcal{F}$  strongly depend on the individual security properties. Therefore,  $\mathcal{T}_2$ , unlike  $\mathcal{T}_1$ , cannot simply be a fixed set of variables, code that need to be added to a functionality, or even a black-box transformation of  $\mathcal{F}'$ . Instead,  $\mathcal{T}_2$  rather constitutes a more abstract guideline on how the functionality  $\mathcal{F}'$  has to be modified. Importantly, such modifications must not alter the behavior when security guarantees hold true since we want to retain the same security guarantees of  $\mathcal{F}$  in those cases:

**Definition 1.** Let  $\mathcal{F}' := \mathcal{T}_1(\mathcal{F})$ . Let  $\mathcal{F}^{\text{acc}} := \mathcal{T}_2(\mathcal{F}')$  be a functionality obtained by introducing additional behavior for capturing broken security properties. We say that  $\mathcal{F}^{\text{acc}}$  is an accountable transformation of  $\mathcal{F}$  if the behavior of  $\mathcal{F}^{\text{acc}}$  and  $\mathcal{F}'$  is identical in all runs until a security property is marked as broken.

Technically, modeling the effects of a broken security property  $p \in \text{Sec}^{\text{acc}} \cup \text{Sec}^{\text{assumption}}$  affecting some ID  $id \in \text{pids}_{\text{judge}} \cup \text{ids}_{\text{assumption}}$  (specifying, e.g., an affected party), generally entails introducing (one or more) conditional clauses of the form “**if** ( $p, id$ ) **is not marked as broken** **then** *<original behavior>* **else** *<new behavior>*”. As the name suggests, *<original behavior>* denotes the original unchanged behavior of the functionality  $\mathcal{F}$ , i.e., the code that enforces  $p$ . The code *<new behavior>* then defines what “breaking  $p$ ” actually means, typically by giving more power to the adversary.

For example, if  $\mathcal{F} := \mathcal{F}_{\text{sig}}$  is an ideal signature functionality and  $p = \text{unforgeability}$  (for, say,  $id = \text{public}$ , modeling public unforgeability), then *<original behavior>* is a check during signature validation

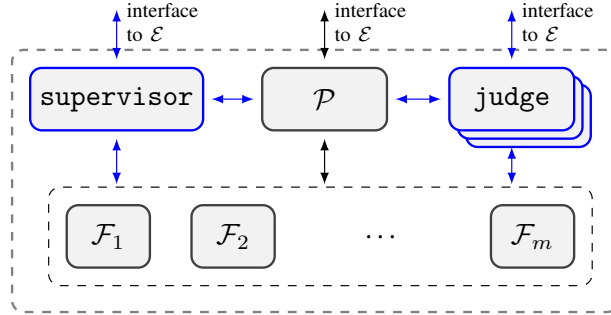


Fig. 4: Example of machines and connections in an accountable real protocol  $\mathcal{P}$  with ideal (possibly accountable) subroutines  $\mathcal{F}_1, \dots, \mathcal{F}_m$ . Blue components are added by AUC.

that, if it detects forgery of a signature, returns `false` irrespective of the actual result of signature validation. Breaking this property would simply disable this forgery check, i.e.,  $\langle \text{new behavior} \rangle$  is empty. Thus, if a signature is forged if unforgeability is broken, the transformed version of  $\mathcal{F}_{\text{sig}}$  might actually return `true`.

More complex conditional statements can be used as well to capture advanced relationships of properties. For example, consider privacy in MPC protocols. The statement “*if there are at least  $t$  PIDs  $pid_i$  of local judges such that  $(\text{privacy}, pid_i)$  is marked as broken then  $\langle \text{leak secret information} \rangle$* ” captures a threshold relationship of local properties: in order to break privacy, the adversary has to mark privacy as broken for at least  $t$  parties.

Observe that all transformations  $\mathcal{T}_2$  following the form described above always satisfy Definition 1. That is, the resulting ideal functionalities  $\mathcal{F}^{\text{acc}} := \mathcal{T}_2(\mathcal{T}_1(\mathcal{F}))$  are indeed accountable transformations of  $\mathcal{F}$  as per Definition 1.

## 2.4 AUC in Real Protocols

As illustrated in Figure 4, modeling accountability in a real protocol  $\mathcal{P}$  using AUC mainly entails adding and specifying a supervisor and several judges that correspond to the ones in the ideal protocol. Again, not always all of these components are needed and can be omitted if not used.

**Judges.** The concrete definitions of judges in real protocols formalize (i) the judging algorithms used for computing verdicts, (ii) inputs and hence evidence needed for obtaining the verdict, (iii) which parties are supposed to provide which information as evidence, and (iv) to whom evidence is provided, namely, the party that is running the judge. We exemplify modeling of real judges for the most common types of public, local, and mandated judges: Intuitively, a *public judge* computes verdicts solely based on publicly available/verifiable data such as information from a public bulletin board or data that individual parties are willing to provide and therefore publish even to external observers. Since an honest party can always locally execute the public judge, the public judge is modeled to be incorruptible. To make formally explicit in the security analysis that also an attacker might run the public judge, public judges should generally publish all collected information via their network interface to the environment. For the same reason, if a public judge is allowed to interact with other parts of the protocol, e.g., to verify a signature in an ideal signature functionality  $\mathcal{F}_{\text{sig}}$ , then typically the judge should provide a network interface for the environment to be able to perform the same interactions.<sup>3</sup> We provide a concrete example of a public judge in Section 3.1. *Local judges*, typically one per (main) protocol party, use public information and also (private) data that protocol participants are willing to share with the party running the local judge but might not want to fully publish. In addition, since a local judge is run by a party herself, it also takes as input the entire internal state and possibly history of that corresponding party, including any private information, to compute a verdict. A local judge should be considered to be corrupted iff its corresponding protocol participant is corrupted. We provide concrete examples of local judges in Sections 3.2 and 3.3. *Mandated judges* reflect the judging procedure of an external mandated auditor, which takes as input publicly available data and also (private) data that protocol participants are willing to share with the auditor. Whether (some of the) mandated judges are corruptible depends on the setting that shall be modeled.

<sup>3</sup>This modeling of public judges is similar in spirit to modeling random oracles. A random oracle represents a public function, which is captured by being incorruptible but providing outputs via an additional network interface to the environment.

When evaluating evidence, real judges often obtain or compute additional information that might be useful for and can even be required by higher-level protocols. Such information can be shared via judicial reports (we discuss the need for this novel concept, including example use cases, in Section 2.7).

**Supervisor.** A supervisor in a real protocol forwards the corruption status of internal protocol participants and specifies when exactly the assumptions underlying an assumption-based property for some party/ID must be considered broken.

This generally involves gathering data from other parts of the protocol. For example, consider a property  $p \in \mathcal{Sec}^{\text{assumption}}$  that relies on an *honest majority* assumption and thus affects everyone (i.e.,  $id = \text{public}$ ), e.g., consistency in Bitcoin [48]. That is, a protocol only ensures  $p$  if more than half of the protocol participants are honest. To determine whether the protocol still ensures  $p$ , the supervisor would check whether a majority of the protocol participants is still honest. If the majority is lost, the supervisor would indicate that the assumptions for  $p$  are no longer met and thus the protocol might no longer guarantee  $p$ .

As we discuss in Section 2.7, our novel concept of a supervisor is necessary to be able to capture a wide range of assumption-based properties as well as real protocols with arbitrary internal structures, such as client-server protocols (e.g., [14, 55]). Our case studies in Sections 3.1 and 3.2 provide concrete examples and indeed require a supervisor.

## 2.5 Composable Security Analysis in AUC

A security analysis in AUC consists of a realization proof where one shows that the protocol  $\mathcal{P}^{\text{acc}}$  at hand indeed realizes (in the UC sense) the accountable ideal functionality  $\mathcal{F}^{\text{acc}}$ . Among others, this formally proves that  $\mathcal{P}^{\text{acc}}$  enjoys the desired security properties, including accountability properties. Since AUC works within existing UC models and uses the standard realization notion, one can now, as usual, build an (accountable or traditional) higher-level protocol  $\mathcal{Q}^{\text{acc}}$  using  $\mathcal{F}^{\text{acc}}$  as a subroutine (written  $(\mathcal{Q}^{\text{acc}} \mid \mathcal{F}^{\text{acc}})$ ) and prove that it is secure, i.e., realizes some  $\mathcal{F}'^{\text{acc}}$ . The composition theorem of the underlying UC model then immediately implies that also the composed protocol  $(\mathcal{Q}^{\text{acc}} \mid \mathcal{P}^{\text{acc}})$  using  $\mathcal{P}^{\text{acc}}$  instead of  $\mathcal{F}^{\text{acc}}$  still realizes  $\mathcal{F}'^{\text{acc}}$ , i.e., achieves all desired strict, assumption-based, and/or accountability properties.

In an accountable higher-level protocol  $\mathcal{Q}^{\text{acc}}$ , the judges can and typically will obtain verdicts and judicial reports from the subroutine judges in  $\mathcal{F}^{\text{acc}}$  resp.  $\mathcal{P}^{\text{acc}}$ . This information can then be used by higher-level judges to compute their own verdicts and reports (cf. the BFT example in Section 2.7 and our case studies in Sections 3.1 and 3.3 for higher-level judges that rely on lower level judges). We note that we do not impose any restrictions on which lower-level judges a higher-level judge may access. For example, often a higher-level local judge of some party will only use local subroutine judges belonging to the same party. This models that the party executes all of its judges from all protocol layers iteratively and can re-use the results of previous computations, such as verdicts identifying misbehavior in subroutines, in the following computations. However, a higher-level local judge can also, e.g., obtain the verdicts of lower level judges of different parties, modeling that one party uses the (claimed) results of other parties. While these results might not necessarily be trustworthy, a higher-level judge might, e.g., be able to aggregate and perform majority voting to obtain a fair result, or achieve a weaker level of accountability that considers the possibility of the subroutine judge lying in their verdict.

## 2.6 Deterrence Analysis

In addition to the UC security analysis, protocol designers should perform a cost-benefit/deterrence analysis of the (possibly composed) real protocol to determine whether honest parties are willing to take part in the protocol and whether accountability indeed deters rational adversaries from misbehavior. This analysis can be performed using any standard approach from the literature. As a simple example, based on the approach by Asharov et al. [6] we describe one possible concrete mechanism for this purpose. For a rational acting honest party  $p_i$ , one considers the utility/profit  $U_{hP}^i$  for running the protocol, the cost  $U_{hE}^i$  for disclosing private data as evidence to judges, and the loss due to (falsely) accusation  $U_{hL}^i$  (by a corrupted judge). For a rational but maliciously acting party  $p_i$ ,  $U_{mP}^i$  is the potential utility/profit from misbehaving,  $U_{mL}^i$  is the cost for misbehaving (e.g., reputation loss or contractual penalties after being detected), and  $U_{mE}^i$  is the cost for

providing (possibly maliciously crafted) evidence to judges. Now, accountability provides a suitable security mechanism in a practical deployment if

$$U_{hP}^i - U_{hE}^i - U_{hL}^i \geq 0 \text{ and} \quad (1)$$

$$U_{hP}^i - U_{hE}^i - U_{hL}^i \geq U_{mP}^i - U_{mE}^i - U_{mL}^i, \quad (2)$$

i.e., honest parties benefit from and hence are willing to take part in the protocol (Eq. (1)) and malicious parties are deterred from misbehaving as they stand to gain more when honestly following the protocol (Eq. (2)).

Often, the result of such an analysis will be obvious. For example, if CAs in a PKI are detected misbehaving, they typically have to close business [92], i.e.,  $U_{mL}^i$  will be much higher than  $U_{mP}^i$ . We also note that utilities/costs might depend on the context/higher-level protocol that a subroutine will be deployed in. For example, the potential profit  $U_{mP}^i$  for tempering with an accountable bulletin board is very large if it is used as subroutine in an e-voting protocol for a major political election. But  $U_{mP}^i$  might be negligible if the bulletin board is just one out of several redundant backups in a distributed cloud storage protocol. In such cases a deterrence analysis cannot be performed for individual components but should rather be performed after fully composing the entire real protocol.

## 2.7 Discussion

AUC is the first general purpose accountability framework for UC models. It builds on and extends existing concepts but also introduces entirely new concepts that are required to construct such a general framework. Here we discuss these concepts and relate them to existing literature.

**Accountability properties:** AUC formalizes a property-based interpretation of accountability, i.e., security properties might be broken as long as misbehaving parties can be identified and blamed. This is a standard interpretation that is widely established and used in the area of formal protocol security analyses [33, 42–45, 52, 57, 62, 63, 67, 68, 79]. There are other (often informal) interpretations of accountability, also outside of the field security (cf. [46]). A relatively close one in the domain of security requires that any object/message/action can be connected to its originator, e. g., by requiring all parties to sign (cf., e. g., [25, 27]). This interpretation is orthogonal to the property-based interpretation. For example, in e-voting it is important that one cannot trace back ballots to individual voters but one can still achieve property-based accountability for such protocols, namely correctness/verifiability of the election result (see, e. g., [69]). Conversely, even if the election servers sign all their messages, it might not be possible to verify from those messages whether all votes were counted correctly. That is, the protocol might not provide accountability for correctness of the election result.

**Verdicts:** The game-based model of Küsters et al. [67] defines verdicts as boolean formulas to be able to express different levels of accountability. We transfer this concept to the setting of universal composability. This allows AUC to capture a wide range of common accountability levels, from strong individual accountability to very weak levels where, say, a verdict only says that someone misbehaved but gives no further information on who exactly is at fault.

**Judges, Relationships, and Judge Dependent Properties:** Judges are an integral part of the specification of protocols since they define the routines, including inputs, that are used by that protocol to detect misbehavior. A formal analysis can then verify for a given detection routine whether it provides (property-based) accountability. Hence, judges are used (at least implicitly) by all works that formally analyze (property-based) accountability (e. g., [6, 13, 34, 37, 54, 62, 63, 67, 79, 87]).

There is a division in the existing literature concerning relationships of security properties: On the one hand, prior general frameworks for accountability such as [62, 63, 67, 79] consider only a single (public/local) judge running at the same time. Hence, these general frameworks cannot actually capture relationships of properties where different parties and/or different types of judges are affected at the same time. This includes, e.g., all local judges obtaining the same verdict or a security property holding true until at least a threshold of  $t$  judges (possibly a mixture of local, mandated, and/or a public one) detect misbehavior. On the other hand, there are works that have already considered and analyzed such relationships by proposing specialized security models that hard code a certain relationship for a specific property and setting. This includes, for example, [6, 8, 14, 55, 85] which require for *identifiable abort* in MPC protocols that honest parties agree on the culprit, i.e., their local judges compute *the same* verdict.

AUC is the first general accountability framework that can capture such relations. At its core, this is enabled by considering *multiple concurrent judges and judge types* as well as *judge dependent properties*. Appendix G illustrates this feature.

**Judicial reports:** This novel concept introduced by AUC allows judges from different protocol layers to exchange information, which is often required to perform a modular analysis.

For example, consider a publicly accountable BFT consensus algorithm used as a subroutine by a higher-level protocol  $\mathcal{P}^{\text{acc}}$ . As long as the (lower-level) public judge  $J_{\text{BFT}}$  in the BFT algorithm does not detect misbehavior, she will typically be able to compute the consensus established among the parties from the available evidence. In  $\mathcal{P}^{\text{acc}}$ , misbehaving parties might be able to break security properties by deviating from the consensus established in the BFT subroutine. Hence, for a higher-level public judge  $J_{\mathcal{P}^{\text{acc}}}$  to detect who deviated from the consensus he must first learn the correct consensus. But  $J_{\mathcal{P}^{\text{acc}}}$  cannot compute this on his own since, by UC composition, he does not see the internals of the BFT subroutine but only gets restricted access via a limited interface. AUC solves this issue and enables a modular analysis by extending this interface via judicial reports:  $J_{\text{BFT}}$ , who has full access within the BFT subroutine and who can thus compute the unique consensus if it exists, can provide this information as part of a report to  $J_{\mathcal{P}^{\text{acc}}}$ . See also our case study in Section 3.1 for a composed protocol whose modular analysis requires judicial reports.

**Supervisor:** AUC is the first framework that allows for modeling and analyzing the combination of assumption-based and accountability-based security for the same properties in a UC model. It is also the first framework that supports modeling and analyzing accountability of real protocols with arbitrary internal parties without assuming any specific internal structure. Both of these features are enabled by, among others, the novel concept of a supervisor.

To formalize the guarantees of assumption-based properties the behavior of the ideal functionality has to change depending on whether the assumption currently holds true. Note that for many assumptions an ideal functionality cannot actually determine whether they are still met as they often depend on internals that only exist in the realization. For example, liveness properties of blockchains typically assume that an internal network subroutine in the real protocol has a bounded message delivery delay; such a network subroutine does not necessarily exist in the ideal functionality. The supervisor solves this issue: firstly, the realization of the supervisor can specify the exact conditions under which an assumption holds true while having full access to all information of the real protocol. By allowing the environment to query whether assumptions are currently broken, the supervisor then further ensures that the assumption will be marked as broken in the ideal functionality if and only if the conditions specified in the real world are no longer met. Hence, e. g., an ideal blockchain functionality would enforce liveness iff the real blockchain supervisor determines that the real network still has a bounded message delay. Our case study in Section 3.1 underlines the need for this aspect of supervisors.

Verdicts in (possibly composed) real protocols might have to blame internal protocol parties, e. g., servers in client-server protocols [50, 85], where clients are main parties that are available to higher-level protocols while servers are purely internal subroutines. Since only main parties but not any of the internal parties of the real protocol also exist in the corresponding ideal protocol, we use the supervisor to ensure that the simulator can mark internal parties as corrupted in the ideal functionality iff they are corrupted in the real protocol (the same property is already guaranteed for the main parties by the underlying model). This mechanism is necessary to ensure that verdicts which are fair in the ideal functionality are also fair in the realization. Our case studies in Sections 3.1 to 3.3 rely on this use of the supervisor.

**Composition:** The combination of the above concepts is what allows AUC, for the first time, to capture UC compositions of arbitrary accountability-based protocols. This is further illustrated by our case studies in Sections 3.1 and 3.3, which are the first modular accountability analyzes for these settings and protocols.

Since AUC works within an arbitrary UC model, AUC inherits and preserves all features of the underlying composition theorems. In particular, if the underlying model supports composition with global state (e. g., [9, 17, 18, 71]), then it is possible to make some or all of the subroutine judges within a composed protocol globally available to the environment and arbitrary other concurrent protocols.<sup>4</sup> Whether it is sensible to consider globally available or rather private subroutine judges depends, as with most other global functionalities, on the protocol and the context that is modeled. For example, local judges within composed MPC protocols (cf. Appendix G)

<sup>4</sup>We note that, as observed by [9, 17, 71], the same UC security proof of the subroutine already implies composition with and without globally available subroutine judges.

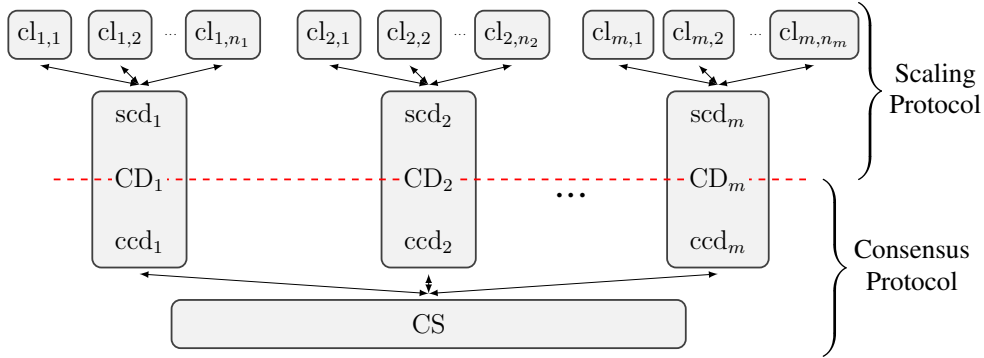


Fig. 5: Consensus scaling: a central service CS establishes consensus. Consensus is distributed via distributors  $CD_1, \dots, CD_m$ . The CD connect to CS via their CS-clients  $ccd_1, \dots, ccd_m$ . The CDs distribute consensus via server components  $scd_1, \dots, scd_m$  to their clients  $cl_{1,1}, cl_{1,2}, \dots$

should typically be modeled as private: the subroutine is supposed to be used only within a single context, namely the MPC protocol, and there is no reason for Alice to share the verdicts and judicial reports of her internal subroutines with arbitrary other parties, protocols, and the environment. In contrast, one might consider modeling judges belonging to a PKI subroutine (cf. Section 3.2) to be globally available within a composed protocol if the same PKI is supposed to be shared by multiple different protocols.

We note that the main difference between global and private subroutine judges is whether only one or multiple higher-level protocols can be composed with the same subroutine. Hence, when only a single specific higher-level protocol has to be considered, then an analysis with a private subroutine is generally already sufficient.

### 3. CASE STUDIES

In this section, we exemplify the usage and features of AUC via three case studies: an accountable scaling protocol, a simplified version of the Web PKI including CTLs, and a key exchange protocol based on an accountable PKI. We provide for each case study a brief deterrence analysis in the appendix. Additionally – as a sanity check – we cast existing accountability definitions from the UC MPC literature into AUC in Appendix G.

#### 3.1 Scaling Accountable Consensus

In this case study, we analyze a common situation from practice, namely, a scaling protocol built on top of a consensus service (CS), cf. Figure 5. The purpose of this scaling protocol is to increase throughput and hence support a larger number of clients. This is achieved by introducing an additional layer of intermediate servers (called CD in Figure 5) that regularly obtain the established ordered sequence of messages/the consensus from the underlying CS, cache the result, and then use this cache to answer incoming requests from clients. This scaling approach is commonly used, e.g., by the prominent Hyperledger Fabric blockchain [3, 50], the Hashgraph consensus service [11], and content delivery networks [39].

More specifically, in this case study our goal is to scale a *consensus service* that provides *public and individual accountability*, takes inputs from clients, appends them to a globally unique ordered list/state, and gives all clients access to this global state. In particular, if a client does not obtain a prefix of the (same) global state, then a judge, based on public evidence, can blame a misbehaving party. We want to show that the composition of a suitable scaling protocol on top of this service retains all properties of the underlying consensus service, most notably public individual accountability. Note that consensus might fail due to misbehavior of parties in CS but could also be introduced by the scaling layer.

This case study exemplifies several features and aspects of the AUC framework, including (i) accountability and assumption-based security properties, (ii) public accountability, (iii) individual accountability, (iv) composition using judicial reports, which are required to express this case study, and (v) complex protocol structures, namely a client-server protocol with *internal*, potentially malicious parties/servers. Full details, including formal specifications of all UC machines, functionalities, and a full security proof are available in Appendix D.

**The scaling protocol.** We consider a scaling protocol that has the structure depicted in Figure 5 and works as described in what follows. While this scaling protocol is a custom one that we chose to illustrate AUC, the general construction is similar to Hyperledger Fabric and Hashgraph. Hence, this case study can serve as a basis for the first UC accountability analysis of these protocols in future work.

*Transaction submission:* When a client, say  $cl_{1,1}$ , submits a transaction (after receiving that transaction from some higher-level protocol), she adds her identity, signs the resulting transaction, and then sends (via an unprotected network) the signed transaction to her *consensus distributor* (CD), here  $CD_1$ . The CD is divided into two components: a server part called SCD which interacts with higher-level clients and a client part CCD which interacts with the CS. An SCD verifies the signature of incoming transactions from higher-level clients and then starts acting as a client CCD towards the CS to add the signed transaction to the globally ordered state. This involves running the code of a client, e.g.  $ccd_1$ , as specified by the underlying consensus protocol. Depending on the consensus protocol, this client code might, e.g., also add its own signature to the signed transaction.

*Accessing the global state:* A SCD, say  $scd_1$ , regularly calls the client code  $ccd_1$  to obtain a current copy of the global state from the CS.  $scd_1$  then signs (with the key of  $CD_1$ ) and caches this global state. Whenever a client, such as  $cl_{1,1}$ , queries its SCD for the global state, the SCD, here  $scd_1$ , responds by returning the most recent signed cached copy of the global state, hence, reducing the load of the CS.  $cl_{1,1}$  accepts and outputs the message to a higher-level protocol if the signature by  $CD_1$  on the whole state is valid.

**Security properties.** We consider two security properties for both the underlying and the scaled consensus protocols, formalized by an ideal functionality (see below):

*Public individual accountability w.r.t. consistency:* Clients (both of the scaling protocol and the CS) obtain a prefix of the same global state or can identify an individual misbehaving party. This definition follows the game-based accountability property formalized for Hyperledger Fabric in [50] but using AUC takes it to the composable UC setting. As mentioned, we want to show this property for the composition of the scaling protocol and the consensus protocol.

*(Assumption-based) liveness:* Liveness states that transactions submitted by honest clients (of the scaling protocol and CS) will be part of the global state after at most  $\delta$  time units assuming a network with bounded message delay (cf., e.g., [48, 50, 78, 80, 89]). To illustrate assumption-based properties, our analysis extends to the case that, at some point in the run, the underlying assumption of a network with bounded message delay might no longer hold true.

**The ideal accountable consensus functionality  $\mathcal{F}_{cp}^{acc}$ .** To define  $\mathcal{F}_{cp}^{acc}$ , we start by considering a (non-accountable) ideal consensus service functionality  $\mathcal{F}_{cp}$  that is essentially a simplified version of established ideal ledger functionalities, e.g., [10, 51], tailored towards the special case of consensus establishment.  $\mathcal{F}_{cp}$  enforces consistency and liveness as preventive security properties. We then apply AUC to obtain an accountable version  $\mathcal{F}_{cp}^{acc}$  (cf., Figure 9 and 10 in Appendix D) that captures the above security properties. More specifically, we set  $\text{Sec}^{acc} = \{\text{consistency}\}$  and  $\text{Sec}^{assumption} = \{\text{liveness}\}$ . Note that  $\text{liveness} \notin \text{Sec}^{acc}$ , and hence,  $\mathcal{F}_{cp}^{acc}$  will not require judges to blame anybody (e.g., the network) if liveness fails.

We start by explaining  $\mathcal{F}_{cp}$  and then the AUC transformation to derive  $\mathcal{F}_{cp}^{acc}$ .  $\mathcal{F}_{cp}$  itself consists of an unbounded number of clients who offer a read and write interface to higher-level protocols. These clients can write transactions to and read from a single globally ordered list/state in  $\mathcal{F}_{cp}$ . Upon writing a new transaction,  $\mathcal{F}_{cp}$  models network traffic by allowing the simulator to determine when and in which order these transactions are appended to the global state.  $\mathcal{F}_{cp}$  guarantees that all incoming transactions by honest clients will be appended to the global state after at most  $\delta$  time units. More formally,  $\mathcal{F}_{cp}$  models (absolute) preventive liveness by disallowing the attacker from advancing time as long as there are pending transactions that have not been added to the state since  $\delta$  time units. Whenever  $\mathcal{F}_{cp}$  receives a read request from an honest client from a higher-level protocol,  $\mathcal{F}_{cp}$  allows  $\mathcal{A}$  to determine the prefix of the global state that will then be returned to the client. For read requests received by corrupted/malicious clients, i.e., clients that are under full control of the adversary,  $\mathcal{F}_{cp}$  always allows  $\mathcal{A}$  to freely determine the output of  $\mathcal{F}_{cp}$ .

To derive  $\mathcal{F}_{cp}^{acc}$  from  $\mathcal{F}_{cp}$ , we implement the AUC transformation step  $\mathcal{T}_2$  as follows. (i)  $\mathcal{F}_{cp}^{acc}$  includes one incorruptible public judge (in a protocol session), i.e.,  $\text{pids}_{\text{judge}} = \{\text{public}\}$ , and considers only assumptions that affect all parties (in a session), i.e.,  $\text{ids}_{\text{assumption}} = \{\text{public}\}$ . (ii) As soon as liveness is marked broken (for  $id = \text{public}$ ),  $\mathcal{F}_{cp}^{acc}$  no longer enforces that messages are added to the global state within  $\delta$  time units. (iii) If (public) consistency is broken,  $\mathcal{F}_{cp}^{acc}$  allows  $\mathcal{A}$  to freely determine the output of  $\mathcal{F}_{cp}^{acc}$  in turn for a fair

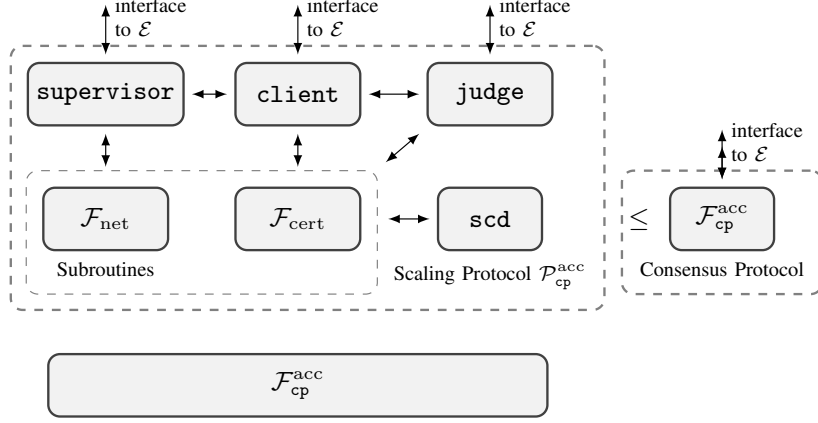


Fig. 6: Illustration of Theorem 1. All machines are also connected to  $\mathcal{A}$ .

verdict matching the customization in  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}}$ . The subroutine  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}}$  forces  $\mathcal{A}$  to provide a verdict which implies individual accountability, i.e., all parties in the verdict can rightfully be blamed for misbehavior. As long as there is no verdict reached yet in  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ ,  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}}$ 's public judicial report returns a view on the global state which contains at least the longest prefix that was read by an honest client so far (everything else after this prefix can be chosen freely by  $\mathcal{A}$ ). If there has been a verdict, the report is empty.

By this construction,  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  indeed models the desired security properties (*individual public*) *accountability w.r.t consistency and assumption-based liveness*.

**Protocol Model.** We model the scaling protocol from Figure 5 via a hybrid protocol as depicted in Figure 6. Specifically, the `client` machine models the code run by the clients  $cl_{i,j}$  and the internal machine `scd` models code of the SCD. The ideal subroutine  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  models an ideal accountable consensus protocol used by the scaling protocol, i.e., it abstracts the code of the CCD and the code of the consensus service CS, all of which are specified in a realization of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  (we discuss a possible realization at the end of this section). In a run of the protocol, there can be an unbounded number of instances of `client` and `scd`, each modeling one party in one protocol session. These parties additionally have access to an ideal signature functionality  $\mathcal{F}_{\text{cert}}$  (which includes a PKI) and an ideal network functionality  $\mathcal{F}_{\text{net}}$  with bounded message delay  $\delta'$ . The adversary is allowed to break the assumption of a bounded network delay in  $\mathcal{F}_{\text{net}}$  by sending a special message.

Observe that the CD consists of two components, SCD and CCD, modeled via separate machines, i.e., the server component `scd` and the client component in  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ , but should be considered corrupted as soon as just one of those machines misbehaves. We capture this expected property by considering a party running `scd` to be corrupted also if the corresponding client in  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  is corrupted, i.e., we use the corruption state of `scd` to represent the corruption status of the combined CD. This idea allows for capturing individual accountability also in a composed protocol, where the same party takes part in multiple parts of a protocol and hence needs to be split into multiple machines. Without this mechanism the judge would be required, by individual accountability and fairness, to identify whether SCD or CCD has misbehaved. In addition to the above, since we identify a party (running `client` or `scd`) with its signature key, we also consider parties to be corrupted if their signature key is corrupted.

Two important aspects of the protocol model are the specifications of the public `judge` and the `supervisor`. The `judge` collects the evidence from clients and from the lower-level `judge`. That is, clients provide all sequences of messages/states that they received from SCDs (include SCD's signature over the state) to the `judge` as evidence. This `judge` also queries the public `judge` of the subroutine  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  to get the most recent verdicts and judicial report from  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ . If the subroutine `judge` returns a non-empty verdict, then the `judge` outputs the same verdict since a misbehaving party was already found in the subroutine and hence consensus might no longer hold true. Otherwise, the `judge` verifies that evidence provided by clients is (i) correctly signed by a CD/SCD and (ii) the provided state is a prefix of the current judicial report, i.e., the correct consensus as determined by the public `judge` in the subroutine consensus protocol. If the first check fails, the client's evidence is not valid and is discarded. If the second check fails, the SCD violated consistency as it signed and forwarded a sequence of messages that differs from the established consensus. Hence, in this case the `judge`



blames an individual SCD. In all other cases no misbehavior was detected and the verdict remains empty. For judicial reports, judge simply forwards the judicial report from the consensus protocol. The public judge follows the modeling as explained in Section 2.4 and, e.g., she reveals all gathered evidence to the network adversary.

The supervisor determines whether the assumptions needed for liveness still hold true by querying (i)  $\mathcal{F}_{\text{net}}$  to check whether the bounded network delay is guaranteed and (ii) the supervisor of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  to check whether liveness assumptions for the subroutine are still met (via `IsAssumptionBroken?`). If any of these checks fail, the supervisor returns that liveness assumptions no longer hold true. Hence, only in this case is the simulator in the ideal world allowed to actually break liveness.

**UC Security Result.** Our security result (cf., Figure 6) states that the scaling protocol  $\mathcal{P}_{\text{cp}}^{\text{acc}}$  using an ideal accountable consensus service subroutine  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  is still an accountable consensus service, i.e., all security guarantees are retained and hence the scaled protocol is also a realization of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ :

**Theorem 1.** *Let  $\mathcal{P}_{\text{cp}}^{\text{acc}}$  and  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  be as described above. Then,*

$$(\mathcal{P}_{\text{cp}}^{\text{acc}} \mid \mathcal{F}_{\text{cp}}^{\text{acc}}) \leq \mathcal{F}_{\text{cp}}^{\text{acc}}.$$

We provide the formal proof for Theorem 1 in Appendix D.

**Discussion.** Using the iUC composition theorem, the ideal subroutine  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  of  $\mathcal{P}_{\text{cp}}^{\text{acc}}$  can be replaced with an arbitrary realization while retaining security results. The perhaps simplest realization consists of clients with access to a consensus service run by a single party, analogous to what is shown in Figure 5 (if CS were considered one party), where the consensus party signs its outputs to provide accountability. By Theorem 1, one can also realize the subroutine  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  via another copy of  $\mathcal{P}_{\text{cp}}^{\text{acc}}$ , i.e., one can iterate the above scaling approach to add additional scaling layers. Security of such a protocol with multiple scaling layers is then directly implied by Theorem 1 and the composition theorem. This nicely demonstrates one of the advantages of composition for accountability properties.

These kinds of composition results are enabled by several features offered by the AUC framework, most notably our novel concept of judicial reports. Indeed, if the public judge in the subroutine  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  had been unable to share his knowledge (i.e., a consistent view on the global state of the subroutine) via a judicial report with the higher-level judge in the scaling protocol, then the judge would be unable to decide, given two inconsistent views, whether CS or CD (CDs in the case of multiple layers) have provided inconsistent views. Without judicial reports, it would thus be impossible to prove individual accountability modularly.

As mentioned at the begin of this section, this case study illustrates the possibilities of AUC. Most notably, composability of accountability-based protocols enabled by judicial reports, the supervisor concept, handling of assumption-based and accountability properties concurrently, and blaming of internal parties (CD).

### 3.2 An Accountable PKI for the Web Based on CTLs

In our second case study, we analyze accountability of a Web PKI based on Certificate Transparency Logs (CTLs) [32, 72]. For the sake of presentation, we consider a slightly simplified version that does not include certificate revocation. We show that such a PKI with CTLs indeed achieves the expected property of *certificate transparency* [72], i.e., accountability w.r.t. certificate correctness. That is, if someone obtains a certified public key for Alice from the PKI, then either this key was indeed registered by Alice or Alice can identify and blame a misbehaving CA for issuing a wrong certificate based on the information provided by CTLs. In Section 3.3, we analyze and prove the security of a standard key exchange (KE) protocol composed with this accountable PKI protocol.

This case study along with the KE protocol uses several features of AUC, including some that were not yet illustrated in Section 3.1, such as (i) local accountability, (ii) individual and group-based accountability levels, and (iii) composition, including the case where higher-level judges use verdicts from lower level judges belonging to different parties. Altogether, using AUC, we are able to perform the first UC analysis of a CTL-based PKI and protocols based thereupon.

In what follows, we present our case study with full formal specifications available in Appendix E.

**Protocol Description.** As depicted in Figure 7, the roles in a CTL-based PKI protocol are (i) clients, (ii) CAs, and (iii) CTLs. Clients request CAs to issue a certificate for them and query CAs for certificates of other clients.

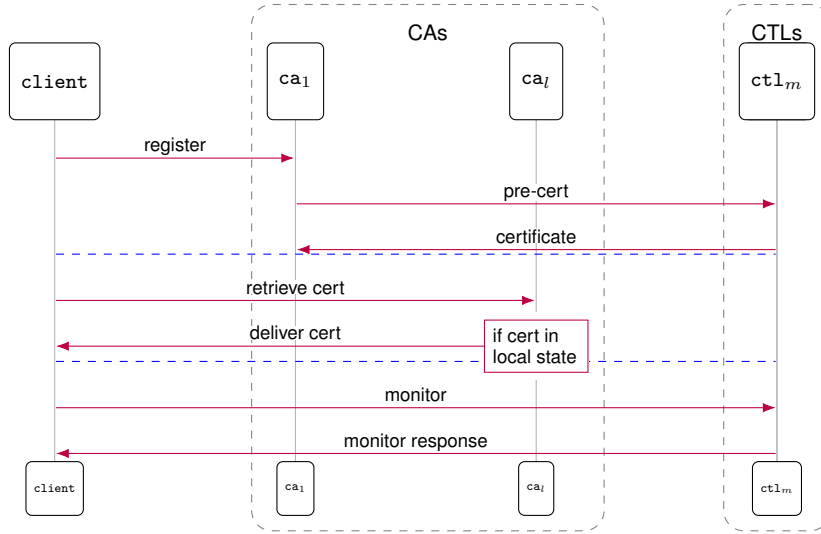


Fig. 7: CTL-based PKI protocol

Here we consider clients that certify at most one key.<sup>5</sup> More specifically, to certify a new key a client sends a registration request, containing its  $pid$  and the public key via an authenticated channel to a CA. When a CA receives a registration request, it checks that it did not issue a certificate for  $pid$  so far. If the request passes the check, the CA generates a pre-certificate containing the original request and a signature of the CA. The CA forwards the pre-certificate to potentially several CTLs. When a CTL receives a pre-certificate, it verifies the CA’s signature. If the signature is valid, the CTL finalizes the certificate and also signs the certificate. CTLs store the certificates they signed/issued and allow clients to monitor these certificates. The underlying idea is that clients can detect identity theft by retrieving certificate lists from CTLs regularly and checking whether there are any certificates in their name that they did not request; we call these *maliciously created certificates* in what follows. Finally, a client  $pid$  can ask to obtain the certificate for another client  $pid'$  from a CA if such a certificate was issued by that CA.

**Security Goal.** Informally, a CTL-based PKI protocol such as the one presented here is supposed to achieve the security goal of *accountability w.r.t. certificate correctness* (typically called *certificate transparency*): honest parties detect maliciously generated certificates for their own identity after some bounded time delay. We denote this *local* accountability property in what follows by “correctCert”. As long as correctCert holds true for a dedicated client, this means that the client actually requested the available certificate and there is no maliciously generated certificate available in the PKI.

**The ideal PKI functionality  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ .** We formalize the property just sketched starting with a non-accountable ideal PKI functionality  $\mathcal{F}_{\text{PKI}}$  that is analogous to Canetti et al.’s ideal functionality  $\mathcal{G}_{\text{BB}}$  [27] and the ideal CA  $\mathcal{F}_{\text{CA}}$  from [17]. We then apply AUC to obtain an accountable version  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  that captures accountability w.r.t. certificate correctness.

Clients are the main roles of  $\mathcal{F}_{\text{PKI}}$ ; CTLs and CAs are internal parties of potential realizations of  $\mathcal{F}_{\text{PKI}}$ .  $\mathcal{F}_{\text{PKI}}$  allows (honest) clients to register one certificate for their own identity at some CA. The adversary  $\mathcal{A}$  decides when and whether a registration is successful. If it succeeds,  $\mathcal{F}_{\text{PKI}}$  issues the certificate (consisting of the party’s  $pid$ , a string, meant to be the party’s public key, and the name of the issuing CA) and adds the certificate to its state. When parties query  $\mathcal{F}_{\text{PKI}}$  for a certificate of an honest  $pid'$  issued by a certain CA,  $\mathcal{A}$  is free to choose when and whether  $\mathcal{F}_{\text{PKI}}$  answers the request. If  $\mathcal{A}$  instructs  $\mathcal{F}_{\text{PKI}}$  to respond,  $\mathcal{F}_{\text{PKI}}$  provides the unique certificate for  $pid'$  (as stored in  $\mathcal{F}_{\text{PKI}}$ ’s state) or it outputs  $\perp$  if there is no certificate recorded at that CA. For corrupted clients  $pid'$ ,  $\mathcal{F}_{\text{PKI}}$  does not provide any guarantees but lets  $\mathcal{A}$  freely determine the response.

<sup>5</sup>This is not an actual restriction. Higher-level protocols can certify multiple keys for a single identity  $pid'$  by setting, e.g.,  $pid = (pid', keyid)$ .

We now apply AUC to  $\mathcal{F}_{\text{PKI}}$  to derive  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  which additionally captures local accountability w.r.t. certificate correctness. We include a local judge  $((\text{local}, \text{pid}, \text{client}), \text{sid}, \text{judge})$  for every client  $(\text{pid}, \text{sid}, \text{client})$  in  $\mathcal{F}_{\text{PKI}}$ , i.e.,  $\text{pids}_{\text{judge}} \subset \{\text{local}\} \times \{0, 1\}^* \times \{\text{client}\}$ . We set  $\text{Sec}^{\text{acc}} = \{\text{correctCert}\}$ . This allows the adversary to indicate that certificate correctness is broken for some client  $\text{pid}$  as long as the adversary provides a verdict to the corresponding local judge. We require that these verdicts blame individual CAs, i.e., the affected client  $\text{pid}$  can identify those CAs that misissued a certificate in their name. This is enforced by instantiating  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}}$  to check that verdicts are of the form  $\bigwedge_{i=1}^n \text{dis}(CA_i)$ , where  $CA_i$  are parties with the (internal) ca role. If  $\text{correctCert}$  is broken for a  $\text{pid}$ , then  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  treats  $\text{pid}$  in the same way as corrupted parties for the purpose of retrieving certificates, i.e., the adversary can return arbitrary certificates issued for  $\text{pid}$ .

**Security model.** We model the CTL-based PKI as a real protocol  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  which implements the previously mentioned roles and operations. We model dynamic sets of clients, CAs, and CTLs. Clients and CAs can be dynamically corrupted, whereas CTLs act as trust anchor and are hence incorruptible.<sup>6</sup>  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  uses an ideal signature functionality  $\mathcal{F}_{\text{sig}}$  for signatures. It further uses three ideal functionalities  $\mathcal{F}_{\text{init}}^{\text{CA}}$ ,  $\mathcal{F}_{\text{psync-net}}$ ,  $\mathcal{F}_{\text{auth}}$  to capture setup assumptions:  $\mathcal{F}_{\text{init}}^{\text{CA}}$  distributes public keys of CAs and CTLs, which are assumed to be known to all parties.  $\mathcal{F}_{\text{psync-net}}$  models a network with bounded message delay, i.e., messages are delivered within  $\delta$  time units. This functionality is used by clients during certificate monitoring to guarantee that a response from the CTL is received in a timely fashion, thereby ensuring that clients can detect maliciously created certificates after some bounded time.  $\mathcal{F}_{\text{psync-net}}$  further provides a clock to all parties, capturing that parties are aware of the current time. Finally,  $\mathcal{F}_{\text{auth}}$  models an ideal authenticated channel which is used during certificate registration, modeling that a CA has some means to identify a client.

We model that a client  $\text{pid}$  regularly monitors CTL certificate lists, namely after at most  $\delta$  time units (any other publicly known bound can be used as well). Hence, by the bounded message delay enforced by  $\mathcal{F}_{\text{psync-net}}$ , we have that  $\text{pid}$  will detect a maliciously created certificate registered at some CTL after at most  $3 \cdot \delta$ . If the client  $\text{pid}$  detects such a certificate, then, from its point of view, the CA is at fault for signing a certificate that was not requested by her.<sup>7</sup> Hence, the local judge of  $\text{pid}$  blames such CAs via the verdict  $\bigwedge_{i=1}^n \text{dis}(CA_i)$ , where  $CA_i$  are exactly those CAs that signed maliciously generated certificates for  $\text{pid}$ . Since a maliciously created certificate is detected by  $\text{pid}$  only after at most  $3 \cdot \delta$  time units, any other party  $\text{pid}'$  that retrieves the certificate before that point in time cannot be sure whether a certificate is genuine or whether  $\text{pid}$  did not yet see and hence did not have the opportunity to complain about the certificate. We therefore model that clients, during certificate retrieval from some CA, accept certificates only with an age of at least  $3 \cdot \delta$  time units. Such a certificate is either correct or  $\text{pid}'$  has already noticed and complained about the malicious certificate, as required by accountability w.r.t. certificate correctness.

Since we consider only local accountability, there is no public judge in  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$ . We also do not use judicial reports in this protocol. As we do not consider assumption-based security properties, we set  $\text{id}_{\text{assumption}} = \emptyset$ . The supervisor in  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  is responsible only for forwarding the corruption status of the internal CAs. This ensures that a simulator in the ideal world can blame a CA in a verdict only if the CA is corrupted in the real world.

**UC Security Result.** We obtain the following result.

**Theorem 2.** *Let  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  be the real PKI protocol and  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  be the ideal accountable PKI functionality as described above. Then,*

$$\mathcal{P}_{\text{PKI}}^{\text{acc}} \leq \mathcal{F}_{\text{PKI}}^{\text{acc}}.$$

We provide a formal proof of Theorem 2 in Appendix E.

### 3.3 A Key Exchange Based on an Accountable PKI

We now analyze a standard authenticated key exchange protocol, the so-called “ISO protocol”, an authenticated version of the Diffie-Hellman key exchange with digital signatures based on the ISO/IEC 9798-3

<sup>6</sup>Trusting a CTL is indeed necessary as a malicious CTL can simply hide certificates during monitoring. This trust can be distributed among several CTLs by requiring  $t \in \mathbb{N}$  CTLs to validate and sign new certificates. In this case, one can obtain a security result if at most  $t - 1$  CTLs are malicious. Our analysis carries over to this setting.

<sup>7</sup>Observe that only  $\text{pid}$  itself can be sure that the CA is at fault. A third party cannot determine whether the CA or  $\text{pid}$  has misbehaved (e.g., by requesting a certificate but then blaming an honest CA). This observation becomes important in the composed protocol analyzed in Section 3.3.




Fig. 8: The ISO protocol for mutually authenticated Diffie-Hellman key exchange between two parties A and B.

standard [56] (see Figure 8). UC security of this protocol has already been studied in various settings [17, 19, 20, 26, 27, 66] but always based on the assumption that the underlying PKI is perfect, i. e., the adversary cannot register certificates for honest parties. In contrast, we base our analysis on  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  which can then be realized by  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  (Section 3.2) using the composition theorem. We thus provide the first analysis of the ISO protocol based on a PKI that may fail, but provides accountability when it does. This not only illustrates the features of AUC highlighted at the beginning of Section 3.2. This also shows that even protocols which traditionally consider only preventive security, such as key exchanges, can benefit from AUC. In the main body, we explain the main aspects of this case study; full details, including formal specifications and proofs, are provided in Appendix F.

**Ideal accountable key exchange.** A signature-based authenticated key exchange can only provide security as long as the underlying public signature key/certificate is trustworthy. We capture this intuition with AUC: We start with a standard ideal key exchange functionality  $\mathcal{F}_{\text{KE}}$  [17, 20] and apply AUC to obtain an accountable version  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ . In  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  we consider the local accountability property  $\text{authenticity} \in \text{Sec}^{\text{acc}}$  which determines if a party, say Alice, can still expect to authenticate her intended session partner, say Bob, or whether the certificate for Bob might be incorrect due to a fault in the PKI. If  $\text{authenticity}$  is marked as broken for Alice by the adversary  $\mathcal{A}$  (in exchange for a verdict), then  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  acts just as  $\mathcal{F}_{\text{KE}}$  does in case one of the parties in the session is corrupted and hence no authentication can be guaranteed, i. e., it leaks the session key, if any, and allows  $\mathcal{A}$  to determine the output for parties that have not yet finished the KE. We require verdicts to be of the form  $\text{dis}(p_{\text{main}}) \vee v$ , where  $p_{\text{main}}$  is a main party, i. e., initiator or responder in this key exchange session, and  $v$  is a verdict containing only internal parties. This is a group based accountability level which captures that Alice in the real protocol, if Bob complains about a maliciously generated certificate, cannot decide whether Bob is lying or the PKI has actually misbehaved (cf. Footnote 7).

**Protocol model.**  $\mathcal{P}_{\text{iso}}^{\text{acc}}$  is a straightforward model of the ISO protocol shown in Figure 8 derived from [17] consisting of the KE protocol part  $\mathcal{P}_{\text{pke}}^{\text{acc}}$  and the ideal subroutine  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  which is used for public key distribution. As required by AUC,  $\mathcal{P}_{\text{iso}}^{\text{acc}}$  contains local judges and a supervisor. As in Section 3.2, the supervisor only forwards the corruption status of internal parties including those that are part of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ , namely  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ 's clients, CAs and CTLs. The main idea of the local judge of Alice is to request the verdict of Bobs local judge in  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . If this judge does not complain, then, by definition of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ , any certificate of Bob that Alice retrieves must have been registered by Bob. If Bobs subroutine judge complains and returns a verdict  $v$  to Alice, then Alice's judge returns the overall verdict  $\text{dis}(p_{\text{Bob}}) \vee v$  (and vice versa for Bob's judge), capturing the aforementioned insecurity that Alice cannot decide whether Bob is lying or whether  $v$  is actually a fair verdict identifying a misbehaving PKI.

**UC Security results.** We can show the following:

**Theorem 3.** *Let  $\mathcal{P}_{\text{KE}}^{\text{acc}}$ ,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ , and  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  be as described above and assume that the DDH assumption holds true. Then,*

$$(\mathcal{P}_{\text{KE}}^{\text{acc}} | \mathcal{F}_{\text{PKI}}^{\text{acc}}) \leq \mathcal{F}_{\text{KE}}^{\text{acc}}.$$

This immediately implies by UC composition that the key exchange remains secure when based on the real accountable CTL-based PKI  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$ :

**Corollary 4.** *Let  $\mathcal{P}_{\text{KE}}^{\text{acc}}$ ,  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$ , and  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  be as described above. Then,*

$$(\mathcal{P}_{\text{KE}}^{\text{acc}} | \mathcal{P}_{\text{PKI}}^{\text{acc}}) \leq \mathcal{F}_{\text{KE}}^{\text{acc}}.$$

*Proof.* Follows from Theorem 2, Theorem 3, and the composition theorem of the underlying UC model.  $\square$

#### 4. RELATED WORK AND CONCLUSION

As already discussed in Section 2.7, AUC focuses on property-based accountability, it generalizes and extends concepts from the literature and also introduces new concepts, such as judicial reports and supervisors to provide a general accountability framework.

**Property-based Accountability in UC.** To the best of our knowledge, the only other works that formalize and use the concept of property-based accountability in a UC model are those on MPC protocols, e.g., [13, 14, 24, 34, 35, 55, 74, 81]. As discussed in the introduction and in Appendix G, these works are specialized to the case of MPC and hence do not serve as general accountability frameworks. Most of these works consider composition of accountability properties with higher-level protocols to be out of scope. A notable exception is the very recent work of Baum et al. [14]. However, Baum et al. focus on the composability of verifiability in MPC protocols adhering to a specific structure. Baum et al. do not provide a general purpose accountability framework which can be used for arbitrary protocols.

**Other simulation-based approaches.** There are also a number of non-UC simulation-based formalizations of accountability properties, e.g., [5, 6, 8, 34, 37, 54, 86, 96]. Just as for the UC approaches mentioned above, these works analyze and are tailored towards MPC protocols and thus do not serve as general accountability frameworks. Furthermore, since they are not based on a UC model, they provide only weaker compositional properties, if any.

The covert adversaries model [6, 8] is perhaps the most prominent line of work in this category. The covert adversaries model formalizes accountability w.r.t. correctness (in the sense of identifiable abort) and w.r.t. privacy. AUC, even when restricted to the special case of MPC, and covert adversaries are incomparable due to different simulation paradigms. Both approaches can formalize accountability w.r.t. to correctness and privacy of MPC protocols (cf. Appendix G). On the one hand, AUC offers stronger composability that, unlike covert adversaries, also includes parallel composition. On the other hand, covert adversaries provide the additional concept of a deterrence factor  $\varepsilon$  to also model cases where a malicious party breaking security might remain undetected by a judge with (potentially non-negligible) probability  $\varepsilon$ . AUC models only the case that this probability is negligible. While it would be straightforward to add the same concept to AUC, we did not do so as it does not appear to offer any benefit within UC models. Indeed, it seems that all covert adversaries protocols that have been analyzed for a non-negligible  $\varepsilon$  use protocol rewinding within their simulators. This technique is not available to UC simulators since it prevents parallel composition, i.e., such protocols are not UC secure anyway. We leave further exploration of this aspect for future work.

**Game-based accountability.** There are many works that formalize accountability within a game-based setting, e.g., [42–46, 62, 63, 67, 79]. Some of these works are closely related to AUC in that they also consider highly general frameworks for accountability, e.g., [42, 45, 62, 67]. The main difference between AUC and these works is that AUC is the first general accountability framework for UC models, thereby providing particularly strong security statements while also offering the benefit of modular protocol analysis and composition. We, however, note that there are aspects in existing game-based accountability frameworks that AUC does not handle yet, such as *causality* [62]. It is an interesting challenge for future work to investigate whether and how these aspects can also be captured in a general accountability framework for UC.

Altogether, AUC lifts some of the work on game-based accountability frameworks to the UC setting, generalizes and unifies existing work on UC accountability, and also introduces several new concepts to make it a general purpose framework for accountability in UC.

#### 5. ACKNOWLEDGMENTS

This research was partially funded by the Ministry of Science of Baden-Württemberg, Germany, for the Doctoral Program “Services Computing”.<sup>8</sup> This work was also supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through grant KU 1434/13-1.

<sup>8</sup><http://www.services-computing.de/>

## REFERENCES

- [1] B. Adida, “Helios: Web-based Open-Audit Voting,” in *Proceedings of the 17th USENIX Security Symposium*, P. C. van Oorschot, Ed. USENIX Association, 2008, pp. 335–348.
- [2] J. F. Almansa, I. Damgård, and J. B. Nielsen, “Simplified Threshold RSA with Adaptive and Proactive Security,” in *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, ser. Lecture Notes in Computer Science, vol. 4004. Springer, 2006, pp. 593–611.
- [3] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. D. Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolic, S. W. Cocco, and J. Yellick, “Hyperledger fabric: a distributed operating system for permissioned blockchains,” in *Proceedings of the Thirteenth EuroSys Conference, EuroSys 2018, Porto, Portugal, April 23-26, 2018*. ACM, 2018, pp. 30:1–30:15.
- [4] Apache Software Foundation, “Apache Kafka,” <https://kafka.apache.org/>, 2017, (Accessed on 04/01/2019).
- [5] G. Asharov, Y. Lindell, T. Schneider, and M. Zohner, “More Efficient Oblivious Transfer Extensions,” *J. Cryptol.*, vol. 30, no. 3, pp. 805–858, 2017.
- [6] G. Asharov and C. Orlandi, “Calling Out Cheaters: Covert Security with Public Verifiability,” in *Advances in Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, ser. Lecture Notes in Computer Science, vol. 7658. Springer, 2012, pp. 681–698.
- [7] N. Asokan, V. Shoup, and M. Waidner, “Asynchronous protocols for optimistic fair exchange,” in *Proceedings of the IEEE Symposium on Research in Security and Privacy*. IEEE Computer Society, 1998, pp. 86–99.
- [8] Y. Aumann and Y. Lindell, “Security Against Covert Adversaries: Efficient Protocols for Realistic Adversaries,” in *Proceedings of the 4th Theory of Cryptography Conference, (TCC 2007)*, ser. Lecture Notes in Computer Science, S. P. Vadhan, Ed., vol. 4392. Springer, 2007, pp. 137–156.
- [9] C. Badertscher, R. Canetti, J. Hesse, B. Tackmann, and V. Zikas, “Universal Composition with Global Subroutines: Capturing Global Setup Within Plain UC,” in *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part III*, ser. Lecture Notes in Computer Science, vol. 12552. Springer, 2020, pp. 1–30.
- [10] C. Badertscher, U. Maurer, D. Tschudi, and V. Zikas, “Bitcoin as a Transaction Ledger: A Composable Treatment,” in *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, ser. Lecture Notes in Computer Science, vol. 10401. Springer, 2017, pp. 324–356.
- [11] L. Baird and A. Luykx, “The Hashgraph Protocol: Efficient Asynchronous BFT for High-Throughput Distributed Ledgers,” in *2020 International Conference on Omni-layer Intelligent Systems, COINS 2020, Barcelona, Spain, August 31 - September 2, 2020*. IEEE, 2020, pp. 1–7.
- [12] J. Baron, K. E. Defrawy, J. Lampkins, and R. Ostrovsky, “Communication-Optimal Proactive Secret Sharing for Dynamic Groups,” in *Applied Cryptography and Network Security - 13th International Conference, ACNS 2015, New York, NY, USA, June 2-5, 2015, Revised Selected Papers*, ser. Lecture Notes in Computer Science, vol. 9092. Springer, 2015, pp. 23–41.
- [13] C. Baum, I. Damgård, and C. Orlandi, “Publicly Auditable Secure Multi-Party Computation,” in *Security and Cryptography for Networks - 9th International Conference, SCN 2014, Amalfi, Italy, September 3-5, 2014. Proceedings*, ser. Lecture Notes in Computer Science, vol. 8642. Springer, 2014, pp. 175–196.
- [14] C. Baum, E. Orsini, P. Scholl, and E. Soria-Vazquez, “Efficient Constant-Round MPC with Identifiable Abort and Public Verifiability,” in *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 12171. Springer, 2020, pp. 562–592.
- [15] A. Boudguiga, N. Bouzerna, L. Granboulan, A. Olivereau, F. Quesnel, A. Roger, and R. Sirdey, “Towards Better Availability and Accountability for IoT Updates by Means of a Blockchain,” in *2017 IEEE European Symposium on Security and Privacy Workshops, EuroS&P Workshops 2017, Paris, France, April 26-28, 2017*. IEEE, 2017, pp. 50–58.
- [16] V. Buterin and V. Griffith, “Casper the Friendly Finality Gadget,” *CoRR*, vol. abs/1710.09437, 2017.
- [17] J. Camenisch, S. Krenn, R. Küsters, and D. Rausch, “iUC: Flexible Universal Composability Made Simple,” in *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part III*, ser. Lecture Notes in Computer Science, vol. 11923. Springer, 2019, pp. 191–221, the full version is available at <http://eprint.iacr.org/2019/1073>.
- [18] R. Canetti, Y. Dodis, R. Pass, and S. Walfish, “Universally Composable Security with Global Setup,” in *Theory of Cryptography, Proceedings of TCC 2007*, ser. Lecture Notes in Computer Science, S. P. Vadhan, Ed., vol. 4392. Springer, 2007, pp. 61–85.
- [19] R. Canetti and J. Herzog, “Universally Composable Symbolic Analysis of Mutual Authentication and Key-Exchange Protocols,” in *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006*, ser. Lecture Notes in Computer Science, S. Halevi and T. Rabin, Eds., vol. 3876. Springer, 2006, pp. 380–403.
- [20] R. Canetti and H. Krawczyk, “Universally Composable Notions of Key Exchange and Secure Channels,” in *Advances in Cryptology - EUROCRYPT 2002, International Conference on the Theory and Applications of Cryptographic Techniques, Proceedings*, ser. Lecture Notes in Computer Science, vol. 2332. Springer, 2002, pp. 337–351.
- [21] R. Canetti, Y. Lindell, R. Ostrovsky, and A. Sahai, “Universally composable two-party and multi-party secure computation,” in *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC 2002)*. ACM Press, 2002, pp. 494–503.
- [22] R. Canetti, “Universally Composable Security: A New Paradigm for Cryptographic Protocols,” in *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS 2001)*. IEEE Computer Society, 2001, pp. 136–145.
- [23] —, “Universally Composable Security,” *J. ACM*, vol. 67, no. 5, pp. 28:1–28:94, 2020.
- [24] R. Canetti, A. Cohen, and Y. Lindell, “A Simpler Variant of Universally Composable Security for Standard Multiparty Computation,” in *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 9216. Springer, 2015, pp. 3–22.
- [25] R. Canetti, K. Hogan, A. Malhotra, and M. Varia, “A Universally Composable Treatment of Network Time,” in *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*. IEEE Computer Society, 2017, pp. 360–375.
- [26] R. Canetti and H. Krawczyk, “Security Analysis of IKE’s Signature-Based Key-Exchange Protocol,” in *Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference*, ser. Lecture Notes in Computer Science, M. Yung, Ed., vol. 2442. Springer, 2002, pp. 143–161.

- [27] R. Canetti, D. Shahaf, and M. Vald, “Universally Composable Authentication and Key-Exchange with Global PKI,” in *Public-Key Cryptography - PKC 2016 - 19th IACR International Conference on Practice and Theory in Public-Key Cryptography, Taipei, Taiwan, March 6-9, 2016, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 9615. Springer, 2016, pp. 265–296.
- [28] M. Castro and B. Liskov, “Practical byzantine fault tolerance and proactive recovery,” *ACM Trans. Comput. Syst.*, vol. 20, no. 4, pp. 398–461, 2002.
- [29] C. Chang and Y. Chang, “Efficient anonymous auction protocols with freewheeling bids,” *Comput. Secur.*, vol. 22, no. 8, pp. 728–734, 2003.
- [30] M. Ciampi, Y. Lu, and V. Zikas, “Collusion-Preserving Computation without a Mediator,” Cryptology ePrint Archive, Tech. Rep. 2020/497, 2020.
- [31] D. Ó. Coileáin and D. O’Mahony, “Accounting and Accountability in Content Distribution Architectures: A Survey,” *ACM Comput. Surv.*, vol. 47, no. 4, pp. 59:1–59:35, 2015.
- [32] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk, “Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile,” RFC 5280, Internet Engineering Task Force, may 2008. [Online]. Available: <http://www.ietf.org/rfc/rfc5280.txt>
- [33] V. Cortier, D. Galindo, R. Küsters, J. Müller, and T. Truderung, “SoK: Verifiability Notions for E-Voting Protocols,” in *IEEE 37th Symposium on Security and Privacy (S&P 2016)*. IEEE Computer Society, 2016, pp. 779–798.
- [34] R. K. Cunningham, B. Fuller, and S. Yakoubov, “Catching MPC Cheaters: Identification and Openability,” in *Information Theoretic Security - 10th International Conference, ICITS 2017, Hong Kong, China, November 29 - December 2, 2017, Proceedings*, ser. Lecture Notes in Computer Science, vol. 10681. Springer, 2017, pp. 110–134.
- [35] E. Cuvelier and O. Pereira, “Verifiable Multi-party Computation with Perfectly Private Audit Trail,” in *Applied Cryptography and Network Security - 14th International Conference, ACNS 2016, Guildford, UK, June 19-22, 2016, Proceedings*, ser. Lecture Notes in Computer Science, vol. 9696. Springer, 2016, pp. 367–385.
- [36] E. Cuvelier, O. Pereira, and T. Peters, “Election Verifiability or Ballot Privacy: Do We Need to Choose?” in *Computer Security - ESORICS 2013 - 18th European Symposium on Research in Computer Security, Egham, UK, September 9-13, 2013, Proceedings*, ser. Lecture Notes in Computer Science, vol. 8134. Springer, 2013, pp. 481–498.
- [37] I. Damgård, C. Orlandi, and M. Simkin, “Black-Box Transformations from Passive to Covert Security with Public Verifiability,” in *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 12171. Springer, 2020, pp. 647–676.
- [38] G. D’Angelo, S. Ferretti, and M. Marzolla, “A Blockchain-based Flight Data Recorder for Cloud Accountability,” in *Proceedings of the 1st Workshop on Cryptocurrencies and Blockchains for Distributed Systems, CRYBLOCK@MobiSys 2018, Munich, Germany, June 15, 2018*. ACM, 2018, pp. 93–98.
- [39] J. Dille, B. M. Maggs, J. Parikh, H. Prokop, R. K. Sitaraman, and W. E. Weihl, “Globally Distributed Content Delivery,” *IEEE Internet Comput.*, vol. 6, no. 5, pp. 50–58, 2002.
- [40] Ethereum Foundation, “Ethereum enterprise,” <https://www.ethereum.org/enterprise/>, 2019, (Accessed on 11/13/2019).
- [41] C. Farkas, G. Ziegler, A. Meretei, and A. Lörcincz, “Anonymity and accountability in self-organizing electronic communities,” in *Proceedings of the 2002 ACM Workshop on Privacy in the Electronic Society, WPES 2002, Washington, DC, USA, November 21, 2002*. ACM, 2002, pp. 81–90.
- [42] J. Feigenbaum, “Privacy, Anonymity, and Accountability in Ad-Supported Services,” in *Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25-28, 2012*. IEEE Computer Society, 2012, pp. 9–10.
- [43] J. Feigenbaum, J. A. Hendler, A. D. Jaggard, D. J. Weitzner, and R. N. Wright, “Accountability and deterrence in online life,” in *Web Science 2011, WebSci ’11, Koblenz, Germany - June 15 - 17, 2011*. ACM, 2011, pp. 7:1–7:7.
- [44] J. Feigenbaum, A. D. Jaggard, and R. N. Wright, “Towards a formal model of accountability,” in *2011 New Security Paradigms Workshop, NSPW ’11, Marin County, CA, USA, September 12-15, 2011*. ACM, 2011, pp. 45–56.
- [45] —, “Open vs. closed systems for accountability,” in *Proceedings of the 2014 Symposium and Bootcamp on the Science of Security, HotSoS 2014, Raleigh, NC, USA, April 08 - 09, 2014*. ACM, 2014, p. 4.
- [46] —, “Accountability in Computing: Concepts and Mechanisms,” *Found. Trends Priv. Secur.*, vol. 2, no. 4, pp. 247–399, 2020.
- [47] E. Funk, J. Riddell, F. Ankel, and D. Cabrera, “Blockchain technology: a data framework to improve validity, trust, and accountability of information exchange in health professions education,” *Academic Medicine*, vol. 93, no. 12, pp. 1791–1794, 2018.
- [48] J. A. Garay, A. Kiayias, and N. Leonardos, “The Bitcoin Backbone Protocol: Analysis and Applications,” in *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 9057. Springer, 2015, pp. 281–310.
- [49] S. Goldwasser and S. Park, “Public Accountability vs. Secret Laws: Can They Coexist?: A Cryptographic Proposal,” in *Proceedings of the 2017 on Workshop on Privacy in the Electronic Society, Dallas, TX, USA, October 30 - November 3, 2017*. ACM, 2017, pp. 99–110.
- [50] M. Graf, R. Küsters, and D. Rausch, “Accountability in a Permissioned Blockchain: Formal Analysis of Hyperledger Fabric,” in *IEEE European Symposium on Security and Privacy, EuroS&P 2020, Genoa, Italy, September 7-11, 2020*. Los Alamitos, CA, USA: IEEE, 2020, pp. 236–255.
- [51] M. Graf, D. Rausch, V. Ronge, C. Egger, R. Küsters, and D. Schröder, “A Security Framework for Distributed Ledgers,” in *CCS ’21: 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, November 15 - 19, 2021*. New York City, USA: ACM, 2021, pp. 1043–1064.
- [52] A. Haeberlen, P. Kouznetsov, and P. Druschel, “Peerreview: practical accountability for distributed systems,” in *Proceedings of the 21st ACM Symposium on Operating Systems Principles 2007, SOSP 2007*, T. C. Bressoud and M. F. Kaashoek, Eds. ACM, 2007, pp. 175–188.
- [53] D. Hofheinz and V. Shoup, “GNUC: A New Universal Composability Framework,” *J. Cryptology*, vol. 28, no. 3, pp. 423–508, 2015.
- [54] C. Hong, J. Katz, V. Kolesnikov, W. Lu, and X. Wang, “Covert Security with Public Verifiability: Faster, Leaner, and Simpler,” in *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part III*, ser. Lecture Notes in Computer Science, vol. 11478. Springer, 2019, pp. 97–121.

- [55] Y. Ishai, R. Ostrovsky, and V. Zikas, "Secure Multi-Party Computation with Identifiable Abort," in *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 8617. Springer, 2014, pp. 369–386.
- [56] "ISO/IEC IS 9798-3, Entity authentication mechanisms — Part 3: Entity authentication using asymmetric techniques," 1993.
- [57] R. Jagadeesan, A. Jeffrey, C. Pitcher, and J. Riely, "Towards a theory of accountability and audit," in *ESORICS*. Springer, 2009.
- [58] J. Kamto, L. Qian, J. Fuller, J. Attia, and Y. Qian, "Key Distribution and management for power aggregation and accountability in Advance Metering Infrastructure," in *IEEE Third International Conference on Smart Grid Communications, SmartGridComm 2012, Tainan, Taiwan, November 5-8, 2012*. IEEE, 2012, pp. 360–365.
- [59] G. O. Karame, E. Androulaki, M. Roeschlin, A. Gervais, and S. Capkun, "Misbehavior in Bitcoin: A Study of Double-Spending and Accountability," *ACM Trans. Inf. Syst. Secur.*, vol. 18, no. 1, pp. 2:1–2:32, 2015.
- [60] A. Kiayias, H. Zhou, and V. Zikas, "Fair and Robust Multi-party Computation Using a Global Transaction Ledger," in *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 9666. Springer, 2016, pp. 705–734.
- [61] T. H. Kim, L. Huang, A. Perrig, C. Jackson, and V. D. Gligor, "Accountable key infrastructure (AKI): a proposal for a public-key validation infrastructure," in *22nd International World Wide Web Conference, WWW '13, Rio de Janeiro, Brazil, May 13-17, 2013*. International World Wide Web Conferences Steering Committee / ACM, 2013, pp. 679–690.
- [62] R. Künnemann, I. Esiyok, and M. Backes, "Automated Verification of Accountability in Security Protocols," in *32nd IEEE Computer Security Foundations Symposium, CSF 2019, Hoboken, NJ, USA, June 25-28, 2019*. IEEE, 2019, pp. 397–413.
- [63] R. Künnemann, D. Garg, and M. Backes, "Accountability in the Decentralised-Adversary Setting," in *34th IEEE Computer Security Foundations Symposium, CSF 2021*,. IEEE, 2021.
- [64] R. Küsters, "Simulation-Based Security with Inexhaustible Interactive Turing Machines," in *Proceedings of the 19th IEEE Computer Security Foundations Workshop (CSFW-19 2006)*. IEEE Computer Society, 2006, pp. 309–320, see [71] for a full and revised version.
- [65] R. Küsters, J. Liedtke, J. Müller, D. Rausch, and A. Vogt, "Ordinos: A Verifiable Tally-Hiding E-Voting System," in *IEEE European Symposium on Security and Privacy, EuroS&P 2020, Genoa, Italy, September 7-11, 2020*. IEEE, 2020, pp. 216–235.
- [66] R. Küsters and D. Rausch, "A Framework for Universally Composable Diffie-Hellman Key Exchange," in *IEEE 38th Symposium on Security and Privacy (S&P 2017)*. IEEE Computer Society, 2017, pp. 881–900.
- [67] R. Küsters, T. Truderung, and A. Vogt, "Accountability: Definition and Relationship to Verifiability," in *Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS 2010)*. ACM, 2010, pp. 526–535, the full version is available at <http://eprint.iacr.org/2010/236>.
- [68] R. Küsters, T. Truderung, and A. Vogt, "Verifiability, Privacy, and Coercion-Resistance: New Insights from a Case Study," in *32nd IEEE Symposium on Security and Privacy (S&P 2011)*. IEEE Computer Society, 2011, pp. 538–553.
- [69] —, "Clash Attacks on the Verifiability of E-Voting Systems," in *33rd IEEE Symposium on Security and Privacy (S&P 2012)*. IEEE Computer Society, 2012, pp. 395–409.
- [70] R. Küsters and M. Tuengerthal, "Composition Theorems Without Pre-Established Session Identifiers," in *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS 2011)*, Y. Chen, G. Danezis, and V. Shmatikov, Eds. ACM, 2011, pp. 41–50.
- [71] R. Küsters, M. Tuengerthal, and D. Rausch, "The IITM model: a simple and expressive model for universal composability," *Journal of Cryptology*, vol. 33, no. 4, pp. 1461–1584, 2020.
- [72] B. Laurie, A. Langley, and E. Kasper, "Certificate Transparency," RFC 6962, jun 2013. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc6962.txt>
- [73] H. Leibowitz, A. Herzberg, and E. Syta, "Provable Security for PKI Schemes," Cryptology ePrint Archive, Tech. Rep. 2019/807, 2019.
- [74] Y. Lindell and B. Pinkas, "Secure Two-Party Computation via Cut-and-Choose Oblivious Transfer," *J. Cryptol.*, vol. 25, no. 4, pp. 680–722, 2012.
- [75] J. Liu, Y. Xiao, and J. Gao, "Achieving Accountability in Smart Grid," *IEEE Syst. J.*, vol. 8, no. 2, pp. 493–508, 2014.
- [76] S. Matsumoto and R. M. Reischuk, "Certificates-as-an-insurance: Incentivizing accountability in ssl/tls," in *Proceedings of the NDSS Workshop on Security of Emerging Network Technologies (SENT'15)*, 2015.
- [77] U. Maurer, "Constructive Cryptography - A New Paradigm for Security Definitions and Proofs," in *Theory of Security and Applications - Joint Workshop, TOSCA 2011, Saarbrücken, Germany, March 31 - April 1, 2011, Revised Selected Papers*, ser. Lecture Notes in Computer Science, vol. 6993. Springer, 2011, pp. 33–56.
- [78] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song, "The Honey Badger of BFT Protocols," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*. ACM, 2016, pp. 31–42.
- [79] K. Morio and R. Künnemann, "Verifying Accountability for Unbounded Sets of Participants," in *34th IEEE Computer Security Foundations Symposium, CSF 2021*,. IEEE, 2021.
- [80] R. Pass, L. Seeman, and A. Shelat, "Analysis of the Blockchain Protocol in Asynchronous Networks," in *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part II*, ser. Lecture Notes in Computer Science, vol. 10211, 2017, pp. 643–673.
- [81] A. Patra and D. Ravi, "Beyond Honest Majority: The Round Complexity of Fair and Robust Multi-party Computation," in *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part I*, ser. Lecture Notes in Computer Science, vol. 11921. Springer, 2019, pp. 456–487.
- [82] R3, "R3 Corda master documentation," <https://docs.corda.net/docs/corda-os/4.4.html>, 2020, (Accessed on 04/24/2020).
- [83] H. V. Ramasamy, A. Agbaria, and W. H. Sanders, "A Parsimonious Approach for Obtaining Resource-Efficient and Trustworthy Execution," *IEEE Trans. Dependable Secur. Comput.*, vol. 4, no. 1, pp. 1–17, 2007.
- [84] K. Ramchen, C. Culnane, O. Pereira, and V. Teague, "Universally Verifiable MPC and IRV Ballot Counting," in *Financial Cryptography and Data Security - 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18-22, 2019, Revised Selected Papers*, ser. Lecture Notes in Computer Science, vol. 11598. Springer, 2019, pp. 301–319.
- [85] M. Rivinius, P. Reisert, D. Rausch, and R. Küsters, "Publicly accountable robust multi-party computation," in *S&P '22*. IEEE, 2022.



- [86] B. Schoenmakers and M. Veeningen, “Universally Verifiable Multiparty Computation from Threshold Homomorphic Cryptosystems,” in *Applied Cryptography and Network Security - 13th International Conference, ACNS 2015, New York, NY, USA, June 2-5, 2015, Revised Selected Papers*, ser. Lecture Notes in Computer Science, vol. 9092. Springer, 2015, pp. 3–22.
- [87] P. Scholl, M. Simkin, and L. Siniscalchi, “Multiparty Computation with Covert Security and Public Verifiability,” *Cryptology ePrint Archive*, Tech. Rep. 2021/366, 2021.
- [88] A. Shamis, P. Pietzuch, B. Canakci, M. Castro, C. Fournet, E. Ashton, A. Chamayou, S. Clebsch, A. Delignat-Lavaud, M. Kerner *et al.*, “IA-CCF: Individual accountability for permissioned ledgers,” in *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, 2022, pp. 467–491.
- [89] C. Stathakopoulou, T. David, and M. Vukolic, “Mir-BFT: High-Throughput BFT for Blockchains,” *CoRR*, vol. abs/1906.05552, 2019.
- [90] Y. S. Tan, R. K. L. Ko, and G. Holmes, “Security and Data Accountability in Distributed Systems: A Provenance Survey,” in *10th IEEE International Conference on High Performance Computing and Communications & 2013 IEEE International Conference on Embedded and Ubiquitous Computing, HPCC/EUC 2013, Zhangjiajie, China, November 13-15, 2013*. IEEE, 2013, pp. 1571–1578.
- [91] The Guardian, “Steve jobs suggests: get rid of the drm on online music,” <https://www.theguardian.com/technology/blog/2007/feb/06/stevejobssugg>, 2007, (Accessed on 07/09/2021).
- [92] VASCO Data Security International, Inc., “news: Vasco announces bankruptcy filing by diginotar b.v.” [https://web.archive.org/web/20110923180445/http://www.vasco.com/company/press\\_room/news\\_archive/2011/news\\_vasco\\_announces\\_bankruptcy\\_filing\\_by\\_diginotar\\_bv.aspx](https://web.archive.org/web/20110923180445/http://www.vasco.com/company/press_room/news_archive/2011/news_vasco_announces_bankruptcy_filing_by_diginotar_bv.aspx), 09 2011, (Accessed on 06/13/2022).
- [93] WIRED, “Shocker: Apple’s drm-free music not so easily stolen,” <https://www.wired.com/2007/06/apples-drmfree-/>, 2007, (Accessed on 07/09/2021).
- [94] Z. Xiao, N. Kathiresshan, and Y. Xiao, “A survey of accountability in computer networks and distributed systems,” *Secur. Commun. Networks*, vol. 9, no. 4, pp. 290–315, 2016.
- [95] J. Yin, J. Martin, A. Venkataramani, L. Alvisi, and M. Dahlin, “Separating agreement from execution for byzantine fault tolerant services,” in *Proceedings of the 19th ACM Symposium on Operating Systems Principles 2003, SOSOP 2003, Bolton Landing, NY, USA, October 19-22, 2003*. ACM, 2003, pp. 253–267.
- [96] B. Zeng, C. Tartary, P. Xu, J. Jing, and X. Tang, “A Practical Framework for t-Out-of-n Oblivious Transfer With Security Against Covert Adversaries,” *IEEE Trans. Inf. Forensics Secur.*, vol. 7, no. 2, pp. 465–479, 2012.
- [97] Z. Zhao, M. Naseri, and Y. Zheng, “Secure quantum sealed-bid auction with post-confirmation,” *Optics Communications*, vol. 283, no. 16, pp. 3194–3197, 2010.
- [98] J. Zhou and D. Gollmann, “A fair non-repudiation protocol,” in *Proceedings of the IEEE Symposium on Research in Security and Privacy*. IEEE Computer Society, 1996, pp. 55–61.
- [99] —, “An Efficient Non-repudiation Protocol,” in *10th Computer Security Foundations Workshop (CSFW '97), June 10-12, 1997, Rockport, Massachusetts, USA*. IEEE Computer Society, 1997, pp. 126–132.
- [100] —, “Evidence and non-repudiation,” *Journal of Network and Computer Applications*, vol. 20, no. 3, pp. 267–281, 1997.

## APPENDIX

### A. A Brief Introduction to the iUC Framework

This section provides a brief introduction to the iUC framework, which underlies all results in this paper. The iUC framework [17] is a highly expressive and user-friendly model for universal composability. It allows for the modular analysis of different types of protocols in various security settings.

The iUC framework uses interactive Turing machines as its underlying computational model. Such interactive Turing machines can be connected to each other to be able to exchange messages. A set of machines  $\mathcal{Q} = \{M_1, \dots, M_k\}$  is called a *system*. In a run of  $\mathcal{Q}$ , there can be one or more instances (copies) of each machine in  $\mathcal{Q}$ . One instance can send messages to another instance. At any point in a run, only a single instance is active, namely, the one to receive the last message; all other instances wait for input. The active instance becomes inactive once it has sent a message; then the instance that receives the message becomes active instead and can perform arbitrary computations. The first machine to run is the so-called *master*. The master is also triggered if the last active machine did not output a message. In iUC, the environment (see next) takes the role of the master. In the iUC framework a special user-specified **CheckID** algorithm is used to determine which instance of a protocol machine receives a message and whether a new instance is to be created (see below).

To define the universal composability security experiment (cf. [17]), one distinguishes between three types of systems: protocols, environments, and adversaries. As is standard in universal composability models, all of these types of systems have to meet a polynomial runtime notion. Intuitively, the security experiment in any universal composability model compares a protocol  $\mathcal{P}$  with another protocol  $\mathcal{F}$ , where  $\mathcal{F}$  is typically an ideal specification of some task, called *ideal protocol* or *ideal functionality*. The idea is that if one cannot distinguish  $\mathcal{P}$  from  $\mathcal{F}$ , then  $\mathcal{P}$  must be “as good as”  $\mathcal{F}$ . More specifically, the protocol  $\mathcal{P}$  is considered secure (written  $\mathcal{P} \leq \mathcal{F}$ ) if for all adversaries  $\mathcal{A}$  controlling the network of  $\mathcal{P}$  there exists an (ideal) adversary  $\mathcal{S}$ , called *simulator*, controlling the network of  $\mathcal{F}$  such that  $\{\mathcal{A}, \mathcal{P}\}$  and  $\{\mathcal{S}, \mathcal{F}\}$  are indistinguishable for all environments  $\mathcal{E}$ . Indistinguishability means that the probability of the environment outputting 1 in runs of the system  $\{\mathcal{E}, \mathcal{A}, \mathcal{P}\}$  is negligibly close to the probability of outputting 1 in runs of the system  $\{\mathcal{E}, \mathcal{S}, \mathcal{F}\}$  (written  $\{\mathcal{E}, \mathcal{A}, \mathcal{P}\} \equiv \{\mathcal{E}, \mathcal{S}, \mathcal{F}\}$ ). The environment can also subsume the role of the network attacker  $\mathcal{A}$ , which yields an

equivalent definition in the iUC framework. We usually show this equivalent but simpler statement in our proofs, i.e., that there exists a simulator  $\mathcal{S}$  such that  $\{\mathcal{E}, \mathcal{P}\} \equiv \{\mathcal{E}, \mathcal{S}, \mathcal{F}\}$  for all environments.

A protocol  $\mathcal{P}$  in the iUC framework is specified via a system of machines  $\{M_1, \dots, M_l\}$ ; the framework offers a convenient template for the specification of such systems. Each machine  $M_i$  implements one or more roles of the protocol, where a role describes a piece of code that performs a specific task. For example, a (real) protocol  $\mathcal{P}_{\text{sig}}$  for digital signatures might contain a `signer` role for signing messages and a `verifier` role for verifying signatures. In a run of a protocol, there can be several instances of every machine, interacting with each other (and the environment) via I/O interfaces and interacting with the adversary (and possibly the environment subsuming a network attacker) via network interfaces. An instance of a machine  $M_i$  manages one or more so-called *entities*. An entity is identified by a tuple  $(pid, sid, role)$  and describes a specific party with party ID (PID)  $pid$  running in a session with session ID (SID)  $sid$  and executing some code defined by the role  $role$  where this role has to be (one of) the role(s) of  $M_i$  according to the specification of  $M_i$ . Entities can send messages to and receive messages from other entities and the adversary using the I/O and network interfaces of their respective machine instances. More specifically, the I/O interfaces of both machines need to be connected to each other (because one machine specifies the other as a subroutine) to enable communication between entities of those machines.

Roles of a protocol can be either public or private. The I/O interfaces of private roles are only accessible by other (entities belonging to) roles of the same protocol, whereas I/O interfaces of public roles can also be accessed by other (potentially unknown) protocols/the environment. Hence, a private role models some internal subroutine that is protected from access outside of the protocol, whereas a public role models some publicly accessible operation that can be used by other protocols. One uses the syntax “ $(\text{pubrole}_1, \dots, \text{pubrole}_n \mid \text{privrole}_1, \dots, \text{privrole}_n)$ ” to uniquely determine public and private roles of a protocol. Two protocols  $\mathcal{P}$  and  $\mathcal{Q}$  can be combined to form a new more complex protocol as long as their I/O interfaces connect only via their public roles. In the context of the new combined protocol, previously private roles remain private while previously public roles may either remain public or be considered private, as determined by the protocol designer. The set of all possible combinations of  $\mathcal{P}$  and  $\mathcal{Q}$ , which differ only in the set of public roles, is denoted by  $\text{Comb}(\mathcal{Q}, \mathcal{P})$ .

An entity in a protocol might become corrupted by the adversary, in which case it acts as a pure message forwarder between the adversary and any connected higher-level protocols as well as subroutines. In addition, an entity might also consider itself (implicitly) corrupted while still following its own protocol because, e.g., a subroutine has been corrupted. Corruption of entities in the iUC framework is highly customizable; one can, for example, prevent corruption of certain entities during a protected setup phase.

The iUC framework supports the modular analysis of protocols via a so-called composition theorem:

**Corollary 5** (Concurrent composition in iUC; informal). *Let  $\mathcal{P}$  and  $\mathcal{F}$  be two protocols such that  $\mathcal{P} \leq \mathcal{F}$ . Let  $\mathcal{Q}$  be another protocol such that  $\mathcal{Q}$  and  $\mathcal{F}$  can be connected. Let  $\mathcal{R} \in \text{Comb}(\mathcal{Q}, \mathcal{P})$  and let  $\mathcal{I} \in \text{Comb}(\mathcal{Q}, \mathcal{F})$  such that  $\mathcal{R}$  and  $\mathcal{I}$  agree on their public roles. Then  $\mathcal{R} \leq \mathcal{I}$ .*

By this theorem, one can first analyze and prove the security of a subroutine  $\mathcal{P}$  independently of how it is used later on in the context of a more complex protocol. Once we have shown that  $\mathcal{P} \leq \mathcal{F}$  (for some other, typically ideal protocol  $\mathcal{F}$ ), we can then analyze the security of a higher-level protocol  $\mathcal{Q}$  based on  $\mathcal{F}$ . Note that this is simpler than analyzing  $\mathcal{Q}$  based on  $\mathcal{P}$  directly as ideal protocols provide absolute security guarantees while typically also being less complex, reducing the potential for errors in proofs. Once we have shown that the combined protocol, say,  $(\mathcal{Q} \mid \mathcal{F})$  realizes some other protocol, say,  $\mathcal{F}'$ , the composition theorem and transitivity of the  $\leq$  relation then directly implies that this also holds true if we run  $\mathcal{Q}$  with an implementation  $\mathcal{P}$  of  $\mathcal{F}$ . That is,  $(\mathcal{Q} \mid \mathcal{P})$  is also a secure realization of  $\mathcal{F}'$ . Please note that the composition theorem does not impose any restrictions on how the protocols  $\mathcal{P}$ ,  $\mathcal{F}$ , and  $\mathcal{Q}$  look like internally. For example, they might have disjoint sessions, but they could also freely share some state between sessions, or they might be a mixture of both. They can also freely share some of their subroutines with the environment, modeling so-called globally available state. This is unlike most other models for universal composability, such as the UC model, which impose several conditions on the structure of protocols for their composition theorem.

## B. Notation in Pseudo Code

ITMs in our paper are specified in pseudo code. Most of our pseudo code notation follows the notation introduced by Camenisch et al. [17]. To ease readability of our figures, we provide a brief overview over the used notation here.

The description in the main part of the ITMs consists of blocks of the form **recv**  $\langle msg \rangle$  **from**  $\langle sender \rangle$  **to**  $\langle receiver \rangle$  **s.t.**  $\langle condition \rangle$ : $\langle code \rangle$  where  $\langle msg \rangle$  is an input pattern,  $\langle sender \rangle$  is the receiving interface (I/O or NET),  $\langle receiver \rangle$  is the dedicated receiver of the message and  $\langle condition \rangle$  is a condition on the input.  $\langle code \rangle$  is the (pseudo) code of this block. The block is executed if an incoming message matches the pattern and the condition is satisfied. More specifically,  $\langle msg \rangle$  defines the format of the message  $m$  that invokes this code block. Messages contain local variables, state variables, strings, and maybe special characters. To compare a message  $m$  to a message pattern  $msg$ , the values of all global and local variables (if defined) are inserted into the pattern. The resulting pattern  $p$  is then compared to  $m$ , where uninitialized local variables match with arbitrary parts of the message. If the message matches the pattern  $p$  and meets  $\langle condition \rangle$  of that block, then uninitialized local variables are initialized with the part of the message that they matched to and  $\langle code \rangle$  is executed in the context of  $\langle receiver \rangle$ ; no other blocks are executed in this case. If  $m$  does not match  $p$  or  $\langle condition \rangle$  is not met, then  $m$  is compared with the next block. Usually a **recv from** block ends with a **send to** clause of form **send**  $\langle msg \rangle$  **to**  $\langle sender \rangle$  where  $\overline{msg}$  is a message that is send via output interface  $sender$ .

If an ITM invokes another ITM, e. g., as a subroutine, ITMs may expect an immediate response. In this case, in a **recv from** block, a **send to** statement is directly followed by a **wait for** statement. We write **wait for**  $\langle msg \rangle$  **from**  $\langle sender \rangle$  **s.t.**  $\langle condition \rangle$  to denote that the ITM stays in its current state and discards all incoming messages until it receives a message  $m$  matching the pattern  $\overline{msg}$  and fulfilling the **wait for** condition. Then the ITM continues the run where it left of, including all values of local variables.

To clarify the presentation and distinguish different types of variables, constants, strings, etc. we follow the naming conventions of Camenisch et al. [17]:

1. (Internal) state variables are denoted by sans-serif fonts.
2. Local (i.e., ephemeral) variables are denoted in *italic font*.
3. Keywords are written in **bold font** (e. g., for operations such as **sending** or **receiving**).
4. Commands, procedure, function names, strings and constants are written in teletype.

To increase readability, we use the following notation:

- For a set of tuples  $K$ ,  $K.add(\_)$  adds the tuple to  $K$ .
- For a string  $S$ ,  $S.add(\_)$  concatenates the given string to  $S$ .
- For a verdicts  $v_1$  and  $v_2$ , we define  $v_1.add(v_2) := v_1 \wedge v_2$ .
- $K.remove(\_)$  removes always the first appearance of the given element/string from the list/tuple/set/string  $K$ .

We use the following additional nomenclature from [17]:

- $(pid_{cur}, sid_{cur}, role_{cur})$  denotes the currently active entity and  $(pid_{call}, sid_{call}, role_{call})$  denotes the entity which called the currently active ITM.
- The macro **corr** $(pid, sid, role)$  is simply a shortcut to invoke the ITM of  $(pid, sid, role)$  and query it for its corruption status.
- The macro **init** $(pid, sid, role)$  triggers the initialization of  $(pid, sid, role)$  and returns the activation to the calling ITM.

## C. Reusing Existing Security Results with AUC

In this section, we briefly discuss how existing security results can be reused when transforming an ideal functionality with AUC. Therefore, we consider an ideal functionality  $\mathcal{F}$  which ensures some security properties  $Sec$  and its accountable transformation  $\mathcal{F}^{acc}$  (cf. Definition 1) which may replace some of  $\mathcal{F}$ 's security properties with the corresponding assumption-based and/or accountability properties, i. e.,  $Sec^{abs} \cup Sec^{assumption} \cup Sec^{acc} = Sec$ . According to Definition 1,  $\mathcal{F}$  and  $\mathcal{F}^{acc}$  behave identical as long as no security property is marked as broken. Thus, an existing security proof remains valid until this case occurs. Thus, one often can reuse the existing security proof as is (up to the point where some property breaks). As AUC is a non-back-box transformation, one needs to enhance proofs in order to capture the behavior when transforming  $\mathcal{F}$  to  $\mathcal{F}^{acc}$ .

We further note that accountability/assumption-based and preventive security properties can be heavily intertwined. Therefore one may also need to adjust the existing simulator to correctly handle broken properties of  $\mathcal{F}_{cp}^{acc}$ .

We essentially illustrate this approach in our case study in Section 3.3.

#### D. Scaling Accountable Consensus (Full Details)

In this section, we provide full details regarding the scalable accountable consensus case study presented in Section 3.1. In particular, we provide (i) a full specification of  $\mathcal{F}_{cp}^{acc}$ , (ii) a full specification of the scaling protocol, and (iii) provide a formal proof for Theorem 1.

**D.1 The Accountable Consensus Functionality  $\mathcal{F}_{cp}^{acc}$ :** In this section, we present a full specification of the ideal accountable consensus service  $\mathcal{F}_{cp}^{acc}$  in Figure 9 and 10 including its subroutine  $\mathcal{F}_{judgeParams}^{acc-cp}$  in Figure 11. For technical details and notation specific to the iUC framework, see the brief summary of the iUC framework in Section A.

$\mathcal{F}_{cp}^{acc}$  models one instance of a consensus protocol per *sid*.  $\mathcal{F}_{cp}^{acc}$ 's duty is to allow clients to write to a stream of messages (also called the state of  $\mathcal{F}_{cp}^{acc}$ ) and read from this stream. For write operations,  $\mathcal{F}_{cp}^{acc}$  expects messages of the form (Submit, *msg*) where *msg* denotes some arbitrary bit string. Submitted transactions/messages do not enter the state of  $\mathcal{F}_{cp}^{acc}$  immediately. Firstly, they get an unambiguous temporary ID and are then stored in a buffer for later processing - this is necessary to capture network artifacts. Besides the submitted message, the buffer also contains the round, resp. point in time (provided by  $\mathcal{F}_{cp}^{acc}$  internal clock, see below), when the message was added to the buffer. This will later be used to enforce liveness. The submitted message as well as its temporary ID are leaked to  $\mathcal{A}$ . Note that the Submit interface is accessible via I/O and NET. The adversary  $\mathcal{A}$  is allowed to read the buffer via the getBuffer command. Further,  $\mathcal{A}$  is allowed to permute the content of the buffer, i.e.,  $\mathcal{A}$  may interchange the order/IDs in the buffer. To add the first element from buffer to  $\mathcal{F}_{cp}^{acc}$ 's state,  $\mathcal{A}$  can use the Update command. When an entry from buffer is moved to state,  $\mathcal{F}_{cp}^{acc}$  assigns an incremental counter as ID to the message and removes the message from the buffer. That is,  $\mathcal{F}_{cp}^{acc}$  establishes a total order over the transactions in state.

When a party (*pid*, *sid*, *role*) wants to read  $\mathcal{F}_{cp}^{acc}$ 's current state, it sends Read via I/O to  $\mathcal{F}_{cp}^{acc}$ . Analogously to submit requests,  $\mathcal{F}_{cp}^{acc}$  buffers the request in *buffer<sub>read</sub>* including an unambiguous ID *id* and the current time.  $\mathcal{A}$  is then expected to trigger the delivery of the response to the request via (Deliver, *id*). If the accountability property consistency is already broken,  $\mathcal{A}$  can freely determine  $\mathcal{F}_{cp}^{acc}$ 's responses to the read request. To do so,  $\mathcal{F}_{cp}^{acc}$  queries  $\mathcal{A}$  for the response and forwards the output to the requestor. If consistency still holds true,  $\mathcal{F}_{cp}^{acc}$  requests  $\mathcal{A}$  for the prefix of the state it should provide to the requestor. More specifically,  $\mathcal{A}$  provides a pointer *N* to the prefix of the state and a potential receiver (*pid<sub>r</sub>*, *sid<sub>cur</sub>*, *role*) to  $\mathcal{F}_{cp}^{acc}$ . If there exists an entry in *buffer<sub>read</sub>* matching *id*,  $\mathcal{F}_{cp}^{acc}$  ignores the receiver provided by  $\mathcal{A}$  and sends the prefix of the state to the requesting party stored in *buffer<sub>msg</sub>* (for *id*) and deletes the read request from *buffer<sub>read</sub>*. If there is no corresponding message to *id* in *buffer<sub>read</sub>*,  $\mathcal{F}_{cp}^{acc}$  pushes the prefix of the state to (*pid<sub>r</sub>*, *sid<sub>cur</sub>*, *role<sub>r</sub>*).

As already mentioned,  $\mathcal{F}_{cp}^{acc}$  also includes an (internal) clock which allows measuring time to capture liveness properties. Parties and  $\mathcal{A}$  can query the current time of  $\mathcal{F}_{cp}^{acc}$  via the GetCurRound command. Further,  $\mathcal{A}$  is allowed to increase the round, resp. current time, in  $\mathcal{F}_{cp}^{acc}$  by one via UpdateRound. In case that liveness still holds true,  $\mathcal{F}_{cp}^{acc}$  checks that there does not exist an entry (*id*, *msg*, *r*) in buffer such that it holds true  $round - r > \delta$  where *round* is the current round in  $\mathcal{F}_{cp}^{acc}$  and  $\delta$  is the liveness parameter of  $\mathcal{F}_{cp}^{acc}$ . If liveness is already broken, there are no restrictions on time updates.

In  $\mathcal{F}_{cp}^{acc}$ , we use the standard corruption behavior for dynamic corruption without secure erasures of iUC. In particular,  $\mathcal{F}_{cp}^{acc}$  only acts as message forwarder for corrupted parties. Further, we model a dynamic set of participants in  $\mathcal{F}_{cp}^{acc}$ .

$\mathcal{F}_{cp}^{acc}$  also relies on the concepts of AUC in order to capture and handle accountability properties. We expect  $\mathcal{F}_{cp}^{acc}$  to be used with the parameters  $Sec^{acc} = \{consistency\}$  and  $Sec^{assumption} = \{liveness\}$ . For the security properties,  $\mathcal{F}_{judgeParams}^{acc-cp}$  enforces that  $\mathcal{A}$  has to trade in a (public) verdict for  $pid_j = public$ , i.e., we expect the parameter  $pids_{judge}$  to be  $\{public\}$ , which ensures individual accountability if he breaks consistency. To capture that liveness may break for all participating parties, we expect  $ids_{assumption}$  to be  $\{public\}$ .  $\mathcal{F}_{judgeParams}^{acc-cp}$  does not provide additional leakage to  $\mathcal{A}$  if he breaks liveness. As judicial report  $\mathcal{F}_{judgeParams}^{acc-cp}$  outputs a prefix of  $\mathcal{F}_{cp}^{acc}$ 's state including an extension provided by  $\mathcal{A}$ . In particular, the prefix needs to contain at least

Description of the ideal and accountable consensus service  $\mathcal{F}_{cp}^{acc} = (\text{client}, \text{judge}, \text{supervisor})$ :

<b>Participating roles:</b> {client, judge, supervisor}	
<b>Corruption model:</b> dynamic corruption without secure erasures	
<b>Protocol parameters:</b>	
- $\delta \in \mathbb{N}$	{The expected liveness guarantee for in time units}
- $\text{Sec}^{acc} \subset \{0, 1\}^*$	{Accountability properties}
- $\text{Sec}^{assumption} \subset \{0, 1\}^*$	{Assumption-based security properties}
- $\text{pids}_{judge} \subset \{0, 1\}^*$	{set of judge entities/(P)IDs in the protocol (which are often directly related to some protocol participants)}
- $\text{ids}_{assumption} \subset \{0, 1\}^*$	{set of entities/IDs where properties are ensured via assumptions}

Description of  $M_{cp}^{acc}$ :

<b>Implemented role(s):</b> {client}	
<b>Subroutines:</b> $\mathcal{F}_{judgeParams} : \text{judgeParams}$	
<b>Internal state:</b>	
- state $\subset \mathbb{N} \times \{0, 1\}^*$ , state = $\emptyset$	{The set of totally ordered transactions/messages}
- buffer $\subset \mathbb{N} \times \{0, 1\}^* \times \mathbb{N}$ , buffer = $\epsilon$	{The buffer of submitted transactions including submission round}
- buffer <sub>read</sub> $\subset \mathbb{N}^2 \times (\{0, 1\}^*)^3$ , buffer <sub>read</sub> = $\emptyset$	{Buffer for handling Read requests, entries of form (id, round, (pid, sid, role))}
- counter, counter <sub>B</sub> , counter <sub>R</sub> $\in \mathbb{N}$ , counter = 0, counter <sub>B</sub> = 0, counter <sub>R</sub> = 0	{The counter for ordered transactions, to order tx in buffer, and read requests in buffer <sub>read</sub> }
- round $\in \mathbb{N}$ , round = 0	{The time in $\mathcal{F}_{cp}^{acc}$ }
- maxHonestOutput $\in \mathbb{N}$ , maxHonestOutput = 0	{The longest prefix of the state queried by an honest party}
- corruptedIntParties $\in \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \setminus (\text{Roles}_{\mathcal{F}}^a \cup \{\text{judge}, \text{supervisor}\})$ , initially $\emptyset$	{The set of corrupted internal parties (pid, sid, role)}
- brokenAssumptions : $\text{Sec}^{assumption} \times \text{ids}_{assumption} \rightarrow \{\text{true}, \text{false}\}$	{Stores broken security assumptions per id, initially false $\forall$ entries}
- brokenProps : $(\text{Sec}^{assumption} \cup \text{Sec}^{acc}) \times (\text{pids}_{judge} \cup \text{ids}_{assumption}) \rightarrow \{\text{true}, \text{false}\}$	{Stores broken security properties per judge/id, initially false $\forall$ entries}
- verdicts : $\text{pids}_{judge} \rightarrow \{0, 1\}^*$	{Verdicts per $p \in \text{pids}_{judge}$ , initially $\epsilon$ }
<b>CheckID</b> (pid, sid, role):	
Accept all messages for the same sid. For messages to (pid, sid, judge) accept only if pid = public.	
<b>Corruption behavior:</b>	
- <b>AllowCorruption</b> (pid, sid, role):	
Do not allow corruption of (pid, sid, supervisor).	
if role = judge:	
send (Corrupt, (pid, sid, judge), internalState)	
to (pid, sid, $\mathcal{F}_{judgeParams} : \text{judgeParams}$ )	
wait for b; return b	{ $\mathcal{F}_{judgeParams}$ decides whether judges can be corrupted}
- <b>DetermineCorrStatus</b> (pid, sid, role):	
if role = judge:	
send (CorruptionStatus?, (pid, sid, judge), internalState) to (pid, sid, $\mathcal{F}_{judgeParams} : \text{judgeParams}$ )	{ $\mathcal{F}_{judgeParams}$ may determine a judge's corruption status}
wait for b; return b	
- <b>AllowAdvMessage</b> (pid, sid, role, pid <sub>receiver</sub> , sid <sub>receiver</sub> , role <sub>receiver</sub> , m)	
Do not allow sending messages to $\mathcal{F}_{judgeParams}$ .	{ $\mathcal{A}$ is not allowed to invoke $\mathcal{F}_{judgeParams}$ in the name of corrupted parties.}
<b>Main:</b>	
<b>recv</b> (Submit, msg):	{Transaction submission}
counter <sub>B</sub> $\leftarrow$ counter <sub>B</sub> + 1; buffer.add([counter <sub>B</sub> , msg, round])	
send (Submit, counter <sub>B</sub> , msg) to NET	{Leak submitted data to $\mathcal{A}$ }
<b>recv</b> Read from I/O:	{Read request}
counter <sub>R</sub> $\leftarrow$ counter <sub>R</sub> + 1; buffer <sub>read</sub> .add(counter <sub>R</sub> , round, (pid <sub>cur</sub> , sid <sub>cur</sub> , role <sub>cur</sub> ))	{Record read request including current time}
send (Read, counter <sub>R</sub> , round, (pid <sub>cur</sub> , sid <sub>cur</sub> , role <sub>cur</sub> )) to NET	{Leak information to $\mathcal{A}$ }
<b>recv</b> (Deliver, id) from NET:	{ $\mathcal{A}$ triggers delivery of response to read request or injects a response without a read request}
if brokenProps[consistency, public] = true:	{Check whether $\mathcal{F}_{cp}^{acc}$ still provides consistency}
send responsively Read to NET	{If consistency is "broken" $\mathcal{A}$ is allowed to determine the output of $\mathcal{F}_{acc}$ }
wait for (Read, state, (pid, sid, role))	
if $\exists (id, round, (pid, sid, role)) \in \text{buffer}_{msg}$ :	{Check that ID exists}
maxHonestOutput $\leftarrow$  state	{Record longest output of the state so far}
buffer <sub>read</sub> .remove([id, round, (pid, sid, role)]); send (Read, state) to (pid, sid, role)	
else:	{If consistency holds, output the current state}
send responsively Read to NET (*)	{Allow $\mathcal{A}$ to determine the prefix of state as response}
wait for (Read, N, (pid <sub>r</sub> , sid <sub>cur</sub> , role <sub>r</sub> ))	{ $\mathcal{A}$ decides on the prefix of the state. If there is no read request for id, $\mathcal{A}$ also determines the recipient.}
if  state  < N:	
Go to (*)	
if $\exists (id, round, (pid, sid, role)) \in \text{buffer}_{msg}$ :	{Check that ID exists}
buffer <sub>read</sub> .remove([id, round, (pid, sid, role)]); Let state be the prefix of state up to and including entry N.	
send (Read, state) to (pid, sid, role)	{Send state to requestor}
else:	
send (Read, state) to (pid <sub>r</sub> , sid <sub>cur</sub> , role <sub>r</sub> )	{Send state the entity determined by $\mathcal{A}$ }
<b>recv</b> Read from NET:	{ $\mathcal{A}$ is allowed to access the full state of $\mathcal{F}_{cp}^{acc}$ .}
reply (Read, state)	
<b>recv</b> getBuffer from NET:	{ $\mathcal{A}$ is allowed to query the current buffer}
reply (getBuffer, buffer)	
<b>recv</b> (Permute, $\pi$ ) from NET:	{ $\mathcal{A}$ is allowed to change the order in the buffer}
require: $\pi$ is a permutation of the IDs in buffer.	
for all (ctr, msg) $\in$ buffer do: buffer.remove([ctr, msg]); buffer.add([\pi(ctr), msg])	
reply ack	

<sup>a</sup>Roles <sub>$\mathcal{F}$</sub>  = {client} here.

Fig. 9: The ideal and accountable consensus service  $\mathcal{F}_{cp}^{acc}$  including accountability (Part 1).

Description of  $M_{cp}^{acc}$  (continued):

**Main:**

**recv** Update from NET: {A may trigger when buffer is included in state}  
 Let  $(i, msg) \in \text{buffer}$  the first entry in buffer;  $\text{counter} \leftarrow \text{counter} + 1$ ;  $\text{state.add}([\text{counter}, msg])$ ;  $\text{buffer.remove}(i, msg)$   
{Remove entry from buffer}

**reply** ack

**recv** UpdateRound from NET: {A triggers round update}

**if**  $(\exists (\_, \_, r) \in \text{buffer},$   
**s.t.**  $\text{round} - r > \delta \vee \exists (\_, r, \_) \in \text{buffer}_{read},$   
**s.t.**  $\text{round} - r > \delta) \wedge \text{brokenProps}[\text{liveness}, \text{public}] = \text{false}$ :

**reply** (UpdateRound, false,  $\epsilon$ ) {Reject round update if there are old messages in queue but liveness still holds}

**else:**

$\text{round} \leftarrow \text{round} + 1$

**reply** (UpdateRound, true,  $\epsilon$ )

**recv** GetCurRound: {Allow access to the "global clock"}

**reply** (GetCurRound, round)

Include static code from the AUC transformation  $\mathcal{T}_1(\cdot)$  here, i.e., include additional code from Figure 2 and 3 here.

Fig. 10: The ideal and accountable consensus service  $\mathcal{F}_{cp}^{acc}$  including accountability (Part 2).

Description of  $\mathcal{F}_{judgeParams}^{acc-cp} = (\text{judgeParams})$ :

**Participating roles:** {judgeParams}  
**Corruption model:** incorruptible

Description of  $M_{judgeParams}^{acc-aCS}$ :

**Implemented role(s):** {judgeParams}

**Internal state:**  
 $\text{ptr} \in \mathbb{N}, \text{ptr} = 0$  {Pointer which stores the length of the prefix of the state reported as judicial report}

**CheckID**( $pid, sid, role$ ):  
 Accept all messages with the same  $sid$ .

**Main:**

**recv** (BreakAccProp,  $verdict, toBreak, internalState$ ) from I/O:  
**if**  $verdict[\text{public}]$  ensures individual accountability)  $\wedge$   
 all other entries in  $verdict$  map to  $\epsilon \wedge$   
 $toBreak = \{(\text{consistency}, \text{public})\}$ : {Handle violation of accountability w.r.t. consistency}  
**reply** (BreakAccProp, true,  $\epsilon$ )

**else:**

**reply** (BreakAssumption, false,  $\epsilon$ )

**recv** (BreakAssumption,  $toBreak, internalState$ ) from I/O: {Do not generate leakage when breaking assumptions}

**reply** (BreakAssumption,  $\epsilon$ )

**recv** (GetJudicialReport,  $msg, internalState$ ) from I/O: {Generate judicial report}

**send responsively** (GetJudicialReport,  $msg, internalState$ ) to NET (\*) {Forward request to A an wait for a pointer to the state including an extension of this state}

**wait for** (GetJudicialReport,  $ptr, ext$ )

**if**  $ptr < \text{maxHonestOutput} \vee ptr > |state| \vee$   
 $ext \notin \mathbb{N} \times \{0, 1\}^*$  {Require valid input from A, prefix must at least contain honest output and extension needs to be valid}  
 $ext \neg$  start with index  $ptr + 1 \vee$   
 further elements  $\neg$  consecutive enumerated:  
 go to (\*)

$ptr \leftarrow ptr$   
 Let  $state_{ptr}$  the prefix of the  $state$  from the beginning including the  $ptr$ ' component unite with  $ext$ .  
**reply** (GetJudicialReport,  $state_{ptr}$ ) {Return state as report}

**recv** (Corrupt, (public,  $sid, judge$ ),  $internalState$ ) from I/O:  $\{\mathcal{F}_{judgeParams}^{acc-cp}$  declines corruption requests for the public judge  
**reply** false {The public judge is incorruptible}

**recv** (CorruptionStatus?(public,  $sid, judge$ ),  $internalState$ ) from I/O:  $\{\mathcal{F}_{judgeParams}^{acc-cp}$  asks for the public judge's corruption  
**reply** false {The public judge is incorruptible}

Fig. 11: The judge parameter functionality  $\mathcal{F}_{judgeParams}^{acc-cp}$  for  $\mathcal{F}_{cp}^{acc}$ .

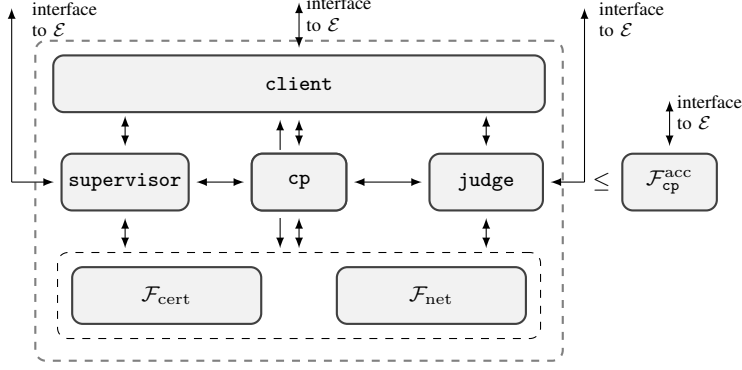


Fig. 12: Possible realization relation of the consensus service. The system  $\mathcal{E}$  denotes the environment, modeling, as usual in UC setting, arbitrary higher level protocols. All machines are additionally connected to the adversary.

the prefix of the state that was already outputted to honest parties. As we model a public judge here, the judge is incorruptible. Thus,  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}}$ , answers false to `Corrupt` and `CorruptionStatus?` queries.

In what follows, we will assume that we already have a realization for  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  which we scale via the scaling protocol. Figure 12 gives an overview over a centralized consensus service CS (including relevant subroutines) in the client-server model which can be used to realize  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  and includes (public) accountability w.r.t. consistency and assumption-based liveness. Importantly, the consensus service CS has to sign answers to read request from its clients in this case. Thus, the (public) judge can gather undeniable evidence whether CS misbehaved according based on CS signatures and can ensure accountability w.r.t. consistency. As already explained in Section 3.1, communication via  $\mathcal{F}_{\text{net}}$  allows including breakable liveness guarantees. Of course,  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  can also be realized by a distributed system. For example, the accountable version of Apache Kafka [4] as presented in [50] is a good candidate to realize  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ .

**D.2 The Scaling Protocol:** In this section, we provide some additional details regarding the consensus scaling protocol  $\mathcal{P}_{\text{cp}}^{\text{acc}}$  we use in our example introduced in Section 3.1. We provide formal definitions for  $\mathcal{P}_{\text{cp}}^{\text{acc}}$  in Figures 13 to 18.

**Remark:** To simplify presentation in Section 3.1, we introduced  $\mathcal{P}_{\text{cp}}^{\text{acc}}$  as  $(\text{client}, \text{supervisor}, \text{judge} \mid \text{scd}, \mathcal{F}_{\text{net}}, \mathcal{F}_{\text{cp}}^{\text{acc}}, \mathcal{F}_{\text{cert}})$ . Formally, the scaling protocol is defined as:

$$\mathcal{P}_{\text{cp}}^{\text{acc}} = (\mathcal{P}_{\text{client}}^{\text{CD}} : \text{client}, \mathcal{P}_{\text{sv}}^{\text{CD}} : \text{supervisor}, \mathcal{P}_{\text{judge}}^{\text{CD}} : \text{judge} \mid \mathcal{P}_{\text{CD}} : \text{scd}, \mathcal{F}_{\text{net}} : \text{net}, \mathcal{F}_{\text{cp}}^{\text{acc}} : \text{client}, \mathcal{F}_{\text{cert}}).$$

**Client  $\mathcal{P}_{\text{client}}^{\text{CD}}$ .** Clients of the scaling protocol  $\mathcal{P}_{\text{client}}^{\text{CD}}$  (cf. Figure 13) are the connection point between higher-level protocols/the environment and the scaling protocol. An instance of  $\mathcal{P}_{\text{client}}^{\text{CD}}$  models one party accessing the scaling protocol. The environment may queue messages/transactions via  $\mathcal{P}_{\text{client}}^{\text{CD}}$  for consensus and read (from) the established totally ordered sequence of transactions. If an instance  $(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \text{role}_{\text{cur}})$  of  $\mathcal{P}_{\text{client}}^{\text{CD}}$  receives a `Submit` command, it firstly queries  $\mathcal{A}$  which CD it should use to submit the transaction. Then, it assigns a unique ID to the message and signs  $(id, tx, \text{pid}_{\text{cur}})$  via  $\mathcal{F}_{\text{cert}}$  and sends the submit request  $(id, tx, \text{pid}_{\text{cur}}, \sigma)$  (where  $\sigma$  is the signature from above) via  $\mathcal{F}_{\text{net}}$  to the determined CD.

If the environment triggers a read request at an instance  $(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \text{role}_{\text{cur}})$  of  $\mathcal{P}_{\text{client}}^{\text{CD}}$ ,  $(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \text{role}_{\text{cur}})$  first queries  $\mathcal{A}$  which CD to use to process the request. Then, it forwards the request via  $\mathcal{F}_{\text{net}}$  to the CD.  $\mathcal{P}_{\text{client}}^{\text{CD}}$  also receives the response to the query via  $\mathcal{F}_{\text{net}}$ . To pull the response from the CD from  $\mathcal{F}_{\text{net}}$ ,  $\mathcal{A}$  triggers  $\mathcal{P}_{\text{client}}^{\text{CD}}$  with the `Pull` command.  $\mathcal{P}_{\text{client}}^{\text{CD}}$  then queries  $\mathcal{F}_{\text{net}}$  for new messages and verifies that the state in the delivered message is validly signed. In this case,  $\mathcal{P}_{\text{client}}^{\text{CD}}$  forwards the message to the judge  $\mathcal{P}_{\text{judge}}^{\text{CD}}$ , waits for its reactivation, and then forwards the response to the environment.

**The consensus distributors  $\mathcal{P}_{\text{CD}}$ .** Though the CD are the core of the scaling protocol, they are internal/private parties which use an internal/private consensus service to achieve consistency/consensus. Internal/private means that the interfaces of  $\mathcal{P}_{\text{CD}}$  are not accessible to higher-level protocols. An instance of  $\mathcal{P}_{\text{CD}}$  (cf. Figure 14) models one dedicated consensus distributor. Communication to/from the environment is always via  $\mathcal{P}_{\text{client}}^{\text{CD}}$ .

Messages from  $\mathcal{P}_{\text{client}}^{\text{CD}}$  to  $\mathcal{P}_{\text{CD}}$  are sent via  $\mathcal{F}_{\text{net}}$  and vice versa. If  $\mathcal{A}$  triggers  $\mathcal{P}_{\text{CD}}$  via the Pull message,  $\mathcal{P}_{\text{CD}}$  pulls for a new message at  $\mathcal{F}_{\text{net}}$ . (i) If the pulled message contains a Submit command,  $\mathcal{P}_{\text{CD}}$  verifies the signature of the client submitting the transaction. If the signature verifies, it queues the submit request at  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  for consensus. In particular,  $\mathcal{P}_{\text{CD}}$  forwards the request to its client component of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ . (ii) If the pulled message is a Read request,  $\mathcal{P}_{\text{CD}}$  creates a signature over state and replies with the state including the created signature via  $\mathcal{F}_{\text{net}}$  to the requestor.

If  $\mathcal{A}$  wants  $\mathcal{P}_{\text{CD}}$  to update its state, it sends Update to  $\mathcal{P}_{\text{CD}}$ . Thereupon,  $\mathcal{P}_{\text{CD}}$  sends a Read request to  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ .  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  replies with a Read message including a state update.  $\mathcal{P}_{\text{CD}}$  verifies that the state update has a valid format and that it is an extension of its current state. If both checks succeed,  $\mathcal{P}_{\text{CD}}$  overwrites its current state with the provided update.

**The public judge  $\mathcal{P}_{\text{judge}}^{\text{CD}}$ .** The public judge  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  (cf. Figure 15) collects information (i) from clients  $\mathcal{P}_{\text{client}}^{\text{CD}}$  and (ii) from the lower-level consensus service  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ . In contrast to clients which push evidence data to  $\mathcal{P}_{\text{judge}}^{\text{CD}}$ , the judge pulls data from the lower-level public judge of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  on every activation during preprocessing. More specifically, as long as  $\text{verdicts} = \varepsilon$ ,  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  queries  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ 's judge for the most recent verdicts and a judicial report and overwrites its current (empty) verdicts and judicialReport with the received data. The judge only processes evidence in the case that there is no verdict so far.

If  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  receives evidence from  $\mathcal{P}_{\text{client}}^{\text{CD}}$  which contains some CD's state (we assume that CD is identified with  $(pid, sid, role)$ ) including its  $pid$  and a signature over the data.  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  checks that the signature verifies (which implicitly includes a check that  $pid$  is indeed a CD). If the signature does not verify, the reported evidence is ignored. Otherwise,  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  checks that the reported state is a prefix of the most recent judicialReport. If this check fails, this means that  $(pid, sid, role)$  violated consistency. Thus,  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  will blame  $(pid, sid, role)$  for its misbehavior.

On request GetVerdict,  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  outputs the verdict stored in verdicts. If the environment invokes the GetJudicialReport interface,  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  outputs the judicial report it gathered from  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ 's judge which is stored in judicialReport.

As  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  models a public judge, it also provides the GetEvidence interface to  $\mathcal{A}$ .  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  provides all collected evidence ( $W$ ), the current judicial report, and its full transcript to  $\mathcal{A}$  if  $\mathcal{A}$  invokes this interface. This models that all data used by  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  is indeed public (as  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  makes it public).  $\mathcal{P}_{\text{judge}}^{\text{CD}}$ 's VerResult interface further allows  $\mathcal{A}$  to verify signatures of the parties of  $\mathcal{P}_{\text{cp}}^{\text{acc}}$ .

**The public supervisor  $\mathcal{P}_{\text{sv}}^{\text{CD}}$ .** The supervisor  $\mathcal{P}_{\text{sv}}^{\text{CD}}$  (cf. Figure 16) is responsible for determining whether the assumption-based security property liveness still holds true and to make the corruption status of internal parties/parties from lower-level protocols accessible to the environment. If the environment queries  $\mathcal{P}_{\text{sv}}^{\text{CD}}$  via the (IsAssumptionBroken?, liveness, public),  $\mathcal{P}_{\text{sv}}^{\text{CD}}$  asks  $\mathcal{F}_{\text{net}}$  via getDeliveryStatus whether liveness still holds true and also queries  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ 's public supervisor and asks whether liveness still holds true for  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ . If both checks succeed,  $\mathcal{P}_{\text{sv}}^{\text{CD}}$  responds that liveness still holds true for the scaling protocol. Otherwise, it reports the property as broken.

If the environment queries for the corruption status of internal parties of the protocol via corruptInt?,  $\mathcal{P}_{\text{sv}}^{\text{CD}}$  checks whether the matching instance of  $\mathcal{P}_{\text{CD}}$  is corrupted or whether the party at  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ , resp. the internal party at  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  is corrupted and outputs the corruption status accordingly.

**The network functionality  $\mathcal{F}_{\text{net}}$ .** The ideal network functionality  $\mathcal{F}_{\text{net}}$  (cf. Figure 18) with breakable liveness mainly provides a Send and a Pull interface to parties. There should be one instance of  $\mathcal{F}_{\text{net}}$  per  $sid$  modeling one separated network for a protocol. In order to measure whether liveness is broken,  $\mathcal{F}_{\text{net}}$  includes an in-built clock (see below). When an entity sends a message via  $\mathcal{F}_{\text{net}}$  to another entity,  $\mathcal{F}_{\text{net}}$  firstly buffers the message including a unique ID, sender, receiver, and the round  $\mathcal{F}_{\text{net}}$  received the message in  $\text{buffer}_{\text{msg}}$ . Additionally, the message (including all metadata) is leaked to  $\mathcal{A}$ .

In order to change the sequence of messages in  $\text{buffer}_{\text{msg}}$ ,  $\mathcal{A}$  is allowed to permute the messages in  $\text{buffer}_{\text{msg}}$  via Permute analogously to the Permute command in  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ .

If a party  $(pid, sid, role)$  wants to access its messages, it can query  $\mathcal{F}_{\text{net}}$  via Pull. In this case,  $\mathcal{F}_{\text{net}}$  looks for the message with the smallest ID addressed to  $(pid, sid, role)$ , deletes this message from  $\text{buffer}_{\text{msg}}$ , and forwards the message to  $(pid, sid, role)$ .

To increase the current time stored in round,  $\mathcal{A}$  can use the UpdateRound command. In the case that liveness still holds true, i.e.,  $\text{inTime} = \text{true}$ ,  $\mathcal{F}_{\text{net}}$  checks whether the round received of all messages in  $\text{buffer}_{\text{msg}}$  is



less than  $\delta$  time units in the past (where  $\delta$  is the liveness parameter of  $\mathcal{F}_{\text{net}}$ ). If the check succeeds,  $\mathcal{F}_{\text{net}}$  increases round by one. Otherwise, the round update is rejected. If liveness does not hold true anymore,  $\mathcal{F}_{\text{net}}$  accepts all UpdateRound requests of  $\mathcal{A}$ .

Via BreakLiveness,  $\mathcal{A}$  can break the liveness guarantees of  $\mathcal{F}_{\text{net}}$ . The other subroutines can query for the time of  $\mathcal{F}_{\text{net}}$  via GetCurRound and they can also query via getDeliveryStatus whether liveness holds true or not. (Note that  $\mathcal{F}_{\text{net}}$  is private in the considered protocol, thus not accessible to the environment).

**The ideal signature functionality  $\mathcal{F}_{\text{cert}}$  with in-built CA.** In the scaling protocol model, we use  $\mathcal{F}_{\text{cert}}$  (cf. Figure 17)– an adapted version of the standard ideal signature functionality  $\mathcal{F}_{\text{sig}}$  (cf. Figure 31) as presented in [17] (which we use later on as well). The major difference between  $\mathcal{F}_{\text{cert}}$  and  $\mathcal{F}_{\text{sig}}$  is that  $\mathcal{F}_{\text{cert}}$  includes an “in-built” CA. This simplifies the handling of signatures in this model.

This concludes the description of  $\mathcal{P}_{\text{cp}}^{\text{acc}}$ . For a description of the corruption behavior, we point to Section 3.1.

**D.3 UC Security Analysis:** We now present the full proof of Theorem 1 in detail. To ease notation, we use  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  for the ideal consensus service subroutine on the real side/hybrid setting and call the instance of the ideal consensus service we want to realize  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$ .

**Theorem 6.** *Let  $\delta \in \mathbb{N}$  be the upper time boundary for message delivery and  $\Sigma = (\text{gen}(1^n), \text{sig}, \text{ver})$  be an EUF-CMA secure signature scheme. Let  $\mathcal{P}_{\text{cp}}^{\text{acc}}$  be as defined above. Let  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  as described above with  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}}$  as subroutine with parameters  $\text{Sec}^{\text{acc}} = \{\text{consistency}\}$ ,  $\text{Sec}^{\text{assumption}} = \{\text{liveness}\}$ ,  $\text{pids}_{\text{judge}} = \{\text{public}\}$ , and  $\text{ids}_{\text{assumption}} = \{\text{public}\}$ . Then:*

$$\mathcal{P}_{\text{cp}}^{\text{acc}} \leq (\mathcal{F}_{\text{cp}}^{\text{acc}'} \mid \mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}})$$

*Proof.* We firstly define a responsive simulator  $\mathcal{S}$  such that the real world running the protocol  $\mathcal{R} := \mathcal{P}_{\text{cp}}^{\text{acc}} = (\mathcal{P}_{\text{client}}^{\text{CD}} : \text{client}, \mathcal{P}_{\text{sv}}^{\text{CD}} : \text{supervisor}, \mathcal{P}_{\text{judge}}^{\text{CD}} : \text{judge} \mid \mathcal{P}_{\text{CD}} : \text{scd}, \mathcal{F}_{\text{net}} : \text{net}, \mathcal{F}_{\text{cp}}^{\text{acc}} : \text{client}, \mathcal{F}_{\text{cert}})$  is indistinguishable from the ideal world running  $\{\mathcal{S}, \mathcal{I}\}$ , with the protocol  $\mathcal{I} := (\mathcal{F}_{\text{cp}}^{\text{acc}'} \mid \mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}})$ , for every ppt environment  $\mathcal{E}$ .

The simulator  $\mathcal{S}$  is defined as follows: as common,  $\mathcal{S}$  is a single machine. It is connected to  $\mathcal{I}$  and the environment  $\mathcal{E}$  via their network interfaces. In a run, there is only a single instance of the machine  $\mathcal{S}$  that accepts and processes all incoming messages. The simulator  $\mathcal{S}$  internally simulates the realization  $\mathcal{R}$ , including its behavior on the network interface connected to the environment, and uses this simulation to compute responses to incoming messages. For ease of presentation, we will refer to this internal simulation by  $\mathcal{R}'$ . More precisely, the simulation runs as follows:

*Network communication from/to the environment*

- Messages that  $\mathcal{S}$  receives on the network connected to the environment (and which are hence meant for  $\mathcal{R}$ ) are forwarded to the internal simulation  $\mathcal{R}'$ .
- Any messages sent by  $\mathcal{R}'$  on its network interface (that are hence meant for the environment) are forwarded to the environment  $\mathcal{E}$ .

*Corruption handling*

- The simulator  $\mathcal{S}$  keeps the corruption status of entities in  $\mathcal{R}'$  and  $\mathcal{I}$  synchronized. That is, whenever an entity of  $\mathcal{P}_{\text{client}}^{\text{CD}}, \mathcal{P}_{\text{CD}}, \mathcal{F}_{\text{cp}}^{\text{acc}}$ , or an internal entity of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  in  $\mathcal{R}'$  starts to consider itself corrupted, the simulator first corrupts the corresponding (internal) entity of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  in  $\mathcal{I}$  before continuing its simulation. Note that corruption of internal entities, i.e., of  $\mathcal{P}_{\text{CD}}$ , or (non-internal) parties of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  is mapped to a corruption of an internal party in  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  in  $\mathcal{I}$ .
- Incoming Messages from corrupted (non-internal) entities of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  in  $\mathcal{I}$  are forwarded on the network to the environment in the name of the corresponding entity  $\mathcal{P}_{\text{client}}^{\text{CD}}$  in  $\mathcal{R}'$ . Conversely, whenever a corrupted entity of  $\mathcal{P}_{\text{client}}^{\text{CD}}$  wants to output a message to a higher-level protocol,  $\mathcal{S}$  instructs the corresponding entity of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  to output the same message to the higher-level protocol.
- For indirectly corrupted parties:<sup>9</sup> Note that these parties are (directly) corrupted in  $\mathcal{I}$ . For these instances,  $\mathcal{I}$  forwards the inputs from  $\mathcal{S}$ . An indirectly corrupted party in  $\mathcal{R}'$  produces an output on an I/O tape then  $\mathcal{S}$  instructs  $\mathcal{I}$  to forward this message to the intended receiver.

<sup>9</sup>Indirectly corrupted parties are not directly corrupted parties, i.e.,  $\mathcal{A}$  did not send Corrupt to the party so far, but (in our case) DetermineCorrStatus outputs true. That is, the party considers itself to be corrupted, e.g., due to a corrupted subrouting. For example, an instance of  $\mathcal{P}_{\text{client}}^{\text{CD}}$  considers itself corrupted if its accompanied signer session at  $\mathcal{F}_{\text{cert}}$  is corrupted.

### Transaction submission

Whenever an honest entity  $entity = (pid, sid, role)$  of  $\mathcal{F}_{cp}^{acc'}$  receives a request  $(Submit, msg)$  to submit a new transaction  $msg$ ,  $\mathcal{F}_{cp}^{acc'}$  buffers the message for later adding it to the state and leaks the full message including the ID of the message in  $\mathcal{F}_{cp}^{acc'}$ 's buffer to  $\mathcal{S}$ .  $\mathcal{S}$  uses the leaked message to simulate the input of  $msg$  to  $\mathcal{P}_{client}^{CD}$  in  $\mathcal{R}'$ .

In the case that a corrupted party submits a message (which can be extracted from  $\mathcal{F}_{net}$ ),  $\mathcal{S}$  submits the transaction via the Submit command to  $\mathcal{F}_{cp}^{acc'}$ .

### Read requests

Whenever an honest entity  $(pid, sid, role)$  receives a request Read to read from the global state,  $\mathcal{F}_{cp}^{acc'}$  buffers this read request and leaks the request to  $\mathcal{S}$  and waits for a trigger to process the read request. Upon receiving this request,  $\mathcal{S}$  uses the leaked message to simulate the input of this message to  $\mathcal{P}_{client}^{CD}$  in  $\mathcal{R}'$  and stores (i) the ID provided by  $\mathcal{F}_{cp}^{acc'}$  for further processing and (ii) which entity triggered the read request.

If  $\mathcal{A}$  triggers a simulated entity  $(pid, sid, role)$  of  $\mathcal{P}_{client}^{CD}$  to output the response to a read request via Pull,  $\mathcal{S}$  simulates the input of Pull in  $\mathcal{R}'$ . If  $(pid, sid, role)$  wants to output a state via a  $(Read, state)$  message on its I/O connection in  $\mathcal{R}'$ ,  $\mathcal{S}$  extract the connected request ID  $id$  matching the request (see above) and sends  $(Deliver, id)$  to  $\mathcal{F}_{cp}^{acc'}$ . If  $brokenProps[consistency, public] = true$  in  $\mathcal{F}_{cp}^{acc'}$ ,<sup>10</sup>  $\mathcal{S}$  simply forwards  $(Read, state, (pid_r, sid_{cur}, role_r))$  from  $\mathcal{R}'$  including the stored receiver as response to  $\mathcal{F}_{cp}^{acc'}$  Read request. Otherwise,  $\mathcal{S}$  compares  $state$  with state from  $\mathcal{F}_{cp}^{acc'}$ <sup>11</sup> and extracts the pointer  $N$ , such that state up to the  $N^{\text{th}}$  entry matches  $state$ .  $\mathcal{S}$  then forwards  $(Read, N, (pid_r, sid_{cur}, role_r))$  to  $\mathcal{F}_{cp}^{acc'}$  where  $(pid_r, sid_{cur}, role_r)$  is the recorded requestor of read request.

### State updates

If there occurs a state update in  $\mathcal{R}'$  in the simulation of  $\mathcal{F}_{cp}^{acc}$ , i.e., the variable state of  $\mathcal{F}_{cp}^{acc}$  is updated,  $\mathcal{S}$  triggers a state update at  $\mathcal{F}_{cp}^{acc'}$ . Therefore, it firstly retrieves the buffer of submitted messages/transactions which are not part of the state so far via `getBuffer` and the state' from  $\mathcal{F}_{cp}^{acc'}$  via `Read`. It checks, which entries of state from the simulation are missing in state' in  $\mathcal{F}_{cp}^{acc'}$ . It then derives a permutation over the received buffer such that the entries missing in state' have the same order as in state starting with the lowest ID in the buffer. Then,  $\mathcal{S}$  sends this permutation via `Permute` to  $\mathcal{F}_{cp}^{acc'}$ .  $\mathcal{F}_{cp}^{acc'}$  activates  $\mathcal{S}$  again with an `ack` message. Thereupon,  $\mathcal{S}$  sends `Update` to  $\mathcal{F}_{cp}^{acc'}$  and is activated again via `ack`.  $\mathcal{S}$  repeats sending the `Update` message until state and state' are equal.

### Public judge and supervisor

- If during the simulation of  $\mathcal{P}_{judge}^{CD}$  in  $\mathcal{R}'$ ,  $\mathcal{P}_{judge}^{CD}$  renders a verdict, i.e.,  $verdicts \neq \varepsilon$ ,  $\mathcal{S}$  extracts the verdict and forwards it to  $\mathcal{F}_{cp}^{acc'}$ . More specifically,  $\mathcal{S}$  sends  $(BreakAccProp, verdict, \{(consistency, public)\})$  to  $(public, sid_{cur}, \mathcal{F}_{cp}^{acc} : judge)$  where  $verdict$  is a map from  $pids_{judge} \rightarrow \{0, 1\}^*$  where only the entry  $verdict[public] = verdicts$  (contains the verdict from  $\mathcal{P}_{judge}^{CD}$ ) and all other entries are mapped to  $\varepsilon$ .
- If  $\mathcal{F}_{cp}^{acc'}$  is queried for a judicial report, it allows  $\mathcal{S}$  to determine some details of the judicial report which is derived from  $\mathcal{F}_{cp}^{acc'}$ 's state. If  $\mathcal{S}$  receives  $(GetJudicialReport, msg, internalState)$  from  $\mathcal{F}_{cp}^{acc'}$ , it extracts the judicial report  $report$  (as currently produced) from  $\mathcal{F}_{cp}^{acc'}$ . It extracts the necessary pointer  $ptr$  to map  $report$  to state from  $\mathcal{F}_{cp}^{acc}$  and then sends  $(GetJudicialReport, ptr, \varepsilon)$  back to  $\mathcal{F}_{cp}^{acc'}$ .
- If  $\mathcal{S}$  receives `GetEvidence` or `VerResult` via NET, it forwards this input to its simulated version of  $\mathcal{P}_{judge}^{CD}$  and sends the output of the simulated machine back via NET. For the transcript provided as response to `GetEvidence`,  $\mathcal{S}$  maps data such as if the simulated version of  $\mathcal{P}_{judge}^{CD}$  would be directly connected to NET, etc.
- As soon as liveness breaks (i) in  $\mathcal{F}_{net}$  or (ii) in  $\mathcal{F}_{cp}^{acc}$  in  $\mathcal{R}'$ ,  $\mathcal{S}$  breaks liveness at  $\mathcal{F}_{cp}^{acc'}$ . In the case of (i), as soon as  $inTime = false$  in  $\mathcal{F}_{net}$ ,  $\mathcal{S}$  sends  $(BreakAssumption, \{(liveness, public)\})$  to  $(public, sid_{cur}, \mathcal{F}_{cp}^{acc} : supervisor)$ . In Case (ii), if the  $brokenProps[liveness, public] = true$  at  $\mathcal{F}_{cp}^{acc}$ ,  $\mathcal{S}$  sends the same message as above to  $\mathcal{F}_{cp}^{acc'}$ .

### Further details

- $\mathcal{S}$  keeps the clocks/rounds of  $\mathcal{R}'$  and  $\mathcal{F}_{cp}^{acc'}$  synchronous. That is,  $\mathcal{S}$  sends `UpdateRound` to  $\mathcal{F}_{cp}^{acc'}$  whenever a round update in the simulated  $\mathcal{F}_{net}$  is performed and before continuing the simulation.

This concludes the description of the simulator. It is easy to see that (i)  $\{\mathcal{S}, \mathcal{I}\}$  is environmentally bounded<sup>12</sup>

<sup>10</sup>Note that  $\mathcal{S}$  explicitly breaks this property. Thus,  $\mathcal{S}$  knows the value of  $brokenProps[consistency, public]$  in  $\mathcal{F}_{cp}^{acc'}$ .

<sup>11</sup>This is possible, as  $\mathcal{S}$  determines the state via `Update` commands.

<sup>12</sup>This is the polynomial runtime notion employed by the iUC framework.

and (ii)  $\mathcal{S}$  is a responsive simulator for  $\mathcal{I}$ , i.e., restricting messages from  $\mathcal{I}$  are answered immediately as long as  $\{\mathcal{S}, \mathcal{I}\}$  runs with a responsive environment. We now argue that  $\mathcal{R}$  and  $\{\mathcal{S}, \mathcal{I}\}$  are indeed indistinguishable for any (responsive) environment  $\mathcal{E} \in \text{Env}(\mathcal{R})$ .

Let  $\mathcal{E} \in \text{Env}(\mathcal{R})$  be an arbitrary but fixed environment. First, observe that  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  provides  $\mathcal{S}$  with full information about all submit and read requests performed by higher-level protocols. Hence, the simulated protocol  $\mathcal{R}'$  within  $\mathcal{S}$  obtains the same inputs and thus performs identical to the real world  $\mathcal{R}$ . Therefore, the network behavior simulated by  $\mathcal{S}$  towards the environment is indistinguishable from the network behavior of  $\mathcal{R}$ . Note that state changes triggered via network interface are synchronized between  $\mathcal{R}'$  and  $\mathcal{I}$ . Together with the state synchronization during I/O interaction (see below), the simulator can keep the states of  $\mathcal{R}'$  and  $\mathcal{I}$  in synchronization. Moreover, we can also conclude that the corruption status of (internal) entities in the real and ideal world is synchronized. Since the simulator has full control over corrupted entities, which are handled via the internal simulation  $\mathcal{R}'$ , this implies that the I/O behavior of corrupted entities of  $\mathcal{R}/\mathcal{I}$  towards higher-level protocols/the environment is also identical in the real and ideal world. Note that purely internal (private) parties have no interface to the environment. Thus, the only way to potentially distinguish the real and ideal world is the I/O behavior of honest entities of  $\mathcal{R}/\mathcal{I}$  towards higher-level protocols.

We will now go over all possible interactions with honest entities on the I/O interface and argue, by induction, that all of those interactions result in identical behavior towards the environment, i.e., are also indistinguishable. At the start of a run, there were no interactions on the I/O interface with honest parties yet. In the following, assume that all I/O interactions so far have resulted in the same behavior visible towards the environment in both the real and ideal world.

**Submission requests:** Submission requests do not directly result into an output to the environment. But, they might affect the output of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  at a later point in time as they have direct impact of the state of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$ , resp.  $\mathcal{R}$ . Thus, we now show that submit requests behave “identical”, i.e., we have to argue that these changes are “synchronized” between  $\mathcal{I}$  and  $\mathcal{R}'$ . In particular, the buffered set of transactions in  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  is a subset of the buffered set of transactions in  $\mathcal{R}'$  (which are buffered (i) as Submit message in  $\mathcal{F}_{\text{net}}$  and (ii) in the buffer of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ ).

Observe that, upon receiving a submission request,  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  behaves similar to  $\mathcal{R}$ : in  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$ , the submitted transaction is directly stored into the buffer. In  $\mathcal{R}'$ , the submit message is forwarded via  $\mathcal{F}_{\text{net}}$  to  $\mathcal{P}_{\text{CD}}$ . So, the “buffers” of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  and  $\mathcal{R}'$  stay synchronized when defining the buffer of  $\mathcal{R}'$  as the set of Submit messages in  $\mathcal{F}_{\text{net}}$ ’s buffer union the buffer of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ . Note that as soon as submit request of corrupted parties are visible to  $\mathcal{S}$  (as specified above),  $\mathcal{S}$  also submits these transactions to  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$ . Thus,  $\mathcal{S}$  is indeed able to keep the buffers in  $\mathcal{I}$  and  $\mathcal{R}$  synchronized.

**State updates:** As buffers in  $\mathcal{R}'$  and  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  are synchronized, we can conclude that the states of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  and  $\mathcal{R}'$  stay synchronized during updates as this boils down to mirror the state from  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  to  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  (as  $\mathcal{S}$  does). (i) The buffer of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  is a subset of the buffer of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$ . (ii) Submitted transaction which are not in the intersection of the states of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  and  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  are currently in the buffer of  $\mathcal{F}_{\text{net}}$ . (iii) By construction, both provide the same security guarantees. Thus,  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  provides at most less guarantees than  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  (in case that properties are broken at  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$ ). Thus,  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  will accept every state update of  $\mathcal{S}$  and both,  $\mathcal{I}$  and  $\mathcal{R}'$  stay synchronized. Note that  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  also provides  $\mathcal{S}$  the possibility to forward answers to read requests which have not been requested. As this is possible in real and ideal world, both worlds stay indistinguishable.

**Read requests:** Observe that, upon receiving a read request,  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  stores the read request in a buffer and leaks the full request to the simulator including an ID and the receiver of the response. Note that the procedure is analogously to a submission requests. A read request does not directly result into an output to the environment. With the same argumentation as above follows that  $\mathcal{S}$  is able to keep the “read buffers” of  $\mathcal{I}$  and  $\mathcal{R}'$  synchronized.

**Deliver response to read request:** If  $\mathcal{S}$  is triggered to output a response to a read request (via Pull), there are two cases to consider: (i) `brokenProps[consistency, public] = true` or (ii) `brokenProps[consistency, public] = false`. In the first case,  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  will simply forward the input from  $\mathcal{S}$  to the requestor. As the output is extracted from  $\mathcal{R}'$  we can conclude that  $\mathcal{R}$  and  $\mathcal{I}$  are indistinguishable. In the second case,  $\mathcal{S}$  can extract the response from  $\mathcal{R}'$ . Honest clients in  $\mathcal{R}'$  receive a validly signed prefix of the state of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  (which is in synchronization with  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$ ’s state). Thus,  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  will accept the Deliver command of  $\mathcal{S}$  and the outputs in both worlds is identical. Thus, both worlds are indistinguishable.

**Verdicts:** Firstly, note that  $\mathcal{S}$  sends a `BreakAccProp` message to break public consistency when consistency in  $\mathcal{R}'$  breaks, resp. when there is a new verdict available in  $\mathcal{R}'$ 's public judge.<sup>13</sup> Secondly, note that this request matches the rules in  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}}$ : the verdict  $\mathcal{S}$  provides ensures individual accountability and  $\mathcal{S}$  solely tries to break public consistency. Thus,  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  accepts the verdict  $\mathcal{S}$  provides. Further note that we assume the used signature scheme to be EUF-CMA secure. Thus, the probability that an honest party receives and accepts a state update that includes a forged signature (and forwards this to a judge) is negligible. Further, as verifier is incorruptible, Thus, we can conclude that  $\mathcal{S}$  can always keep the verdicts and whether consistency is broken synchronized in  $\mathcal{I}$  and  $\mathcal{R}'$  and the output on a `GetVerdict` request is equal in  $\mathcal{R}$  and  $\mathcal{I}$ . Thus, both worlds are/stay indistinguishable.

**Judicial Reports:** Observe that the states of  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  in  $\mathcal{R}'$  and  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  are synchronized (see above). As the judicial report is derived as a prefix of the state,  $\mathcal{S}$  always provides a pointer to the full state of both  $\mathcal{R}'$  and  $\mathcal{I}$  to determine the judicial report of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$ . Thus, delivered judicial reports are equal in both worlds and  $\mathcal{R}$  and  $\mathcal{I}$  stay indistinguishable.

**GetEvidence and VerResult:** We note that the environment cannot use the `GetEvidence` or the `VerResult` interface of  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  to distinguish between real and ideal world. The activation order of the simulated  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  matches the real world order and also messages to  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  match in both worlds perfectly. As  $\mathcal{S}$  maps the transcript such that it matches the real world execution the transcript (and provided state) is indistinguishable in both worlds. Similarly,  $\mathcal{S}$  fully simulates all interaction with  $\mathcal{F}_{\text{cert}}$  (in the same order in both worlds). Thus, requests in both worlds will lead to the same output.

**Supervisor:** The same holds true for `BreakAssumption` requests to the supervisor. The message of  $\mathcal{S}$  matches the rules and thus will be accepted. This allows  $\mathcal{S}$  to keep both worlds synchronized regarding the assumption-based security property liveness. Thus, both worlds will answer requests to the supervisor identically, i. e., both worlds remain indistinguishable.

**Time requests:** As the simulator updates the internal clock of  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  every time  $\mathcal{F}_{\text{net}}$ 's clock in  $\mathcal{R}'$  is increased, both worlds always output the same value for the current time. Note that, even if  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  enforces *liveness*, any round update requests of  $\mathcal{S}$  will indeed be accepted: as  $\mathcal{S}$  keeps the liveness guarantees of  $\mathcal{I}$  and  $\mathcal{R}'$  synchronized,  $\mathcal{S}$  will not ask for a time update that violates the liveness check in  $\mathcal{F}_{\text{cp}}^{\text{acc}'}$  as long as liveness still holds true.

We also note that the (additional) public evidence  $\mathcal{A}$ , resp.  $\mathcal{E}$ , may gather at the  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  (via `GetEvidence`) does not enable  $\mathcal{E}$  to distinguish between both worlds. As  $\mathcal{R}'$  gets the same input as  $\mathcal{I}/\mathcal{R}$ , the output of `GetEvidence` in  $\mathcal{R}'$  is the same as in the real world  $\mathcal{R}$ . We emphasize that the judicial report of  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  already includes most of the information  $\mathcal{E}$  receives via `GetEvidence`. The transcript is the same in  $\mathcal{I}$  and  $\mathcal{R}$  as the order of activations of and the inputs to  $\mathcal{P}_{\text{judge}}^{\text{CD}}$ , resp. its simulated version, are identically. Additionally, `GetEvidence` provides signatures to  $\mathcal{E}$ . However, as signatures do not contain private information and are ideally generated in both worlds, it is not possible to distinguish between  $\mathcal{R}$  and  $\mathcal{I}$  based on the additional signature data.

Altogether,  $\mathcal{R}$  and  $\{\mathcal{S}, \mathcal{I}\}$  behave identical in terms of behavior visible to the environment  $\mathcal{E}$  and thus are indistinguishable. □

**D.4 Deterrence Analysis:** Common consensus services, e. g., Hashgraph [11], charge fees for their service. Clients also benefit from a consensus service as they do not have to establish a consensus service on their own. Thus, we conclude  $U_{hP}^i > 0$  for all parties involved in the consensus service. As  $\mathcal{P}_{\text{cp}}^{\text{acc}}$  provides public accountability and verdicts are fair (by definition), we can conclude that the public judge will not falsely accuse an honest party. Further, all data used in the judging procedure is public. Thus, we can assume that  $U_{hE}^i = U_{hL}^i = 0$ . To deter parties from misbehavior, penalties for violating consistency need to be sufficiently high. Thus, we assume that  $U_{mL}^i \gg 0$ . In particular, penalties need to negate potential profit due to malicious behavior, i. e., it needs to hold true  $U_{mP}^i - U_{mL}^i < U_{hP}^i$ . The exact value of  $U_{mP}^i$  and hence the minimally required fees depend on the context that uses the consensus protocol. As the judging procedure would identify faked evidence, we assume  $U_{mE}^i = 0$ . Thus, Equation 1 and 2 hold true.

<sup>13</sup>We remark that  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  does not update her verdict.

Description of the client  $\mathcal{P}_{\text{client}}^{\text{CD}} = (\text{client})$  of the consensus distributors:

**Participating roles:**  $\{\text{client}\}$

**Corruption model:** *Dynamic corruption without secure erasures*

Description of  $M_{\text{client}}$ :

**Implemented role(s):**  $\{\text{client}\}$

**Subroutines:**  $\mathcal{F}_{\text{cert}} : \text{verifier}, \mathcal{F}_{\text{cert}} : \text{signer}, \mathcal{F}_{\text{net}} : \text{net}, \mathcal{P}_{\text{judge}}^{\text{CD}} : \text{judge}$

**Internal state:**

- $\text{owner} \in (\{0, 1\}^*)^3$  {The entity  $(pid, sid, role)$  that has access to this instance of  $\mathcal{P}_{\text{client}}^{\text{CD}}$ }
- $\text{serial} \in \mathbb{N}, \text{serial} = 0$  {The serial for "agreed" transactions}

**CheckID** $(pid, sid, role)$ :

Accept all messages for the same  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})$ .

**Corruption behavior:**

**DetermineCorrStatus** $(pid, sid, role)$ :

```

if corrupted = true: {Checks whether a party itself is corrupted.}
  return true
  corrRes  $\leftarrow$  corr(pidcur, sidcur, signer) {Request corruption status at  $\mathcal{F}_{\text{cert}}$ }
return corrRes {Return whether  $\mathcal{F}_{\text{cert}}$  instance is corrupted}

```

**Initialization:**

$\text{owner} \leftarrow (pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}})$  {Record that only  $(pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}})$  has access to this instance of  $\mathcal{P}_{\text{client}}^{\text{CD}}$ }

**Main:**

**recv** (Submit,  $tx$ ) **from** I/O **s.t.**  $(pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}}) = \text{owner}$ : {Transaction submission with specified CD}  
**send responsively** (Submit,  $tx$ ) **to** (NET) { $\mathcal{A}$  decides which CD  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})$  will use}

**wait for** (Submit,  $tx, pid_{\text{CD}}$ )

$\text{serial} \leftarrow \text{serial} + 1$

$\sigma' \leftarrow \text{sign}(\text{serial}, tx, pid_{\text{cur}})$

**send** (Send,  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}}), (pid_{\text{CD}}, sid_{\text{CD}}, scd), (\text{Submit}, \text{serial}, tx, pid_{\text{cur}}, \sigma)$ ) **to**  $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{net}} : \text{net})$  {Send transaction via  $\mathcal{F}_{\text{net}}$  to  $\mathcal{P}_{\text{CD}}$ }

**recv** Read **from** I/O **s.t.**  $(pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}}) = \text{owner}$ :

**send responsively** Read **to** (NET) {Read request from IO}

**wait for** (Read,  $pid_{\text{CD}}$ )

**send** (Send,  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}}), (pid_{\text{CD}}, sid_{\text{CD}}, scd), (\text{Read}, pid_{\text{cur}})$ ) **to**  $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{net}} : \text{net})$  {Forward request via  $\mathcal{F}_{\text{net}}$ }

**recv** Pull **from** NET:

**send** Pull **to**  $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{net}} : \text{net})$

**wait for** (Pull,  $msg$ )

**if**  $msg = (\text{Read}, state, pid_{\text{CD}}, \sigma)$ :

$b \leftarrow \text{verifySig}(pid_{\text{CD}}, state, \sigma)$

**if**  $b \wedge state \subset state$ :

$state \leftarrow state$

**send** (Evidence,  $pid_{\text{CD}}, state, \sigma$ ) **to**  $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{P}_{\text{judge}}^{\text{CD}} : \text{judge})$

**wait for** ack

**send** (Read,  $state$ ) **to** owner

{Output response of read request to owner}

**Procedures and Functions:**

**function** sign( $msg$ ) :

**send** (sign,  $msg$ ) **to**  $(pid_{\text{cur}}, (pid_{\text{cur}}, sid_{\text{cur}}), \mathcal{F}_{\text{cert}} : \text{signer})$

**wait for** (sign,  $\sigma$ )

**return**  $\sigma$

{Sign message at  $\mathcal{F}_{\text{cert}}$ }

**function** verifySig( $pid, msg, \sigma$ ) :

**send** (VerResult,  $msg, \sigma$ ) **to**  $(pid_{\text{cur}}, (pid, sid_{\text{cur}}), \mathcal{F}_{\text{cert}} : \text{verifier})$

**wait for** (VerResult,  $result$ )

**return**  $result$

{Verify signature at  $\mathcal{F}_{\text{cert}}$ }

Fig. 13: The consensus distributors client  $\mathcal{P}_{\text{client}}^{\text{CD}}$ .

Description of the consensus distributors  $\mathcal{P}_{CD} = (\text{scd})$ :

<b>Participating roles:</b> $\{\text{scd}\}$ <b>Corruption model:</b> <i>Dynamic corruption without secure erasures</i>
Description of $M_{\text{scd}}$ :
<b>Implemented role(s):</b> $\{\text{scd}\}$ <b>Subroutines:</b> $\mathcal{F}_{\text{cp}}^{\text{acc}} : \text{client}, \mathcal{F}_{\text{cert}} : \text{verifier}, \mathcal{F}_{\text{cert}} : \text{signer}, \mathcal{F}_{\text{net}} : \text{net}, \mathcal{P}_{\text{judge}}^{\text{CD}} : \text{judge}$ <b>Internal state:</b> - $\text{state} \subset \mathbb{N} \times \{0, 1\}^*, \text{state} = \emptyset$ <span style="float: right;">{The set of totally ordered transactions/messages}</span> <b>CheckID(<math>pid, sid, role</math>):</b> Accept all messages for the same $sid$ . <b>Corruption behavior:</b> <b>DetermineCorrStatus(<math>pid, sid, role</math>):</b> <b>if</b> corrupted = true: <span style="float: right;">{Checks whether a party itself is corrupted.}</span> <b>return</b> true $\text{corrRes}_1 \leftarrow \text{corr}(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{cert}} : \text{signer})$ <span style="float: right;">{Request corruption status at <math>\mathcal{F}_{\text{cert}}</math>}</span> $\text{corrRes}_2 \leftarrow \text{corr}(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{cp}}^{\text{acc}} : \text{client})$ <span style="float: right;">{Request corruption status at <math>\mathcal{F}_{\text{cp}}^{\text{acc}}</math>}</span> <b>return</b> $\text{corrRes}_1 \vee \text{corrRes}_2$ <span style="float: right;">{Return whether <math>\mathcal{F}_{\text{cert}}</math> or <math>\mathcal{F}_{\text{cp}}^{\text{acc}}</math> instance is corrupted}</span> <b>Main:</b> <b>recv Pull from NET:</b> <span style="float: right;">{A triggers <math>pid_{\text{cur}}</math> to query <math>\mathcal{F}_{\text{net}}</math>}</span> <b>send</b> Pull <b>to</b> $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{net}} : \text{net})$ <b>wait for</b> (Pull, $msg$ ) <b>if</b> $msg = (\text{Submit}, \text{serial}, tx, pid, \sigma)$ : <span style="float: right;">{Received message is a Submit}</span> <b>if</b> $b \leftarrow \text{verifySig}(pid, (\text{serial}, tx, pid), \sigma)$ : <span style="float: right;">{Check valid signature}</span> <b>if</b> $b$ : <span style="float: right;">{If signature is valid, submit data to consensus service}</span> <b>send</b> (Submit, ( $\text{serial}, tx, pid, \sigma$ )) <b>to</b> $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{cp}}^{\text{acc}} : \text{client})$ <span style="float: right;">{Forward tx to consensus service}</span> <b>else if</b> $msg = (\text{Read}, pid)$ : <span style="float: right;">{In case of a Read request}</span> $\sigma \leftarrow \text{sign}(\text{state})$ <span style="float: right;">{Sign current state}</span> <b>send</b> (Send, $(pid_{\text{cur}}, sid_{\text{cur}}, \text{role}_{\text{cur}}), (pid, sid_{\text{cur}}, \text{client}), (\text{Read}, \text{state}, pid_{\text{cur}}, \sigma)$ ) <b>to</b> $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{net}} : \text{net})$ <span style="float: right;">{Reply with full state including signature}</span> <span style="float: right;">{Request update at consensus service}</span> <span style="float: right;">{Forward request to consensus service}</span> <span style="float: right;">{State update from consensus service}</span> <b>recv</b> (Read, $state$ ) <b>from</b> $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{cp}}^{\text{acc}} : \text{client})$ : <b>require:</b> $state$ is a set with entries of form $(ctr, msg)$ , where $ctr \in \mathbb{N}, msg \in \{0, 1\}^*$ . <b>if</b> $\text{state} \subset state$ : $state \leftarrow state$ <b>Procedures and Functions:</b> <b>function</b> $\text{sign}(msg)$ : <span style="float: right;">{Sign message at <math>\mathcal{F}_{\text{cert}}</math>}</span> <b>send</b> (sign, $msg$ ) <b>to</b> $(pid_{\text{cur}}, (pid_{\text{cur}}, sid_{\text{cur}}), \mathcal{F}_{\text{cert}} : \text{signer})$ <b>wait for</b> (sign, $\sigma$ ) <b>return</b> $\sigma$ <b>function</b> $\text{verifySig}(pid, msg, \sigma)$ : <span style="float: right;">{Verify signature at <math>\mathcal{F}_{\text{cert}}</math>}</span> <b>send</b> (VerResult, $msg, \sigma$ ) <b>to</b> $(pid_{\text{cur}}, (pid, sid_{\text{cur}}), \mathcal{F}_{\text{cert}} : \text{verifier})$ <b>wait for</b> (VerResult, $result$ ) <b>return</b> $result$

Fig. 14: The consensus distributor ITM  $\mathcal{P}_{CD}$  which scales consensus.

### E. Accountable Key Exchange based on an accountable PKI (Full Details)

In this section, we provide full details regarding the accountable PKI case study as presented in Section 3.2. This section is structured as follows: we firstly present the details regarding the accountable PKI as introduced in Section 3. In particular, we provide a detailed introduction to the ideal accountable PKI functionality  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . Afterwards, we explain the protocol  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  and provide a formal specification of the model. We conclude this part with a security analysis of the accountable PKI.

**E.1 An Accountable Ideal PKI Functionality  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ :** We now present the full specification of the ideal accountable PKI functionality  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  including the formal specification of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  in Figure 19 including its subroutine  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}}$  in Figure 20.  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  extends Canetti et al.'s  $\mathcal{G}_{\text{BB}}$  as presented in [27].

One session of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  models one instance of a PKI including several different CAs and CTLs. From a high-level perspective,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  offers the possibility to register certificates for identities at a dedicated CA and to retrieve certificates of identities from CAs. We model that parties use a local PKI client which is responsible for the communication with the PKI, i.e., the client is connected via I/O to the environment modeling higher-level protocols which access/use the PKI. To interact with the PKI, honest parties, resp. clients identified by  $(pid, sid, \text{client})$ , can register a certificate for their own identity. To register a certificate, clients send a message of the form (Register,  $pk, pid_{\text{CA}}$ ) (where  $pid_{\text{CA}}$  is the PID of the CA that should attest the entities identity and  $pk$  be a string, typically the public key of the client) to  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ .  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  does not immediately process the

Description of  $\mathcal{P}_{\text{judge}}^{\text{CD}} = (\text{judge})$ :

**Participating roles:**  $\{\text{judge}\}$   
**Corruption model:** *incorruptible*

Description of  $M_{\text{judge}}$ :

**Implemented role(s):**  $\{\text{judge}\}$   
**Subroutines:**  $\mathcal{F}_{\text{cert}}$  : verifier,  $\mathcal{F}_{\text{cp}}^{\text{acc}}$  : judge  
**Internal state:**  
- counter  $\in \mathbb{N}$ , counter = 0 {Evidence counter}  
-  $W \subset \mathbb{N} \times (\{0, 1\}^*)^2$  {Reported evidence (enumerated) per party, entries of form (ctr, pid, state)}  
- judicialReport  $\in \{0, 1\}^*$ , judicialReport =  $\varepsilon$  {The judicial report from  $\mathcal{F}_{\text{cp}}^{\text{acc}}$ }  
- verdicts  $\in \{0, 1\}^*$ , verdicts =  $\varepsilon$  {Recorded verdict}  
**CheckID**( $pid, sid, role$ ):  
Accept all messages with the same  $sid$  addressed to (public,  $sid$ , judge).  
**MessagePreprocessing:**  
**if** verdicts =  $\varepsilon$ :  
  **send** GetVerdict **to** (public,  $sid_{\text{cur}}, \mathcal{F}_{\text{cp}}^{\text{acc}}$  : judge) {Query for verdicts regarding consensus service}  
  **wait for** (GetVerdict, verdict)  
  verdicts  $\leftarrow$  verdict  
**if** verdicts =  $\varepsilon$ : {Update judicial report}  
  **send** GetJudicialReport **to** (public,  $sid_{\text{cur}}, \mathcal{F}_{\text{cp}}^{\text{acc}}$  : judge)  
  **wait for** (GetJudicialReport, judicialReport)  
  judicialReport  $\leftarrow$  judicialReport  
**Main:**  
**recv** (Evidence,  $pid_{\text{CD}}, state, \sigma$ ) **from** I/O: {Evidence from clients}  
   $b \leftarrow$  verifySig( $pid_{\text{CD}}, state, \sigma$ )  
  **if**  $b$ :  
    **if**  $state \not\subset$  judicialReport:  
      verdicts.add(dis( $pid_{\text{CD}}, sid_{\text{cur}}, \mathcal{P}_{\text{CD}} : \text{scd}$ )) {Honest CD will always forward state which is a subset of a judicial report}  
    **reply** ack  
  **recv** GetVerdict **from** I/O:  
  **reply** (GetVerdict, verdicts)  
  **recv** (GetJudicialReport,  $msg$ ) **from** I/O:  
  **reply** (GetJudicialReport, judicialReport) {Return judicial report (from lower-level) to requestor}  
  **recv** GetEvidence **from** NET:  
  **reply** (GetEvidence,  $W$ , judicialReport, transcript<sup>a</sup>) {A may query the public judge for the evidence it gathered including all details}  
  **recv** VerResult( $pid, msg, \sigma$ ) **from** NET:  
  result  $\leftarrow$  verifySig( $pid, msg, \sigma$ ) {A verify signatures via  $\mathcal{P}_{\text{judge}}^{\text{CD}}$ 's interface to  $\mathcal{F}_{\text{cert}}$ }  
  **reply** (VerResult, result)  
**Procedures and Functions:**  
**function** verifySig( $pid, msg, \sigma$ ) : {Verify signature at  $\mathcal{F}_{\text{cert}}$ }  
  **send** (VerResult,  $msg, \sigma$ ) **to** ( $pid_{\text{cur}}, (pid, sid_{\text{cur}}), \mathcal{F}_{\text{cert}}$  : verifier)  
  **wait for** (VerResult, result)  
  **return** result  
<sup>a</sup>transcript is a special variable in iUC, which, informally speaking, contains a transcript of all messages sent and received by the current machine instance.

Fig. 15: The judging functionality  $\mathcal{P}_{\text{judge}}^{\text{CD}}$  for  $\mathcal{P}_{\text{cp}}^{\text{acc}}$  (the consensus scaling protocol).

request, it stores the register message including the  $pid$  of the party which requests the certificate in a buffer  $\text{buffer}_S$  for later processing. This allows to model asynchronous network. Also,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  leaks the full request to  $\mathcal{A}$ . If an honest party ( $pid_r, sid, \text{client}$ ) wants to access a certificate of a party  $pid$ , the process is similar: the party sends (Retrieve,  $pid, pid_{\text{CA}}$ ) to access the certificate of identity  $pid$  at CA  $pid_{\text{CA}}$ . Retrieve requests are also stored in a buffer  $\text{buffer}_R$  including the requestor ( $pid_r, sid, \text{client}$ ) for processing later on.  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  also leaks the full request to  $\mathcal{A}$ .

As  $\mathcal{A}$  has full control over the network,  $\mathcal{A}$  needs to inform  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  when it should accept a certificate and add it to its state. Therefore,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  allows  $\mathcal{A}$ : (i) to drop entries from  $\text{buffer}_S$  and (ii) to add buffered certificates to  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ 's state. In Case (i),  $\mathcal{A}$  sends a (Drop $_S$ ,  $pid, pk, pid_{\text{CA}}$ ) message to  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . If the entry ( $pid, pk, pid_{\text{CA}}$ ) exists in  $\text{buffer}_S$ ,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  removes the entry from  $\text{buffer}_S$  and replies (Drop $_S$ , ack) to  $\mathcal{A}$ . Otherwise,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  declines the request and responds (Drop $_S$ , nack) to  $\mathcal{A}$ . In Case (ii),  $\mathcal{A}$  sends the command (Update,  $buffer$ ) to  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ , where  $buffer$  is expected to be a subset of  $\text{buffer}_S$ . If  $buffer \subset \text{buffer}_S$ ,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  adds the elements from  $buffer$  to the state and removes them from the  $\text{buffer}_S$ .  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  then replies to  $\mathcal{A}$  that it accepted the update.

Similarly,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  handles retrieve requests.  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  allows the adversary  $\mathcal{A}$  (i) to drop entries from  $\text{buffer}_R$  and

Description of  $\mathcal{P}_{sv}^{CD} = (\text{supervisor})$ :

Participating roles: {supervisor}  
Corruption model: *incorruptible*

Description of  $M_{\text{supervisor}}$ :

Implemented role(s): {supervisor}  
Subroutines:  $\mathcal{F}_{cp}^{acc}$ : supervisor,  $\mathcal{F}_{net}$ : net  
CheckID( $pid, sid, role$ ):  
Accept all messages with the same  $sid$ .  
Main:  
**recv** (IsAssumptionBroken?,  $prop, id$ ):  
  **if**  $id = \text{public} \wedge prop = \text{liveness}$ : {The model includes public liveness as assumption-based security property}  
  **send** getDeliveryStatus **to** ( $pid_{cur}, sid_{cur}, \mathcal{F}_{net}$ : net) {Query  $\mathcal{F}_{net}$  whether liveness still holds}  
  **wait for** (getDeliveryStatus,  $inTime$ )  
  **send** (IsAssumptionBroken?, liveness, public) **to** ( $pid_{cur}, sid_{cur}, \mathcal{F}_{cp}^{acc}$ : supervisor) {Query  $\mathcal{F}_{cp}^{acc}$ 's supervisor whether liveness still holds}  
  **wait for** (IsAssumptionBroken?,  $b$ )  
  **if**  $inTime \wedge \neg b$ : {Liveness holds true}  
   **reply** (IsAssumptionBroken?, false)  
  **else:**  
   **reply** (IsAssumptionBroken?, true)  
  **else:**  
   **reply** (IsAssumptionBroken?, false)  
**recv** (corruptInt?, ( $pid, sid, role$ )) **from** I/O **s.t.**  $role \notin \{\text{client}, \text{judge}, \text{supervisor}\}$ :  
   $corrRes \leftarrow \text{false}; b \leftarrow \text{false}$  {Check corruption depending on protocol layer}  
  **if**  $role = \text{scd}$ : {Handle internal parties of  $\mathcal{P}_{cp}^{acc}$ }  
    $corrRes \leftarrow \text{corr}((pid, sid, role))$  {Check whether party is corrupted if directly accessible}  
  **else:** {Handle internal parties from lower protocol levels}  
   **send** (corruptInt, ( $pid, sid, role$ )) **to** ( $\text{public}, sid_{cur}, \mathcal{F}_{cp}^{acc}$ : supervisor) {Check if internal party at  $\mathcal{F}_{cp}^{acc}$  is corrupted}  
  **wait for** (corruptInt,  $b$ )  
  **if**  $corrRes \vee b$ :  
   **reply** (corruptInt, true)  
  **else:**  
   **reply** (corruptInt, false)  
**recv** (BreakAssumption,  $toBreak, internalState$ ) **from** I/O: {Leakage request}  
  **reply** (BreakAssumption,  $\varepsilon$ ) {Provide empty leakage}

Fig. 16: The supervisor  $\mathcal{P}_{sv}^{CD}$  for  $\mathcal{P}_{CS}$ .

(ii) to trigger the response to a retrieve request. Case (i) is analogously to Case (i) from the paragraph above. Here,  $\mathcal{A}$  sends a ( $\text{Drop}_R, pid, pid_{CA}, pid_r, role_r$ ) message to  $\mathcal{F}_{PKI}^{acc}$ . If the entry ( $pid, pid_{CA}, pid_r, role_r$ ) exists in  $\text{buffer}_R$ ,  $\mathcal{F}_{PKI}^{acc}$  removes the entry from  $\text{buffer}_R$  and replies ( $\text{Drop}_R, \text{ack}$ ) to  $\mathcal{A}$ . Otherwise,  $\mathcal{F}_{PKI}^{acc}$  declines the request and responds ( $\text{Drop}_R, \text{nack}$ ) to  $\mathcal{A}$ . In Case (ii),  $\mathcal{A}$  sends ( $\text{Deliver}, pid, pid_{CA}, pid_r, role_r, pk_A$ ) to  $\mathcal{F}_{PKI}^{acc}$ . There are now two cases to consider: (i) correctCert is broken for ( $pid, sid, \text{client}$ ) or the party is corrupted and (ii) correctCert is not broken for ( $pid, sid, \text{client}$ ) and the party is not corrupted. In the latter case,  $\mathcal{F}_{PKI}^{acc}$  removes the entry from  $\text{buffer}_R$  and sends the requested output (extracted from its local state) via ( $\text{Retrieve}, pid, pid_{CA}, \text{state}[pid, pid_{CA}]$ ) to ( $pid_r, sid, role_r$ ) (and ignores further data provided by  $\mathcal{A}$ ). Otherwise,  $\mathcal{F}_{PKI}^{acc}$  allows  $\mathcal{A}$  to determine  $pid$ 's certificate. That is,  $\mathcal{F}_{PKI}^{acc}$  forwards the public key  $pk_A$  provided by  $\mathcal{A}$  via ( $\text{Retrieve}, pid, pid_{CA}, pk_A$ ) to ( $pid_r, sid, role_r$ ).

To cover accountability properties, we expect  $\mathcal{F}_{PKI}^{acc}$  to be used with the parameters  $\text{Sec}^{acc} = \{\text{correctCert}\}$  and  $\text{Sec}^{\text{assumption}} = \emptyset$ . For correctCert, we require verdicts which state that one or several CAs misbehaved.  $\mathcal{F}_{PKI}^{acc}$  includes local judges – one local judge ((local,  $pid, \text{client}$ ),  $sid, \text{judge}$ ) per client identity ( $pid, sid, \text{client}$ ), i.e.,  $\text{pids}_{\text{judge}}$  is expected to be  $\{\text{local}\} \times \{0, 1\}^* \times \{\text{client}\}$ . As judicial report  $\mathcal{F}_{\text{judgeParams}}^{acc-PKI}$  outputs an empty report. Further, the functionality  $\mathcal{F}_{\text{judgeParams}}^{acc-PKI}$  specifies that verdicts need to be of form  $\bigwedge_{i=1}^n \text{dis}(pid_{CA}^i, sid_{cur}, ca), n \in \mathbb{N}$ , where ( $pid_{CA}^i, sid_{cur}, ca$ ) are CA identities in  $\mathcal{P}_{PKI}^{acc}$ . As the model does not include assumption-based security properties, we expect  $\text{ids}_{\text{assumption}} = \emptyset$  and specify that  $\mathcal{F}_{\text{judgeParams}}^{acc-PKI}$  always outputs  $\varepsilon$  as response to a BreakAssumption request. Note that (i) local judges act as pure message forwarder for  $\mathcal{A}$  if the associated party is corrupted and (ii)  $\mathcal{A}$  may corrupt a local judge only if its accompanied party, resp. client is already corrupted.

**E.2 An Accountable PKI based on CTLs:** In what follows, we provide the full specification of our PKI protocol, resp. its model  $\mathcal{P}_{PKI}^{acc}$ , as presented in Section 3. We also provide a formal description of all components



Description of the protocol  $\mathcal{F}_{\text{cert}} = (\text{signer}, \text{verifier})$ :

<b>Participating roles:</b> $\{\text{signer}, \text{verifier}\}$	
<b>Corruption model:</b> <i>incorruptible</i>	{See text below}
<b>Protocol parameters:</b>	
- $p \in \mathbb{Z}[x]$ .	{Polynomial that bounds the runtime of the algorithms provided by the adversary.}
- $\eta \in \mathbb{N}$	{The security parameter.}
- <b>sig</b>	{Signing algorithm, outputs a signature $\sigma$ on input $(\text{msg}, \text{sk})$ . The generated signature has a length of $\eta$ bits}
- <b>ver</b>	{Signature verifying algorithm, outputs verification result on input $(\text{msg}, \sigma, \text{pk})$ }
- <b>gen</b>	{Key generation algorithm, outputs $(\text{pk}, \text{sk})$ on input $1^\eta$ }

Description of  $M_{\text{signer}, \text{verifier}}$ :

<b>Implemented role(s):</b> $\{\text{signer}, \text{verifier}\}$	
<b>Internal state:</b>	
- $(\text{pk}, \text{sk}) \in (\{0, 1\}^* \cup \{\perp\})^2 = (\perp, \perp)$ .	{Key pair.}
- $\text{pidowner} \in \{0, 1\}^* \cup \{\perp\} = \perp$ .	{Party ID of the key owner.}
- $\text{msglist} \subset \{0, 1\}^* = \emptyset$ .	{Set of recorded messages.}
- $\text{corrupted} \in \{\text{true}, \text{false}\} = \text{false}$ .	{Is signature key corrupted?}
<b>CheckID</b> $(\text{pid}, \text{sid}, \text{role})$ :	
Check that $\text{sid} = (\text{pid}', \text{sid}')$ :	
If this check fails, output <b>reject</b> .	
Otherwise, accept all entities with the same SID.	{A single instance manages all parties and roles in a single session. A session models one signature key pair belonging to party $\text{pid}'$ .
<b>Corruption behavior:</b>	
- <b>DetermineCorrStatus</b> $(\text{pid}, \text{sid}, \text{role})$ : Return <b>corrupted</b> .	
<b>Initialization:</b>	
$(\text{pk}, \text{sk}) \xleftarrow{\$} \text{Gen}(1^\eta)$	{Generate public/secret key pair}
Parse $\text{sid}_{\text{cur}}$ as $(\text{pid}, \text{sid})$ .	
$\text{pidowner} \leftarrow \text{pid}$ .	
<b>Main:</b>	
<b>recv</b> $(\text{Sign}, \text{msg})$ from I/O to $(\text{pidowner}, \_, \text{signer})$ :	
$\sigma \leftarrow \text{sig}^{(p)}(\text{msg}, \text{sk})$ .	
<b>add</b> $\text{msg}$ to $\text{msglist}$ .	
<b>reply</b> $(\text{Signature}, \sigma)$ .	{Record msg for verification and return signature.}
<b>recv</b> $(\text{Verify}, \text{msg}, \sigma)$ from I/O to $(\_, \_, \text{verifier})$ :	
$b \leftarrow \text{ver}^{(p)}(\text{msg}, \sigma, \text{pk})$ .	{Verify signature.}
<b>if</b> $b = \text{true} \wedge \text{msg} \notin \text{msglist} \wedge \text{corrupted} = \text{false}$ :	
<b>reply</b> $(\text{VerResult}, \text{false})$ .	{Prevent forgery.}
<b>else:</b>	
<b>reply</b> $(\text{VerResult}, b)$ .	{Return verification result.}
<b>recv</b> <b>corruptSigKey</b> from NET:	
$\text{corrupted} \leftarrow \text{true}$ .	{Allow network attacker to corrupt signature keys.}
<b>reply</b> $(\text{corruptSigKey}, \text{ok})$ .	

Fig. 17: The ideal signature functionality  $\mathcal{F}_{\text{cert}}$ .

of the protocol in Figure 22 to 31. Figure 21 provides an overview over the different components of the PKI protocol.

**Remark:** The full formal definition of  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  as presented in Section 3.1 is  $\mathcal{P}_{\text{PKI}}^{\text{acc}} = (\mathcal{P}_{\text{client}}^{\text{CA}} : \text{client}, \mathcal{P}_{\text{sv}}^{\text{CA}} : \text{supervisor}, \mathcal{P}_{\text{judge}}^{\text{CA}} : \text{judge} \mid \mathcal{P}^{\text{CA}} : \text{ca}, \mathcal{P}_{\text{CTL}} : \text{ctl}, \mathcal{F}_{\text{psync-net}} : \text{psync-net}, \mathcal{F}_{\text{sig}}, \mathcal{F}_{\text{auth}} : \text{auth})$ .

We remark that one PKI consisting of clients, several CAs, and several CTLs is modeled as one instance ( $\text{sid}$ ) of  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$ .

**CA clients**  $\mathcal{P}_{\text{client}}^{\text{CA}}$ . The CA client  $\mathcal{P}_{\text{client}}^{\text{CA}}$  (cf. Figure 22 and 23) handles the connection between the PKI and the environment.  $\mathcal{P}_{\text{client}}^{\text{CA}}$ 's main purpose is to request at most one certificate at the PKI and to query for certificates of other parties. Additionally,  $\mathcal{P}_{\text{client}}^{\text{CA}}$  also monitors the PKI, resp. its CTLs and surveils whether there exists solely certificates for its identity requested by itself.

If the environment instructs an instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$ , namely the instance for  $(\text{pid}, \text{sid}, \text{client})$ , to register a string  $\text{pk}$  as its public key at a certain CA  $(\text{pid}_{\text{CA}}, \text{sid}_{\text{cur}}, \text{ca})$ , it sends  $(\text{Register}, \text{pk}, \text{pid}_{\text{CA}})$  via an authenticated channel  $\mathcal{F}_{\text{auth}}$  to  $\mathcal{P}_{\text{client}}^{\text{CA}}$ . An honest instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  registers at most one certificate. If the environment calls the register command once again,  $\mathcal{P}_{\text{client}}^{\text{CA}}$  does not process the request.

If the environment wants to retrieve a certificate of  $\text{pid}$  issued by  $\text{pid}_{\text{CA}}$ , it instructs its instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  to do so via  $(\text{Retrieve}, \text{pid}, \text{pid}_{\text{CA}})$ .  $\mathcal{P}_{\text{client}}^{\text{CA}}$  stores the request in  $\text{requests}$  and forwards the request via NET, i.e., unprotected, to  $(\text{pid}_{\text{CA}}, \text{sid}, \text{ca})$  (with the same session ID). Note that there is neither a guarantee that the request arrives at the CA nor that it arrives unchanged. When  $\mathcal{P}_{\text{client}}^{\text{CA}}$  receives the answer to the retrieve

Description of the protocol  $\mathcal{F}_{\text{net}} = (\text{net})$ :

**Participating roles:**  $\{\text{net}\}$   
**Corruption model:** *incorruptible*  
**Protocol parameters:**  
 -  $\delta \in \mathbb{N}$

*{The expected liveness guarantee for  $\mathcal{F}_{\text{net}}$ }*

Description of  $M_{\text{net}}$ :

**Implemented role(s):**  $\{\text{net}\}$

**Internal state:**

- $\text{buffer}_{\text{msg}} \subset \mathbb{N}^2 \times (\{0, 1\}^*)^7$ ,  $\text{buffer}_{\text{msg}} = \emptyset$  *{Buffer for messages consisting of tuples  $(ctr, \text{round}, \text{sender}, \text{receiver}, \text{content})$ }*
- $\text{counter} \in \mathbb{N}$ ,  $\text{counter} = 0$  *{Counter for messages}*
- $\text{round} \in \mathbb{N}$ ,  $\text{round} = 0$  *{Current round/time unit}*
- $\text{inTime} \in \{\text{true}, \text{false}\}$ ,  $\text{inTime} = \text{true}$  *{Indicator whether messages were within  $\delta$  time units}*

**CheckID( $pid, sid, role$ ):**

Accept all messages with the same  $sid$ .

**Main:**

**recv** (Send,  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})$ ,  $(pid, sid_{\text{cur}}, role)$ ,  $msg$ ) **from** I/O:  $\{(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})$  sends a message to  $(pid, sid_{\text{cur}}, role)$   
 $\text{counter} \leftarrow \text{counter} + 1$   
 $\text{buffer}_{\text{msg}}.add([\text{counter}, \text{round}, (pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}}), (pid, sid_{\text{cur}}, role), msg])$  *{Record message}*  
**send** (Send, counter,  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})$ ,  $(pid, sid_{\text{cur}}, role)$ ,  $msg$ ) **to** NET *{Forward leakage and identifier to  $\mathcal{A}$ }*

**recv** Pull **from** I/O: *{Parties can pull messages}*  
 Let  $(counter, (pid, sid_{\text{cur}}, role), (pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}}), msg) \in \text{buffer}_{\text{msg}}$ , such that there exists no smaller  $counter$   
 with recipient  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})$ .  
**if**  $(counter, (pid, sid_{\text{cur}}, role), (pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}}), msg)$  as above exists:  
 $\text{buffer}_{\text{msg}}.remove([\text{counter}, (pid, sid_{\text{cur}}, role), (pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}}), msg])$  *{Remove message from delivery queue}*  
**reply** (Pull,  $msg$ ) *{Deliver message to receiver}*  
**else:**  
**reply** (Pull,  $\perp$ )

**recv** (Permute,  $\pi$ ) **from** NET: *{ $\mathcal{A}$  is allowed to change the order in the  $\text{buffer}_{\text{msg}}$ }*  
**require:**  $\pi$  is a permutation of the IDs in  $\text{buffer}_{\text{msg}}$ .  
**for all**  $(ctr, msg) \in \text{buffer}_{\text{msg}}$  where  $msg$  represents all other content of  $\text{buffer}_{\text{msg}}$  besides  $ctr$  **do:**  
 $\text{buffer}_{\text{msg}}.remove([ctr, msg])$   
 $\text{buffer}_{\text{msg}}.add([\pi(ctr), msg])$

**recv** UpdateRound **from** NET: *{ $\mathcal{A}$  triggers round update}*  
**if**  $\exists (c, \text{round}, \_, \_) \in \text{buffer}_{\text{msg}}$ , **s.t.**  $\text{round} - \text{round} > \delta \wedge \text{inTime} = \text{true}$ :  
**reply** (UpdateRound, false,  $\epsilon$ )  
**else:**  
 $\text{round} \leftarrow \text{round} + 1$   
**reply** (UpdateRound, true,  $\epsilon$ )  
**reply** (UpdateRound, true,  $\epsilon$ )

**recv** (GetCurRound): *{ $\mathcal{A}$  and  $\mathcal{E}$  are allowed to query the current round.}*  
**reply** (GetCurRound, round)

**recv** BreakLiveness **from** NET: *{ $\mathcal{A}$  breaks liveness}*  
 $\text{inTime} \leftarrow \text{false}$

**recv** getDeliveryStatus **from** I/O: *{ $\mathcal{E}$  may query whether messages were delivered "in time"}*  
**reply** (getDeliveryStatus, inTime)

Fig. 18: The simplified network model  $\mathcal{F}_{\text{net}}$  with breakable  $\delta$ -liveness.

request via the message (Retrieve,  $(ctr, (serial, pid, pk, pid_{\text{CA}}), pid_{\text{CTL}}, \tau, \sigma)$ ) on the NET interface,  $\mathcal{P}_{\text{client}}^{\text{CA}}$  firstly request the current time  $\tau'$  at  $\mathcal{F}_{\text{psync-net}}$ . Then,  $\mathcal{P}_{\text{client}}^{\text{CA}}$  updates its information regarding the CA  $pid_{\text{CA}}$  and queries  $\mathcal{F}_{\text{init}}^{\text{CA}}$  (see below) for the public key of  $pid_{\text{CA}}$  and the CTLs that the CAs uses including their public keys (this is done via the message (GetCaCtls,  $pid_{\text{CA}}$ )). When  $\mathcal{P}_{\text{client}}^{\text{CA}}$  receives the answer (GetCaCtls,  $pk_{\text{CA}}, (pid_1, \dots, pid_n), (pk_1, \dots, pk_n)$ ), it stores the CA's public key  $pk_{\text{CA}}$  in  $pk_{\text{CA}}[pid_{\text{CA}}]$ . Further, it stores the CTLs connected to the CA in  $\text{CTLs}[pid_{\text{CA}}]$  and their public keys in  $pk_{\text{CTL}}[pid_{\text{CTL}}]$ . After  $\mathcal{P}_{\text{client}}^{\text{CA}}$  gathered the trusted information from  $\mathcal{F}_{\text{init}}^{\text{CA}}$ , it checks whether (i)  $\mathcal{P}_{\text{client}}^{\text{CA}}$  indeed requested the certificate for  $pid$  at  $pid_{\text{CA}}$ , (ii) the provided certificate is older than  $3 \cdot \delta$  time units, (iii)  $pid_{\text{CTL}}$  is a CTL for  $pid_{\text{CA}}$ , (iv) the certificate is validly signed by the CTL  $pid_{\text{CTL}}$ , and (v) the certificate is also validly signed by the CA  $pid_{\text{CA}}$ . If the checks succeed,  $\mathcal{P}_{\text{client}}^{\text{CA}}$  removes the request from its request queue and sends the certificate to the requestor.

The adversary is responsible for starting the certificate monitoring process of  $\mathcal{P}_{\text{client}}^{\text{CA}}$ .<sup>14</sup> When  $\mathcal{A}$  triggers the monitoring process by sending Monitor,  $\mathcal{P}_{\text{client}}^{\text{CA}}$  updates the set of CTLs at  $\mathcal{F}_{\text{init}}^{\text{CA}}$  (by using the GetCTLs). Let  $\text{CTL}$  be the set of all CTLs provided by  $\mathcal{F}_{\text{init}}^{\text{CA}}$ .  $\mathcal{P}_{\text{client}}^{\text{CA}}$  then queries all CTLs via (Monitor,  $pid_{\text{cur}}, \text{CTL}$ ) via  $\mathcal{F}_{\text{psync-net}}$ . Note that  $\mathcal{P}_{\text{client}}^{\text{CA}}$  and  $\mathcal{F}_{\text{psync-net}}$  are connected via a (direct and thus secure) I/O interface. When

<sup>14</sup>We note that security guarantees are incorporated into  $\mathcal{F}_{\text{psync-net}}$  which we will discuss later on.

Description of the ideal and accountable PKI  $\mathcal{F}_{\text{PKI}}^{\text{acc}} = (\text{client}, \text{judge}, \text{supervisor})$ :

<b>Participating roles:</b> $\{\text{client}, \text{judge}, \text{supervisor}\}$ <b>Corruption model:</b> <i>dynamic without secure erasures</i> <b>Protocol parameters:</b> <ul style="list-style-type: none"> <li>- <math>\text{Sec}^{\text{acc}} \subset \{0, 1\}^*</math> <span style="float: right;">{Accountability properties}</span></li> <li>- <math>\text{Sec}^{\text{assumption}} \subset \{0, 1\}^*</math> <span style="float: right;">{Assumption-based security properties}</span></li> <li>- <math>\text{pids}_{\text{judge}} \subset \{0, 1\}^*</math> <span style="float: right;">{set of judge entities/(P)IDs in the protocol (which are often directly related to some protocol participants)}</span></li> <li>- <math>\text{ids}_{\text{assumption}} \subset \{0, 1\}^*</math> <span style="float: right;">{set of entities/IDs where properties are ensured via assumptions}</span></li> </ul>
---

Description of  $M_{\text{client}}^{\text{acc-PKI}}$ :

<b>Implemented role(s):</b> $\{\text{client}, \text{judge}, \text{supervisor}\}$ <b>Subroutines:</b> $\mathcal{F}_{\text{judgeParams}} : \text{judgeParams}$ <b>Internal state:</b> <ul style="list-style-type: none"> <li>- <math>\text{state} : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \text{state} = \emptyset</math> <span style="float: right;">{The registered certificates, resp. identities (as mapping), initially <math>\perp</math> for all entries}</span></li> <li>- <math>\text{buffer}_S \subset (\{0, 1\}^*)^3, \text{buffer}_S = \varepsilon, \text{buffer}_R \subset (\{0, 1\}^*)^4, \text{buffer}_R = \emptyset</math> <span style="float: right;">{Requested certificates/buffer of form <math>(\text{pid}, \text{pid}_{\text{CA}}, \text{pid}_r, \text{role}_r)</math>}</span></li> <li>- <math>\text{corruptedIntParties} \in \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \setminus (\text{Roles}_{\mathcal{F}} \cup \{\text{judge}, \text{supervisor}\})</math>, initially <math>\emptyset</math> <span style="float: right;">{The set of corrupted internal parties <math>(\text{pid}, \text{sid}, \text{role})</math>}</span></li> <li>- <math>\text{brokenAssumptions} : \text{Sec}^{\text{assumption}} \times \text{ids}_{\text{assumption}} \rightarrow \{\text{true}, \text{false}\}</math> <span style="float: right;">{Stores broken security assumptions per id, initially false <math>\forall</math> entries}</span></li> <li>- <math>\text{brokenProps} : (\text{Sec}^{\text{assumption}} \cup \text{Sec}^{\text{acc}}) \times (\text{pids}_{\text{judge}} \cup \text{ids}_{\text{assumption}}) \rightarrow \{\text{true}, \text{false}\}</math> <span style="float: right;">{Stores broken security properties per judge/id, initially false <math>\forall</math> entries}</span></li> <li>- <math>\text{verdicts} : \text{pids}_{\text{judge}} \rightarrow \{0, 1\}^*</math> <span style="float: right;">{Verdicts per <math>p \in \text{pids}_{\text{judge}}</math>, initially <math>\varepsilon</math>}</span></li> </ul> <b>CheckID</b> $(\text{pid}, \text{sid}, \text{role})$ : Accept all messages for the same $\text{sid}$ . <b>Corruption behavior:</b> <ul style="list-style-type: none"> <li>- <b>AllowCorruption</b><math>(\text{pid}, \text{sid}, \text{role})</math>:  Do not allow corruption of <math>(\text{pid}, \text{sid}, \text{supervisor})</math>.  if <math>\text{role} = \text{judge}</math>:  <b>send</b> <math>(\text{Corrupt}, (\text{pid}, \text{sid}, \text{judge}), \text{internalState})</math>  <b>to</b> <math>(\text{pid}, \text{sid}, \mathcal{F}_{\text{judgeParams}} : \text{judgeParams})</math> <span style="float: right;">{<math>\mathcal{F}_{\text{judgeParams}}</math> decides whether judges can be corrupted}</span>  <b>wait for b; return b</b></li> <li>- <b>DetermineCorrStatus</b><math>(\text{pid}, \text{sid}, \text{role})</math>:  if <math>\text{role} = \text{judge}</math>: <span style="float: right;">{<math>\mathcal{F}_{\text{judgeParams}}</math> may determine a judge's corruption status}</span>  <b>send</b> <math>(\text{CorruptionStatus?}, (\text{pid}, \text{sid}, \text{judge}), \text{internalState})</math>  <b>to</b> <math>(\text{pid}, \text{sid}, \mathcal{F}_{\text{judgeParams}} : \text{judgeParams})</math>  <b>wait for b; return b</b></li> <li>- <b>AllowAdvMessage</b><math>(\text{pid}, \text{sid}, \text{role}, \text{pid}_{\text{receiver}}, \text{sid}_{\text{receiver}}, \text{role}_{\text{receiver}}, m)</math>  Do not allow sending messages to <math>\mathcal{F}_{\text{judgeParams}}</math>. <span style="float: right;">{<math>\mathcal{A}</math> is not allowed to invoke <math>\mathcal{F}_{\text{judgeParams}}</math> in the name of corrupted parties.}</span></li> </ul> <b>Message preprocessing:</b> if message is addressed to $((\text{local}, \text{pid}, \text{role}), \text{sid}, \text{judge})$ and $(\text{pid}, \text{sid}, \text{role})$ is corrupted: Forward message to $\mathcal{A}$ <span style="float: right;">{Forward request to corrupted local judges to <math>\mathcal{A}</math>}</span> if Receive $(\text{Fwd}, \text{msg}, (\text{pid}', \text{sid}', \text{role}'))$ to $((\text{local}, \text{pid}, \text{role}), \text{sid}, \text{judge})$ via NET and $(\text{pid}, \text{sid}, \text{role})$ is corrupted: <b>send</b> $\text{msg}$ <b>to</b> $(\text{pid}, \text{sid}, \text{role})$ <span style="float: right;">{<math>\mathcal{A}</math> corrupted local judge acts as message forwarder for <math>\mathcal{A}</math>}</span> <b>Main:</b> <b>recv</b> $(\text{Register}, \text{pk}, \text{pid}_{\text{CA}})$ <b>from</b> I/O: <span style="float: right;">{Register certificate.}</span> $\text{buffer}_S.\text{add}([\text{pid}_{\text{cur}}, \text{pk}, \text{pid}_{\text{CA}}])$ <span style="float: right;">{Record registration attempt}</span> <b>send</b> $(\text{Register}, \text{pk}, \text{pid}_{\text{cur}}, \text{pid}_{\text{CA}})$ <b>to</b> NET <span style="float: right;">{Leak request to <math>\mathcal{A}</math>}</span> <b>recv</b> $(\text{Retrieve}, \text{pid}, \text{pid}_{\text{CA}})$ <b>from</b> I/O: <span style="float: right;">{Query for certificate}</span> $\text{buffer}_R.\text{add}([\text{pid}, \text{pid}_{\text{CA}}, \text{pid}_{\text{call}}, \text{role}_{\text{call}}])$ <b>send</b> $(\text{Retrieve}, \text{pid}, \text{pid}_{\text{CA}}, \text{pid}_{\text{call}}, \text{role}_{\text{call}})$ <b>to</b> NET <span style="float: right;">{Leak request to <math>\mathcal{A}</math>}</span> <b>recv</b> $(\text{Drop}_S, \text{pid}, \text{pk}, \text{pid}_{\text{CA}})$ <b>from</b> NET: <span style="float: right;">{<math>\mathcal{A}</math> is allowed to drop the register requests}</span> if $(\text{pid}, \text{pk}, \text{pid}_{\text{CA}}) \in \text{buffer}_S$ : $\text{buffer}_S.\text{remove}([\text{pid}, \text{pk}, \text{pid}_{\text{CA}}])$ <b>reply</b> $(\text{Drop}_S, \text{ack})$ <b>else:</b> <b>reply</b> $(\text{Drop}_S, \text{nack})$ <b>recv</b> $(\text{Deliver}, \text{pid}, \text{pid}_{\text{CA}}, \text{pid}_r, \text{role}_r, \text{pk}_{\mathcal{A}})$ <b>from</b> NET: <span style="float: right;">{<math>\mathcal{A}</math> may trigger when retrieve requests are answered.}</span> <b>require:</b> $(\text{pid}, \text{pid}_{\text{CA}}, \text{pid}_r, \text{role}_r) \in \text{buffer}_R$ <span style="float: right;">{If correctCert is broken, <math>\mathcal{A}</math> may inject malicious certificates.}</span> if $\text{brokenProps}[\text{correctCert}, (\text{pid}, \text{sid}_{\text{cur}}, \text{client})] = \text{false} \wedge (\text{pid}, \text{sid}_{\text{cur}}, \text{client}) \notin \text{CorruptionSet}$ : $\text{buffer}_R.\text{remove}([\text{pid}, \text{pid}_{\text{CA}}, \text{pid}_r, \text{role}_r])$ <b>send</b> $(\text{Retrieve}, \text{pid}, \text{pid}_{\text{CA}}, \text{state}[\text{pid}, \text{pid}_{\text{CA}}])$ <b>to</b> $(\text{pid}_r, \text{sid}_{\text{cur}}, \text{role}_r)$ <b>else:</b> $\text{buffer}_R.\text{remove}([\text{pid}, \text{pid}_{\text{CA}}, \text{pid}_r, \text{role}_r])$ <b>send</b> $(\text{Retrieve}, \text{pid}, \text{pid}_{\text{CA}}, \text{pk}_{\mathcal{A}})$ <b>to</b> $(\text{pid}_r, \text{sid}_{\text{cur}}, \text{role}_r)$ <span style="float: right;">{<math>\mathcal{A}</math> determines the certificate if correctCert is broken for pid or pid is corrupted}</span> <b>recv</b> $(\text{Drop}_R, \text{pid}, \text{pid}_{\text{CA}}, \text{pid}_r, \text{role}_r)$ <b>from</b> NET: <span style="float: right;">{<math>\mathcal{A}</math> is allowed to drop the read requests}</span> if $(\text{pid}, \text{pid}_{\text{CA}}, \text{pid}_r, \text{role}_r) \in \text{buffer}_R$ : $\text{buffer}_R.\text{remove}([\text{pid}, \text{pid}_{\text{CA}}, \text{pid}_r, \text{role}_r])$ <b>reply</b> $(\text{Drop}_R, \text{ack})$ <b>else:</b> <b>reply</b> $(\text{Drop}_R, \text{nack})$ <b>recv</b> $(\text{Update}, \text{buffer})$ <b>from</b> NET <b>s.t.</b> $\text{buffer} \subset \text{buffer}_S$ : <span style="float: right;">{<math>\mathcal{A}</math> may trigger when <math>\text{buffer}_S</math> is included in state}</span> <b>for all</b> $(\text{pid}, \text{pk}, \text{pid}_{\text{CA}}) \in \text{buffer}$ <b>do:</b> if $\text{state}[\text{pid}, \text{pid}_{\text{CA}}] = \perp$ : $\text{state}[\text{pid}, \text{pid}_{\text{CA}}] \leftarrow \text{pk}$ $\text{buffer}_S.\text{remove}[\text{buffer}]$ <span style="float: right;">{Update buffer}</span> <b>reply</b> $(\text{Update}, \text{ack})$ Include static code from the AUC transformation $\mathcal{T}_1(\cdot)$ here, i.e., include additional code from Figure 2 and 3 here. <sup>a</sup> $\text{Roles}_{\mathcal{F}} = \{\text{client}\}$ here.
--

Fig. 19: An ideal and accountable PKI  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  derived from  $\mathcal{G}_{\text{BB}}$  [27].

Description of  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}} = (\text{judgeParams})$ :

**Participating roles:**  $\{\text{judgeParams}\}$   
**Corruption model:** *incorruptible*

Description of  $M_{\text{judgeParams}}^{\text{acc-PKI}}$ :

**Implemented role(s):**  $\{\text{judgeParams}\}$

**CheckID**( $pid, sid, role$ ):

Accept all messages with the same  $sid$ .

**Main:**

**recv** (BreakAccProp, verdict, toBreak, internalState) **from** I/O: {No restrictions on verdict}  
**if** All entries in *verdict* are of form  $\bigwedge_{i=1}^n \text{dis}(pid_{ca}^i, sid_{cur}, ca), n \in \mathbb{N}$  and all  $pid_{ca}^i \in \{0, 1\}^*$ :  
**if**  $toBreak \subset \{\text{correctCert}\} \times (\{0, 1\}^*)^3$ :  
    **reply** (Cheat, true,  $\varepsilon$ )  
**else:**  
    **reply** (Cheat, false,  $\varepsilon$ )  
**else:**  
    **reply** (Cheat, false,  $\varepsilon$ )

**recv** (BreakAssumption, toBreak, internalState) **from** I/O: {Do not generate leakage when breaking assumptions}  
**reply** (BreakAssumption,  $\varepsilon$ )

**recv** (GetJudicialReport, msg, internalState) **from** I/O: {Generate judicial report}  
**reply** (GetJudicialReport,  $\varepsilon$ ) {Return an empty report}

**recv** (Corrupt, (id, sid, judge), internalState) **from** I/O: { $\mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}}$  allows to corrupt a local judge iff the accompanied client is corrupted}  
**if**  $id = (\text{local}, pid, \text{client}) \wedge (pid, sid, \text{client}) \in \text{CorruptionSet}$ :  
    **reply** true  
**else:**  
    **reply** false

**recv** (CorruptionStatus?, (id, sid, judge), internalState) **from** I/O: { $\mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}}$  interprets a local judge as corrupted iff the accompanied client is corrupted}  
**if**  $(id, sid, \text{judge}) \in \text{CorruptionSet}$ :  
    **reply** true  
**else if**  $id = (\text{local}, pid, \text{client}) \wedge (pid, sid, \text{client}) \in \text{CorruptionSet}$ :  
    **reply** true  
**else:**  
    **reply** false

Fig. 20:  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}}$  for  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ .

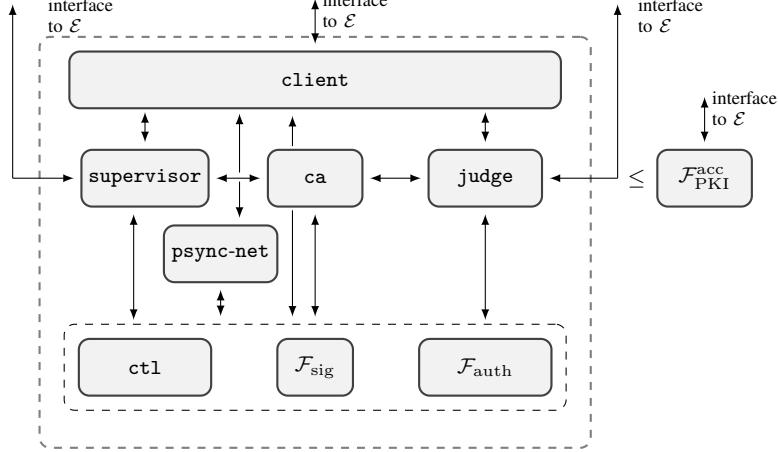


Fig. 21: Realization relation of a CA/CTL example stated in Theorem 2. The system  $\mathcal{E}$  denotes the environment, modeling, as usual in UC setting, arbitrary higher level protocols. All machines are additionally connected to the network adversary who is responsible for delivering network messages.

$\mathcal{F}_{\text{psync-net}}$  returns an answer of a CTL (containing all certificates issued for the requested  $pid$  recorded at that CTL),  $\mathcal{P}_{\text{client}}^{\text{CA}}$  pulls updates public keys and data of CAs and CTLs at  $\mathcal{F}_{\text{init}}^{\text{CA}}$  (see above). Afterwards, it verifies whether the received certificates are validly signed (by a CA and a CTL). If one of the certificates does not match the public key  $pk$  stored in  $\mathcal{P}_{\text{client}}^{\text{CA}}$ ,  $\mathcal{P}_{\text{client}}^{\text{CA}}$  reports the falsely issued certificate to its local judge  $\mathcal{P}_{\text{judge}}^{\text{CA}}$ . *Corruption:* The adversary can dynamically corrupt every instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$ . In this case,  $\mathcal{A}$  has full control

over the input and output interface of the instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  and can act on its behalf. Further note that an instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  considers itself corrupted if its session at  $\mathcal{F}_{\text{sig}}$  is corrupted or if the used authenticated channel session at  $\mathcal{F}_{\text{auth}}$  is corrupted.

Further note that:

- Every new instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  registers itself automatically at  $\mathcal{F}_{\text{psync-net}}$  when activated for the first time.
- $\mathcal{A}$  is responsible to trigger an instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  to register itself for the usage of authenticated channels (via  $\mathcal{F}_{\text{auth}}$ ).

*Remark:* Note that we do not model that clients provide “self-signed certificates” as part of their certificate requests to a CA, i.e., they do sign their public key in combination with their identity (and potentially some metadata) and forward this certificate to the CA. We omit self-signed certificates in our model due to simplicity. To incorporate this into our model, one can extend  $\mathcal{P}_{\text{client}}^{\text{CA}}$  with a parameter (which is intended to include the signature scheme) for creating a self-signed certificate. Before requesting a certificate,  $\mathcal{P}_{\text{client}}^{\text{CA}}$  then needs to invoke this parameter with the input data and has to enrich the certificate request with the output produced by the parameter. To cover this case, one also needs to adapt  $\mathcal{P}^{\text{CA}}$ . It needs to be enhanced with a new parameter to check the validity of the client’s self-signed certificate.  $\mathcal{P}^{\text{CA}}$  only process certificate request if the parameter (on input of the certificate request) accepts the certificate request.

**Certification authorities  $\mathcal{P}^{\text{CA}}$ .** The main purpose of CAs, resp.  $\mathcal{P}^{\text{CA}}$  (cf. Figure 24), is to issue and to distribute certificates. There exists one instance of  $\mathcal{P}^{\text{CA}}$  per CA. The issuance of certificates is divided into two steps: (i) a CA generates a so-called pre-certificate and forwards the pre-certificate to a CTL. (ii) The CTL attests the correct generation of the (pre-)certificate and provides a finalized certificate to the CA. More specifically,  $\mathcal{P}^{\text{CA}}$  works as follows: On the first activation of an instance of  $\mathcal{P}^{\text{CA}}$ , it queries the  $\mathcal{F}_{\text{init}}^{\text{CA}}$  for initialization.  $\mathcal{F}_{\text{init}}^{\text{CA}}$  initializes  $\mathcal{P}^{\text{CA}}$ ’s instance at  $\mathcal{F}_{\text{sig}}$ , i.e.,  $\mathcal{P}^{\text{CA}}$ ’s signing and public key, and provides its public key  $pk$  to  $\mathcal{P}^{\text{CA}}$ .  $\mathcal{F}_{\text{init}}^{\text{CA}}$  also informs  $\mathcal{P}^{\text{CA}}$  which CTLs it will/has to use (including their public keys).

As clients send their certificate requests via an authenticated channel  $\mathcal{F}_{\text{auth}}$ ,  $\mathcal{A}$  is responsible for triggering  $\mathcal{P}^{\text{CA}}$  to pull certificate requests from  $\mathcal{F}_{\text{auth}}$ . When  $\mathcal{P}^{\text{CA}}$  receives such request (of form  $(\text{Register}, pid, pk)$ ), it creates a pre-certificate: it adds an ID/serial number (an incremental counter) and its own identity to  $(pid, pk)$ , and signs  $(id, (pid, pk), pid_{\text{cur}})$ . It then forwards the so-called pre-certificate  $(id, (pid, pk), \sigma_{\text{CA}})$  for finalization via NET to its connected CTLs. Note that  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  does not enforce liveness properties here, i.e., we do not enforce that these requests are actually delivered to a CTL.

As  $\mathcal{A}$  has full control over the network, we model that  $\mathcal{A}$  triggers  $\mathcal{P}^{\text{CA}}$ ’s update process. If  $\mathcal{P}^{\text{CA}}$  receives a  $(\text{Read}, pid_{\text{CTL}}, state, \sigma_{\text{CTL}})$  message via NET from one of its CTLs (as stored in CTLs),  $\mathcal{P}^{\text{CA}}$  checks whether (i) the message is validly signed by a CTL  $pid_{\text{CTL}}$  from CTLs, (ii) the provided  $state$  is an extension of the state previously delivered by  $pid_{\text{CTL}}$  (which is stored in  $\text{certs}[pid_{\text{CTL}}]$ ), and (iii) all certificates in  $state$  are validly signed by  $pid_{\text{CTL}}$ . If these checks succeed,  $\mathcal{P}^{\text{CA}}$  stores  $state$  (i.e., the certificates provided by  $pid_{\text{CTL}}$ ) as update in  $\text{certs}[pid_{\text{CTL}}]$ .

When clients want to retrieve a certificate from  $\mathcal{P}_{\text{CS}}$ , they send  $(\text{Retrieve}, pid, pid_{\text{CA}})$  to the instance of  $\mathcal{P}^{\text{CA}}$ , namely  $(pid_{\text{CA}}, sid, ca)$ . Note that these requests are delivered via A/NET and can thus be tampered. When receiving such a request,  $\mathcal{P}^{\text{CA}}$  checks in its storage  $\text{certs}$  whether it issued a pre-certificate for  $pid$ . If it issued a pre-certificate,  $\mathcal{P}^{\text{CA}}$  looks for the finalized certificate in the certificates provided by the CTLs. If  $\mathcal{P}^{\text{CA}}$  can find a finalized certificate in  $\text{certs}_{\text{CTL}}$  (for any CTL), it forwards the first certificate matching the search criteria (via NET) to the requesting party.

*Corruption:* Analogously to  $\mathcal{P}_{\text{client}}^{\text{CA}}$ ,  $\mathcal{A}$  can dynamically corrupt every instance of  $\mathcal{P}^{\text{CA}}$ . In case of a corruption,  $\mathcal{A}$  gains full control over the input and output interface of the instance of  $\mathcal{P}^{\text{CA}}$  and can act on its behalf. Further note that an instance of  $\mathcal{P}^{\text{CA}}$  considers itself corrupted if its session at  $\mathcal{F}_{\text{sig}}$  is corrupted or if the used authenticated channel session at  $\mathcal{F}_{\text{auth}}$  is corrupted.

Moreover,  $\mathcal{A}$  is responsible for triggering an instance of  $\mathcal{P}^{\text{CA}}$  to register itself for the usage of authenticated channels (via  $\mathcal{F}_{\text{auth}}$ ).

**Certificate Transaction Logs  $\mathcal{P}_{\text{CTL}}$ .** CTLs validate and finalize (pre-)certificates. Further, they allow clients to check whether there exist falsely/maliciously issued certificates for their identities. We provide the specification of  $\mathcal{P}_{\text{CTL}}$  in Figure 25. One instance of  $\mathcal{P}_{\text{CTL}}$  models one CTL. Similarly to  $\mathcal{P}^{\text{CA}}$ ,  $\mathcal{P}_{\text{CTL}}$  queries the initialization functionality  $\mathcal{F}_{\text{init}}^{\text{CA}}$  for initialization on its first activation.  $\mathcal{F}_{\text{init}}^{\text{CA}}$  initializes  $\mathcal{P}_{\text{CTL}}$  instance at  $\mathcal{F}_{\text{sig}}$  and provides its public key  $pk$  to  $\mathcal{P}_{\text{CTL}}$ .  $\mathcal{F}_{\text{init}}^{\text{CA}}$  also informs  $\mathcal{P}_{\text{CTL}}$  for which CAs it will finalize pre-certificates.

If an instance of  $\mathcal{P}_{\text{CTL}}$  identified by  $(pid_{\text{CTL}}, sid, ct1)$  receives a `Submit` request to finalize a pre-certificate (via `NET`),  $\mathcal{P}_{\text{CTL}}$  checks whether the pre-certificate is of form  $(id, (pid_r, pk_r), pid_{\text{CA}}, \sigma_{\text{CA}})$ . If the current instance of  $\mathcal{P}_{\text{CTL}}$  finalizes pre-certificates of  $pid_{\text{CA}}$ , i.e.,  $pid_{\text{CA}}$  is stored in CAs, and the signature  $\sigma_{\text{CA}}$  verifies,  $\mathcal{P}_{\text{CTL}}$  starts the certificate finalization process. Therefore, it firstly queries  $\mathcal{F}_{\text{psync-net}}$  for the current time  $\tau$ . Then,  $\mathcal{P}_{\text{CTL}}$  adds an own ID (incremental counter), its own identity, and the current timestamp  $\tau$  to the pre-certificate and then signs  $(id_{\text{CTL}}, (id, (pid_r, pk_r), pid_{\text{CA}}, \sigma_{\text{CA}}), pid_{\text{CTL}}, \tau)$ . Afterwards, it stores the finalized certificate  $((id_{\text{CTL}}, (id, (pid_r, pk_r), pid_{\text{CA}}, \sigma_{\text{CA}}), pid_{\text{CTL}}, \sigma_{\text{CTL}}, \tau)$  in its state.

If a party queries for  $\mathcal{P}_{\text{CTL}}$ 's state via `Read`,  $\mathcal{P}_{\text{CTL}}$  signs its state and replies to the request with  $(\text{Read}, pid_{\text{cur}}, \text{state}, \sigma_{\text{CTL}})$ .

To allow clients to monitor issued certificates,  $\mathcal{P}_{\text{CTL}}$  provides the `Monitor` interface/command. If  $\mathcal{P}_{\text{CTL}}$  receives  $(\text{Monitor}, ctr, pid)$  from a client  $pid$  (via  $\mathcal{F}_{\text{psync-net}}$ ), it extracts all certificates issued for  $pid$  from its state and sends the certificates (including the used message identifier  $ctr$ ) back to  $\mathcal{F}_{\text{psync-net}}$ .

*Corruption:* CTLs are trusted anchors in our model and thus,  $\mathcal{P}_{\text{CTL}}$  is incorruptible. Further note that signature keys of CTLs, i.e., their `signer` instances at  $\mathcal{F}_{\text{sig}}$ , as well as their signature verifying function, i.e., their `verifier` instances at  $\mathcal{F}_{\text{sig}}$ , cannot be corrupted. This is handled/ensured via  $\mathcal{F}_{\text{init}}^{\text{CA}}$  and our slightly adapted variant of  $\mathcal{F}_{\text{sig}}$ .

*Remark:* Note that one can enhance our model to a setting where only a threshold of CTLs per CA needs to be honest. For the sake of presentation, we decided to present the simpler variant with incorruptible CTLs.

**The initialization functionality  $\mathcal{F}_{\text{init}}^{\text{CA}}$ .** The initialization functionality  $\mathcal{F}_{\text{init}}^{\text{CA}}$  (cf. Figure 26) models a trusted setup/distribution of CAs/CTLs identities and their public keys. There is one instance of  $\mathcal{F}_{\text{init}}^{\text{CA}}$  per  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$ . In particular,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  includes (i) initialization of CAs, resp.  $\mathcal{P}^{\text{CA}}$ , (ii) initialization of CTLs, resp.  $\mathcal{P}_{\text{CTL}}$ , and (iii) initialization of judges  $\mathcal{P}_{\text{judge}}^{\text{CA}}$ . It further provides a interface for all participants in  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  to query for CA, resp. CTL, data including their public keys and is used to ensure that clients and their accompanied local judges are initialized simultaneously.

If an instance of  $\mathcal{P}^{\text{CA}}$  asks for initialization,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  queries  $\mathcal{A}$  for the CTLs  $\mathcal{P}^{\text{CA}}$  should use. If one of the provided CTLs is not already registered,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  initializes its session at  $\mathcal{F}_{\text{sig}}$ . Note that  $\mathcal{F}_{\text{init}}^{\text{CA}}$  enforces that all used signature keys for CTLs are not corrupted. When storing which CTLs “work” for a dedicated CA,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  checks that every CTLs session at  $\mathcal{F}_{\text{sig}}$ /signature key is not corrupted. Otherwise,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  declines the initialization of  $\mathcal{P}^{\text{CA}}$  and queries  $\mathcal{A}$  for a new set of CTLs for  $\mathcal{P}^{\text{CA}}$ . If the registration of the CTLs succeeds,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  also initializes  $\mathcal{P}^{\text{CA}}$  signature key at  $\mathcal{F}_{\text{sig}}$  if necessary. All data is stored in  $\mathcal{F}_{\text{init}}^{\text{CA}}$  state: each CA's  $pid$ , resp. CTL, is added to  $pid_{\text{CA}}$ , resp.  $pid_{\text{CTL}}$ . The map  $pk_{\text{CA}}$ , resp.  $pk_{\text{CTL}}$ , stores the public key of each registered CA, resp. CTL. The map  $\text{CTLs}$ ,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  stores for each CA which CTLs are assigned to it. Finally,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  sends (i) the CA's public key, (ii) the CTLs the CA uses, and (iii) the CTLs' public keys back to the instance of  $\mathcal{P}^{\text{CA}}$  which asked for initialization.

Similarly,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  initializes an instance of  $\mathcal{P}_{\text{CTL}}$ . This mainly includes to provide the CTL (i) its public key, (ii) the CAs for which it finalizes pre-certificates, and (iii) the relevant CAs public keys (to the dedicated instance of  $\mathcal{P}_{\text{CTL}}$ ). The handling is similar to the initialization of a CA. If the CTL/CA is not registered at  $\mathcal{F}_{\text{init}}^{\text{CA}}$  so far,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  creates a signature key for it at  $\mathcal{F}_{\text{sig}}$ . Note that  $\mathcal{F}_{\text{init}}^{\text{CA}}$  ensures that signature keys of CTLs are not corrupted – if a newly registered CTL's  $\mathcal{F}_{\text{sig}}$  session is corrupted,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  aborts the registration for this CTL.

When a judge ( $\mathcal{P}_{\text{judge}}^{\text{CA}}$ ) queries  $\mathcal{F}_{\text{init}}^{\text{CA}}$  for initialization,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  simply forwards all its information regarding the PKI session to the judge. This includes, (i) the identities (PIDs) of the registered CAs, resp. CTLs, (ii) their public keys, and (iii) the information which CA and CTLs work together.

Besides initializing parties, parties can also query  $\mathcal{F}_{\text{init}}^{\text{CA}}$  with (i) `GetCaCtIs`, (ii) `GetCAs`, and (iii) `GetCTLs`. If a party queries for  $(\text{GetCaCtIs}, pid_{\text{CA}})$ ,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  replies with the public key of the CA  $pid_{\text{CA}}$ , the (PIDs of the) CTLs of the CA, and the CTLs' public keys. In Case (ii),  $\mathcal{F}_{\text{init}}^{\text{CA}}$  provides all registered (PIDs of) CAs and their public keys to the requestor. In (iii),  $\mathcal{F}_{\text{init}}^{\text{CA}}$  provides all registered (PIDs of) CTLs and their public keys to the requestor.

Further, we use the initialization functionality  $\mathcal{F}_{\text{init}}^{\text{CA}}$  to ensures that clients and their judges are initialized “at the same time”. This simplifies the proof later on, as we do not have to take care of several edge cases in cases of non-initialized parties.<sup>15</sup>

<sup>15</sup>Note that an ITM in iUC which is initialized via a corruption status request always responds that the ITM is not corrupted so far. We want to avoid this edge case in our modeling.

**The certificate monitoring enforcing**  $\mathcal{F}_{\text{psync-net}}$ . The network functionality  $\mathcal{F}_{\text{psync-net}}$  (cf. Figure 27) ensures that honest parties can check in “near-time” whether there are malicious certificates issued for their identities. Especially for this purpose,  $\mathcal{F}_{\text{psync-net}}$  models partially synchronous network between clients and CTLs. There is one instance of  $\mathcal{F}_{\text{psync-net}}$  per PKI session.

On every activation,  $\mathcal{F}_{\text{psync-net}}$  updates the list of available CTLs at  $\mathcal{F}_{\text{init}}^{\text{CA}}$  (via `GetCTLs`) and checks whether all registered honest clients (stored in  $\text{pid}_{\text{client}}^h$ ) are still honest (via querying all clients in  $\text{pid}_{\text{client}}^h$  for their corruption status). If previously honest clients have been corrupted,  $\mathcal{F}_{\text{psync-net}}$  removes them from  $\text{pid}_{\text{client}}^h$  (and also from monitoring in `monReg` and `requestQueue`, see below).

When honest parties start a monitoring process, i. e., they query every registered CTL for the certificates published for their identity, they send a `Monitor` message to  $\mathcal{F}_{\text{psync-net}}$  including their identity/*pid* and (the PIDs of) the CTLs which they want to monitor.  $\mathcal{F}_{\text{psync-net}}$  stores this request in `requestQueue`. More specifically, it stores one entry per CTL in `requestQueue` including (i) an unambiguous identifier for each entry and (ii) the time  $\tau$  (according to  $\mathcal{F}_{\text{psync-net}}$ 's internal clock) when  $\mathcal{F}_{\text{psync-net}}$  received the monitoring request. Additionally,  $\mathcal{F}_{\text{psync-net}}$  stores in `monReg` that  $\mathcal{P}_{\text{client}}^{\text{CA}}$  started a monitoring process at time  $\tau$ .  $\mathcal{F}_{\text{psync-net}}$  then leaks the complete `requestQueue` to  $\mathcal{A}$ .

Similar to the initialization of the monitoring process at  $\mathcal{P}_{\text{client}}^{\text{CA}}$ ,  $\mathcal{A}$  is expected to trigger further processing of the monitoring at  $\mathcal{F}_{\text{psync-net}}$ , i. e., (i) forwarding a monitoring request to a CTL and (ii) deliver the response of a CTL back to the monitoring party. If the adversary triggers  $\mathcal{F}_{\text{psync-net}}$  via (`Deliver, ctr`),  $\mathcal{F}_{\text{psync-net}}$  firstly checks whether a request for the *ctr* exists in `requestQueue` and which of the two cases above needs to be handled. If the request was not handled by the dedicated CTL so far,  $\mathcal{F}_{\text{psync-net}}$  adds in `requestQueue` a (current) timestamp to entry *ctr* and forwards the monitor request (including the identifier *ctr*) to the CTL. In Case (ii), the `Deliver` message of  $\mathcal{A}$  triggers the delivery of the (stored) response of the CTL (see below) to the monitoring party. Before  $\mathcal{F}_{\text{psync-net}}$  delivers the response, it updates the entry with ID *ctr* in `requestQueue` and adds a (current) timestamp as delivery time.

When a CTL (immediately) responds to the request from Case (i) above, it sends (`Monitor, ctr, resp`) to  $\mathcal{F}_{\text{psync-net}}$ .  $\mathcal{F}_{\text{psync-net}}$  records the response in the dedicated entry for *ctr* in `requestQueue`. Note that (i) CTLs are modeled to be incorruptible and (ii) CTLs and  $\mathcal{F}_{\text{psync-net}}$  have a direct I/O connection for handling monitoring request, i. e., responses from CTLs to  $\mathcal{F}_{\text{psync-net}}$  will be delivered “immediately” and cannot be manipulated by  $\mathcal{A}$ . Thus,  $\mathcal{A}$  is not activated when a CTL generates a response to a monitoring request and sends it to  $\mathcal{F}_{\text{psync-net}}$ . Therefore, we do not store an additional timestamp to determine when the CTL replies to a response. This timestamp is always equal to the timestamp stored in Case (i) from above.

The adversary is allowed to update the internal clock of  $\mathcal{F}_{\text{psync-net}}$  via the `UpdateRound` command.  $\mathcal{F}_{\text{psync-net}}$  accepts a time update (and increases the time by one) iff (i) there is no monitoring process initially triggered more than  $\delta$  time units ago but without response from a CTL, (ii) there is no response from a monitoring process waiting for delivery for more then  $\delta$  time units (both properties can be derived from `requestQueue`), and (iii) every honest party started a monitoring process at least  $\delta$  time units ago (stored in `monReg`).  $\mathcal{F}_{\text{psync-net}}$  also allows parties to read from its in-build clock (via `GetCurRound`).

The network functionality  $\mathcal{F}_{\text{psync-net}}$  is incorruptible.

**The local judge**  $\mathcal{P}_{\text{judge}}^{\text{CA}}$ . The (local) judge  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  (cf. Figure 28) handles when a client finds a maliciously issued certificate for its identity.

On every activation,  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  updates/pulls the list of CAs, resp. CTLs, which are currently available in the PKI, including their public keys (via `InitMejudge` from  $\mathcal{F}_{\text{init}}^{\text{CA}}$ ).

In case that a client finds a certificate during monitoring issued for her identify which she did not request, the client forwards the certificate to her local judge.  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  verifies that the certificate is valid, i. e., CA and CTL signature verify. As this indicates that the issuing CA misbehaved (at least from the current clients perspective, as clients requests certificates via an authenticated channel at CA),  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  adds  $\text{dis}(\text{pid}_{\text{CA}}, \text{sid}_{\text{cur}}, \mathcal{P}^{\text{CA}} : \text{ca})$  (where  $(\text{pid}_{\text{CA}}, \text{sid}_{\text{cur}}, \mathcal{P}^{\text{CA}} : \text{ca})$  denotes the CA that attested the reported maliciously generated certificate) to her (local) verdict. Note that external observers or the environment cannot trust the verdict of  $\mathcal{P}_{\text{judge}}^{\text{CA}}$ . From a higher-level perspective, one cannot be sure which of the two parties, the CA or the client, indeed misbehaved in this case. A malicious client (in turn a malicious judge) could have requested the certificate at an honest CA but then blames the CA that she did not request the certificate. Vice versa, CAs can simply generate certificates for any client.  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  also marks `correctCert` broken for its associated client.

$\mathcal{P}_{\text{judge}}^{\text{CA}}$  implements the `GetVerdict` interface which simply outputs the stored local verdict and a interface to access a judicial report. On receiving `GetJudicialReport`,  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  does not provide information and outputs  $\varepsilon$ .

Corrupting a local judge as long its accompanied client party is still honest is not allowed. This models that a clients validation routine still works honestly/correctly as long as the accompanied client is honest/works correctly. Note that  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  is a pure message forwarder to/from  $\mathcal{A}$  in case that the associated client is corrupted.

**The public supervisor  $\mathcal{P}_{\text{sv}}^{\text{CA}}$ .** As we do not model assumption-based security properties in this case study, the purpose of the supervisor  $\mathcal{P}_{\text{sv}}^{\text{CA}}$  (cf. Figure 29) is to provide the environment access to the corruption status of internal protocol participants (here: CAs). Thus, on request (`corruptInt?`,  $(pid, sid, role)$ ) where  $role$  needs to be `ca`,  $\mathcal{P}_{\text{sv}}^{\text{CA}}$  queries  $(pid, sid, role)$  for its corruption status and forwards the result to the requestor.

Note that we include a “dummy” interface for `IsAssumptionBroken?`. On every request,  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  simply replies “false”.

**The ideal authenticated channel functionality  $\mathcal{F}_{\text{auth}}$ .** The authenticated channel functionality  $\mathcal{F}_{\text{auth}}$  (cf. Figure 30) is derived from Canetti’s message authentication functionality [22] in combination with the secure channel functionality  $\mathcal{F}_{\text{sec-channel}}$  from [70]. In our case study, we use  $\mathcal{F}_{\text{auth}}$  to authenticate clients to CAs. Thus, CAs can be sure that a client indeed registers a certificate for its identity.

One instance of  $\mathcal{F}_{\text{auth}}$  is meant to handle all authenticated channels in an instance/session of  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$ . The whole session of  $\mathcal{F}_{\text{auth}}$  can be corrupted (as long as there are no channels initialized yet). Parties “register” at  $\mathcal{F}_{\text{auth}}$  via `Establish`. The adversary may trigger the finalization the party’s registration. Parties can send messages to specified receivers via  $\mathcal{F}_{\text{auth}}$ .  $\mathcal{F}_{\text{auth}}$  buffers the messages and  $\mathcal{A}$  may trigger the delivery of the message. Note that  $\mathcal{A}$  has full control over the messages send to/delivered from  $\mathcal{F}_{\text{auth}}$  when  $\mathcal{F}_{\text{auth}}$  is corrupted. Furthermore,  $\mathcal{A}$  may drop messages which are queued for delivery in  $\mathcal{F}_{\text{auth}}$ .

**The ideal signature functionality  $\mathcal{F}_{\text{sig}}$ .** The ideal signature functionality  $\mathcal{F}_{\text{sig}}$  (cf. Figure 31) is mainly the same as presented in [17] and matches common approaches for ideal signature functionalities.

As we need to ensure in  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  that CTL signatures are incorruptible,  $\mathcal{F}_{\text{sig}}$  is additionally connected to  $\mathcal{F}_{\text{init}}^{\text{CA}}$  to determine whether a party is a CTL.  $\mathcal{F}_{\text{sig}}$  declines corruption of CTLs.

In contrast to [17], we do not allow  $\mathcal{A}$  to corrupt `verifier` instances of  $\mathcal{F}_{\text{sig}}$  as one typically expects that signature verification is a local process and not externalized and thus honest as long the considered party is honest. Indeed, we need to ensure for the model that at least `verifier` instance of (local) judges need to be honest as long as the accompanied judge is honest. Otherwise, the security proof would fail.

**E.3 UC Security Analysis:** We now present the formal variant of Theorem 2 and its security proof in detail.

**Theorem 7.** *Let  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  be as defined above and let  $\delta \in \mathbb{N}$  be the upper time boundary for message delivery in  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$ , resp.  $\mathcal{F}_{\text{psync-net}}$ , and  $\Sigma = (\text{gen}(1^\eta), \text{sig}, \text{ver})$  be an EUF-CMA secure signature scheme.*

*Let  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  be as described above with  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}}$  as subroutine with parameters  $\text{Sec}^{\text{acc}} = \{\text{correctCert}\}$ ,  $\text{Sec}^{\text{assumption}} = \emptyset$ ,  $\text{pids}_{\text{judge}} = \{\text{local}\} \times \{0, 1\}^* \times \{\text{client}\}$ , and  $\text{ids}_{\text{assumption}} = \emptyset$ . Then:*

$$(\mathcal{P}_{\text{PKI}}^{\text{acc}}) \leq (\mathcal{F}_{\text{PKI}}^{\text{acc}} \mid \mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}})$$

We remark that the following proof repeats many general points/techniques from the proof of Theorem 6, e.g., regarding the simulator, in verbatim. This is mainly due to the fact that both protocols – the scaling protocol and the PKI protocol – rely on some kind of global state. In the scaling protocol, the global state is a sequence of totally ordered transactions. In the PKI protocol, the global state consists of a set of certificates.

*Proof.*

Firstly, we define a responsive simulator  $\mathcal{S}$  such that the real world running the protocol  $\mathcal{R} := \mathcal{P}_{\text{PKI}}^{\text{acc}} = (\mathcal{P}_{\text{client}}^{\text{CA}} : \text{client}, \mathcal{P}_{\text{sv}}^{\text{CA}} : \text{supervisor}, \mathcal{P}_{\text{judge}}^{\text{CA}} : \text{judge} \mid \mathcal{P}^{\text{CA}} : \text{ca}, \mathcal{P}_{\text{CTL}} : \text{ctl}, \mathcal{F}_{\text{sig}}, \mathcal{F}_{\text{psync-net}} : \text{psync-net}, \mathcal{F}_{\text{auth}} : \text{auth})$  is indistinguishable from the ideal world running  $\{\mathcal{S}, \mathcal{I}\}$ , with the protocol  $\mathcal{I} := (\mathcal{F}_{\text{PKI}}^{\text{acc}} \mid \mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}})$ , for every ppt environment  $\mathcal{E}$ .

The simulator  $\mathcal{S}$  is defined as follows: as common,  $\mathcal{S}$  is a single machine. It is connected to  $\mathcal{I}$  and the environment  $\mathcal{E}$  via their network interfaces. In a run, there is only a single instance of the machine  $\mathcal{S}$  that accepts and processes all incoming messages. The simulator  $\mathcal{S}$  internally simulates the realization  $\mathcal{R}$ , including its behavior on the network interface connected to the environment, and uses this simulation to compute responses



Description of a CA client  $\mathcal{P}_{\text{client}}^{\text{CA}} = (\text{client})$ :

<b>Participating roles:</b> {client} <b>Corruption model:</b> <i>Dynamic corruption without secure erasures</i> <b>Protocol parameters:</b> - $\delta \in \mathbb{N}$	{Network delay}
--	-----------------

Description of  $M_{\text{client}}$ :

<b>Implemented role(s):</b> {client} <b>Subroutines:</b> $\mathcal{F}_{\text{auth}} : \text{auth}$ , $\mathcal{F}_{\text{sig}} : \text{signer}$ , $\mathcal{F}_{\text{sig}} : \text{verifier}$ , $\mathcal{P}_{\text{judge}}^{\text{CA}} : \text{judge}$ , $\mathcal{F}_{\text{psync-net}} : \text{psync-net}$ , $\mathcal{F}_{\text{init}}^{\text{CA}} : \text{init}$ <b>Internal state:</b> - $\text{pk} \in \{0, 1\}^*$ , $\text{pk} = \perp$ {Public key of the client} - $\text{certCA} \in \{0, 1\}^*$ , $\text{certCA} = \perp$ {The CA which should create $\text{pid}_{\text{cur}}$ 's certificate} - $\text{requests} \subset \{0, 1\}^* \times (\{0, 1\}^*)^3$ , $\text{requests} = \emptyset$ {The recorded requests of $\text{pid}_{\text{cur}}$ including requesting party. - $\text{pk}_{\text{CA}} : \{0, 1\}^* \rightarrow \{0, 1\}^*$ {The pubkeys of the CAs, initially $\perp$ - $\text{CTLs} : \{0, 1\}^* \rightarrow \{0, 1\}^*$ {The CTLs of each CA, initially $\perp$ - $\text{pk}_{\text{CTL}} : \{0, 1\}^* \rightarrow \{0, 1\}^*$ {Pubkeys of CTLs, initially $\perp$ <b>CheckID(<math>\text{pid}</math>, <math>\text{sid}</math>, <math>\text{role}</math>):</b> Accept all messages for the same $\text{sid}$ . <b>Corruption behavior:</b> <b>DetermineCorrStatus(<math>\text{pid}</math>, <math>\text{sid}</math>, <math>\text{role}</math>):</b> <b>if</b> corrupted = true: {Checks whether party itself is corrupted. <b>return</b> true $\text{corrRes}_1 \leftarrow \text{corr}(\text{pid}_{\text{cur}}, (\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}), \text{signer})$ {Request corruption status at $\mathcal{F}_{\text{sig}}$ <b>send</b> CorruptionStatus? <b>to</b> ( $\_$ , $\text{sid}_{\text{cur}}$ , $\mathcal{F}_{\text{auth}} : \text{auth}$ ) {Query $\mathcal{F}_{\text{auth}}$ for its corruption status. <b>wait for</b> (CorruptionStatus?, $\text{corrRes}_2$ ) <b>return</b> ( $\text{corrRes}_1 \vee \text{corrRes}_2$ ) {Return whether $\mathcal{F}_{\text{sig}}$ or $\mathcal{F}_{\text{auth}}$ instances are corrupted}	
<b>EntityInitialization:</b> <b>send</b> Establish <b>to</b> ( $\text{pid}_{\text{cur}}$ , $\text{sid}_{\text{cur}}$ , $\mathcal{F}_{\text{psync-net}} : \text{psync-net}$ ); <b>wait for</b> $\_$ {Register party at $\mathcal{F}_{\text{psync-net}}$ <b>if</b> ITM was not activated via <b>init</b> macro, resp. a <b>InitEntity</b> message: <b>send</b> initPartner <b>to</b> ( $\text{pid}_{\text{cur}}$ , $\text{sid}_{\text{cur}}$ , $\mathcal{F}_{\text{init}}^{\text{CA}} : \text{init}$ ) <b>wait for</b> initPartner	
<b>Main:</b> <b>recv</b> (Register, $\text{pk}$ , $\text{pid}_{\text{CA}}$ ) <b>from</b> I/O: {Client should register certificate <b>if</b> $\text{pk} = \varepsilon$ : {Store public key. $\text{pk} \leftarrow \text{pk}$ , $\text{certCA} \leftarrow \text{pid}_{\text{CA}}$ <b>send</b> (Send, ( $\text{pid}_{\text{CA}}$ , $\text{sid}_{\text{cur}}$ , $\text{ca}$ ), (Register, $\text{pid}_{\text{cur}}$ , $\text{pk}$ ) <b>to</b> ( $\text{pid}_{\text{cur}}$ , $\text{sid}_{\text{cur}}$ , $\mathcal{F}_{\text{auth}} : \text{auth}$ )	
<b>recv</b> (Retrieve, $\text{pid}$ , $\text{pid}_{\text{CA}}$ ) <b>from</b> I/O: {Query for certificate $\text{requests.add}([\text{pid}, \text{pid}_{\text{CA}}], (\text{pid}_{\text{call}}, \text{sid}_{\text{call}}, \text{role}_{\text{call}}))$ <b>send</b> (Retrieve, $\text{pid}$ , $\text{pid}_{\text{CA}}$ ) <b>to</b> NET {Retrieve requests are dispatched via $\mathcal{A}$ , i.e., there are no delivery gura	
<b>recv</b> (Retrieve, ( $\text{ctr}$ , ( $\text{serial}$ , $\text{pid}$ , $\text{pk}$ , $\text{pid}_{\text{CA}}$ , $\sigma_{\text{CA}}$ ), $\text{pid}_{\text{CTL}}$ , $\tau$ , $\sigma_{\text{CTL}}$ )) <b>from</b> NET: {Query trusted CA/CTL information at $\mathcal{F}_{\text{init}}^{\text{CA}}$ <b>send</b> (GetCaCtIs, $\text{pid}_{\text{CA}}$ ) <b>to</b> ( $\text{pid}_{\text{cur}}$ , $\text{sid}_{\text{cur}}$ , $\mathcal{F}_{\text{init}}^{\text{CA}} : \text{init}$ ) <b>wait for</b> (GetCaCtIs, $\text{pk}_{\text{CA}}$ , ( $\text{pid}_1, \dots, \text{pid}_n$ ), ( $\text{pk}_1, \dots, \text{pk}_n$ )) $\text{pk}[\text{pid}_{\text{CA}}] \leftarrow \text{pk}_{\text{CA}}$ {Store trusted CA public key $\text{CTLs}[\text{pid}_{\text{CA}}] \leftarrow (\text{pid}_1, \dots, \text{pid}_n)$ <b>for</b> $i \in \{1, \dots, n\}$ <b>do</b> : {Store CTL information $\text{pk}_{\text{CTL}}[\text{pid}_i] \leftarrow \text{pk}_i$ <b>send</b> GetCurRound <b>to</b> ( $\text{pid}_{\text{cur}}$ , $\text{sid}_{\text{cur}}$ , $\text{psync-net}$ ) {Get current time <b>wait for</b> (GetCurRound, $\tau'$ ) <b>if</b> $\exists ([\text{pid}, \text{pid}_{\text{CA}}], (\text{pid}_c, \text{sid}_c, \text{role}_c)) \in \text{requests} \wedge \text{pid}_{\text{CTL}}$ is in $\text{CTLs}[\text{pid}_{\text{CA}}] \wedge \tau' > 3 \cdot \delta + \tau$ : {Handle answer only if there is a matching request and $\text{pid}_{\text{CTL}}$ is registered as CTL for $\text{pid}_{\text{CA}}$ and certificate is sufficient old to be considered valid $b_1 \leftarrow \text{verifySig}([\text{ctr}, (\text{serial}, \text{pid}, \text{pk}, \text{pid}_{\text{CA}}, \sigma_{\text{CA}}), \text{pid}_{\text{CTL}}, \tau], \sigma_{\text{CTL}}, \text{pk}_{\text{CTL}}[\text{pid}_{\text{CTL}}], \text{pid}_{\text{CTL}})$ {Check CTL signature $b_2 \leftarrow \text{verifySig}([\text{serial}, \text{pid}, \text{pk}, \text{pid}_{\text{CA}}], \sigma_{\text{CA}}, \text{pk}_{\text{CA}}[\text{pid}_{\text{CA}}], \text{pid}_{\text{CA}})$ {Check CA signature <b>if</b> $b_1 \wedge b_2$ : $\text{requests.remove}([\text{pid}, \text{pid}_{\text{CA}}], [\text{pid}_c, \text{sid}_c, \text{role}_c])$ {Remove request <b>send</b> (Retrieve, $\text{pid}$ , $\text{pid}_{\text{CA}}$ , $\text{pk}$ ) <b>to</b> ( $\text{pid}_c$ , $\text{sid}_c$ , $\text{role}_c$ )	
<b>recv</b> Establish <b>from</b> NET: { $\mathcal{A}$ triggers the registration at $\mathcal{F}_{\text{auth}}$ <b>send</b> (Establish, $\varepsilon$ ) <b>to</b> ( $\text{pid}_{\text{cur}}$ , $\text{sid}_{\text{cur}}$ , $\mathcal{F}_{\text{auth}} : \text{auth}$ )	

Fig. 22: The CA client  $\mathcal{P}_{\text{client}}^{\text{CA}}$  (Part 1).

```

Main:
  recv Monitor from NET:                                     {A triggers certificate monitoring}
  send GetCTLs to (pidcur, sidcur,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  : init)      {Request CTLs at  $\mathcal{F}_{\text{init}}^{\text{CA}}$ }
  wait for (GetCTLs, CTL, pidCTL)
  send (Monitor, pidcur, CTL) to (pidcur, sidcur,  $\mathcal{F}_{\text{psync-net}}$  : psync-net) {Send request to  $\mathcal{F}_{\text{psync-net}}$ }

  recv (Monitor, ctr, msg) from NET:                         {Check monitored certificates}
  send GetCAs to (pidcur, sidcur,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  : init)      {Request CAs at  $\mathcal{F}_{\text{init}}^{\text{CA}}$ }
  wait for (GetCAs, CA, pkCA)
  send GetCTLs to (pidcur, sidcur,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  : init)      {Request CTLs at  $\mathcal{F}_{\text{init}}^{\text{CA}}$ }
  wait for (GetCTLs, CTL, pkCTL)
  if msg = {(ctr', [serial, pidcur, pk, pid'CA,  $\sigma_{\text{CA}}$ ], pid'CTL,  $\tau$ ,  $\sigma_{\text{CTL}}$ ), ...}: {Verify signatures}
  for all (ctr', [serial, pidcur, pk, pid'CA,  $\sigma_{\text{CA}}$ ], pid'CTL,  $\tau$ ,  $\sigma_{\text{CTL}}$ ) from above set do:
    b  $\leftarrow$  verifySig([(ctr', [serial, pidcur, pk, pid'CA,  $\sigma_{\text{CA}}$ ], pid'CTL,  $\tau$ ,  $\sigma_{\text{CTL}}$ )] {Check CTL signature}
    if pid'CA  $\in$  CA:
      c  $\leftarrow$  verifySig([serial, pidcur, pk, pid'CA,  $\sigma_{\text{CA}}$ , pkCA[pid'CA], pid'CA]) {Check CA signature}
    else:
      c  $\leftarrow$  false
    if b  $\wedge$  c: {Certificate is validly signed}
      if pk  $\neq$  pk:
        send (EvidenceM, (ctr', [serial, pid, pk, pid'CA,  $\sigma_{\text{CA}}$ ], pid'CTL,  $\tau$ ,  $\sigma_{\text{CTL}}$ )) to ((local, pid, client), sidcur,  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  :
judge) {Forward evidence to  $\mathcal{P}_{\text{judge}}^{\text{CA}}$ }

    // In the model: stop processing after finding the first maliciously generated certificate //

Procedures and Functions:
  function verifySig(msg,  $\sigma$ , pk, pid) : {Verify signature at  $\mathcal{F}_{\text{sig}}$ }
  send (VerResult, msg,  $\sigma$ , pk) to ((pidcur, ct1), (pid, sidcur),  $\mathcal{F}_{\text{sig}}$  : verifier); wait for (VerResult, result)
  return result

```

Fig. 23: The CA client  $\mathcal{P}_{\text{client}}^{\text{CA}}$  (Part 2).

to incoming messages. For ease of presentation, we will refer to this internal simulation by  $\mathcal{R}'$ . More precisely, the simulation runs as follows:

#### Network communication from/to the environment

- Messages that  $\mathcal{S}$  receives on the network interface (connected to the environment) and which are hence meant for  $\mathcal{R}$  are forwarded to the internal simulation  $\mathcal{R}'$ .
- Any messages sent by  $\mathcal{R}'$  on its network interface (that are hence meant for the environment) are forwarded to the environment  $\mathcal{E}$ .

#### Corruption handling

- The simulator  $\mathcal{S}$  keeps the corruption status of entities in  $\mathcal{R}'$  and  $\mathcal{I}$  synchronized. That is, whenever an entity of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  or  $\mathcal{P}^{\text{CA}}$  in  $\mathcal{R}'$  starts to consider itself corrupted, the simulator corrupts the corresponding (internal) entity of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  in  $\mathcal{I}$  before continuing its simulation. Note that corruption of internal entities, i. e., of  $\mathcal{P}^{\text{CA}}$ , is mapped to a corruption of an internal party in  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  in  $\mathcal{I}$ .
- Incoming messages from corrupted (non-internal) entities of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  in  $\mathcal{I}$  are forwarded on the network interface to the environment in the name of the corresponding entity/instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  in  $\mathcal{R}'$ . Conversely, whenever a corrupted entity of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  wants to output a message to a higher-level protocol,  $\mathcal{S}$  instructs the corresponding entity of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  to output the same message to the higher-level protocol.

#### Certificate registration

Whenever an honest entity  $entity = (pid, sid, client)$  of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  receives a request (Register,  $pk, pid_{\text{CA}}$ ) to register a certificate, resp. a public key, for its own identity  $pid$ ,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  buffers the request in  $buffer_{\mathcal{S}}$  and leaks the full message buffer to  $\mathcal{S}$ . The simulator  $\mathcal{S}$  uses the leaked message to simulate the registration of  $pk$  for  $pid$  at  $pid_{\text{CA}}$  in  $\mathcal{R}'$ , i. e., it simulates the input of (Register,  $pk, pid_{\text{CA}}$ ) to the  $(pid, sid, client)$ 's instance of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  in  $\mathcal{R}'$ .

We remark that  $\mathcal{S}$  does not need to handle the registration of certificates from corrupted parties.  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  does not provide guarantees for these certificates and allows  $\mathcal{A}$  to freely determine the output in case that a party requests the certificate of a corrupted party.

#### Retrieve requests

Whenever the environment instructs an honest entity  $(pid, sid, client)$  via (Retrieve,  $pid_{\text{req}}, pid_{\text{CA}}$ ) to request a certificate of  $pid_{\text{req}}$  from  $pid_{\text{CA}}$ ,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  buffers this request (in  $buffer_{\mathcal{R}}$ ) and leaks the full request to  $\mathcal{S}$ . Upon receiving the leakage from  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ ,  $\mathcal{S}$  forwards this message to  $\mathcal{P}_{\text{client}}^{\text{CA}}$  in  $\mathcal{R}'$  and simulates its behavior.

**Participating roles:** {ca}**Corruption model:** *Dynamic corruption without secure erasures*Description of  $M_{\text{ca}}$ :**Implemented role(s):** {ca}**Subroutines:**  $\mathcal{F}_{\text{init}}^{\text{CA}} : \text{init}$ ,  $\mathcal{F}_{\text{sig}} : \text{signer}$ ,  $\mathcal{F}_{\text{sig}} : \text{verifier}$ ,  $\mathcal{P}_{\text{judge}}^{\text{CA}} : \text{judge}$ **Internal state:**

- $\text{pk}_{\text{CA}} \in \{0, 1\}^*$ ,  $\text{pk}_{\text{CA}} = \varepsilon$  {The public key of CA ( $\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \text{role}_{\text{cur}}$ )}
- $\text{CTLs} \subset \{0, 1\}^*$ ,  $\text{CTLs} = \emptyset$  {The CA's CTL's pids}
- $\text{pk}_{\text{CTL}} : \text{CTLs} \rightarrow \{0, 1\}^*$  {CTLs' pubkeys, initially  $\perp$ }
- $\text{serial} \in \mathbb{N}$ ,  $\text{serial} = 0$  {The serial number of certificates issued by entity is a counter}
- $\text{certs} \subseteq \mathbb{N} \times (\{0, 1\}^*)^4$ ,  $\text{certs} = \emptyset$  {CAs issued certificates of form (serial, subject, pk, issuer,  $\sigma$ )}
- $\text{certs}_{\text{CTL}} : \text{CTLs} \rightarrow \{0, 1\}^*$  {The state of each CTL after reading initially  $\perp$ }

**CheckID**( $\text{pid}, \text{sid}, \text{role}$ ):Accept all messages for ( $\text{pid}, \text{sid}, \text{role}$ ) where  $\text{sid} = (\text{pid}, \text{sid}')$  and  $\text{role} = \text{ca}$ .Do not accept an entity ( $\text{pid}, \text{sid}, \text{role}$ ) if  $\text{pid}$  can be parsed to ( $\text{pid}', \text{client}$ ){Avoid local judges like PIDs}**Corruption behavior:****DetermineCorrStatus**( $\text{pid}, \text{sid}, \text{role}$ ):

if corrupted = true:

{Checks whether party itself is corrupted.}

return true

 $\text{corrRes}_1 \leftarrow \text{corr}(\text{pid}_{\text{cur}}, (\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}), \text{signer})$ {Request corruption status at  $\mathcal{F}_{\text{sig}}$ }send CorruptionStatus? to ( $\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \mathcal{F}_{\text{auth}} : \text{auth}$ ){Query  $\mathcal{F}_{\text{auth}}$  for its corruption status.}wait for (CorruptionStatus?,  $\text{corrRes}_2$ ) $\text{corrRes}_3 \leftarrow \text{false}$ {Check that all used verifier instances are not corrupted}for all queries to ( $\text{pid}_{\text{cur}}, (\text{pid}', \text{sid}')$ , verifier) at  $\mathcal{F}_{\text{sig}}$  from transcript do: $\text{corrRes}_3 = \text{corrRes}_3 \vee \text{corr}(\text{pid}_{\text{cur}}, (\text{pid}', \text{sid}'), \text{verifier})$ return ( $\text{corrRes}_1 \vee \text{corrRes}_2 \vee \text{corrRes}_3$ ){Return whether  $\mathcal{F}_{\text{sig}}$  or  $\mathcal{F}_{\text{auth}}$  instances are corrupted}**Initialization:**send InitMeCA to ( $\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \mathcal{F}_{\text{init}}^{\text{CA}} : \text{init}$ ){Request CA setup at  $\mathcal{F}_{\text{init}}^{\text{CA}}$ }wait for (InitMeCA,  $\text{pk}_{\text{CA}}, (\text{pid}_1, \dots, \text{pid}_n), (\text{pk}_1, \dots, \text{pk}_n), \text{pk}_{\text{ctl}}$ ){Store (trusted) public keys and used CTLs} $\text{pk}_{\text{CA}} \leftarrow \text{pk}_{\text{CA}}$  $\text{CTLs} \leftarrow \{\text{pid}_1, \dots, \text{pid}_n\}$ for all  $i \in \{1, \dots, n\}$  do: $\text{pk}_{\text{CTL}}[\text{pid}_i] \leftarrow \text{pk}_i$ **Main:**recv (Received, ( $\text{pid}, \text{sid}_{\text{cur}}, \text{role}$ ), (Register,  $\text{pid}, \text{pk}$ )) from I/O:{pid registers public key/certificate via  $\mathcal{F}_{\text{auth}}$ }if  $b \wedge \nexists (\_, \text{pid}, \_, \_, \_) \in \text{certs}$ :{One certificate per party}serial  $\leftarrow$  serial + 1 $\text{cert}' \leftarrow (\text{serial}, \text{pid}, \text{pk}, \text{pid}_{\text{cur}})$ {Simplified version of a certificate for pid w/o signature} $\sigma' \leftarrow \text{sign}(\text{cert}')$ {Generate signature for certificate}certs.add( $[\text{cert}', \sigma']$ ){Record certificate}send (Submit,  $\text{cert}', \sigma', \text{CTLs}$ ) to NET{Publish certificate at CTLs}recv (Read,  $\text{pid}_{\text{CTL}}, \text{state}, \sigma$ ) from NET s.t.  $\text{pid}_{\text{CTL}} \in \text{CTLs}$ :{Current state of CTL pid} $b \leftarrow \text{verifySig}(\text{state}, \sigma, \text{pk}_{\text{CTL}}[\text{pid}_{\text{CTL}}], \text{pid}_{\text{CTL}})$ {Verify signature}if  $b \wedge \text{certs}_{\text{CTL}}[\text{pid}] \subset \text{state}$ :{ $\mathcal{P}^{\text{CA}}$  only accepts valid states which extend the current state}check  $\leftarrow$  truefor all ( $\text{ctr}, \text{cert}, \text{pid}, \sigma$ )  $\in$  state do:if  $\text{pid} \neq \text{pid}_{\text{CTL}}$ :{A CTL only has its "own" certificates}

break

 $b \leftarrow \text{verifySig}([\text{ctr}, \text{cert}], \sigma, \text{pk}_{\text{CTL}}[\text{pid}], \text{pid})$ if  $\neg b$ :check  $\leftarrow$  check  $\vee$  b{CTL list is valid}

if check:

certs<sub>CTL</sub>[pid]  $\leftarrow$  staterecv (Retrieve,  $\text{pid}, \text{pid}_{\text{cur}}$ ) from NET:{Client queries for certificate}if  $\exists (\text{serial}, \text{pid}, \text{pk}, \text{pid}_{\text{cur}}, \sigma) \in \text{certs}$ :if  $\exists \text{certs}_{\text{CTL}}[\text{pid}_{\text{CTL}}]$  of a CTL  $\text{pid}_{\text{CTL}}$ , s.t. ( $\text{ctr}, (\text{serial}, \text{pid}, \text{pk}, \text{pid}_{\text{cur}}, \sigma), \text{pid}_{\text{CTL}}, \tau, \sigma_{\text{CTL}}$ )  $\in$  certs<sub>CTL</sub>[ $\text{pid}_{\text{CTL}}$ ]:{Let  $\text{pid}_{\text{CTL}}$  be the first hit during search}reply (Retrieve, ( $(\text{ctr}, (\text{serial}, \text{pid}, \text{pk}, \text{pid}_{\text{cur}}, \sigma), \text{pid}_{\text{CTL}}, \tau, \sigma_{\text{CTL}})$ ))

recv Establish from NET:

{A triggers the registration at  $\mathcal{F}_{\text{auth}}$ }send (Establish,  $\varepsilon$ ) to ( $\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \mathcal{F}_{\text{auth}} : \text{auth}$ )**Procedures and Functions:**function verifySig( $\text{msg}, \sigma, \text{pk}, \text{pid}$ ):{Verify signature at  $\mathcal{F}_{\text{sig}}$ }send (VerResult,  $\text{msg}, \sigma, \text{pk}$ ) to ( $(\text{pid}_{\text{cur}}, \text{ctl}), (\text{pid}, \text{sid}_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{verifier}$ )wait for (VerResult,  $\text{result}$ )return  $\text{result}$ function sign( $\text{msg}$ ):{Sign message at  $\mathcal{F}_{\text{sig}}$ }send (sign,  $\text{msg}$ ) to ( $(\text{pid}_{\text{cur}}, \text{ctl}), (\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{signer}$ )wait for (sign,  $\sigma$ )return  $\sigma$ Fig. 24: The model of a CA  $\mathcal{P}^{\text{CA}}$ .

**Participating roles:**  $\{\text{ctl}\}$   
**Corruption model:** *incorruptible*

Description of  $M_{\text{ctl}}$ :

**Implemented role(s):**  $\{\text{ctl}\}$   
**Subroutines:**  $\mathcal{F}_{\text{sig}}$  : verifier,  $\mathcal{F}_{\text{sig}}$  : signer,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  : init  
**Internal state:**  
 -  $\text{state} \subset \mathbb{N} \times \{0, 1\}^* \times \{\text{pid}_{\text{cur}}\} \times \mathbb{N} \times \{0, 1\}^*$ ,  $\text{state} = \emptyset$  {The set of certificates of form  $(\text{ctr}, \text{cert}, \text{pid}_{\text{cur}}, \tau, \sigma_{\text{CTL}})$ }  
 -  $\text{CAs} \subset \{0, 1\}^*$ ,  $\text{CAs} = \emptyset$  {The set of CAs, the CTL will finalize pre-certificates for}  
 -  $\text{pk}_{\text{CTL}} \in \{0, 1\}^*$  {The public key of the CTL}  
 -  $\text{pk}_{\text{CA}} : \text{CAs} \rightarrow \{0, 1\}^*$  {The (trusted) public keys of the CA}  
 -  $\text{counter} \in \mathbb{N}$ ,  $\text{counter} = 0$  {Serial}  
**CheckID**( $\text{pid}, \text{sid}, \text{role}$ ):  
 Accept all messages for the same  $\text{sid}$ .  
**Initialization:**  
**send**  $\text{InitMe}_{\text{CTL}}$  **to**  $(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \mathcal{F}_{\text{init}}^{\text{CA}} : \text{init})$  {Request initialization at  $\mathcal{F}_{\text{init}}^{\text{CA}}$ }  
**wait for**  $(\text{InitMe}_{\text{CTL}}, \text{pk}_{\text{CTL}}, (\text{pid}_1, \dots, \text{pid}_n), (\text{pk}_1, \dots, \text{pk}_n))$   
 $\text{pk}_{\text{CTL}} \leftarrow \text{pk}_{\text{CTL}}$   
 $\text{CAs} \leftarrow \{\text{pid}_1, \dots, \text{pid}_n\}$  {Store (trusted) CAs and their public keys}  
**for all**  $i \in \{1, \dots, n\}$  **do**:  
 $\text{pk}_{\text{CA}}[\text{pid}_i] \leftarrow \text{pk}_i$   
**Main:**  
**recv** (Submit,  $\text{cert}$ ) **from** NET: {Certificate submission}  
**require:**  $\text{cert} = (\text{serial}, \text{pid}_r, \text{pk}_r, \text{pid}_{\text{CA}}, \sigma_{\text{CA}})$ , **s.t.**  $\text{pid}_{\text{CA}} \in \text{CAs}$   
 $b \leftarrow \text{verifySig}(\text{serial}, \text{pid}_r, \text{pk}_r, \text{pid}_{\text{CA}}, \sigma_{\text{CA}}, \text{pk}[\text{pid}_{\text{CA}}], \text{pid}_{\text{CA}})$  {Check valid signature}  
**if**  $b$ : {If signature is valid, record data in state}  
**send**  $\text{GetCurRound}$  **to**  $(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \mathcal{F}_{\text{psync-net}} : \text{psync-net})$  {Query current time at  $\mathcal{F}_{\text{psync-net}}$ }  
**wait for**  $(\text{GetCurRound}, \tau)$   
 $\text{counter} \leftarrow \text{counter} + 1$   
 $\sigma_{\text{CTL}} \leftarrow \text{sign}([\text{counter}, \text{cert}, \text{pid}_{\text{cur}}, \tau])$  {Mark certificate as valid}  
 $\text{state.add}([\text{counter}, \text{cert}, \text{pid}_{\text{cur}}, \tau, \sigma_{\text{CTL}}])$   
**recv** Read: {Read request}  
 $\sigma \leftarrow \text{sign}(\text{state})$  {Sign full state}  
**reply** (Read,  $\text{pid}_{\text{cur}}, \text{state}, \sigma$ ) {Reply with full state including signatures}  
**recv** (Monitor,  $\text{pid}, \text{ctr}$ ) **from** I/O: {Monitor maliciously published certificates}  
 $\text{resp} \leftarrow \varepsilon$   
**for all**  $(\text{ctr}', [\text{serial}, \text{pid}, \text{pk}, \text{pid}_{\text{CA}}, \sigma_{\text{CA}}], \text{pid}_{\text{cur}}, \tau, \sigma_{\text{CTL}})$  **in** state **do**: {Collect certificates for  $\text{pid}$ }  
 $\text{resp.add}(\text{ctr}', [\text{serial}, \text{pid}, \text{pk}, \text{pid}_{\text{CA}}, \sigma_{\text{CA}}], \text{pid}_{\text{cur}}, \tau, \sigma_{\text{CTL}})$   
**reply** (Monitor,  $\text{ctr}, \text{resp}$ ) {Return certificate for  $\text{pid}$  to requestor}  
**Procedures and Functions:**  
**function**  $\text{verifySig}(\text{msg}, \sigma, \text{pk}, \text{pid})$  : {Verify signature at  $\mathcal{F}_{\text{sig}}$ }  
**send**  $(\text{VerResult}, \text{msg}, \sigma, \text{pk})$  **to**  $((\text{pid}_{\text{cur}}, \text{role}_{\text{cur}}), (\text{pid}, \text{sid}_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{verifier})$   
**wait for**  $(\text{VerResult}, \text{result})$   
**return**  $\text{result}$   
**function**  $\text{sign}(\text{msg})$  : {Sign message at  $\mathcal{F}_{\text{sig}}$ }  
**send**  $(\text{sign}, \text{msg})$  **to**  $((\text{pid}_{\text{cur}}, \text{role}_{\text{cur}}), (\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{signer})$   
**wait for**  $(\text{sign}, \sigma)$   
**return**  $\sigma$

Fig. 25: The simplified CTL  $\mathcal{P}_{\text{CTL}}$ .

If  $\mathcal{A}$  triggers a (simulated) entity  $(\text{pid}, \text{sid}, \text{client})$  to respond on a retrieval request by sending (Retrieve,  $\text{cert}$ ) to the entity,  $\mathcal{S}$  simulates the input the message in  $\mathcal{R}'$ . If  $(\text{pid}, \text{sid}, \text{client})$  (in  $\mathcal{R}'$ ) wants to output  $\text{cert}$ , resp. parts of  $\text{cert}$  within a message of form (Retrieve,  $\text{pid}, \text{pid}_{\text{CA}}, \text{pk}$ ) to  $(\text{pid}_c, \text{sid}_c, \text{role}_c)$  at the end of its activation,  $\mathcal{S}$  sends (Deliver,  $\text{pid}, \text{pid}_{\text{CA}}, \text{pid}_c, \text{role}_c, \text{pk}$ ) to  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . This triggers  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  to deliver the retrieval response. Note, that in case of  $\text{brokenProps}[(\text{correctCert}, (\text{pid}, \text{sid}_{\text{cur}}, \text{client}))] = \text{true}$ , the output is forwarded without further checks. If  $\text{correctCert}$  still holds true,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  will only forward the output, if a matching request in  $\text{buffer}_{\mathcal{R}}$  exists.

#### State updates

When a CTL finalizes an *uncorrupted* party's certificate (for which  $\text{correctCert}$  still holds true) in  $\mathcal{R}'$ ,  $\mathcal{S}$  triggers an update at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  to include the finalized certificate in  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . Therefore, it firstly extracts the newly generated certificate from the CTL's state. In particular,  $\mathcal{S}$  extracts the identity  $\text{pid}$  for which the certificate is issued, the used public key  $\text{pk}$  as well as the used CA  $\text{pid}_{\text{CA}}$ .  $\mathcal{S}$  then sends (Update,  $\{(\text{pid}, \text{pk}, \text{pid}_{\text{CA}})\}$ ) to  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ .

**Participating roles:**  $\{\text{init}\}$   
**Corruption model:** *incorruptible*

Description of  $M_{\text{init}}$ :

**Implemented role(s):**  $\{\text{init}\}$   
**Subroutines:**  $\mathcal{F}_{\text{sig}} : \text{signer}$   
**Internal state:**

- $\text{pid}_{\text{CA}} \subset \{0, 1\}^*$ ,  $\text{pid}_{\text{CA}} = \emptyset$  {The pids of the CAs}
- $\text{pk}_{\text{CA}} : \text{pid}_{\text{CA}} \rightarrow \{0, 1\}^*$  {The public keys of CAs, initially  $\perp$  for all entries}
- $\text{pid}_{\text{CTL}} \subset \{0, 1\}^*$ ,  $\text{pid}_{\text{CTL}} = \emptyset$  {The pids of the CTLs}
- $\text{pk}_{\text{CTL}} : \text{pid}_{\text{CTL}} \rightarrow \{0, 1\}^*$  {The public keys of CTLs, initially  $\perp$  for all entries}
- $\text{CTLs} : \text{pid}_{\text{CA}} \rightarrow (\text{pid}_{\text{CTL}} \cup \{\perp\})^*$  {A CAs trusted CTLs, initially  $\perp$ }
- $\text{caller} \in (\{0, 1\}^*)^3 \cup \{\perp\} = \perp$  {Stores the callers during client/judge initialization.}

**CheckID**( $\text{pid}, \text{sid}, \text{role}$ ):  
 Accept all messages for the same  $\text{sid}$ .

**Main:**

```

recv InitMeCA from I/O: {CA request initialization}
send responsively (InitMeCA) to NET (*) {Query  $\mathcal{A}$  for initialization details of  $\text{pid}_{\text{cur}}$ }
wait for (InitMeCA,  $\text{pid}_1, \dots, \text{pid}_n$ )
if  $n < 1$ :
  go to (*)
for all  $i \in \{1, \dots, n\}$  do:
   $\text{corrRes} \leftarrow \text{corr}(\text{pid}_i, (\text{pid}_i, \text{sid}_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{signer})$  {Check whether CTL key is still uncorrupted}
  if  $\text{corrRes}$ :
    go to (*)
  if  $\text{pk}_{\text{CTL}}[\text{pid}_i] = \perp$ :
    send InitSign to  $(\text{pid}_i, (\text{pid}_i, \text{sid}_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{signer})$  {Generate CTL's signing keys}
    wait for (InitSign, success,  $\text{pk}_i$ )
     $\text{pk}_{\text{CTL}}[\text{pid}_i] \leftarrow \text{pk}_i$  {Store public key}
  if  $\text{pk}_{\text{CA}}[\text{pid}_{\text{cur}}] \neq \perp$ : {Establish session at  $\mathcal{F}_{\text{sig}}$  if it is not recorded}
    send InitSign to  $(\text{pid}_{\text{cur}}, (\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{signer})$  {Generate CA's signing keys}
    wait for (InitSign, success,  $\text{pk}_{\text{CA}}$ )
     $\text{pk}_{\text{CA}}[\text{pid}_{\text{cur}}] \leftarrow \text{pk}_{\text{CA}}$  {Store generated public key}
   $\text{CTLs}[\text{pid}_{\text{CA}}] \leftarrow (\text{pid}_1, \dots, \text{pid}_n)$  {Store CA's CTLs}
  reply (InitMeCA,  $\text{pk}_{\text{CA}}[\text{pid}_{\text{cur}}], (\text{pid}_1, \dots, \text{pid}_n), (\text{pk}_{\text{CTL}}[\text{pid}_1], \dots, \text{pk}_{\text{CTL}}[\text{pid}_n])$ ) {Output CA's and CTLs' pubkeys to CA}

recv InitMeCTL from I/O: {CTL request initialization}
if  $\text{corr}(\text{pid}_{\text{call}}, (\text{pid}_{\text{call}}, \text{sid}_{\text{call}}), \mathcal{F}_{\text{sig}} : \text{signer})$ : {CTL key is corrupted, thus this entity cannot act as (trusted) CTL}
  reply (InitMeCTL,  $\perp$ )
send responsively (InitMeCTL) to NET (*) {Query  $\mathcal{A}$  for initialization details of  $\text{pid}_{\text{cur}}$ }
wait for (InitMeCTL,  $\text{pid}_1, \dots, \text{pid}_n$ )
if  $n < 1$ :
  go to (*)
if  $\text{pk}_{\text{CTL}}[\text{pid}_{\text{cur}}] \neq \perp$ : {Establish session at  $\mathcal{F}_{\text{sig}}$  if it is not recorded}
  send InitSign to  $(\text{pid}_{\text{cur}}, (\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{signer})$  {Generate CTL's signing keys}
  wait for (InitSign, success,  $\text{pk}_{\text{CTL}}$ )
   $\text{pk}_{\text{CTL}}[\text{pid}_{\text{cur}}] \leftarrow \text{pk}_{\text{CTL}}$  {Store generated public key}
for all  $i \in \{1, \dots, n\}$  do:
  if  $\text{pk}_{\text{CA}}[\text{pid}_i] = \perp$ :
    send InitSign to  $(\text{pid}_i, (\text{pid}_i, \text{sid}_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{signer})$  {Generate CA's signing keys}
    wait for (InitSign, success,  $\text{pk}_i$ )
     $\text{pk}_{\text{CA}}[\text{pid}_i] \leftarrow \text{pk}_i$  {Store public key}
  reply (InitMeCTL,  $\text{pk}_{\text{CTL}}[\text{pid}_{\text{cur}}], (\text{pid}_1, \dots, \text{pid}_n), (\text{pk}_{\text{CA}}[\text{pid}_1], \dots, \text{pk}_{\text{CA}}[\text{pid}_n])$ ) {Output CA's and CTL's pubkey to CA}

recv InitMeJudge from I/O: { $\mathcal{P}_{\text{judge}}^{\text{CA}}$  requests initialization}
  reply (InitMeJudgeCA,  $\text{pid}_{\text{CA}}, \text{pk}_{\text{CA}}, \text{pid}_{\text{CTL}}, \text{pk}_{\text{CTL}}, \text{CTLs}$ ) {Output current state to  $\mathcal{F}_{\text{judge}}$ }

recv (GetCaCtls,  $\text{pid}_{\text{CA}}$ ) from I/O: {Clients can query for CA/CTL  $\text{pk}$ 's}
require:  $\text{pid}_{\text{CA}} \in \text{pid}_{\text{CA}} \wedge \text{pk}[\text{pid}_{\text{CA}}] \neq \perp$ 
  Let  $\text{pid}_1, \dots, \text{pid}_n$  be the pids from  $\text{CTLs}[\text{pid}_{\text{CA}}]$ 
  reply (GetCaCtls,  $\text{pk}_{\text{CA}}[\text{pid}_{\text{CA}}], (\text{pid}_1, \dots, \text{pid}_n), (\text{pk}_{\text{CTL}}[\text{pid}_1], \dots, \text{pk}_{\text{CTL}}[\text{pid}_n])$ )

recv GetCAs from I/O: {Clients can query for CAs and their  $\text{pk}$ 's}
  reply (GetCaCtls,  $\text{pid}_{\text{CA}}, \text{pk}_{\text{CA}}$ )

recv GetCTLs from I/O: {Clients can query for CTLs and their  $\text{pk}$ 's}
  reply (GetCTLs,  $\text{pid}_{\text{CTL}}, \text{pk}_{\text{CTL}}$ )

recv initPartner from  $(\text{pid}, \text{sid}, \text{role})$ : {Request to initialize accompanied local judge or client}
  caller  $\leftarrow (\text{pid}, \text{sid}, \text{role})$ 
  if  $\text{role} = \text{client}$ :
    init((local,  $\text{pid}$ , client),  $\text{sid}$ , judge) {Trigger initialization of the entities local judge}
  else if  $\text{role} = \text{judge}$ :
    Parse  $\text{pid}$  as (local,  $\text{pid}'$ , client)
    init( $\text{pid}'$ ,  $\text{sid}$ , client) {Trigger initialization of the accompanied client}

recv InitEntityDone from  $(\text{pid}, \text{sid}, \text{role})$ : {Finish initialization of accompanied local judge or client}
  caller  $\leftarrow \text{caller}$ ; caller  $\leftarrow \perp$ 
  send initPartner to caller

```

Fig. 26: The initialization ITM  $\mathcal{F}_{\text{init}}^{\text{CA}}$ .

Description of the certificate monitoring enforcing network functionality  $\mathcal{F}_{\text{psync-net}} = (\text{psync-net})$ :

**Participating roles:**  $\{\text{psync-net}\}$   
**Corruption model:** *incorruptible*  
**Protocol parameters:**  
 -  $\delta \in \mathbb{N}$  {The upper bound for liveness}

Description of  $M_{\text{psync-net}}^{\text{in}}$ :

**Implemented role(s):**  $\{\text{psync-net}\}$   
**Subroutines:**  $\mathcal{P}_{\text{client}}^{\text{CA}} : \text{client}, \mathcal{P}_{\text{CTL}} : \text{ctl}$   
**Internal state:**  
 -  $\tau \in \mathbb{N}, \tau = 0$  {Current time in the  $\mathcal{F}_{\text{psync-net}}$   
 -  $\text{counter} \in \mathbb{N}, \text{counter} = 0$  {Request counter}  
 -  $\text{requestQueue} \subset \mathbb{N} \times \mathbb{N} \times (\{0, 1\}^*)^4 \times \mathbb{N}^2, \text{requestQueue} = \emptyset$  {The set of messages to be delivered via  $\mathcal{F}_{\text{psync-net}}$  of form  $(ctr, t_{\text{rec}}, pid_s, pid_r, msg, resp, t_{\text{del}}, t_{\text{out}})$   
 -  $\text{monReg} : \{0, 1\}^* \rightarrow \mathbb{N}$  {Last recorded check request of party  $pid$ . Initially all set to 0.}  
 -  $\text{pid}_{\text{CTL}} \subset \{0, 1\}^*, \text{pid}_{\text{CTL}} = \emptyset$  {The available CTLs}  
 -  $\text{pid}_{\text{client}}^h \subset \{0, 1\}^*$  {The set of honest clients}  
**CheckID( $pid, sid, role$ ):**  
 Accept all messages for the same  $sid$ .  
**MessagePreprocessing:**  
**send** GetCTLs **to**  $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{init}}^{\text{CA}} : \text{init})$  {Request CTLs at  $\mathcal{F}_{\text{init}}^{\text{CA}}$ }  
**wait for** (GetCTLs,  $CTL, pid_{\text{CTL}}$ )  
 $pid_{\text{CTL}} \leftarrow CTL$  {Update available CTLs}  
**for all**  $pid \in \text{pid}_{\text{client}}^h$  **do:**  
**send** CorruptionStatus? **to**  $(pid, sid_{\text{cur}}, \mathcal{P}_{\text{client}}^{\text{CA}} : \text{client})$   
**wait for** (CorruptionStatus?,  $b$ )  
**if**  $b$ :  
 $\text{pid}_{\text{client}}^h \cdot \text{remove}(pid)$  {Remove corrupted party}  
 Remove entries for  $pid$  from requestQueue and monReg.  
**Main:**  
**recv** UpdateRound **from** NET: {Triggering a clock update increases the time}  
**if**  $\exists (ctr, t_{\text{rec}}, pid_s, pid_r, msg, resp', t_{\text{del}}, t_{\text{out}}) \in \text{requestQueue}, \text{s.t. } t_{\text{del}} = \perp \wedge \text{round} - t_{\text{rec}} > \delta$ :  
**reply** (UpdateRound,  $\text{ack}$ ) {Requests need to delivered in  $\delta$  time units to a CTL}  
**else if**  $\exists (ctr, t_{\text{rec}}, pid_s, pid_r, msg, resp', t_{\text{del}}, t_{\text{out}}) \in \text{requestQueue}, \text{s.t. } t_{\text{del}} \neq \perp \wedge t_{\text{rec}} \neq \perp \wedge t_{\text{out}} = \perp \wedge \text{round} - t_{\text{del}} > \delta$ :  
**reply** (UpdateRound,  $\text{ack}$ ) {Requests need to delivered in  $\delta$  time units to a CTL}  
**else if**  $\exists pid \in \text{monReg} \text{ s.t. } \text{round} - \text{monReg}[pid] > \delta$ :  
**reply** (UpdateRound,  $\text{ack}$ ) {A needs to allow honest parties to frequently trigger monitoring}  
**else:**  
 $\tau \leftarrow \tau + 1$  {Clock update successful}  
**recv** GetCurRound:  
**reply** (GetCurRound,  $\tau$ ) {Handling reads from the clock}  
**recv** (Monitor,  $pid, \{pid_1, \dots, pid_l\}$ ) **from** I/O **s.t.**  $pid_{\text{call}} = pid, \text{role}_{\text{call}} = \text{client}, \{pid_1, \dots, pid_l\} \subset \text{pid}_{\text{CTL}} : \{\text{Monitor CTLs}\}$  {Record monitor requests}  
**for all**  $pid' \in \{pid_1, \dots, pid_l\}$  **do:**  
 $\text{counter} \leftarrow \text{counter} + 1$   
 $\text{requestQueue.add}([\text{counter}, \text{round}, pid, pid', (\text{Monitor}, \text{counter}, pid), \perp, \perp, \perp, \perp])$  {Record requests to CTLs}  
 $\text{monReg}[pid_{\text{cur}}] \leftarrow \text{round}$  {Record last request of party}  
**send** (Monitor, requestQueue) **to** NET {Leak full information to A}  
**recv** (Deliver,  $ctr$ ) **from** NET: {A triggers delivery of message with id  $ctr$  to  $\mathcal{F}_{\text{CTL}}$  or  $\mathcal{P}_{\text{client}}^{\text{CA}}$ }  
**if**  $\exists (ctr, t_{\text{rec}}, pid_s, pid_r, msg, resp, t_{\text{del}}, t_{\text{out}}) \in \text{requestQueue}, \text{s.t. } t_{\text{del}} = \perp$ :  
 $\text{requestQueue.remove}([\text{ctr}, t_{\text{rec}}, pid_s, pid_r, msg, resp, t_{\text{del}}, t_{\text{out}}])$   
 $\text{requestQueue.add}([\text{ctr}, t_{\text{rec}}, pid_s, pid_r, msg, resp, \text{round}, t_{\text{out}}])$  {Record delivery to CTL}  
**send**  $(msg, ctr)$  **to**  $(pid_r, sid_{\text{cur}}, \mathcal{F}_{\text{CTL}} : \text{ctl})$  {Deliver request to CTL}  
**else if**  $\exists (ctr, t_{\text{rec}}, pid_s, pid_r, msg, resp, t_{\text{del}}, t_{\text{out}}) \in \text{requestQueue}, \text{s.t. } t_{\text{del}} \neq \perp \wedge t_{\text{out}} = \perp$ :  
 $\text{requestQueue.remove}([\text{ctr}, t_{\text{rec}}, pid_s, pid_r, msg, resp, t_{\text{del}}, t_{\text{out}}])$  {Record delivery to client}  
 $\text{requestQueue.add}([\text{ctr}, t_{\text{rec}}, pid_s, pid_r, msg, resp, t_{\text{del}}, \text{round}])$  {Deliver answer to client}  
**send**  $(resp)$  **to**  $(pid_s, sid_{\text{cur}}, \mathcal{P}_{\text{client}}^{\text{CA}} : \text{client})$   
**recv** (Monitor,  $ctr, resp$ ) **from** I/O: {Response from CTL on monitor request}  
**if**  $\exists (ctr, t_{\text{rec}}, pid_s, pid_r, msg, resp', t_{\text{del}}, t_{\text{out}}) \in \text{requestQueue}, \text{s.t. } t_{\text{out}} = \perp$ :  
 $\text{requestQueue.remove}([\text{ctr}, t_{\text{rec}}, pid_s, pid_r, msg, resp', t_{\text{del}}, t_{\text{out}}])$   
 $\text{requestQueue.add}([\text{ctr}, t_{\text{rec}}, pid_s, pid_r, msg, resp, t_{\text{del}}, t_{\text{out}}])$  {Record reponse of CTL}  
**send** Monitor, requestQueue **to** NET {Leak full information to A}  
**recv** Establish **from** I/O:  
 $\text{pid}_{\text{client}} \cdot \text{add}(pid_{\text{call}})$  {Client registers at  $\mathcal{F}_{\text{psync-net}}$ }  
 $\text{monReg}[pid_{\text{cur}}] \leftarrow \text{round}$  {Record registration round as starting point in monReg}

Fig. 27: The certificate monitoring enforcing network functionality  $\mathcal{F}_{\text{psync-net}}$  for modeling a partially synchronous network.

Description of  $\mathcal{P}_{\text{judge}}^{\text{CA}} = (\text{judge})$ :

**Participating roles:** {judge}  
**Corruption model:** *Dynamic corruption without secure erasures*

Description of  $M_{\text{judge}}$ :

**Implemented role(s):** {judge}  
**Subroutines:**  $\mathcal{F}_{\text{sig}}$  : verifier,  $\mathcal{F}_{\text{init}}^{\text{CA}}$  : init  
**Internal state:**  
-  $\text{pid}_{\text{CA}} \subset \{0, 1\}^*$  {The pids of the CAs}  
-  $\text{pk}_{\text{CA}} : \text{pid}_{\text{CA}} \rightarrow \{0, 1\}^*$  {The public keys of CAs, initially  $\perp$  for all entries}  
-  $\text{pid}_{\text{CTL}} \subset \{0, 1\}^*$  {The pids of the CTLs}  
-  $\text{pk}_{\text{CTL}} : \text{pid}_{\text{CTL}} \rightarrow \{0, 1\}^*$  {The public keys of CTLs, initially  $\perp$  for all entries}  
-  $\text{CTLs} : \text{pid}_{\text{CA}} \rightarrow (\text{pid}_{\text{CTL}} \cup \{\perp\})^*$  {A CAs trusted CTLs, initially  $\perp$ }  
-  $\text{verdicts} \in \{0, 1\}^*$ ,  $\text{verdicts} = \varepsilon$  {Recorded verdict}  
**CheckID**( $\text{pid}, \text{sid}, \text{role}$ ):  
Accept all messages with the same  $\text{sid}$  addressed to ( $\text{pid}, \text{sid}, \text{judge}$ ) where  $\text{pid}$  can be parsed as ( $\text{local}, \text{pid}', \text{client}$ ).  
**Corruption behavior:**  
- **AllowCorruption**( $\text{pid}, \text{sid}, \text{role}$ ) :  
Parse  $\text{pid}_{\text{cur}}$  as ( $\text{pid}', \text{role}'$ )  
 $\text{corrRes} \leftarrow \text{corr}(\text{pid}', \text{sid}_{\text{cur}}, \text{role}')$   
**return**  $\text{corrRes}$  {Corruption is allowed if accompanied client is corrupted.}  
- **DetermineCorrStatus**( $\text{pid}, \text{sid}, \text{role}$ ) : {Consider local judge corrupted if associated client is corrupted.}  
Parse  $\text{pid}$  as ( $\text{pid}', \text{role}'$ )  
 $\text{corr} \leftarrow \text{corr}(\text{pid}', \text{sid}, \text{role}')$   
**if**  $\text{corr}$ :  
**return** true  
**else:**  
**return** false  
**EntityInitialization:**  
**if** ITM was not activated via **init** macro, resp. a **InitEntity** message:  
**send** **initPartner** **to** ( $\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \mathcal{F}_{\text{init}}^{\text{CA}}$  : init)  
**wait for** **initPartner**  
**MessagePreprocessing:**  
**send** **InitMe<sub>judge</sub>** **to** ( $\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \mathcal{F}_{\text{init}}^{\text{CA}}$  : init) {Update CA/CTL identities and pubkeys}  
**wait for** (**InitMe<sub>judge</sub>**,  $\text{pid}'_{\text{CA}}, \text{pk}_{\text{CA}}, \text{pid}'_{\text{CTL}}, \text{pk}_{\text{CTL}}, \text{CTLs}$ )  
 $\text{pid}_{\text{CA}} \leftarrow \text{pid}'_{\text{CA}}; \text{pk}_{\text{CA}} \leftarrow \text{pk}_{\text{CA}}; \text{pid}_{\text{CTL}} \leftarrow \text{pid}'_{\text{CTL}}; \text{pk}_{\text{CTL}} \leftarrow \text{pk}_{\text{CTL}}; \text{CTLs} \leftarrow \text{CTLs}$   
**Main:**  
**recv** (**Evidence<sub>M</sub>**, ( $\text{ctr}', [\text{serial}, \text{pid}, \text{pk}, \text{pid}'_{\text{CA}}, \sigma_{\text{CA}}], \text{pid}'_{\text{CTL}}, \tau, \sigma_{\text{CTL}}$ )) **from** I/O {Client complains that someone registered a certificate for her identity}  
**s.t.**  $\text{pid}_{\text{call}} = \text{pid}_{\text{cur}}, \text{sid}_{\text{call}} = \text{sid}_{\text{cur}}, \text{role}_{\text{call}} = \text{client}$ :  
 $b \leftarrow \text{verifySig}([\text{ctr}', [\text{serial}, \text{pid}_{\text{cur}}, \text{pk}, \text{pid}'_{\text{CA}}, \sigma_{\text{CA}}], \text{pid}_{\text{CTL}}, \tau], \sigma_{\text{CTL}}, \text{pid}_{\text{CTL}}[\text{pid}'_{\text{CTL}}, \text{pid}'_{\text{CTL}}])$  {Check CTL signature}  
**if**  $\text{pid}'_{\text{CA}} \in \text{pid}_{\text{CA}}$ :  
 $c \leftarrow \text{verifySig}([\text{serial}, \text{pid}_{\text{cur}}, \text{pk}, \text{pid}'_{\text{CA}}], \sigma_{\text{CA}}, \text{pid}_{\text{CA}}[\text{pid}'_{\text{CA}}, \text{pid}'_{\text{CA}}])$  {Check CA signature}  
**else:**  
 $c \leftarrow \text{false}$   
**if**  $b \wedge c$ : {Certificate is valid}  
 $\text{verdicts.add}(\text{dis}(\text{pid}_{\text{CA}}, \text{sid}_{\text{cur}}, \mathcal{P}^{\text{CA}} : \text{ca}))$  {The local judge blames the CA for misbehavior}  
**recv** **GetVerdict**:  
**reply** (**GetVerdict**,  $\text{verdicts}$ ) {Both I/O and NET may send this message.}  
**recv** (**GetJudicialReport**,  $\text{msg}$ ) **from** I/O:  
**reply** (**GetJudicialReport**,  $\varepsilon$ )  
**Procedures and Functions:**  
**function** **verifySig**( $\text{msg}, \sigma, \text{pk}, \text{pid}$ ) : {Verify signature at  $\mathcal{F}_{\text{sig}}$ }  
**send** (**VerResult**,  $\text{msg}, \sigma, \text{pk}$ ) **to** (( $\text{pid}_{\text{cur}}, \text{ctl}$ ), ( $\text{pid}, \text{sid}_{\text{cur}}$ ),  $\mathcal{F}_{\text{sig}}$  : verifier)  
**wait for** (**VerResult**,  $\text{result}$ )  
**return**  $\text{result}$

Fig. 28: The judging functionality  $\mathcal{P}_{\text{judge}}^{\text{CA}}$  for a CA (Part 1).

Description of  $\mathcal{P}_{sv}^{CA} = (\text{supervisor})$ :

<b>Participating roles:</b> {supervisor} <b>Corruption model:</b> <i>incorruptible</i> <b>Protocol parameters:</b> - $\text{Sec}_{\text{assumption}} = \emptyset$	<i>{The set of security properties that may break}</i>
--	--

Description of  $M_{\text{supervisor}}^{CA}$ :

<b>Implemented role(s):</b> {supervisor} <b>Subroutines:</b> $\mathcal{P}^{CA} : \text{ca}, \mathcal{P}_{\text{CTL}} : \text{ctl}$ <b>CheckID</b> ( $pid, sid, role$ ): Accept all messages with the same $sid$ . <b>Main:</b> <b>recv</b> (IsAssumptionBroken?, $msg$ ): <b>reply</b> (IsAssumptionBroken?, false) <b>recv</b> (corruptInt?, ( $pid, sid, role$ )) <b>s.t.</b> $role \notin \{\text{client}, \text{judge}, \text{supervisor}\}$ : <b>if</b> $role \neq \text{ca}$ : <b>reply</b> (corruptInt, false) <b>else:</b> $corrRes \leftarrow \text{corr}(pid, sid_{\text{cur}}, \text{ca})$ <b>reply</b> (corruptInt, $corrRes$ )	<i>{Request corruption status at <math>\mathcal{P}^{CA}</math>}</i>
---	---

Fig. 29: The supervisor  $\mathcal{P}_{sv}^{CA}$  for CA.

### Local judges

When an *uncorrupted* instance  $((pid, role), sid, \text{judge})$  of  $\mathcal{P}_{\text{judge}}^{CA}$  in  $\mathcal{R}'$  renders a verdict, i. e., verdicts  $\neq \varepsilon$ , or its verdicts changes,  $\mathcal{S}$  extracts the verdicts and forwards it to  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . More specifically,  $\mathcal{S}$  sends  $(\text{BreakAccProp}, \text{verdict}, \{(\text{correctCert}, (pid, sid, role))\})$  to  $((\text{local}, pid, role), sid, \mathcal{F}_{\text{PKI}}^{\text{acc}} : \text{judge})$  where  $\text{verdict}$  is a map from  $\text{ids}_{\text{assumption}} \rightarrow \{0, 1\}^*$  where the entry  $\text{verdict}[id] = \text{verdicts}$  (contains the verdict from  $\mathcal{P}_{\text{judge}}^{CA}$ ). All other entries in  $\text{verdict}$  are mapped to  $\varepsilon$ .

If a judge's associated client is corrupted, the local judge is also considered corrupted and serves as pure message forwarder to/from  $\mathcal{A}$ . Thus – as stated above –  $\mathcal{S}$  forwards the messages from/to  $\mathcal{A}$  to/from  $\mathcal{I}$  in this case.

This concludes the description of the simulator. It is easy to see that (i)  $\{\mathcal{S}, \mathcal{I}\}$  is environmentally bounded, and (ii)  $\mathcal{S}$  is a responsive simulator for  $\mathcal{I}$ , i. e., restricting messages from  $\mathcal{I}$  are answered immediately as long as  $\{\mathcal{S}, \mathcal{I}\}$  runs with a responsive environment. We now argue that  $\mathcal{R}$  and  $\{\mathcal{S}, \mathcal{I}\}$  are indeed indistinguishable for any (responsive) environment  $\mathcal{E} \in \text{Env}(\mathcal{R})$ .

In the following part of the proof, let  $\mathcal{E} \in \text{Env}(\mathcal{R})$  be an arbitrary but fixed environment. First, observe that  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  provides  $\mathcal{S}$  with full information about (i) all creations requests for certificate of uncorrupted clients (for which also  $\text{correctCert}$  holds true) and (ii) retrieve requests performed by higher-level protocols/the environment. Hence,  $\mathcal{S}$ 's simulated protocol  $\mathcal{R}'$  obtains the same inputs as  $\mathcal{R}$  (resp.  $\mathcal{I}$ ) and thus performs identical to  $\mathcal{R}$ . We remark that further I/O input, e. g., requests to judges, do not influence the state of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . Thus, we can conclude that the network behavior simulated by  $\mathcal{S}$  towards the environment is indistinguishable from the network behavior of  $\mathcal{R}$ . Moreover, we can also conclude that the corruption status of (internal) entities in the real and ideal world is synchronized. Since the simulator has full control over corrupted entities, which are handled via the internal simulation  $\mathcal{R}'$ , this implies that the I/O behavior of corrupted entities of  $\mathcal{R}/\mathcal{I}$  towards the environment is also identical in the real and ideal world. Note that purely internal (private) parties have no interface to the environment. The only way to potentially distinguish the real and ideal world is the I/O behavior of honest entities of  $\mathcal{R}/\mathcal{I}$  towards higher-level protocols.

We will now go over all possible interactions with honest entities on the I/O interface and argue, by induction, that all of those interactions result in identical behavior towards the environment, i. e., are also indistinguishable. At the start of a run, there were no interactions on the I/O interface with honest parties yet. Thus,  $\mathcal{R}$  and  $\mathcal{I}$  are indistinguishable. In the following, we assume that all I/O interactions so far have resulted in the same behavior visible towards the environment in both the real and ideal world.

In what follows, we call the state which includes the certificates of uncorrupted parties for which  $\text{correctCert}$  still holds true *relevant state*. The relevant state can be extracted from the union of all state variables of all  $\mathcal{P}_{\text{CTL}}$  instances in a session of  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$  in  $\mathcal{R}$ , resp.  $\mathcal{R}'$ , by removing certificates from (i) corrupted parties and (ii) from honest parties whose local judges already store a verdict  $\neq \varepsilon$ . In  $\mathcal{I}$ , the relevant state can be derived from  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ 's



Description of the ideal authenticated channel functionality  $\mathcal{F}_{\text{auth}} = (\text{auth})$ :

<b>Participating roles:</b> $\{\text{auth}\}$ <b>Corruption model:</b> <i>custom</i>
Description of $M_{\text{auth}}$ :
<b>Implemented role(s):</b> $\{\text{auth}\}$ <b>Internal state:</b> <ul style="list-style-type: none"> <li>- <math>\text{queue} : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^*</math> <span style="float: right;">{Queue of messages from entity <math>e_1</math> to entity <math>e_2</math>, initially <math>\perp</math> for all entries}</span></li> <li>- <math>\text{status} : (\{0, 1\}^*)^3 \rightarrow \{\text{inactive}, \text{active}, \text{established}\}</math> <span style="float: right;">{The status of <math>(pid, sid, role)</math>, initially inactive for all entries}</span></li> <li>- <math>\text{corrStatus} \in \{0, 1\}</math> <span style="float: right;">{The corruption status of the <math>\mathcal{F}_{\text{auth}}</math>}</span></li> </ul> <b>CheckID</b> $(pid, sid, role)$ : Accept all messages for the same $sid$ . <b>Main:</b> <div style="display: flex; justify-content: space-between;"> <div style="width: 80%;"> <p><b>recv corrupt from NET:</b>              if <math>\forall</math> entries in status are not equal to established:                <math>\text{corrStatus} \leftarrow \text{true}</math>                <b>reply</b> (corrupt, ack)              else:                <b>reply</b> (corrupt, nack)</p> <p><b>recv CorruptionStatus? from I/O:</b>              if <math>\text{corrStatus} = \text{true}</math>:                <b>reply</b> (CorruptionStatus?, true)              else:                <b>reply</b> (CorruptionStatus?, false)</p> <p><b>recv</b> (Establish, <math>m</math>) <b>from</b> I/O <b>s.t.</b> <math>\text{sid}_{\text{call}} = \text{sid}_{\text{cur}}</math>, <math>\text{status}[(pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}})] = \text{inactive}</math>:                <math>\text{status}[(pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}})] \leftarrow \text{active}</math>                <b>send</b> (Establish, <math>m</math>, <math>(pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}})</math>) <b>to</b> NET <span style="float: right;">{Establish session}</span></p> <p><b>recv</b> (Establish, <math>(pid, sid_{\text{cur}}, role)</math>) <b>from</b> NET <b>s.t.</b> <math>\text{status}[(pid, sid_{\text{cur}}, role)] = \text{active}</math>:                <math>\text{status}[(pid, sid_{\text{cur}}, role)] \leftarrow \text{established}</math>                <b>send</b> (Establish, <math>m</math>, <math>(pid, sid_{\text{cur}}, role)</math>) <b>to</b> <math>(pid, sid_{\text{cur}}, role)</math> <span style="float: right;">{Establish session}</span></p> <p><b>recv</b> (Send, <math>(pid, sid_{\text{cur}}, role)</math>, <math>msg</math>) <b>from</b> I/O <b>s.t.</b> <math>\text{status}[(pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}})] = \text{established} \wedge pid \neq pid_{\text{call}}</math>: <span style="float: right;">{Send message via authenticated channel}</span>              if <math>\text{corrStatus} = \text{false}</math>:                <math>\text{queue}[(pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}}), (pid, sid_{\text{cur}}, role)].\text{add}(msg)</math> <span style="float: right;">{Add msg to the queue of pid}</span>                <b>send</b> (Send, <math>(pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}}), (pid, sid_{\text{cur}}, role), msg</math>) <b>to</b> NET <span style="float: right;">{Leak communication to NET}</span></p> <p><b>recv</b> (Deliver, <math>(pid_1, sid_{\text{cur}}, role_1), (pid_2, sid_{\text{cur}}, role_2), msg</math>) <b>from</b> I/O <b>s.t.</b> <math>\text{status}[(pid_2, sid_{\text{cur}}, role_2)] = \text{established} \wedge pid_1 \neq pid_2</math>: <span style="float: right;">{Deliver message via authenticated channel}</span>              if <math>\text{corrStatus} = \text{false}</math>:                if <math>\text{queue}[(pid_1, sid_{\text{cur}}, role_1), (pid_2, sid_{\text{cur}}, role_2)] = \perp</math>: <span style="float: right;">{There are no queued messages}</span>                  <b>reply</b> (Deliver, <math>\perp</math>) <span style="float: right;">{Return error}</span>                else:                  remove the first message from <math>\text{queue}[(pid_1, sid_{\text{cur}}, role_1), (pid_2, sid_{\text{cur}}, role_2)]</math>, let <math>msg'</math> be this message                  <b>send</b> (Received, <math>(pid_1, sid_{\text{cur}}, role_1), msg'</math>) <b>to</b> <math>(pid_2, sid_{\text{cur}}, role_2)</math> <span style="float: right;">{Deliver first message from queue}</span>                else:                  <b>send</b> (Received, <math>(pid_1, sid_{\text{cur}}, role_1), msg</math>) <b>to</b> <math>(pid_2, sid_{\text{cur}}, role_2)</math> <span style="float: right;">{Deliver message from adversary if corrupted}</span></p> <p><b>recv</b> (Drop, <math>(pid_1, sid_{\text{cur}}, role_1), (pid_2, sid_{\text{cur}}, role_2)</math>) <b>from</b> NET:              if <math>\text{queue}[(pid_1, sid_{\text{cur}}, role_1), (pid_2, sid_{\text{cur}}, role_2)] \neq \perp</math>:                remove the first message from <math>\text{queue}[(pid_1, sid_{\text{cur}}, role_1), (pid_2, sid_{\text{cur}}, role_2)]</math>. <span style="float: right;">{Drop message}</span>                <b>reply</b> (Drop, ack)</p> </div> <div style="width: 15%; font-size: small; vertical-align: top;"> <p>{A (tries to) corrupts current instance of <math>\mathcal{F}_{\text{auth}}</math>              {This models static corruption</p> </div> </div>

Fig. 30: The ideal authenticated channel functionality  $\mathcal{F}_{\text{auth}}$  (cf. [70]).

state variable by excluding (i) certificates of corrupted parties and (ii) certificates where  $\text{correctCert}$  does not hold true for the certificate owner.

**Certificate creation/registration requests:** Certificate creation/registration requests do not directly result in an output to the environment. But, they might affect the output of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  later on as they have direct impact on the relevant state of  $\mathcal{I}/\mathcal{R}'$ , resp.  $\mathcal{R}$ . Thus, we now show that registration requests behave “identical”, i.e., we have to argue that these changes in relevant state are “synchronized” between  $\mathcal{I}$  and  $\mathcal{R}'$ . In particular, the buffered set of relevant certificates registrations, i.e., certificate requested by uncorrupted parties for which  $\text{correctCert}$  still holds true, in  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  is equal to the buffered set of certificate registration from the same class of clients in  $\mathcal{R}'$ . We define *buffered registration request* in  $\mathcal{R}'$  as the following requests/messages: (i) certificate requests stored in uncorrupted instances of  $\mathcal{P}_{\text{client}}^{\text{CA}}$  (in  $\text{pk}$  and  $\text{certCA}$ ) for which  $\text{correctCert}$  still holds true (ii) minus the set of finalized certificates from the different instances of  $\mathcal{P}_{\text{CTL}}$ . In what follows, we will call these certificates the *buffered relevant certificates for registration at  $\mathcal{R}'$* .

Observe that, upon receiving a certificate registration,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  behaves similar to  $\mathcal{R}'$ : in  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ , the submitted certificate registration is directly stored in  $\text{buffer}_S$ . In  $\mathcal{R}'$ , the submit message is stored in  $\mathcal{P}_{\text{client}}^{\text{CA}}$  and forwarded

Description of the protocol  $\mathcal{F}_{\text{sig}} = (\text{signer}, \text{verifier})$ :

<b>Participating roles:</b> $\{\text{signer}, \text{verifier}\}$ <b>Subroutines:</b> $\mathcal{F}_{\text{init}}^{\text{CA}} : \text{init}$ <b>Corruption model:</b> <i>dynamic with secure erasures</i> <b>Protocol parameters:</b> - $p \in \mathbb{Z}[x]$ .	$\left\{ \begin{array}{l} \text{Polynomial that bounds the runtime of the algorithms provided by the} \\ \text{adversary.} \end{array} \right.$
---	---

Description of  $M_{\text{signer}, \text{verifier}}$ :

<b>Implemented role(s):</b> $\{\text{signer}, \text{verifier}\}$ <b>Internal state:</b> - $(\text{sig}, \text{ver}, \text{pk}, \text{sk}) \in (\{0, 1\}^* \cup \{\perp\})^4 = (\perp, \perp, \perp, \perp)$ . - $\text{pidowner} \in \{0, 1\}^* \cup \{\perp\} = \perp$ . - $\text{msglist} \subset \{0, 1\}^* = \emptyset$ . - $\text{KeysGenerated} \in \{\text{ready}, \perp\} = \perp$ . <b>CheckID</b> ( $\text{pid}, \text{sid}, \text{role}$ ): Check that $\text{sid} = (\text{pid}', \text{sid}')$ . If this check fails, output <b>reject</b> . Otherwise, accept all entities with the same SID. <b>Corruption behavior:</b> - <b>AllowCorruption</b> ( $\text{pid}, \text{sid}, \text{role}$ ): Parse $\text{pid}_{\text{cur}}$ to $(\text{pid}', \text{role}')$ <b>send</b> GetCTLs to $(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \mathcal{F}_{\text{init}}^{\text{CA}} : \text{init})$ <b>wait for</b> (GetCaCtls, $\text{pid}_{\text{CTL}}, \text{pk}_{\text{CTL}}$ ) if $\text{pid}' \in \text{pid}_{\text{CTL}} \wedge \text{role}' = \text{ctl}$ : <b>return false</b> <b>else if</b> $\text{role} = \text{verifier}$ : <b>return false</b> <b>else:</b> <b>return true</b> - <b>LeakedData</b> ( $\text{pid}, \text{sid}, \text{role}$ ): If $(\text{pid}, \text{sid}, \text{role})$ determines its initial corruption status, use the default behavior of <b>LeakedData</b> . Otherwise, if $\text{role} = \text{signer}$ and $\text{pid} = \text{pidowner}$ , return <b>KeysGenerated</b> . In all other cases return $\perp$ . <b>Initialization:</b> <b>send responsively</b> InitMe to NET; <b>wait for</b> (Init, $(\text{sig}, \text{ver}, \text{pk}, \text{sk})$ ). $(\text{sig}, \text{ver}, \text{pk}, \text{sk}) \leftarrow (\text{sig}, \text{ver}, \text{pk}, \text{sk})$ . Parse $\text{sid}_{\text{cur}}$ as $(\text{pid}, \text{sid})$ . $\text{pidowner} \leftarrow \text{pid}$ . <b>Main:</b> <b>recv</b> InitSign from I/O to $(\text{pidowner}, \_, \text{signer})$ : $\text{KeysGenerated} \leftarrow \text{ready}$ . <b>reply</b> (InitSign, success, pk). <b>recv</b> (Sign, $\text{msg}$ ) from I/O to $(\text{pidowner}, \_, \text{signer})$ s.t. $\text{KeysGenerated} = \text{ready}$ : $\sigma \leftarrow \text{sig}^{(p)}(\text{msg}, \text{sk})$ . $b \leftarrow \text{ver}^{(p)}(\text{msg}, \sigma, \text{pk})$ . if $\sigma = \perp \vee b \neq \text{true}$ : <b>reply</b> (Signature, $\perp$ ). <b>else:</b> <b>add</b> $\text{msg}$ to msglist. <b>reply</b> (Signature, $\sigma$ ). <b>recv</b> (Verify, $\text{msg}, \sigma, \text{pk}$ ) from I/O to $(\_, \_, \text{verifier})$ : $b \leftarrow \text{ver}^{(p)}(\text{msg}, \sigma, \text{pk})$ . if $\text{pk} \wedge b = \text{true} \wedge \text{msg} \notin \text{msglist} \wedge (\text{pidowner}, \text{sid}_{\text{cur}}, \text{signer}) \notin \text{CorruptionSet}$ : <b>reply</b> (VerResult, false). <b>else:</b> <b>reply</b> (VerResult, $b$ ).	$\left\{ \begin{array}{l} \text{Algorithms and key pair.} \\ \text{Party ID of the key owner.} \\ \text{Set of recorded messages.} \\ \text{Has signer initialized his key?} \end{array} \right.$  $\left\{ \begin{array}{l} \text{A single instance manages all parties} \\ \text{and roles in a single session.} \end{array} \right.$ $\left\{ \begin{array}{l} \text{The highlighted part differs between a standard version} \\ \text{of } \mathcal{F}_{\text{sig}} \text{ and the version considered here} \end{array} \right.$ $\left\{ \begin{array}{l} \text{Query for the list of CTLs of session sid, decline} \\ \text{corruption requests for CTL keys} \end{array} \right.$  $\left\{ \begin{array}{l} \text{Not allowed to corrupt CLT keys.} \\ \text{Corruption of verifier instances not allowed} \end{array} \right.$  $\left\{ \begin{array}{l} \text{Successful initialization. Note that signer can submit} \\ \text{InitSign multiple times, always with the same effect.} \end{array} \right.$  $\left\{ \begin{array}{l} \text{Sign and check that verification succeeds.} \\ \text{Signing or verification test failed.} \end{array} \right.$  $\left\{ \begin{array}{l} \text{Record msg for verification and return signature.} \\ \text{Verify signature.} \\ \text{Prevent forgery.} \\ \text{Return verification result.} \end{array} \right.$
--	--

Fig. 31: The ideal signature functionality  $\mathcal{F}_{\text{sig}}$ .

via  $\mathcal{F}_{\text{auth}}$  to  $\mathcal{P}^{\text{CA}}$ . So, the buffered relevant certificates  $\mathcal{R}'$  also are in a buffer of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ , i. e., the buffered relevant certificates are synchronized between  $\mathcal{R}'$  and  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  (that is, the sets are equal in both worlds). Note that  $\mathcal{S}$  does not have to take care of certificates of corrupted parties as  $\mathcal{A}/\mathcal{S}$  fully determines the output of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  for these parties anyways. Thus,  $\mathcal{S}$  is indeed able to keep the buffered relevant certificates in  $\mathcal{I}$  and  $\mathcal{R}'$  synchronized.

**State updates:** Again, state updates do not lead to a direct output to the environment. However, state updates may influence the output to the environment later on. Thus, we argue here that the relevant state (see above) is synchronized between  $\mathcal{I}$  and  $\mathcal{R}'$  (thus, it will allow for the same output in  $\mathcal{I}$  and  $\mathcal{R}$ ). As the buffered relevant certificates in  $\mathcal{R}'$  and  $\mathcal{I}$  are synchronized (see above), we can conclude that the relevant states of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  and  $\mathcal{R}'$  stay synchronized during state updates. This boils down to mirroring the relevant state (see above for the definition) from the CTLs to  $\mathcal{I}$  (as  $\mathcal{S}$  does). Thus, all relevant certificates which  $\mathcal{S}$  would add to the state in  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  are also buffered in  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . Therefore,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  will accept every state update of  $\mathcal{S}$  and thus the relevant state stay synchronized in  $\mathcal{I}$  and  $\mathcal{R}'$ .

**Retrieve requests:** Observe that, upon a retrieve request,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  stores the read request in  $\text{buffer}_R$  and leaks the full request to the simulator including an ID and the receiver of the response. Note that the procedure is analogously to a certificate registration. A retrieve request does not directly result in an output to the environment. With the same argumentation as above follows that  $\mathcal{S}$  is able to keep the “retrieve buffers”, i. e. the set of requested certificates stored in honest instances of  $\mathcal{P}_{\text{client}}^{\text{CA}}$ , resp.  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ , for honest participants of  $\mathcal{I}$  and  $\mathcal{R}'$  synchronized.

**Deliver responses to retrieve requests:** If  $\mathcal{S}$  is triggered to output a response to a retrieve request, i. e.,  $\mathcal{S}$  receives Retrieve via NET including the necessary data to output,  $\mathcal{S}$  firstly simulates received input in  $\mathcal{R}'$  and determines whether this input actually leads to an output on an I/O interface. In case of an output,  $\mathcal{S}$  forwards the output in a Deliver message to  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . In this case, there are two cases to consider: (i) `correctCert` is broken for the certificate owner `pid` or (ii) `correctCert` still hold true for `pid`. In the first case,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  will simply forward the input from  $\mathcal{S}$  to the requestor as all checks at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  are disabled. As the output is extracted from  $\mathcal{R}'$ , we can conclude that  $\mathcal{R}$  and  $\mathcal{I}$  are indistinguishable. In the second case, the output is honestly generated in  $\mathcal{R}'$  and  $\mathcal{I}$ . As argued above, the relevant states of both worlds are synchronized. Thus,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  will accept the Deliver command of  $\mathcal{S}$  and the outputs in both worlds will be identical.

When a party queries for a certificate of corrupted party,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  allows  $\mathcal{A}/\mathcal{S}$  to freely determine the certificate delivered to the requesting party. We already discussed this case after the description of the simulator and emphasize that in this case,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  simply forwards the output provided by  $\mathcal{S}$  (which matches  $\mathcal{R}'$ ).

Altogether, both worlds are also indistinguishable in the case of retrieve requests.

**Monitoring Process:** Note that the monitoring process only takes places in  $\mathcal{R}'$  and only its output, a verdict if there is any, is mirrored to  $\mathcal{I}$  (see below). Thus, the monitoring process will not cause a case where  $\mathcal{I}$  and  $\mathcal{R}$  are distinguishable.

**Verdicts:** First note, that  $\mathcal{S}$  sends a BreakAccProp message to break `correctCert` when an unambiguous or an not requested certificate of an uncorrupted party in  $\mathcal{R}'$  appears. More specifically,  $\mathcal{S}$  sends a BreakAccProp when there is a new/updated verdict available in an uncorrupted instance of  $\mathcal{P}_{\text{judge}}^{\text{CA}}$ . We remark that these requests match the rules in  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}}$  and thus are always accepted by  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . Further, we assume that the used signature scheme is EUF-CMA secure. Thus, the probability that an honest party receives and accepts a certificate created by  $\mathcal{A}$  via forging signatures (without corrupting the party owning the key or corrupting the key itself) only occurs with negligible probability. Also, the verifier instance of an honest local judge is incorruptible. Thus, we can conclude that  $\mathcal{S}$  can always keep the verdicts and whether consistency is broken synchronized in  $\mathcal{I}$  and  $\mathcal{R}'$ . Thus, the output on a GetVerdict request is equal in  $\mathcal{R}'$  and  $\mathcal{I}$ . Thus, ideal and real world are indistinguishable.

We also remark that  $\mathcal{P}_{\text{client}}^{\text{CA}}$  only forwards certificates to the environment if they are older than  $3 \cdot \delta$  time units. In this case, the monitoring process for the particular certificate is already finished (in  $\mathcal{R}/\mathcal{R}'$ ) and thus the associated local judge will provide a verdict in case of a maliciously created certificate. This triggers  $\mathcal{S}$  to break `correctCert` (including providing a verdict) for the dedicated party and ensures that there is always a verdict available/the status of `correctCert` correctly reflected/synchronized between  $\mathcal{R}'$  and  $\mathcal{I}$ . Thus, also in this special case,  $\mathcal{R}$  and  $\mathcal{I}$  are indistinguishable.

**Judicial Reports:** The judicial report (for uncorrupted instances of  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-PKI}}$ ) in both worlds is always  $\varepsilon$ . Thus,  $\mathcal{R}$  and  $\mathcal{I}$  stay indistinguishable.

**Supervisor:** The same holds true for BreakAssumption requests to the supervisor. The message of  $\mathcal{S}$  matches the rules and thus will be accepted.

Altogether,  $\mathcal{R}$  and  $\{\mathcal{S}, \mathcal{I}\}$  behave identical in terms of behavior visible to the environment  $\mathcal{E}$  and thus are indistinguishable.  $\square$

**E.4 Deterrence Analysis:** As running CA business is profitable in practice and clients benefit from PKIs, we should have that  $U_{hP}^i \gg 0$ . We assume  $U_{hL}^i$  to be negligible as parties falsely accusing another cannot provide undeniable evidence that the party indeed misbehaved. As all relevant data is public in the PKI setting, we assume  $U_{hE}^i = 0$ . In real world scenarios, CAs would typically lose their complete reputation when misbehaving and thus often have to close business afterwards [92]. Thus, we estimate  $U_i^c \ll U_{hP}^i$ . In order to deter CAs from misbehavior, penalties need to be chosen significantly high such that  $U_{mL}^i \gg U_{mP}^i - U_{hP}^i$ . As it is not possible to generate faked undeniable evidence, we assume  $U_{mE}^i = 0$ . Thus, Equation 1 and 2 are met.

#### F. Key Exchange Based on $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ (Full Details)

Before we have a look at the security proof of Theorem 3, we present the ideal functionality  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  and the key exchange protocol specified in ISO 9798-3 [56] in detail. This case study relies on existing analysis of the ISO protocol [20, 27, 66]. More specifically, we derive the considered ideal key exchange functionality  $\mathcal{F}_{\text{KE}}$  from the functionality explained in [20] and use AUC to derive  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  from  $\mathcal{F}_{\text{KE}}$ . In particular, we add the accountability property authenticity to  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ .  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  ensures for “honest” key exchange sessions, that they provide security guarantees as long as non of the parties was involved in a misconduct in  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . In contrast to existing works, which typically rely on an ideal/perfect PKI which does not allow maliciously generated certificates, our model of the ISO protocol  $\mathcal{P}_{\text{iso}}^{\text{acc}}$  uses the PKI  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  for accessing the public signature keys of the involved parties for mutual authentication. Theorem 2 then allows us to replace  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  with its realization  $\mathcal{P}_{\text{PKI}}^{\text{acc}}$ , a realistic PKI with CTLs.

In this section, we will firstly provide the formal specification of the (accountable) ideal key exchange functionality  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  and the model of the ISO protocol  $\mathcal{P}_{\text{iso}}^{\text{acc}}$ . Finally, we will provide a formal proof for Theorem 3.

**F.1 The Accountable Key Exchange Functionality  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ :** We also provide the full formal specification of  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  in Figure 32 and 33 including the specification of  $\mathcal{F}_{\text{judgeParams}}^{\text{KE}}$  in Figure 34. The highlighted parts in the specification of  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  mark the changes during the AUC transformation from  $\mathcal{F}_{\text{KE}}$  to  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ . In what follows, we firstly present  $\mathcal{F}_{\text{KE}}$  and then discuss the differences between  $\mathcal{F}_{\text{KE}}$  and  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ .

The functionality  $\mathcal{F}_{\text{KE}}$  consists of two roles, initiator and responder. Both, initiator and responder, are implemented via a single machine. One instance of this machine models a single key exchange. In  $\mathcal{F}_{\text{KE}}$ , key exchange sessions are predefined via SIDs of form  $(sid', pid_i, pid_r)$  where  $sid'$  is a local session ID for the key exchange,  $pid_i$  is the PID of the initiator and  $pid_r$  is the responder’s PID. If two honest entities finish a key exchange, then  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  ensures that they obtain an ideal session key which is unknown to the adversary.

$\mathcal{F}_{\text{KE}}$  models a key exchange in three “phases” which are reflected in the available commands:

**Public key registration (I/O):** before parties can start a key exchange, they need to register their identities.

**Initialization of key exchange (I/O):** when parties have registered, they can start a key exchange.

**Finish key exchange (Net):**  $\mathcal{F}_{\text{KE}}$  allows the adversary/simulator to determine when a started key exchange finished. The FinishKE message triggers the output of keys to higher-level protocols. In a session consisting of honest initiator and responder, the session key is derived ideally, i. e., chosen uniformly at random from a cyclic group  $(G, n, g)$ . If one of the parties is corrupted, the adversary/the simulator may determine the session key.

In a session without corrupted parties,  $\mathcal{F}_{\text{KE}}$  provides mutual authentication.

We add the following to  $\mathcal{F}_{\text{KE}}$  to derive  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ : if one of the two parties involved in a key exchange does not provide authenticity, i. e., the parties are not corrupted but brokenProps[authenticity, (local, pid, role)] = true for one of both participants,  $\mathcal{A}$  is allowed to determine the session key during key exchange. If  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  already established a session key, the key is leaked to  $\mathcal{A}$ . Additionally,  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ , resp.  $\mathcal{F}_{\text{judgeParams}}^{\text{KE}}$ , restricts the local judges output to be a verdict blaming some internal protocol party and/or the intended key exchange

partner. Note that the local judge considers itself as corrupted if the accompanied initiator, resp. responder entity is corrupted. Also, a local judge can only be (directly) corrupted if the accompanied initiator, resp. responder entity is corrupted. If the local judge is directly corrupted, it acts (as common) as pure message forwarder for  $\mathcal{A}$ . Judicial reports do not contain content, i. e., the judicial report is always  $\varepsilon$ . We do not consider a supervisor in this example/functionality. Formally, the supervisor does not accept break attempts and always outputs an empty leakage.

**F.2 The ISO 9798-3 Protocol as Accountable Key Exchange Protocol:** The ISO protocol [56] or in short the *ISO Protocol* as depicted in Figure 8 is a key exchange protocol which builds mutual authentication on top of a Diffie-Hellman key exchange.

We model the ISO protocol in a modular way using several smaller protocols. The static structure of all protocols, including their I/O connections for direct communication, is shown in Figure 35. In the following, we give a high-level overview over each part of the protocol.

The (accountable) ISO protocol is modeled as a real protocol  $\mathcal{P}_{\text{iso}}^{\text{acc}}$ . We derived  $\mathcal{P}_{\text{iso}}^{\text{acc}}$  mainly from  $\mathcal{P}_{\text{iso}}$  [17]. However, as already stated above, we model the key exchange via an pre-established SID which predetermines the session of the key exchange and initiator and responder in the particular session (cf., e. g., [20]). In what follows, we firstly explain our variant of  $\mathcal{P}_{\text{iso}}$ . Then, we explain how to derive the accountable variant  $\mathcal{P}_{\text{iso}}^{\text{acc}}$  from  $\mathcal{P}_{\text{iso}}$ .

The real protocol  $\mathcal{P}_{\text{iso}}$  consists of three roles, *initiator*, *responder*, and *setup*. The *setup* role models secure generation and distribution of a system parameter, namely, a description of a cyclic group  $(G, n, g)$ . As this parameter must be shared between all runs of a key exchange protocol, *setup* is implemented by a single machine which spawns a single instance that manages all entities and always outputs the same parameter. The roles *initiator* and *responder* implement parties  $A$  and  $B$ , respectively, from Figure 8. Each role is implemented by a separate machine and every instance of those machines manages exactly one entity. Thus, these instances directly correspond to an actual implementation where each run of a key exchange protocol spawns a new program instance. Both, *initiator* and *responder* register their public keys at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  before they actually start the key exchange.

In  $\mathcal{P}_{\text{iso}}$ , we use a standard ideal signature functionality  $\mathcal{F}_{\text{sig}}^{16}$  for signing messages and verifying signatures. As common, the ideal functionality  $\mathcal{F}_{\text{sig}}$  consists of two roles, *signer* and *verifier*, that allows for the corresponding operations. Both roles are implemented by the same machine and instances of that machine manage entities that share the same SID. The SID  $sid$  of an entity is structured as a tuple  $(pid_{\text{owner}}, (sid', pid_i, pid_r))$ , modeling a specific key pair of the party  $pid_{\text{owner}}$ . More specifically, in protocol  $\mathcal{P}_{\text{iso}}$ , every party  $pid$  owns a single key pair per key exchange session, represented by SID  $(pid, (sid', pid_i, pid_r))$ . We use the accountable PKI functionality  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  as PKI. To simplify presentation, we use a fixed session of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  (namely  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  runs in the session with SID  $\varepsilon$ ). This, however, can trivially be extended to the case where multiple PKIs are used by different key exchanges. We use an initialization functionality  $\mathcal{F}_{\text{init}}^{\text{iso}}$  which ensures that *initiator* and *responder* and their judges are initialized “at the same time”. This simplifies the proof later on, as we do not have to take care of several edge cases in cases of non-initialized parties.

The corruption model of  $\mathcal{P}_{\text{iso}}$  allows  $\mathcal{A}$  to corrupt users  $(pid, sid, role)$  (statically) before the key exchange starts. We consider a party corrupted, if (i) a party is directly corrupted, (ii) the instance of the party at  $\mathcal{F}_{\text{sig}}$ , i. e., its “private signature key”, is corrupted, (iii) its verification instance at  $\mathcal{F}_{\text{sig}}$ , used to verify the intended partners signature, is corrupted, and (iv) its session at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  is corrupted. This essentially models a party-wise corruption.

Additionally to  $\mathcal{P}_{\text{iso}}$ ,  $\mathcal{P}_{\text{iso}}^{\text{acc}}$  adds local judges to  $\mathcal{P}_{\text{iso}}$  – one local judge per *signer*, resp. *verifier* instance of  $\mathcal{P}_{\text{iso}}$  – and a dummy supervisor. For the local judges in  $\mathcal{P}_{\text{iso}}^{\text{acc}}$ , we essentially pass-through the accountability properties from  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ . The judge queries for verdicts at its own local judge (instance) at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  and the judge of the intended partner. In case that there is a verdict from  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  from the intended partners judge,  $\mathcal{P}_{\text{judge}}^{\text{iso}}$  adds the intended partner as potentially misbehaving party to the verdict (i. e., adds  $\forall\text{dis}(\text{intended partner})$ ), if the intended partner’s judge provides a verdict from the lower level. This captures that the judge cannot be certain who behaved maliciously on the PKI level: the issuing CA or the intended partner.

<sup>16</sup>In contrast to Appendix E, where we use an adapted variant of  $\mathcal{F}_{\text{sig}}$ . Additionally, we also allow  $\mathcal{A}$  to corrupt the *verifier* instance here.

Description of the protocol  $\mathcal{F}_{\text{KE}}^{\text{acc}} = (\text{initiator}, \text{responder}, \text{supervisor}, \text{judge})$ :

<b>Participating roles:</b> {initiator, responder, judge, supervisor}	
<b>Corruption model:</b> <i>dynamic with secure erasures</i>	
<b>Protocol parameters:</b>	
- groupGen( $1^n$ ).	{Algorithm for generating tuples $(G, n, g)$ describing cyclic groups $G$ of size $n$ with generator $g$ .
- Sec <sup>acc</sup> $\subset \{0, 1\}^*$	{Accountability properties, meant to be set to {authenticity}}
- Sec <sup>assumption</sup> $\subset \{0, 1\}^*$	{Assumption-based security properties, meant to be set to $\emptyset$ }
- pids <sub>judge</sub> $\subset \{0, 1\}^*$	{set of judge entities/(P)IDs in the protocol (which are often directly related to some protocol participants)}
- ids <sub>assumption</sub> $\subset \{0, 1\}^*$	{set of entities/IDs where properties are ensured via assumptions}

Description of  $M_{\text{initiator}, \text{responder}, \text{supervisor}, \text{judge}}$ :

<b>Implemented role(s):</b> {initiator, responder, judge, supervisor}	
<b>Subroutines:</b> $\mathcal{F}_{\text{judgeParams}}^{\text{KE}}$	
<b>Internal state:</b>	
- $(G, n, g) \in (\{0, 1\}^* \cup \{\perp\})^3 = (\perp, \perp, \perp)$	{Global group parameters.}
- state : $(\{0, 1\}^*)^3 \rightarrow \{\perp, \text{registered}, \text{started}, \text{finished}\}$	{Stores the current state of entities in key exchange; initially $\perp$ .
- initiator : $(\{0, 1\}^*)^3$	{The initiator of the key exchange}
- responder : $(\{0, 1\}^*)^3$	{The responder of the key exchange}
- caller : {initiator, responder} $\rightarrow (\{0, 1\}^*)^3 \cup \{\perp\}$	{Stores the calling entity for both entities in key exchange; initially $\perp$ .
- sessionKey : $\{0, 1\}^* \cup \{\perp\}$	{The session key; initially $\perp$ .
- corruptedIntParties $\in \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \setminus (\text{Roles}_{\mathcal{F}}^a \cup \{\text{judge}, \text{supervisor}\})$ , initially $\emptyset$	{The set of corrupted internal parties $(pid, sid, role)$ }
- brokenAssumptions : Sec <sup>assumption</sup> $\times$ ids <sub>assumption</sub> $\rightarrow \{\text{true}, \text{false}\}$	{Stores broken security assumptions per id, initially false $\forall$ entries}
- brokenProps : (Sec <sup>assumption</sup> $\cup$ Sec <sup>acc</sup> ) $\times$ (pids <sub>judge</sub> $\cup$ ids <sub>assumption</sub> ) $\rightarrow \{\text{true}, \text{false}\}$	{Stores broken security properties per judge/id, initially false $\forall$ entries}
- verdicts : pids <sub>judge</sub> $\rightarrow \{0, 1\}^*$	{Verdicts per $p \in$ pids <sub>judge</sub> , initially $\varepsilon$ }
<b>CheckID(<math>pid, sid, role</math>):</b>	
Accept all entities of form $(pid, (sid', pid_i, pid_r), role)$ where $role \in \{\text{initiator}, \text{responder}\}$ and $pid \in \{pid_i, pid_r\}$ or $((\text{local}, pid', role'), (sid', pid_i, pid_r), \text{judge})$ where $role' \in \{\text{initiator}, \text{responder}\}$ and also $pid' \in \{pid_i, pid_r\}$ or $(\_, (sid', pid_i, pid_r), \text{supervisor})$ .	{One session of $\mathcal{F}_{\text{KE}}^{\text{acc}}$ manages a single instance of a key exchange.}
<b>Corruption behavior:</b>	
- <b>AllowCorruption(<math>pid, sid, role</math>):</b>	
Do not allow corruption of $(pid, sid, \text{supervisor})$ .	
<b>if</b> $role = \text{judge}$ :	
<b>send</b> (Corrupt, $(pid, sid, \text{judge})$ , internalState)	
<b>to</b> $(pid, sid, \mathcal{F}_{\text{judgeParams}} : \text{judgeParams})$	{ $\mathcal{F}_{\text{judgeParams}}$ decides whether judges can be corrupted}
<b>wait for</b> $b$ ; <b>return</b> $b$	
- <b>DetermineCorrStatus(<math>pid, sid, role</math>):</b>	
<b>if</b> $role = \text{judge}$ :	
<b>send</b> (CorruptionStatus?, $(pid, sid, \text{judge})$ , internalState)	{ $\mathcal{F}_{\text{judgeParams}}$ may determine a judge's corruption status}
<b>to</b> $(pid, sid, \mathcal{F}_{\text{judgeParams}} : \text{judgeParams})$	
<b>wait for</b> $b$ ; <b>return</b> $b$	
- <b>AllowAdvMessage(<math>pid, sid, role, pid_{\text{receiver}}, sid_{\text{receiver}}, role_{\text{receiver}}, m</math>)</b>	{ $A$ is not allowed to invoke $\mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}}$ }
Do not allow sending messages to $\mathcal{F}_{\text{judgeParams}}$ .	
- <b>LeakedData(<math>pid, sid, role</math>):</b>	
If called while $(pid, sid, role)$ determines its initial corruption status, use the default behavior of <b>LeakedData</b> .	
That is, output the initially received message and the sender of that message.	
Otherwise, return $(\text{caller}[pid, sid, role], \text{sessionKey})$ .	
<b>Initialization:</b>	
Parse $sid_{\text{cur}}$ as $(sid', pid_i, pid_r)$ .	
initiator $\leftarrow (pid_i, (sid', pid_i, pid_r), \text{initiator})$	{Store initiator and responder}
responder $\leftarrow (pid_r, (sid', pid_i, pid_r), \text{responder})$	
$(G, n, g) \leftarrow \text{groupGen}(1^n)$ .	
<b>if</b> $sender = \text{NET} \wedge m = \text{InitGroup}$ :	
<b>send</b> (LeakGroup, $(G, n, g)$ ) <b>to</b> NET.	{Allow adversary to start initialization and then return the generated group. No other actions are performed in this case.}
<b>else</b> :	
<b>send responsively</b> (LeakGroup, $(G, n, g)$ ) <b>to</b> NET;	{Leak group parameters to the adversary. Note that this command forces the adversary to respond before interacting with the protocol in any other way, i.e., the run can continue as expected.}
<b>wait for</b> $\_$ .	
<b>MessagePreprocessing:</b>	
<b>if</b> $role_{\text{cur}} = \text{initiator} \wedge \text{caller}[\text{initiator}] = \perp$ :	
$\text{caller}[\text{initiator}] \leftarrow \text{entity}_{\text{call}}$	
<b>else if</b> $role_{\text{cur}} = \text{responder} \wedge \text{caller}[\text{responder}] = \perp$ :	
$\text{caller}[\text{responder}] \leftarrow \text{entity}_{\text{call}}$	
<b>Main:</b>	
<b>rcv</b> RegisterPubKey <b>from</b> I/O <b>to</b> $(\_, \_, \text{initiator}) \vee (\_, \_, \text{responder})$ <b>s.t.</b> $\text{state}[(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})] = \perp$ :	{Register that initiator and responder want exchange keys}
$\text{state}[(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})] \leftarrow \text{registered}$ .	
<b>send</b> RegisterPubKey <b>to</b> NET.	
<b>rcv</b> InitKE <b>from</b> I/O <b>to</b> $(\_, \_, \text{initiator}) \vee (\_, \_, \text{responder})$ <b>s.t.</b> $\text{state}[(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})] = \text{registered}$ :	{Start key exchange key}
$\text{state}[(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})] \leftarrow \text{started}$ .	
<b>send</b> InitKE <b>to</b> NET.	
<sup>a</sup> Here: $\text{Roles}_{\mathcal{F}} := \{\text{initiator}, \text{responder}\}$ .	

Fig. 32: The accountable ideal key exchange functionality  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  (part 1).

Description of  $M_{\text{initiator, responder, supervisor, judge}}$  (continued):

<b>Main:</b>	<div style="text-align: right; font-size: small; margin-bottom: 5px;"> <math>\left\{ \begin{array}{l} \text{Note that Main continues processing the message } m \text{ that Initialization has} \\ \text{already parsed if Initialization does not end the run.} \end{array} \right.</math> </div> <p> <b>recv</b> (FinishKE, <math>k</math>) <b>from</b> NET <b>to</b> <math>(\_, \_, \text{initiator}) \vee (\_, \_, \text{responder})</math> <b>s.t.</b> <math>\text{state}[(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \text{role}_{\text{cur}})] = \text{started}</math>:  <b>if</b> sessionKey = <math>\perp</math>:              Choose sessionKey <math>\leftarrow</math> G uniformly at random              state[entity<sub>cur</sub>] <math>\leftarrow</math> finished.  <b>if</b> initiator <math>\in</math> CorruptionSet <math>\vee</math> responder <math>\in</math> CorruptionSet <math>\vee</math>              brokenProps[authenticity, (local, initiator, initiator)] = true <math>\vee</math>              brokenProps[authenticity, (local, responder, responder)] = true: <math>\left\{ \begin{array}{l} \mathcal{A} \text{ may choose the key if he has corrupted the partner or} \\ \text{authenticity is broken for one of the involved parties.} \end{array} \right.</math>              <b>send</b> (FinishKE, <math>k</math>) <b>to</b> caller[entity<sub>cur</sub>].  <b>else:</b>              <b>send</b> (FinishKE, sessionKey) <b>to</b> caller[entity<sub>cur</sub>].         </p> <p style="font-size: x-small; margin-top: 10px;">Include static code from the AUC transformation <math>\mathcal{T}_1(\cdot)</math> here, i.e., include additional code from Figure 2 and 3 here. Do not include the public judge here.</p>
--------------	---

Fig. 33: The accountable ideal key exchange functionality  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  (part 2).

Description of  $\mathcal{F}_{\text{judgeParams}}^{\text{KE}} = (\text{judgeParams})$ :

<b>Participating roles:</b> {judgeParams} <b>Corruption model:</b> <i>incorruptible</i>
--

Description of  $M_{\text{judgeParams}}$ :

<b>Implemented role(s):</b> {judgeParams} <b>CheckID</b> ( $\text{pid}, \text{sid}, \text{role}$ ): Accept all messages with the same $\text{sid}$ . <b>Main:</b>	<p> <b>recv</b> (BreakAccProp, <math>\text{verdict}, \text{toBreak}, \text{internalState}</math>) <b>from</b> I/O:              Let initiator, responder be initiator and responder extracted from <math>\text{internalState}</math>. Further, let <math>\text{pid}_{\text{initiator}}</math>, resp. <math>\text{pid}_{\text{responder}}</math> be their PIDs.              <b>if</b> <math>\text{toBreak} = \{\text{authenticity}\} \times (\{\text{local}\} \times \underbrace{(\{\text{pid}_{\text{initiator}}, \text{initiator}\}, \{\text{pid}_{\text{responder}}, \text{responder}\})}_{=: \text{ids} \{s \text{ is defined to be the session ID of the considered session}\}})</math>:         </p> <p>                 <b>for all</b> <math>\text{id} \in \text{ids}</math> <b>do</b>:                      <b>if</b> <math>\text{verdict}</math> for <math>\text{id}</math> does not match conjunction of verdicts of form                          <math>\text{dis}((\text{pid}', \text{sid}', \text{role}')) \vee V</math> or <math>V</math>                          where <math>(\text{pid}', \text{sid}', \text{role}') \in \{\text{initiator}, \text{responder}\}</math> and                          <math>V</math> is a verdict not blaming parties with roles from {initiator, responder, judge, supervisor}:                          Let <math>k</math> be sessionKey extracted from <math>\mathcal{F}_{\text{KE}}^{\text{acc}}</math>'s provided internal state.                          <b>reply</b> (BreakAccProp, true, <math>k</math>) <span style="float: right; font-size: x-small;">{Accept verdict/broken properties, leak session key}</span>                      <b>else:</b>                          <b>reply</b> (BreakAccProp, false, <math>\varepsilon</math>) <span style="float: right; font-size: x-small;">{Decline verdict/broken properties}</span> </p> <p> <b>recv</b> (GetJudicialReport, <math>\text{msg}, \text{internalState}</math>) <b>from</b> I/O:              <b>reply</b> (GetJudicialReport, <math>\varepsilon</math>) <span style="float: right; font-size: x-small;">{Generate judicial report We do not provide a judicial report here}</span> </p> <p> <b>recv</b> (BreakAssumption, <math>\text{toBreak}, \text{internalState}</math>) <b>from</b> I/O:              <b>reply</b> (BreakAssumption, <math>\varepsilon</math>) <span style="float: right; font-size: x-small;">{Do not generate leakage when breaking assumptions}</span> </p> <p> <b>recv</b> (Corrupt, <math>(\text{id}, \text{sid}, \text{judge}), \text{internalState}</math>) <b>from</b> I/O:              <b>if</b> <math>\text{id} = (\text{local}, \text{pid}, \text{role}) \wedge (\text{pid}, \text{sid}, \text{role}) \in \text{CorruptionSet}</math>: <math>\left\{ \begin{array}{l} \mathcal{F}_{\text{judgeParams}}^{\text{KE}} \text{ allows to corrupt a local judge iff the} \\ \text{accompanied initiator or responder is corrupted} \end{array} \right.</math>                  <b>reply</b> true              <b>else:</b>                  <b>reply</b> false         </p> <p> <b>recv</b> (CorruptionStatus?, <math>(\text{id}, \text{sid}, \text{judge}), \text{internalState}</math>) <b>from</b> I/O: <math>\left\{ \begin{array}{l} \mathcal{F}_{\text{judgeParams}}^{\text{KE}} \text{ interprets a local judge as corrupted iff} \\ \text{the accompanied initiator or responder is corrupted} \end{array} \right.</math>              <b>if</b> <math>(\text{id}, \text{sid}, \text{judge}) \in \text{CorruptionSet}</math>:                  <b>reply</b> true              <b>else if</b> <math>\text{id} = (\text{local}, \text{pid}, \text{role}) \wedge (\text{pid}, \text{sid}, \text{role}) \in \text{CorruptionSet}</math> :                  <b>reply</b> true              <b>else:</b>                  <b>reply</b> false         </p>
--	---

Fig. 34: The judge parameter functionality  $\mathcal{F}_{\text{judgeParams}}^{\text{KE}}$  for  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ .

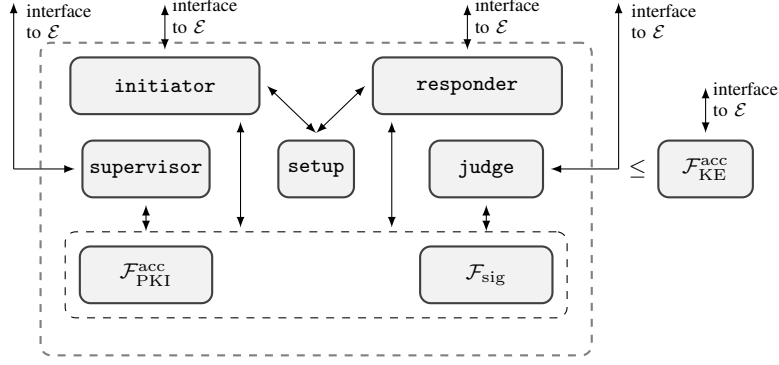


Fig. 35: Realization relation of a key exchange example stated in Theorem 8. The system  $\mathcal{E}$  denotes the environment, modeling, as usual in UC setting, arbitrary higher level protocols. All machines are additionally connected to the network adversary.

As we do not consider assumption-based security properties in this case study, the supervisor  $\mathcal{P}_{sv}^{iso}$  always responds with `false` if a party queries for the status of an assumption. Further,  $\mathcal{P}_{sv}^{iso}$  allows querying for the corruption status of (internal) parties in  $\mathcal{F}_{PKI}^{acc}$ .

**F.3 Security Analysis:** We now provide the formal result that we can use the ISO protocol on top of an accountable PKI to realize the (accountable) key exchange functionality  $\mathcal{F}_{KE}^{acc}$ .

**Theorem 8.** *Let  $\text{groupGen}(1^\eta)$  be an algorithm that outputs descriptions  $(G, n, g)$  of cyclic groups (i. e.,  $G$  is a group of size  $n$  with generator  $g$ ) such that  $n$  grows exponentially in  $\eta$  and the DDH assumption holds true. Let  $\mathcal{P}_{iso}^{acc}$  be the (accountable) ISO protocol as described above, let  $\mathcal{F}_{PKI}^{acc}$  be the accountable ideal PKI functionality and  $\mathcal{F}_{judgeParams}^{acc-PKI}$  the associated subroutine both with parameters  $\text{Sec}^{acc} = \{\text{correctCert}\}$  and  $\text{Sec}^{assumption} = \emptyset$ , and let  $\mathcal{F}_{sig}$  be the (standard) ideal signature functionality, and let  $\mathcal{F}_{KE}^{acc}$  be the (accountable) ideal functionality for authenticated key exchange with  $\text{Sec}^{acc} = \{\text{authenticity}\}$ ,  $\text{Sec}^{assumption} = \emptyset$ ,  $\text{pids}_{judge} = \{\text{local}\} \times (\{0, 1\}^*)^2$ , and  $\text{ids}_{assumption} = \emptyset$ . Then the following holds true:*

$$(\mathcal{P}_{iso}^{acc}, \mathcal{P}_{sv}^{iso}, \mathcal{P}_{judge}^{iso} \mid \mathcal{F}_{PKI}^{acc}, \mathcal{F}_{judgeParams}^{acc-PKI}, \mathcal{F}_{sig}) \leq (\mathcal{F}_{KE}^{acc} \mid \mathcal{F}_{judgeParams}^{KE}).$$

*Proof.* We show that the real key exchange  $\mathcal{R} := (\mathcal{P}_{iso}^{acc}, \mathcal{P}_{sv}^{iso}, \mathcal{P}_{judge}^{iso} \mid \mathcal{F}_{PKI}^{acc}, \mathcal{F}_{judgeParams}^{acc-PKI}, \mathcal{F}_{sig})$  realizes the ideal key exchange  $\mathcal{I} := (\mathcal{F}_{KE}^{acc} \mid \mathcal{F}_{judgeParams}^{KE})$ . As part of this, we define a responsive simulator  $\mathcal{S}$  such that the real world running  $\mathcal{R}$  is indistinguishable from the ideal world running  $\{\mathcal{S}, \mathcal{I}\}$  for every ppt environment  $\mathcal{E}$ .

First note that it is easy to see that both  $\mathcal{R}$  and  $\mathcal{I}$  are environmentally bounded and complete. Now, the simulator  $\mathcal{S}$  is defined as follows: the simulator is a single machine that is connected to  $\mathcal{I}$  and the environment via their network interfaces. In a run, there is only a single instance of the machine  $\mathcal{S}$  that accepts all incoming messages. The simulator  $\mathcal{S}$  internally simulates the *full protocol*  $\mathcal{R}$ , including its behavior on the network interface to the environment. To improve readability, we refer to  $\mathcal{S}$ 's simulated version of the real protocol as  $\mathcal{R}'$ . More precisely, the simulation runs as follows:

- At the start of a run,  $\mathcal{S}$  obtains the group parameters used by  $\mathcal{F}_{KE}^{acc}$ : if the simulator is activated for the first time via the `(LeakGroup, (G, n, g))` message, then he simply saves this message and returns `ok`. Otherwise,  $\mathcal{S}$  sends an `InitGroup` message to (an arbitrary entity of)  $\mathcal{F}_{KE}^{acc}$  to trigger **Initialization** and obtain the group parameters. Note that the environment cannot observe whether the simulator has manually triggered **Initialization** of  $\mathcal{F}_{KE}^{acc}$ . The group parameters are used by  $\mathcal{S}$  as output of the internally simulated  $\mathcal{P}_{iso}^{acc}$  : `setup` role.
- Upon receiving `RegisterPubKey` from an honest party,<sup>17</sup>  $\mathcal{S}$  forwards this message to the dedicated machine of  $\mathcal{R}'$ . We note that  $\mathcal{S}$  can derive all involved parties and their roles from the `SID` (which predefines initiator

<sup>17</sup>We consider an entity to be honest if it outputs `false` upon `CorruptionStatus?` requests. Conversely, we call an entity corrupted if it outputs `true`, even if it was not explicitly corrupted.



and responder of the key exchange session including their local judges). In  $\mathcal{R}'$ , the message then triggers the initialization of the party's signer instance at  $\mathcal{F}_{\text{sig}}$  and registers the generated public key for the dedicated party at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ .

- Upon receiving a `InitKE` from an honest entity,  $\mathcal{S}$  forwards this message – again – in the name of a higher-level protocol to the simulated entity in  $\mathcal{R}'$ . This triggers the start of a key exchange as defined in the ISO protocol in  $\mathcal{R}'$ .
- As soon as an honest initiator in  $\mathcal{R}'$  outputs `(FinishKE, sessionKey)`,  $\mathcal{S}$  sends the message `(FinishKE, sessionKey)` to the initiator identity in  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ .
- As soon as an honest responder in  $\mathcal{R}'$  outputs `(FinishKE, sessionKey)`,  $\mathcal{S}$  distinguishes three cases:
  1. In the case that initiator is corrupted after its output of `FinishKE`,  $\mathcal{I}$  leaks the initiator's ideal key  $k'$  to  $\mathcal{S}$ . Then,  $\mathcal{S}$  sends `(FinishKE, k')` to  $\mathcal{I}$ .
  2. If authenticity is broken for (at least) one of both parties after initiator outputs the session key via `FinishKE`. In this case,  $\mathcal{S}$  learns the (ideal) session key  $k'$  from  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  when providing the verdict to  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  before outputting `FinishKE`. Then,  $\mathcal{S}$  sends `(FinishKE, k')` to  $\mathcal{I}$ .
  3. In all other cases,  $\mathcal{S}$  sends the message `(FinishKE, sessionKey)` to the responder identity in  $\mathcal{F}_{\text{KE}}^{\text{acc}}$ .
- $\mathcal{S}$  forwards all network communication from the environment to corresponding entities in  $\mathcal{R}'$ , and vice versa.
- $\mathcal{S}$  keeps the (internal) corruption status of entities in roles of  $\mathcal{R}'$  and  $\mathcal{I}$  synchronized. In particular, if an entity of  $\mathcal{I}$  asks for its initial corruption status, then the same entity in  $\mathcal{R}'$  is simulated to also do so. Furthermore, as soon as a simulated entity of  $\mathcal{R}'$  considers itself to be corrupted (either explicitly or due to a corrupted subroutine),  $\mathcal{S}$  corrupts the same entity in  $\mathcal{I}$ .
- If a corrupted entity *entity* in a public role of  $\mathcal{R}$  outputs a message on its I/O interface to a higher-level protocol, then  $\mathcal{S}$  instructs the (explicitly) corrupted entity *entity* in  $\mathcal{I}$  to forward the same message on its I/O interface. The same is also done in the other direction. Note that this also includes corrupted local judges.
- Always when an instance  $((\text{local}, \text{pid}, \text{role}), \text{sid}, \text{judge})$  of  $\mathcal{P}_{\text{judge}}^{\text{iso}}$  in  $\mathcal{R}'$  renders a verdict,  $\mathcal{S}$  extracts this verdict and sends `(BreakAccProp, verdict, toBreak)` where *verdict* maps  $(\text{local}, \text{pid}, \text{role})$  to the verdict extracted from the simulation (and all other entries to  $\varepsilon$ ) and *toBreak* =  $\{(\text{authenticity}, (\text{local}, \text{pid}, \text{role}))\}$  to  $((\text{local}, \text{pid}, \text{role}), \text{sid}, \text{judge})$  in  $\mathcal{I}$ .
- As the supervisor does not have to handle assumption-based properties, there is no additional specification for  $\mathcal{S}$  necessary.

This concludes the description of the simulator. It is easy to see that  $\{\mathcal{S}, \mathcal{I}\}$  is environmentally bounded and  $\mathcal{S}$  is responsive for  $\mathcal{I}$ . We note that the behavior of the supervisor is indistinguishable between  $\mathcal{R}$  and  $\mathcal{I}$  as he always provides the same output in both worlds. Now, let  $\mathcal{E}$  be an arbitrary but fixed responsive environment.

We will now go over all possible interactions with honest entities on the I/O interface and argue, by induction, that all of those interactions result in identical behavior towards the environment, i.e.,  $\mathcal{I}$  and  $\mathcal{R}$  are indistinguishable. At the start of a run, there were no interactions on the I/O interface with honest parties yet, i.e.,  $\mathcal{I}$  and  $\mathcal{R}$  are indistinguishable. In the following, the induction assumption is that all I/O interactions to far have resulted in the same behavior visible towards the environment in both the real and ideal world.

We firstly have a look at the I/O output of initiator and responder. We then prove that the output of local judges is also indistinguishable between the  $\mathcal{R}$  and  $\{\mathcal{S}, \mathcal{I}\}$ .

**Initiator and responder.** We note that all I/O calls to honest initiator and responder do not directly lead to an output to  $\mathcal{E}$  in  $\mathcal{I}$  and  $\mathcal{R}'$  (for the corrupted case, see below). Thus,  $\mathcal{I}$  and  $\mathcal{R}$  are indistinguishable. However, note that the phases of the protocol are synchronized in  $\mathcal{I}$  and  $\mathcal{R}'$ . `RegisterPubKey` initializes the signatures key at  $\mathcal{F}_{\text{sig}}$  and also registers the public key at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  in  $\mathcal{R}'$  which does not have an adequate in  $\mathcal{I}$ . `InitKE` starts the key exchange itself in both worlds.

Note that honest initiator or responder do not send `FinishKE` to  $\mathcal{E}$  if their partners public key is not registered at the PKI  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  or if the signatures provided via NET do not verify. Also note that a party initialization leads to an initialization of both parties, initiator and responder and also their local judges. This enforces that the corruption status of the machines is always correctly handled and that we do not have to take care of edge cases due to uninitialized parties.

When the initiator in  $\mathcal{R}'$  wants to output the message `(FinishKE, sessionKey)` to  $\mathcal{E}$ ,  $\mathcal{S}$  triggers  $\mathcal{I}$  via `(FinishKE, sessionKey)` also to output a session key. If initiator and responder are uncorrupted and the property

authenticity holds true for both,  $\mathcal{I}$  outputs a randomly selected group element (potentially  $\neq$  sessionKey) from  $h_{\text{init}}$  on behalf of initiator. As  $\mathcal{E}$  is not aware of the sessionKey (both parties are honest) and as the DDH assumption holds true,  $\mathcal{E}$  can only distinguish with negligible probability between  $\mathcal{I}$  and  $\mathcal{R}'$ .

If both parties are still uncorrupted but authenticity is broken for at least one of both parties, this leads to the case that  $\mathcal{I}$  directly outputs the sessionKey provided by  $\mathcal{S}$ . As the session key is extracted from  $\mathcal{R}'$ , this case is also indistinguishable between  $\mathcal{R}$  and  $\mathcal{I}$ .

In case that initiator considers itself as corrupted in  $\mathcal{R}'$ , responder outputs the sessionKey. Note that this session key matches the one exchanged with initiator thus potentially leaked to  $\mathcal{A}$ . Thus,  $\mathcal{E}$  cannot distinguish between  $\mathcal{I}$  and  $\mathcal{R}$ .

The same holds true for the case, responder is corrupted but initiator is not.

The two cases above also include the cases when authenticity is additionally broken for one party – as this does not change the behavior of  $\mathcal{I}$  in the case that already one of the parties is corrupted.

Note that we have to consider two special cases here:

1. initiator gets corrupted after it outputted the key. However, as  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  leaks the key upon corruption to  $\mathcal{S}$  and  $\mathcal{S}$  forwards this leaked key to the responder,  $\mathcal{E}$  cannot distinguish between  $\mathcal{I}$  and  $\mathcal{R}$ .
2. initiator or responder are both honest but one of them/both lose/s authenticity after the initiator already outputted FinishKE including an ideally generated session key to  $\mathcal{E}$ . However, as  $\mathcal{F}_{\text{KE}}^{\text{acc}}$  leaks the key upon breaking to  $\mathcal{S}$  and  $\mathcal{S}$  forwards this leaked key to the responder,  $\mathcal{E}$  cannot distinguish between  $\mathcal{I}$  and  $\mathcal{R}$ .

**Judges.** Note that  $\mathcal{S}$  updates verdicts in  $\mathcal{I}$  as soon as they occur in  $\mathcal{R}'$ . Further note that GetVerdict and GetJudicialReport do not influence any state in  $\mathcal{R}$  and  $\mathcal{I}$ . We also note that answers to GetJudicialReport are indistinguishable in both worlds.

Further,  $\mathcal{P}_{\text{judge}}^{\text{iso}}$  always accepts the input of  $\mathcal{S}$  as rendered verdicts in  $\mathcal{R}'$  are compliant with the restrictions imposed in  $\mathcal{F}_{\text{judgeParams}}^{\text{KE}}$ . Thus, local judges output the same verdicts in  $\mathcal{I}$  and  $\mathcal{R}'$  and we can conclude that both worlds are indistinguishable.

As local judges act as pure message forwarder when their associated party is corrupted, we can also conclude that  $\mathcal{I}$  and  $\mathcal{R}$  are indistinguishable in this case.

**Supervisor.** Supervisors is always the same in  $\mathcal{R}$  and  $\mathcal{I}$ . Thus, both worlds are indistinguishable regarding their output.

We can also conclude that the corruption status of (internal) entities in the real and ideal world is synchronized. Since the simulator has full control over corrupted entities, which are handled via the internal simulation  $\mathcal{R}'$ , this implies that the I/O behavior of corrupted entities of  $\mathcal{R}/\mathcal{I}$  towards the environment is also identical in the real and ideal world.

Note that the behavior of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  is only visible to  $\mathcal{I}$  via its network interface. As  $\mathcal{S}$  forwards these messages to  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  in  $\mathcal{R}'$  and also provide the output of  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  to  $\mathcal{E}$  via NET,  $\mathcal{E}$  cannot use  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  to distinguish between  $\mathcal{I}$  and  $\mathcal{R}$ .

Overall, we can conclude that  $\mathcal{I}$  and  $\mathcal{R}$  are indistinguishable. □

**F.4 Deterrence Analysis:** Please note that we consider the fully composed protocol here, i. e., including the used PKI. We emphasize that the meaning of a single established key exchange is negligible in contrast to a working PKI. Therefore, being caught cheating in the PKI is the major deterrence to behave honestly in the key exchange. In particular, we conclude that the utilities for PKI participants have the same relations as in Appendix E.4, i. e., Equation 1 and 2 are met for these parties. For the parties involved in the key exchange, we assume that  $U_{hP}^i = U_{hE}^i = U_{hL}^i = U_{mP}^i = U_{mL}^i = U_{mE}^i = 0$  expressing that a party typically does not have a big advantage from executing the key exchange honest or maliciously. Thus, Equation 1 and 2 holds true for all involved parties.

### G. Capturing MPC Accountability Properties via AUC

As already mentioned in the introduction, there are several works that capture accountability properties in a UC model for the special case of MPC protocols (e.g., [13, 14, 24, 34, 35, 55, 74, 81]). These properties include (*publicly identifiable abort* [14, 34, 55]), (*public/universal verifiability* [35, 84, 86]), *auditability* [13],

Description of the protocol  $\mathcal{P}_{\text{iso}}^{\text{acc}} = (\text{initiator}, \text{responder} \mid \text{setup})$ :

<p><b>Participating roles:</b> initiator, responder, setup  <b>Corruption model:</b> static  <b>Protocol parameters:</b>          - groupGen(<math>1^n</math>).</p>	<p>{Algorithm for generating tuples <math>(G, n, g)</math> describing cyclic groups <math>G</math> of size <math>n</math> with generator <math>g</math>.</p>
---	--

Description of  $M_{\text{initiator}}$ :

<p><b>Implemented role(s):</b> initiator  <b>Subroutines:</b> setup, <math>\mathcal{F}_{\text{sig}}</math>, <math>\mathcal{F}_{\text{PKI}}^{\text{acc}}</math>, <math>\mathcal{F}_{\text{init}}^{\text{iso}}</math> : init  <b>Internal state:</b>          - <math>\text{pk} \in \{0, 1\}^*</math>, <math>\text{pk} = \varepsilon</math>, <math>\text{pid}_{\text{CA}} \in \{0, 1\}^*</math>, <math>\text{pid}_{\text{CA}} = \varepsilon</math>          - <math>(G, n, g) \in (\{0, 1\}^* \cup \{\perp\})^3 = (\perp, \perp, \perp)</math>          - <math>\text{state} \in \{\perp, \text{registered}, \text{started}, \text{finished}\} = \perp</math>          - <math>\text{initiator} : (\{0, 1\}^*)^3</math>          - <math>\text{responder} : (\{0, 1\}^*)^3</math>          - <math>\text{caller} \in (\{0, 1\}^*)^3 \cup \{\perp\} = \perp</math>          - <math>\text{sessionKey} \in \{0, 1\}^* \cup \{\perp\} = \perp</math>          - <math>e_{\text{init}} \in \mathbb{Z}_n \cup \{\perp\} = \perp</math>          - <math>h_{\text{resp}} \in G \cup \{\perp\} = \perp</math></p>	<p>{The public key of <math>(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \text{role}_{\text{cur}})</math> and the used CA          {Global group parameters.          {Current state in key exchange.          {The initiator of the key exchange          {The responder of the key exchange          {Stores the initial caller of this entity/instance.          {Stores the session key.          {Secret DH exponent of initiator.          {Public DH key share of responder.</p>
<p><b>CheckID</b>(<math>\text{pid}, \text{sid}, \text{role}</math>):          Accept all entities of form <math>(\text{pid}_i, (\text{sid}', \text{pid}_i, \text{pid}_r), \text{initiator})</math>.</p>	
<p><b>Corruption behavior:</b>          - <b>DetermineCorrStatus</b>(<math>\text{pid}, \text{sid}, \text{role}</math>) : {Entity corrupted if one of its signature keys or the verification subroutine are corrupted.</p>	
<p><math>\text{out} \leftarrow \text{corr}(\text{pid}, (\text{pid}, \text{sid}), \mathcal{F}_{\text{sig}} : \text{signer}) \vee \text{corr}(\text{pid}, (\text{pid}(\text{responder}), \text{sid}), \mathcal{F}_{\text{sig}} : \text{verifier})</math>  <math>\text{out} \leftarrow \text{out} \vee \text{corr}(\text{pid}, \text{sid}, \text{role}), \varepsilon, \mathcal{F}_{\text{PKI}}^{\text{acc}} : \text{client}</math>  <b>return out</b></p>	<p>{ <math>\text{pid}(\text{responder})</math>          { denotes responder's          { PID</p>
<p>- <b>AllowAdvMessage</b>(<math>\text{pid}, \text{sid}, \text{role}, \text{pid}_{\text{receiver}}, \text{sid}_{\text{receiver}}, \text{role}_{\text{receiver}}, m</math>):          If <math>\text{role}_{\text{receiver}} = \mathcal{F}_{\text{sig}} : \text{signer}</math>, return <b>false</b>.<sup>a</sup>          Otherwise output true iff <math>\text{pid} = \text{pid}_{\text{receiver}}</math>.</p>	
<p><b>Initialization:</b>          Parse <math>\text{sid}_{\text{cur}}</math> as <math>(\text{sid}', \text{pid}_i, \text{pid}_r)</math>.  <math>\text{initiator} \leftarrow (\text{pid}_i, (\text{sid}', \text{pid}_i, \text{pid}_r), \text{initiator})</math>.  <math>\text{responder} \leftarrow (\text{pid}_r, (\text{sid}', \text{pid}_i, \text{pid}_r), \text{responder})</math>.  <math>\text{caller} \leftarrow (\text{pid}_{\text{call}}, \text{sid}_{\text{call}}, \text{role}_{\text{call}})</math>.  <b>send</b> GetParameters <b>to</b> <math>(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \text{setup})</math>;  <b>wait for</b> (GetParameters, <math>(G, n, g)</math>).  <math>(G, n, g) \leftarrow (G, n, g)</math>.  <b>if</b> ITM was not activated via <b>init</b> macro, resp. a <b>InitEntity</b> message:  <b>send</b> initPartner <b>to</b> <math>(\text{pid}_{\text{cur}}, \text{sid}_{\text{cur}}, \mathcal{F}_{\text{init}}^{\text{iso}} : \text{init})</math>  <b>wait for</b> initPartner</p>	<p>{Store initiator and responder.          {Get DH parameters</p>
<p><b>Main:</b>          See Figure 37.</p>	
<p><sup>a</sup>In our modeling, the corruption status of <b>signer</b> entities indicates whether the adversary has access to the corresponding signature keys, i.e., whether he can sign his own messages (as in this case the <b>signer</b> entity should be considered compromised). Thus, the adversary is not allowed to access uncorrupted <b>signer</b> entities. If the <b>signer</b> entity is corrupted, then the adversary already knows the secret key and can sign messages on his own, so there is no need to give him access in this case.</p>	

Fig. 36: A real key exchange protocol  $\mathcal{P}_{\text{iso}}^{\text{acc}}$  for realizing  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  (part 1). Note that each instance of  $M_{\text{initiator}}$  and  $M_{\text{responder}}$  corresponds to a single entity.

Description of  $M_{\text{initiator}}$  (continued):

<p><b>Main:</b></p> <p><b>recv</b> RegisterPubKey from I/O s.t. state[(pid<sub>cur</sub>, sid<sub>cur</sub>, role<sub>cur</sub>)] = <math>\perp</math>:  state(pid<sub>cur</sub>, sid<sub>cur</sub>, role<sub>cur</sub>) <math>\leftarrow</math> registered.  <b>send</b> InitSign to (pid<sub>cur</sub>, (pid<sub>cur</sub>, sid<sub>cur</sub>), <math>\mathcal{F}_{\text{sig}}</math> : signer);  <b>wait for</b> (InitSign, success, pk); pk <math>\leftarrow</math> pk;  <b>send responsively</b> (Register, pk) to NET;  <b>wait for</b> (Register, pk, pid<sub>CA</sub>) s.t. pid<sub>CA</sub> <math>\in</math> {0, 1}<sup>*</sup>;  pid<sub>CA</sub> <math>\leftarrow</math> pid<sub>CA</sub>;  <b>send</b> (Register, pk, pid<sub>CA</sub>) to ((pid<sub>cur</sub>, sid<sub>cur</sub>, role<sub>cur</sub>), <math>\varepsilon</math>, <math>\mathcal{F}_{\text{PKI}}^{\text{acc}}</math> : client).<sup>a</sup></p> <p><b>recv</b> InitKE from I/O s.t. state = registered:  state <math>\leftarrow</math> started.  Choose <math>e_{\text{init}} \leftarrow \mathbb{Z}_n</math> uniformly at random, compute <math>h_{\text{init}} = g^{e_{\text{init}}}</math>.  <b>send</b> (Send, (pid<sub>cur</sub>, sid<sub>cur</sub>, role<sub>cur</sub>), h<sub>init</sub>) to NET.</p> <p><b>recv</b> (FinishKE, responder, h<sub>resp</sub>, (g<sup>e<sub>init</sub></sup>, h<sub>resp</sub>, responder), <math>\sigma</math>) from NET s.t. state = started:</p> <p><b>send responsively</b> (RetrieveCA, responder) to NET  <b>wait for</b> (RetrieveCA, responder, pid<sub>CA</sub>);  <b>send</b> (Retrieve, responder, pid<sub>CA</sub>) to ((pid<sub>cur</sub>, sid<sub>cur</sub>, role<sub>cur</sub>), <math>\varepsilon</math>, <math>\mathcal{F}_{\text{PKI}}^{\text{acc}}</math> : client);  <b>wait for</b> (Retrieve, intendedPartner, pid<sub>CA</sub>, pk).  <b>if</b> pk = <math>\perp</math>:    <b>abort</b>.  <b>send</b> (Verify, (g<sup>e<sub>init</sub></sup>, h<sub>resp</sub>, responder), <math>\sigma</math>, pk) to (pid<sub>cur</sub>, (pid(responder), sid<sub>cur</sub>), <math>\mathcal{F}_{\text{sig}}</math> : verifier);  <b>wait for</b> (VerResult, b).  <b>if</b> b = false:    <b>abort</b>.  h<sub>resp</sub> <math>\leftarrow</math> h<sub>resp</sub>; sessionKey <math>\leftarrow</math> h<sub>resp</sub><sup>e<sub>init</sub></sup>; state <math>\leftarrow</math> finished.  <b>send</b> (FinishKE, sessionKey) to caller.</p> <p><b>recv</b> GetLastMessage from NET s.t. state = finished:  m = (h<sub>resp</sub>, g<sup>e<sub>init</sub></sup>, responder).  <b>send</b> (Sign, m) to (pid<sub>cur</sub>, (pid<sub>cur</sub>, sid<sub>cur</sub>), <math>\mathcal{F}_{\text{sig}}</math> : signer);  <b>wait for</b> (Signature, <math>\sigma</math>).  <b>send</b> (send, <math>\sigma</math>) to NET.</p>	<p>{Register initiator's signature and certificate}</p> <p>{Register identity at <math>\mathcal{F}_{\text{PKI}}^{\text{acc}}</math>}</p> <p>{Start KE and send first message.}</p> <p>{Receive second message and output key.}</p> <p>{Query <math>\mathcal{A}</math> for responder's CA;}</p> <p>{Get public verification key of intended partner.}</p> <p>{pid(responder) denoted the PID of responder}</p>
---	---

<sup>a</sup>We consider here one instance of the PKI, namely the one with SID  $\varepsilon$ .

Description of  $M_{\text{setup}}$ :

<p><b>Implemented role(s):</b> setup</p> <p><b>Internal state:</b>  - (G, n, g) <math>\in</math> ({0, 1}<sup>*</sup> <math>\cup</math> {<math>\perp</math>})<sup>3</sup> = (<math>\perp</math>, <math>\perp</math>, <math>\perp</math>)</p> <p><b>CheckID</b>(pid, sid, role): Accept all entities.</p> <p><b>Corruption behavior:</b>  - <b>AllowCorruption</b>(pid, sid, role) : return false.</p> <p><b>Initialization:</b>  (G, n, g) <math>\leftarrow</math> groupGen(1<sup><math>\eta</math></sup>).</p> <p><b>Main:</b></p> <p><b>recv</b> GetParameters from _ :  <b>reply</b> (GetParameters, (G, n, g)).</p>	<p>{Global group parameters.}</p> <p>{The adversary may not corrupt the (honestly generated) setup parameters.}</p> <p>{Everyone may retrieve the group parameters, including the adversary on the network.}</p>
--	--

Fig. 37: A real key exchange protocol  $\mathcal{P}_{\text{iso}}^{\text{acc}}$  for realizing  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  (part 2).

openability [34], and privacy [8]. AUC can capture these accountability properties as special cases. This is not only an important sanity check but also shows that AUC generalizes and unifies existing UC accountability literature. Here we illustrate this for the most common property: identifiable abort. The other properties can be dealt with analogously, see Appendix H for more details.

**Identifiable abort.** The standard definition of a basic ideal MPC functionality  $\mathcal{F}_{\text{MPC}}$  (cf., e.g., [6, 8, 21, 24]) is based on three phases. In the first phase, it takes inputs from  $m$  parties, with inputs of corrupted parties being chosen by the adversary. In the second phase, it acts as a trusted third party that computes some function  $f$  on those inputs. In the final phase, each party receives an output of  $f$  but otherwise obtains no information. Hence,  $\mathcal{F}_{\text{MPC}}$  provides preventive security for correctness of the outputs and for privacy/secretcy of the inputs.

Instead of preventive security for correctness, MPC protocols often rather consider the weaker property of *identifiable abort* [6, 8, 14, 55, 85], which states that either all honest parties obtain a correct output or all honest parties agree on the name of a malicious party who has caused the output to be incorrect and hence the protocol to abort. In other words, identifiable abort is a type of (local) accountability w.r.t. correctness that additionally requires individual accountability and certain relationships between local properties of different parties.

In the literature, *identifiable abort* has been formalized within ideal functionalities  $\mathcal{F}_{\text{MPC}}^{\text{id-ab}}$  by letting the simulator, during the final output phase, decide whether all honest parties obtain their correct output or provide

Description of  $M_{\text{responder}}$ :

**Implemented role(s): responder**

**Subroutines:** setup,  $\mathcal{F}_{\text{sig}}$ ,  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ ,  $\mathcal{F}_{\text{init}}^{\text{iso}}$  : init

**Internal state:**

- $pk \in \{0, 1\}^*$ ,  $pk = \varepsilon$ ,  $pid_{\text{CA}} \in \{0, 1\}^*$ ,  $pid_{\text{CA}} = \varepsilon$  {The public key of  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})$  and the used CA
- $(G, n, g) \in (\{0, 1\}^* \cup \{\perp\})^3 = (\perp, \perp, \perp)$  {Global group parameters.
- $state \in \{\perp, \text{registered}, \text{started}, \text{inSession}, \text{finished}\} = \perp$  {Current state in key exchange.
- $initiator : (\{0, 1\}^*)^3$  {The initiator of the key exchange
- $responder : (\{0, 1\}^*)^3$  {The responder of the key exchange
- $caller \in (\{0, 1\}^*)^3 \cup \{\perp\} = \perp$  {Stores the initial caller of this entity/instance.
- $sessionKey \in \{0, 1\}^* \cup \{\perp\} = \perp$  {Stores the session key.
- $e_{\text{resp}} \in \mathbb{Z}_n \cup \{\perp\} = \perp$  {Secret DH exponent of responder.
- $h_{\text{init}} \in G \cup \{\perp\} = \perp$  {Public DH key share of initiator.

**CheckID( $pid, sid, role$ ):**  
 Accept all entities of form  $(pid_r, (sid', pid_i, pid_r), responder)$ .

**Corruption behavior:**

- **DetermineCorrStatus( $pid, sid, role$ ):** {Consider entity corrupted if one of the signature keys or the verification subroutine is corrupted.  
 $out \leftarrow \text{corr}(pid, (pid, sid), \mathcal{F}_{\text{sig}} : \text{signer}) \vee \text{corr}(pid, (pid(\text{initiator}), sid), \mathcal{F}_{\text{sig}} : \text{verifier})$   
 $out \leftarrow out \vee \text{corr}(pid, sid, role), \varepsilon, \mathcal{F}_{\text{PKI}}^{\text{acc}} : \text{client}$   
 return out
- **AllowAdvMessage( $pid, sid, role, pid_{\text{receiver}}, sid_{\text{receiver}}, role_{\text{receiver}}, m$ ):** {cf. explanation in Figure 36.  
 If  $role_{\text{receiver}} = \mathcal{F}_{\text{sig}} : \text{signer}$ , return false.  
 Otherwise output true iff  $pid = pid_{\text{receiver}}$ .

**Initialization:**

Parse  $sid_{\text{cur}}$  as  $(sid', pid_i, pid_r)$ .  
 $initiator \leftarrow (pid_i, (sid', pid_i, pid_r), initiator)$ . {Store initiator and responder.}  
 $responder \leftarrow (pid_r, (sid', pid_i, pid_r), responder)$ .  
 $caller \leftarrow (pid_{\text{call}}, sid_{\text{call}}, role_{\text{call}})$ .  
**send** GetParameters **to**  $(pid_{\text{cur}}, sid_{\text{cur}}, \text{setup})$ ; {Get DH parameters}  
**wait for** (GetParameters,  $(G, n, g)$ ).  
 $(G, n, g) \leftarrow (G, n, g)$ .  
**if** ITM was not activated via **init** macro, resp. a **InitEntity** message:  
**send** initPartner **to**  $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{init}}^{\text{iso}} : \text{init})$   
**wait for** initPartner

**Main:**

**rcv** RegisterPubKey **from** I/O **s.t.**  $state[(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})] = \perp$ : {Register responder's signature and certificate}  
 $state[(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})] \leftarrow \text{registered}$ .  
**send** InitSign **to**  $(pid_{\text{cur}}, (pid_{\text{cur}}, sid_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{signer})$ ;  
**wait for** (InitSign, success,  $pk$ );  $pk \leftarrow pk$ ;  
**send responsively** (Register,  $pk$ ) **to** NET;  
**wait for** (Register,  $pk, pid_{\text{CA}}$ ) **s.t.**  $pid_{\text{CA}} \in \{0, 1\}^*$ ;  
 $pid_{\text{CA}} \leftarrow pid_{\text{CA}}$ ;  
**send** (Register,  $pk, pid_{\text{CA}}$ ) **to**  $((pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}}), \varepsilon, \mathcal{F}_{\text{PKI}}^{\text{acc}} : \text{client})$ .<sup>a</sup> {Register identity at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ }

**rcv** InitKE **from** I/O **s.t.**  $state = \text{registered}$ : {Start KE.}  
 $state \leftarrow \text{started}$ .  
**send** (InitKE,  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})$ ) **to** NET. {Notify network that the key exchange has started and the responder is ready to receive the first message.}

**rcv** (Receive, initiator,  $h_{\text{init}}$ ) **from** NET **s.t.**  $state = \text{started}$ : {Receive first message, send second message.}  
 $h_{\text{init}} \leftarrow h_{\text{init}}$   
 Choose  $e_{\text{resp}} \leftarrow \mathbb{Z}_n$  uniformly at random, compute  $h_{\text{resp}} = g^{e_{\text{resp}}}$ .  
**send** (Sign,  $(h_{\text{init}}, g^{e_{\text{resp}}}, initiator)$ ) **to**  $(pid_{\text{cur}}, (pid_{\text{cur}}, sid_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{signer})$ ;  
**wait for** (Signature,  $\sigma$ ).  
 $state \leftarrow \text{inSession}$   
**send** (Send,  $(pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}}), g^{e_{\text{resp}}}, \sigma$ ) **to** NET.

**rcv** (Receive,  $\sigma$ ) **from** NET **s.t.**  $state = \text{inSession}$ : {Receive third message, output key.}  
**send responsively** (RetrieveCA, initiator) **to** NET {Query  $\mathcal{A}$  for initiator's CA;}  
**wait for** (RetrieveCA, initiator,  $pid_{\text{CA}}$ );  
**send** (Retrieve, initiator,  $pid_{\text{CA}}$ ) **to**  $((pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}}), \varepsilon, \mathcal{F}_{\text{PKI}}^{\text{acc}} : \text{client})$ ;  
**wait for** (Retrieve, initiator,  $pid_{\text{CA}}, pk$ ). {Get public verification key of intended partner.}  
**if**  $pk = \perp$ :  
 abort.  
**send** (Verify,  $(g^{e_{\text{resp}}}, h_{\text{init}}, (pid_{\text{cur}}, sid_{\text{cur}}, role_{\text{cur}})), \sigma, pk$ ) **to**  $(pid_{\text{cur}}, (initiator, sid_{\text{cur}}), \mathcal{F}_{\text{sig}} : \text{verifier})$ ;  
**wait for** (VerResult,  $b$ ).  
**if**  $b = \text{false}$ :  
 abort.  
 $sessionKey \leftarrow h_{\text{init}}^{e_{\text{resp}}}$ ;  $state \leftarrow \text{finished}$ .  
**send** (FinishKE, sessionKey) **to** caller.

<sup>a</sup>We consider here one instance of the PKI, namely the one with SID  $\varepsilon$ .

Fig. 38: A real key exchange protocol  $\mathcal{P}_{\text{iso}}^{\text{acc}}$  for realizing  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  (part 3).

Description of the initialization machine  $\mathcal{F}_{init}^{iso} = (\text{init})$ :

<b>Participating roles:</b> $\{\text{init}\}$ <b>Corruption model:</b> <i>incorruptible</i>	
Description of $M_{init}$ :	
<b>Implemented role(s):</b> $\{\text{init}\}$ <b>Subroutines:</b> $\mathcal{P}_{iso}^{acc}$ : initiator, $\mathcal{P}_{iso}^{acc}$ : responder <b>Internal state:</b> <ul style="list-style-type: none"> <li>- <math>\text{init} : (\{0, 1\}^*)^3 \rightarrow \{\text{true}, \text{false}\}</math> <span style="float: right;">{Initialization status of the machine, initially false}</span></li> <li>- <math>\text{initiator} : (\{0, 1\}^*)^3</math> <span style="float: right;">{The initiator of the key exchange}</span></li> <li>- <math>\text{responder} : (\{0, 1\}^*)^3</math> <span style="float: right;">{The responder of the key exchange}</span></li> <li>- <math>\text{judge}_{initiator} : (\{0, 1\}^*)^3</math> <span style="float: right;">{The initiator's local judge}</span></li> <li>- <math>\text{judge}_{responder} : (\{0, 1\}^*)^3</math> <span style="float: right;">{The responder's local judge}</span></li> <li>- <math>\text{caller} \in (\{0, 1\}^*)^3 \cup \{\perp\} = \perp</math> <span style="float: right;">{Stores the callers.}</span></li> </ul>	
<b>CheckID(<math>pid, sid, role</math>):</b> Accept all messages for the same $sid$ .	
<b>Initialization:</b> <ul style="list-style-type: none"> <li>Parse <math>\text{sid}_{cur}</math> as <math>(sid', pid_i, pid_r)</math>.</li> <li><math>\text{initiator} \leftarrow (pid_i, (sid', pid_i, pid_r), \text{initiator})</math>. <span style="float: right;">{Store initiator and responder}</span></li> <li><math>\text{responder} \leftarrow (pid_r, (sid', pid_i, pid_r), \text{responder})</math>.</li> <li><math>\text{judge}_{initiator} \leftarrow ((\text{local}, pid_i, \text{initiator}), (sid', pid_i, pid_r), \text{judge})</math>.</li> <li><math>\text{judge}_{responder} \leftarrow ((\text{local}, pid_r, \text{responder}), (sid', pid_i, pid_r), \text{judge})</math>. <span style="float: right;">{Also store their judges entities}</span></li> <li><math>\text{caller} \leftarrow (pid_{call}, sid_{call}, role_{call})</math></li> </ul>	
<b>MessagePreprocessing:</b> <ul style="list-style-type: none"> <li><b>if</b> <math>(pid_{call}, sid_{call}, role_{call}) = \text{initiator}</math>:</li> <li style="padding-left: 20px;"><math>\text{init}[\text{initiator}] \leftarrow \text{true}</math></li> <li><b>else if</b> <math>(pid_{call}, sid_{call}, role_{call}) = \text{responder}</math>:</li> <li style="padding-left: 20px;"><math>\text{init}[\text{responder}] \leftarrow \text{true}</math></li> <li><b>else if</b> <math>(pid_{call}, sid_{call}, role_{call}) = \text{judge}_{initiator}</math>:</li> <li style="padding-left: 20px;"><math>\text{init}[\text{judge}_{initiator}] \leftarrow \text{true}</math></li> <li><b>else if</b> <math>(pid_{call}, sid_{call}, role_{call}) = \text{judge}_{responder}</math>:</li> <li style="padding-left: 20px;"><math>\text{init}[\text{judge}_{responder}] \leftarrow \text{true}</math></li> </ul>	
<b>Main:</b> <ul style="list-style-type: none"> <li><b>recv</b> <math>\text{initPartner}</math> <b>from</b> responder <math>\vee</math> initiator: <span style="float: right;">{Request to initialize}</span></li> <li style="padding-left: 20px;"><b>if</b> <math>\text{init}[\text{initiator}] = \perp</math>:</li> <li style="padding-left: 40px;"><math>\text{init}(\text{initiator})</math> <span style="float: right;">{Trigger initialization of initiator}</span></li> <li style="padding-left: 20px;"><b>else if</b> <math>\text{init}[\text{responder}] = \perp</math>:</li> <li style="padding-left: 40px;"><math>\text{init}(\text{responder})</math> <span style="float: right;">{Trigger initialization of responder}</span></li> <li style="padding-left: 20px;"><b>else if</b> <math>\text{init}[\text{judge}_{initiator}] = \perp</math>:</li> <li style="padding-left: 40px;"><math>\text{init}(\text{judge}_{initiator})</math> <span style="float: right;">{Trigger initialization of judge}_{initiator}</span></li> <li style="padding-left: 20px;"><b>else if</b> <math>\text{init}[\text{judge}_{responder}] = \perp</math>:</li> <li style="padding-left: 40px;"><math>\text{init}(\text{judge}_{responder})</math> <span style="float: right;">{Trigger initialization of judge}_{responder}</span></li> <li style="padding-left: 20px;"><b>else:</b></li> <li style="padding-left: 40px;"><b>send</b> <math>\text{initPartner}</math> <b>to</b> caller <span style="float: right;">{All parties are initialized}</span></li> <li><b>recv</b> <math>\text{InitEntityDone}</math> <b>from</b> initiator <math>\vee</math> responder <math>\vee</math> <math>\text{judge}_{initiator} \vee \text{judge}_{responder}</math>:</li> <li style="padding-left: 20px;"><b>if</b> <math>\text{init}[\text{initiator}] = \perp</math>:</li> <li style="padding-left: 40px;"><math>\text{init}(\text{initiator})</math> <span style="float: right;">{Trigger initialization of initiator}</span></li> <li style="padding-left: 20px;"><b>else if</b> <math>\text{init}[\text{responder}] = \perp</math>:</li> <li style="padding-left: 40px;"><math>\text{init}(\text{responder})</math> <span style="float: right;">{Trigger initialization of responder}</span></li> <li style="padding-left: 20px;"><b>else if</b> <math>\text{init}[\text{judge}_{initiator}] = \perp</math>:</li> <li style="padding-left: 40px;"><math>\text{init}(\text{judge}_{initiator})</math> <span style="float: right;">{Trigger initialization of judge}_{initiator}</span></li> <li style="padding-left: 20px;"><b>else if</b> <math>\text{init}[\text{judge}_{responder}] = \perp</math>:</li> <li style="padding-left: 40px;"><math>\text{init}(\text{judge}_{responder})</math> <span style="float: right;">{Trigger initialization of judge}_{responder}</span></li> <li style="padding-left: 20px;"><b>else:</b></li> <li style="padding-left: 40px;"><b>send</b> <math>\text{initPartner}</math> <b>to</b> caller <span style="float: right;">{All parties are initialized}</span></li> </ul>	

Fig. 39: The initialization ITM  $\mathcal{F}_{init}^{iso}$ .

Description of  $\mathcal{P}_{\text{judge}}^{\text{iso}} = (\text{judge})$ :

**Participating roles:** {judge}  
**Corruption model:** *incorruptible*

Description of  $M_{\text{judge}}$ :

**Implemented role(s):** {judge}  
**Subroutines:**  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  : judge,  $\mathcal{P}_{\text{iso}}^{\text{acc}}$   
**Internal state:**  
- initiator :  $(\{0, 1\}^*)^3$  {The initiator of the key exchange}  
- responder :  $(\{0, 1\}^*)^3$  {The responder of the key exchange}  
**CheckID**( $pid, sid, role$ ):  
Accept all entities of form  $((\text{local}, pid_i, \text{initiator}), (sid', pid_i, pid_r), \text{judge})$  and  $((\text{local}, pid_r, \text{responder}), (sid', pid_i, pid_r), \text{judge})$ .  
**Initialization:**  
Parse  $sid_{\text{cur}}$  as  $(sid', pid_i, pid_r)$ .  
initiator  $\leftarrow (pid_i, (sid', pid_i, pid_r), \text{initiator})$ . {Store initiator and responder}  
responder  $\leftarrow (pid_r, (sid', pid_i, pid_r), \text{responder})$ .  
**if** ITM was not activated via **init** macro, resp. a **InitEntity** message:  
  **send** **initPartner** **to**  $(pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{init}}^{\text{iso}} : \text{init})$   
  **wait for** **initPartner**  
**Corruption behavior:**  
- **AllowCorruption**( $pid, sid, role$ ) :  
  Parse  $pid$  as  $(pid', role')$   
   $corr \leftarrow \text{corr}(pid', sid, role')$   
  **if**  $corr$ :  
    **return true**  
  **else:**  
    **return false**  
- **DetermineCorrStatus**( $pid, sid, role$ ) : {Consider local judge corrupted if associated client is corrupted.}  
  Parse  $pid$  as  $(pid', role')$   
   $corr \leftarrow \text{corr}(pid', sid, role')$   
  **if**  $corr$ :  
    **return true**  
  **else:**  
    **return false**  
**Main:**  
**recv** **GetVerdict** **from** I/O: {Forward local verdict from  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ }  
  **send** **InitMe** **to** initiator  
  **wait for** \_  
  **send** **InitMe** **to** responder  
  **wait for** \_  
  Parse  $pid_{\text{cur}}$  as  $(pid', role')$ .  
  **send** **GetVerdict** **to**  $((\text{initiator}, \text{client}), \varepsilon, \mathcal{F}_{\text{PKI}}^{\text{acc}} : \text{judge})$  {Query for the initiator's status at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ }  
  **wait for** (GetVerdict,  $v_1$ )  
  **send** **GetVerdict** **to**  $((\text{responder}, \text{client}), \varepsilon, \mathcal{F}_{\text{PKI}}^{\text{acc}} : \text{judge})$  {Query for the responder's status at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ }  
  **wait for** (GetVerdict,  $v_2$ )  
  **if**  $v_1 \neq \varepsilon \wedge role' = \text{responder}$ :  
     $v_1.\text{add}(\text{dis}(\text{initiator}))$  {The judge cannot be shure whether the intended partners verdict is trustful}  
  **if**  $v_2 \neq \varepsilon \wedge role' = \text{initiator}$ :  
     $v_2.\text{add}(\text{dis}(\text{responder}))$  {The judge cannot be shure whether the intended partners verdict is trustful}  
  Let *verdict* be  
  1.  $v_1$ , if  $role' = \text{responder}$   
  2.  $v_2$ , if  $role' = \text{initiator}$   
  **reply** (GetVerdict, *verdict*)  
**recv** (GetJudicialReport, *msg*) **from** I/O: {Forward judicial report from lower level.}  
  **reply** (GetJudicialReport,  $\varepsilon$ )

Fig. 40: The judging functionality  $\mathcal{P}_{\text{judge}}^{\text{iso}}$  for  $\mathcal{P}_{\text{iso}}^{\text{acc}}$ .

Description of  $\mathcal{P}_{sv}^{iso} = (\text{supervisor})$ :

**Participating roles:** {supervisor}  
**Corruption model:** *incorruptible*

Description of  $M_{\text{supervisor}}$ :

**Implemented role(s):** {supervisor}  
**Subroutines:**  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$  : supervisor  
**CheckID**( $pid, sid, role$ ):  
 Accept all messages with the same  $sid$ .  
**Main:**

```

recv (IsAssumptionBroken?,  $msg, id$ ):
  reply (IsAssumptionBroken?, false)
recv (corruptInt?,  $pid, sid, role$ ) s.t.  $role \notin \{\text{initiator, responder, judge, supervisor}\}$ :
  if  $role = \text{client}$ :
     $corrRes \leftarrow \text{corr}(pid, sid, \mathcal{F}_{\text{PKI}}^{\text{acc}} : \text{client})$  {Request corruption status at  $\mathcal{F}_{\text{PKI}}^{\text{acc}}$ }
    reply (corruptInt,  $corrRes$ )
  else:
    send (corruptInt?,  $pid, sid, role$ ) to ( $pid_{\text{cur}}, sid_{\text{cur}}, \mathcal{F}_{\text{PKI}}^{\text{acc}} : \text{supervisor}$ )
    wait for (corruptInt?,  $corrRes$ )
    reply (corruptInt,  $corrRes$ )
  
```

Fig. 41: The supervisor  $\mathcal{P}_{sv}^{iso}$  for  $\mathcal{P}_{iso}^{\text{acc}}$ .

the name  $pid$  of a malicious party that caused the protocol to abort (cf., [14, 34, 55]). Depending on this choice,  $\mathcal{F}_{\text{MPC}}^{\text{id-ab}}$  either returns the outputs of  $f$  or a special message (abort,  $pid$ ) to the honest parties.

**Capturing identifiable abort using AUC.** We can easily capture the same properties as  $\mathcal{F}_{\text{MPC}}^{\text{id-ab}}$  by applying AUC to the basic functionality  $\mathcal{F}_{\text{MPC}}$ . We set  $\text{Sec}^{\text{acc}} = \{\text{correctness}\}$  and, as part of the transformation  $\mathcal{T}_2$ , redefine the behavior of  $\mathcal{F}_{\text{MPC}}$  to output abort, instead of the actual function output, iff *correctness* is broken for that party. To also capture the same relationships and exact accountability level required by identifiable abort, we instantiate  $\mathcal{F}_{\text{judgeParams}}$  to impose the following additional requirements whenever the simulator tries to break *correctness*: (i) no honest party has already obtained an output, (ii) if *correctness* is broken for one honest party, it must be broken for all others at the same time, and (iii) all verdicts for local judges of honest parties are identical and of the form  $\text{dis}(A)$  for some party  $A$ .

The resulting functionality  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$  models the exact same properties as  $\mathcal{F}_{\text{MPC}}^{\text{id-ab}}$ , with the only syntactical difference being that  $\mathcal{F}_{\text{MPC}}^{\text{id-ab}}$  includes the verdict as part of the protocol output while in  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$  the verdict is obtained separately from a judge.

#### H. Further Details on MPC Accountability Properties via AUC

In this section, we detail our presentation from Appendix G and show that AUC can be seen as a generalization of the existing accountability related MPC/UC literature. Therefore, we firstly illustrate in detail how the common MPC property of (internal) identifiable abort translates to AUC. We use a standard MPC functionality  $\mathcal{F}_{\text{MPC}}$  without in-build identifiable abort and depict how identifiable abort is typically captured and how one can capture it with AUC. We then briefly discuss that this approach also works for other common properties, such as *verifiability*.

**The ideal MPC functionality  $\mathcal{F}_{\text{MPC}}$ .** For common ideal MPC functionalities (cf., e.g., [8, 21, 24]) one assumes that there are  $m$  parties involved in the computation, each party  $p_i$  contributes some secret input  $x_i$ , and the parties want to distributedly compute the output of a function  $f$ . The ideal MPC functionality  $\mathcal{F}_{\text{MPC}}$  models a trusted third party which executes the computation on behalf of the participating users. Within  $\mathcal{F}_{\text{MPC}}$ , the adversary  $\mathcal{A}$  is typically allowed to (statically) corrupt parties. The execution of the computation via  $\mathcal{F}_{\text{MPC}}$  is separated in phases. All parties are expected/forced to be in the same phase of the protocol: (i) *Input*: Each honest party inputs  $x_i$  to  $\mathcal{F}_{\text{MPC}}$ . Parties controlled by the adversary could input tainted data  $x'_i$ . (ii) *Computation*:  $\mathcal{F}_{\text{MPC}}$  evaluates  $f$  given the inputs from the parties. (iii) *Finalization*:  $\mathcal{F}_{\text{MPC}}$  provides the output of the computation to the parties.  $\mathcal{F}_{\text{MPC}}$  typically guarantees the correct computation of the function  $f$ . We present a formal specification of  $\mathcal{F}_{\text{MPC}}$  in Figure 42 where the **highlighted** code denotes the differences between  $\mathcal{F}_{\text{MPC}}$  and  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$ .

**Identifiable abort in UC literature.** With identifiable abort (cf., e.g., [14, 34, 55]), one typically protects *correctness* of a MPC protocol and in particular the correctness of the protocol's output. This models that



parties abort the protocol and blame a misbehaving party if, e.g., (i) a party identifies a misbehaving party which tries to distort the computation or (ii) a (corrupted) party aborts the protocol. We focus on internal/local identifiable abort, i.e., we interpret identifiable abort as a *local* accountability property. This is often the case in MPC protocols as verdicts often rely on information that is only accessible to protocol participants and thus do not convince outsiders of the protocol.

An adapted ideal MPC functionality  $\mathcal{F}'_{\text{MPC}}$  which covers identifiable abort typically differs from  $\mathcal{F}_{\text{MPC}}$  in the following details: (i) Honest parties do not receive the output of the computation in case of an abort. (ii) To abort, a corrupted party  $p$  inputs the string `abort` to  $\mathcal{F}'_{\text{MPC}}$ , this informs  $\mathcal{F}'_{\text{MPC}}$  that  $\mathcal{A}$ , resp. party  $p$ , aborted the execution of the protocol. (iii)  $\mathcal{F}'_{\text{MPC}}$  then typically informs honest parties that  $p$  aborted the protocol. This “individual verdict” on  $p$  allows the honest parties in the MPC protocol to exclude misbehaving parties from the next computation attempt.

**Identifiable Abort via AUC.** In contrast to Appendix G, we do not consider the traditional interpretation of MPC protocols which assumes that all parties do all phases of the MPC protocol synchronously – we call this case synchronous MPC. We focus here on what we call the “asynchronous” case, where parties are not forced to be in the same phase of the protocol. However, we provide a full specification of the synchronous variant of  $\mathcal{F}_{\text{MPC}}$  which we call  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$  in Figure 42 and 43 which include the properties of identifiable abort, called non-aborting, verifiability, called correctness, and privacy. In what follows, we focus on the asynchronous variant of  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$  and include identifiable abort to the functionality. We call this functionality  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$  (cf. Figure 44 and 45). In  $\mathcal{F}_{\text{MPC}}$ , the following major adaptations are necessary to derive  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$  (besides mapping messages): (i) Apply the AUC transformation  $\mathcal{T}_1(\cdot)$  to  $\mathcal{F}_{\text{MPC}}$ , (ii) add one local judge per participant/node to  $\mathcal{F}_{\text{MPC}}$ , i.e.,  $\text{pids}_{\text{judge}} = \{\text{local}\} \times \{0, 1\}^*$  and every set member is meant to be of form  $(\text{local}, \text{pid}, \text{role})$  ( $\text{ids}_{\text{assumption}} = \emptyset$  as we do not consider assumption-based properties here), (iii) add the property correctness to  $\text{Sec}^{\text{acc}}$  in  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$  and  $\mathcal{F}_{\text{judgeParams}}^{\text{async-MPC}}$ <sup>18</sup> (iv) Allow  $\mathcal{A}$  to abort the computation using the `BreakAccProp` command.  $\mathcal{A}$ 's request needs to mark correctness broken for honest parties and needs to include a (fair) individual verdict for each of the parties. (v)  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$  outputs `abort` to a party if correctness is broken for this party. Parties can query their judge to retrieve the verdict which includes who aborted the protocol. (vi) We adapt  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$ , resp.  $\mathcal{F}_{\text{judgeParams}}^{\text{async-MPC}}$ , such that they ensure that (a) aborting is not possible after the first honest party received output from  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$  and vice versa, (b) that no party receives an output if one party received `abort`. (vii)  $\mathcal{F}_{\text{judgeParams}}^{\text{async-MPC}}$  enforces that (local) verdicts are available via the internal local judges, and that (viii) all internal judges output the same verdict.

Clearly, the AUC formulation of identifiable abort can be mapped to the traditional definition of identifiable abort: As soon as  $\mathcal{A}$  signals that a corrupted party will abort the protocol, all honest parties will be informed that the protocol aborted and which party caused the abort. This statement holds true for  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$  as well as for  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$ . Of course one can enhance our approach above analogously to achieve *public* identifiable abort (cf., e.g., [30, 60]).

**Verifiability.** Verifiability in MPC is also a correctness property [35, 84, 86]: in the case of public or universal verifiability, every party including protocol outsiders should be able to verify that the MPC's output was computed correctly or honestly. Due to the privacy guarantees of MPC protocols, public verifiability typically only provides a weak level of accountability as one cannot identify misbehaving parties individually.

The common formalization of verifiability is analogous to the formalization of identifiable abort. For public verifiability (cf. [6]), one sometimes expects some *certificate of misbehavior* created by a judge which allows outsiders to verify that the blamed party indeed misbehaved.

To capture verifiability with AUC, one uses the same techniques as for identifiable abort. The major difference between both approaches are (i) verifiability is a public property, thus restrictions are incorporated into a public judge instead of local judges, (ii)  $\mathcal{F}_{\text{judgeParams}}^{\text{async-MPC}}$  does not impose restrictions on the verdict, and (iii) one may use the judicial report defined at  $\mathcal{F}_{\text{judgeParams}}^{\text{async-MPC}}$  to output a certificate of misbehavior.

For other MPC accountability properties, such as *auditability* [13], *openability* [34] or *privacy* [8],  $\mathcal{A}$  also signals the ideal functionality that a property is broken and the ideal functionality then acts appropriately which typically signals that a party or a group of parties misbehaved. The approach above works analogously for other accountability properties from MPC. Thus, AUC can be seen as generalization of existing work regarding accountability in MPC/UC literature.

<sup>18</sup>We note that `correctness` captures identifiable abort in  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$  while it captures verifiability in  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$ . For the sake of presentation for  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$  we want to keep naming in  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$  as close to literature as possible.

Description of the protocol  $\mathcal{F}_{\text{MPC}}^{\text{acc}} = (\text{node}, \text{judge}, \text{supervisor})$

**Participating roles:** {node, judge, supervisor}  
**Corruption model:** static  
**Protocol parameters:**  
-  $n \in \mathbb{N}$  {Input length for the MPC function}  
-  $f : (\{0, 1\}^*)^{|P|} \rightarrow (\{0, 1\}^*)^{|P|}$ , **s.t.**  $f = (f_1, \dots, f_{|P|})$  {the MPC function}  
-  $\text{Sec}^{\text{acc}} \subset \{0, 1\}^*$  {Accountability properties, expected here: {privacy, correctness, non-aborting}}  
-  $\text{Sec}^{\text{assumption}} \subset \{0, 1\}^*$  {Assumption-based security properties, expected here:  $\emptyset$ }  
-  $\text{pids}_{\text{judge}} \subset \{0, 1\}^*$  {set of judge entities/(P)IDs in the protocol (which are often directly related to some protocol participants)}  
-  $\text{ids}_{\text{assumption}} \subset \{0, 1\}^*$  {set of entities/IDs where properties are ensured via assumptions}

Description of  $\mathcal{M}_{\text{node}}$ :

**Implemented role(s):** {node}  
**Subroutines:**  $\mathcal{F}_{\text{judgeParams}}^{\text{MPC}} : \text{judgeParams}$   
**Internal state:**  
-  $P \subset \{0, 1\}^*$ ,  $P \neq \emptyset$ ,  $|P| < \infty$  {Parties involved in the protocol}  
- **inputs** :  $P \rightarrow \{0, 1\}^n$  {The inputs for the computation, initially always set to  $\perp$ }  
- **outputs** :  $P \rightarrow \{0, 1\}^*$  {The outputs of the computation, initially always set to  $\perp$ }  
- **state**  $\in \{\text{uninitialized}, \text{computation}, \text{finalized}\}$ , **state** = uninitialized {State of the MPC computation}  
- **corruptedIntParties**  $\in \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \setminus (\text{Roles}_{\mathcal{F}}^a \cup \{\text{judge}, \text{supervisor}\})$ , initially  $\emptyset$  {The set of corrupted internal parties (pid, sid, role)}  
- **brokenAssumptions** :  $\text{Sec}^{\text{assumption}} \times \text{ids}_{\text{assumption}} \rightarrow \{\text{true}, \text{false}\}$  {Stores broken security assumptions per id, initially false  $\forall$  entries}  
- **brokenProps** :  $(\text{Sec}^{\text{assumption}} \cup \text{Sec}^{\text{acc}}) \times (\text{pids}_{\text{judge}} \cup \text{ids}_{\text{assumption}}) \rightarrow \{\text{true}, \text{false}\}$  {Stores broken security properties per judge/id, initially false  $\forall$  entries}  
- **verdicts** :  $\text{pids}_{\text{judge}} \rightarrow \{0, 1\}^*$  {Verdicts per  $p \in \text{pids}_{\text{judge}}$ , initially  $\varepsilon$ }  
**CheckID**(pid, sid, role):  
Accept all messages if pid  $\in P$  and with same sid. {A single instance manages all parties in a single session.}  
**Corruption behavior:**  
- **AllowCorruption**(pid, sid, role):  
Do not allow corruption of (pid, sid, supervisor).  
**if** role = judge:  
**send** (Corrupt, (pid, sid, judge), internalState)  
**to** (pid, sid,  $\mathcal{F}_{\text{judgeParams}}$  : judgeParams) { $\mathcal{F}_{\text{judgeParams}}$  decides whether judges can be corrupted}  
**wait for** b; **return** b  
- **DetermineCorrStatus**(pid, sid, role):  
**if** role = judge:  
**send** (CorruptionStatus?(pid, sid, judge), internalState)  
**to** (pid, sid,  $\mathcal{F}_{\text{judgeParams}}$  : judgeParams) { $\mathcal{F}_{\text{judgeParams}}$  may determine a judge's corruption status}  
**wait for** b; **return** b  
- **AllowAdvMessage**(pid, sid, role, pid<sub>receiver</sub>, sid<sub>receiver</sub>, role<sub>receiver</sub>, m) { $\mathcal{A}$  is not allowed to invoke  $\mathcal{F}_{\text{judgeParams}}^{\text{acc-cp}}$ }  
Do not allow sending messages to  $\mathcal{F}_{\text{judgeParams}}$ .  
**Initialization:**  
**send responsively** InitMe to NET { $\mathcal{A}$  determines the involved parties}  
**wait for** (Init, P) **s.t.**  $P \subset \{0, 1\}^*$   
 $P \leftarrow P$   
**MessagePreprocessing:**  
**if** state = uninitialized  $\wedge$  number of inputs that are  $\neq \perp$  is  $|P| \wedge$  brokenProps is false for all entries:  
state  $\leftarrow$  computation {Update state/computation phase}  
**Main:**  
**recv** (Compute, x) **from** I/O **s.t.** pid<sub>cur</sub>  $\in P \wedge$  state  $\in \{\text{uninitialized}\}$ : {Input from uncorrupted users}  
**inputs**[pid<sub>cur</sub>]  $\leftarrow$  x {Record user input}  
**recv** (Compute, pid, x) **from** NET **s.t.** pid in CorruptionSet  $\wedge$  pid  $\in P$ , state  $\in \{\text{uninitialized}, \text{computation}\}$ : { $\mathcal{A}$  input for corrupted parties}  
**inputs**[pid<sub>cur</sub>]  $\leftarrow$  x {Record user input}  
**recv** Finalize **from** NET:  
**if** state = computation  $\wedge \forall$  pid  $\in P$  : **inputs**[pid]  $\neq \perp \wedge$  state = uninitialized: {That is, there was no cheating/corrupting}  
Let  $w_i$  **s.t.** (pid<sub>i</sub>, w<sub>i</sub>)  $\in$  **inputs**,  $i \in \{1, \dots, |P|\}$   
 $(y_1, \dots, y_{|P|}) \leftarrow f(w_1, \dots, w_{|P|})$  {If there was no "interrupt" from  $\mathcal{A}$ , f is computed as expected}  
**for**  $i = 1$  to  $|P|$  **do**:  
**outputs**[pid<sub>i</sub>]  $\leftarrow$  y<sub>i</sub>  
state  $\leftarrow$  finalized  
Let **output** contain every **outputs**[pid] where pid is in CorruptionSet.  
**reply** (Finalize, output) {Send output of corrupted parties to  $\mathcal{A}$ }  
**recv** GetResult **from** I/O **s.t.** pid<sub>cur</sub>  $\in P$ :  
**if** brokenProps[non-aborting, (local, pid<sub>cur</sub>, role<sub>cur</sub>):  
**reply** (GetResult, abort, verdicts) {We are in the identifiable abort case}  
**else if** brokenProps[privacy, (local, pid<sub>cur</sub>, role<sub>cur</sub>):  
**reply** (GetResult, corrupted, verdicts) {Output abort information and verdict}  
{We are in the case that  $\mathcal{A}$  broke privacy} {Output privacy break and verdict}  
**else if** brokenProps[correctness, (local, pid<sub>cur</sub>, role<sub>cur</sub>):  
**reply** (GetResult, correctness, verdicts) {We are in the case that  $\mathcal{A}$  broke correctness of the calculation}  
**reply** (GetResult, outputs[pid<sub>cur</sub>]) {Output correctness break and verdict}  
**reply** (GetResult, outputs[pid<sub>cur</sub>])

<sup>a</sup>Here:  $\text{Roles}_{\mathcal{F}} := \{\text{node}\}$ .

Fig. 42: The MPC functionality  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$ .

Description of  $\mathcal{F}_{\text{judgeParams}}^{\text{MPC}} = (\text{judgeParams})$ :

**Participating roles:**  $\{\text{judgeParams}\}$   
**Corruption model:** *incorruptible*

Description of  $M_{\text{judgeParams}}$ :

**Implemented role(s):**  $\{\text{judgeParams}\}$   
**CheckID**( $pid, sid, role$ ):  
 Accept all messages with the same  $sid$ .  
**Main:**

```

recv (BreakAccProp, verdict, toBreak, internalState) from I/O:
  if toBreak = {non-aborting}  $\times$  ({local}  $\times$   $\{0, 1\}^{2*}$ )  $\wedge$  state  $\in$  {uninitialized, compute}:    {A aborts computation}
    reply (BreakAccProp, true,  $\epsilon$ )
  if toBreak = {correctness}  $\times$  ({local}  $\times$   $\{0, 1\}^{2*}$ )  $\wedge$  state = compute:    {This models a cheat attempt}
    reply (BreakAccProp, true,  $\epsilon$ )
  if toBreak = {privacy}  $\times$  ({local}  $\times$   $\{0, 1\}^{2*}$ )  $\wedge$  state = compute:    {This models an attempt to break input privacy}
    Let  $\bar{\omega}$  be the vector of all values from inputs.
    if eval(verdict) = true:    {Need to check whether verdict evaluates to true as leakage forwarded to A anyways}
      reply (BreakAccProp, true,  $\bar{\omega}$ )
    else:
      reply (BreakAccProp, false,  $\epsilon$ )
  else:
    reply (BreakAccProp, false,  $\epsilon$ )

recv (GetJudicialReport, msg, internalState) from I/O:    {Generate judicial report}
  reply (GetJudicialReport,  $\epsilon$ )    {Return state variable as report}

recv (Corrupt, (id, sid, judge), internalState) from I/O:
  if id = (local, pid, role)  $\wedge$  (pid, sid, role)  $\in$  CorruptionSet:    { $\mathcal{F}_{\text{judgeParams}}^{\text{KE}}$  allows to corrupt a local judge iff the}
    reply true    {accompanied initiator or responder is corrupted}
  else:
    reply false

recv (CorruptionStatus?, (id, sid, judge), internalState) from I/O:    { $\mathcal{F}_{\text{judgeParams}}^{\text{KE}}$  interprets a local judge as corrupted iff}
  if (id, sid, judge)  $\in$  CorruptionSet:    {the accompanied initiator or responder is corrupted}
    reply true
  else if id = (local, pid, role)  $\wedge$  (pid, sid, role)  $\in$  CorruptionSet :
    reply true
  else:
    reply false
  
```

Fig. 43: The judge parameter functionality  $\mathcal{F}_{\text{judgeParams}}^{\text{MPC}}$  for  $\mathcal{F}_{\text{MPC}}^{\text{acc}}$ .

Description of the protocol  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}} = (\text{node}, \text{judge}, \text{supervisor})$

**Participating roles:** {node, judge, supervisor}  
**Corruption model:** static  
**Protocol parameters:**  
-  $n \in \mathbb{N}, f : (\{0, 1\}^*)^{|P|} \rightarrow (\{0, 1\}^*)^{|P|}$ , **s.t.**  $f = (f_1, \dots, f_{|P|})$  {Input length for the MPC function and the function itself}  
-  $\text{Sec}^{\text{acc}} \subset \{0, 1\}^*$  {Accountability properties}  
-  $\text{Sec}^{\text{assumption}} \subset \{0, 1\}^*$  {Assumption-based security properties}  
-  $\text{pids}_{\text{judge}} \subset \{0, 1\}^*$  {set of judge entities/(P)IDs in the protocol (which are often directly related to some protocol participants)}  
-  $\text{ids}_{\text{assumption}} \subset \{0, 1\}^*$  {set of entities/IDs where properties are ensured via assumptions}

Description of  $M_{\text{node}}$ :

**Implemented role(s):** {node}  
**Subroutines:**  $\mathcal{F}_{\text{judgeParams}}^{\text{MPC}}$  : judgeParams  
**Internal state:**  
-  $P \subset \{0, 1\}^*, P \neq \emptyset, |P| < \infty$  {Parties involved in the protocol}  
-  $\text{inputs} : P \rightarrow \{0, 1\}^n$  {The inputs for the computation, initially always set to  $\perp$ }  
-  $\text{outputs} : P \rightarrow \{0, 1\}^*$  {The outputs of the computation, initially always set to  $\perp$ }  
-  $\text{state} : P \rightarrow \{\text{uninitialized}, \text{computation}, \text{finalized}, \text{finishedOutput}\}$  {State of the MPC computation per party, initially uninitialized}  
-  $\text{corruptedIntParties} \in \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^* \setminus (\text{Roles}_{\mathcal{F}}^a \cup \{\text{judge}, \text{supervisor}\})$ , initially  $\emptyset$  {The set of corrupted internal parties (pid, sid, role)}  
-  $\text{brokenAssumptions} : \text{Sec}^{\text{assumption}} \times \text{ids}_{\text{assumption}} \rightarrow \{\text{true}, \text{false}\}$  {Stores broken security assumptions per id, initially false  $\forall$  entries}  
-  $\text{brokenProps} : (\text{Sec}^{\text{assumption}} \cup \text{Sec}^{\text{acc}}) \times (\text{pids}_{\text{judge}} \cup \text{ids}_{\text{assumption}}) \rightarrow \{\text{true}, \text{false}\}$  {Stores broken security properties per judge/id, initially false  $\forall$  entries}  
-  $\text{verdicts} : \text{pids}_{\text{judge}} \rightarrow \{0, 1\}^*$  {Verdicts per  $p \in \text{pids}_{\text{judge}}$ , initially  $\varepsilon$ }  
**CheckID**(pid, sid, role):  
Accept all messages if pid  $\in P$  and with same sid. {A single instance manages all parties in a single session.}  
**Additional Corruption behavior:**  
- **AllowCorruption**(pid, sid, role):  
Do not allow corruption of (pid, sid, supervisor).  
**if** role = judge:  
**send** (Corrupt, (pid, sid, judge), internalState) **to** (pid, sid,  $\mathcal{F}_{\text{judgeParams}}$  : judgeParams) { $\mathcal{F}_{\text{judgeParams}}$  decides whether judges can be corrupted}  
**wait for b; return b**  
- **DetermineCorrStatus**(pid, sid, role):  
**if** role = judge:  
**send** (CorruptionStatus?, (pid, sid, judge), internalState) **to** (pid, sid,  $\mathcal{F}_{\text{judgeParams}}$  : judgeParams) { $\mathcal{F}_{\text{judgeParams}}$  may determine a judge's corruption status}  
**wait for b; return b**  
- **AllowAdvMessage**(pid, sid, role, pid<sub>receiver</sub>, sid<sub>receiver</sub>, role<sub>receiver</sub>, m)  
Do not allow sending messages to  $\mathcal{F}_{\text{judgeParams}}$ . { $\mathcal{A}$  is not allowed to invoke  $\mathcal{F}_{\text{judgeParams}}$  in the name of corrupted parties.}  
**Initialization:**  
**send responsively** InitMe **to** NET { $\mathcal{A}$  determines the involved parties}  
**wait for** (Init, P) **s.t.**  $P \subset \{0, 1\}^*$   
 $P \leftarrow P$   
**Main:**  
**rcv** (Compute, x) **from** I/O **s.t.** pid<sub>cur</sub>  $\in P \wedge \text{state} \in \{\text{uninitialized}\}$ : {Input from uncorrupted users}  
state[pid<sub>cur</sub>]  $\leftarrow$  computation; inputs[pid<sub>cur</sub>]  $\leftarrow$  x {Update phase of pid<sub>cur</sub> and record user input}  
**rcv** (Compute, pid, x) **from** NET **s.t.** pid in CorruptionSet  $\wedge$  pid  $\in P$ , state  $\in \{\text{uninitialized}, \text{computation}\}$ : { $\mathcal{A}$  input for corrupted parties}  
state[pid<sub>cur</sub>]  $\leftarrow$  computation {Update phase of pid<sub>cur</sub>}  
inputs[pid<sub>cur</sub>]  $\leftarrow$  x {Record user input}  
**rcv** Finalize **from** NET:  
**if**  $\forall \text{pid} \in P : \text{state}[\text{pid}] = \text{computation} \wedge \forall \text{pid} \in P : \text{inputs}[\text{pid}] \neq \perp$ : {Trigger computation of f}  
Let  $w_i$  **s.t.** (pid<sub>i</sub>, w<sub>i</sub>)  $\in$  inputs,  $i \in \{1, \dots, |P|\}$   
 $(y_1, \dots, y_{|P|}) \leftarrow f(w_1, \dots, w_{|P|})$  {If there was no "interrupt" from  $\mathcal{A}$ , f is computed as expected}  
**for**  $i = 1$  to  $|P|$  **do**:  
outputs[pid<sub>i</sub>]  $\leftarrow$  y<sub>i</sub>  
**if** There is a brokenProps of type correctness true:  
output  $\leftarrow$   $\varepsilon$  {The protocol aborted, no one receives output}  
**else**:  
Let output contain every outputs[pid] where pid is in CorruptionSet.  
**reply** (Finalize, output) {Send output of corrupted parties to  $\mathcal{A}$ }  
**rcv** (Finalize, P) **from** NET **s.t.**  $P \subset P$ : { $\mathcal{A}$  may advance participants from phase computation to finalized}  
**for all**  $p \in P$  **do**:  
**if** state[p] = computation: state[p]  $\leftarrow$  finalized  
**rcv** GetResult **from** I/O **s.t.** pid<sub>cur</sub>  $\in P$ : {Access only for "registered" parties}  
**if** brokenProps[(correctness, (local, pid<sub>cur</sub>, role<sub>cur</sub>))]: {Handle identifiable abort}  
**reply** (GetResult, abort) {Output abort information}  
**else if** state[pid<sub>cur</sub>]  $\in \{\text{uninitialized}, \text{computation}\}$ : {Party did not finish computation so far}  
**reply** (GetResult,  $\perp$ )  
**else if** There exists an entry brokenProps[(correctness, (local, pid<sub>cur</sub>, role<sub>cur</sub>)) = true: {If abort is recorded at another party}  
**reply** (GetResult,  $\perp$ ) {Suppress output in case of an abort of another party}  
**else if** state[pid<sub>cur</sub>] = finalized:  
state[pid<sub>cur</sub>]  $\leftarrow$  finishedOutput  
**reply** (GetResult, outputs[pid<sub>cur</sub>]) {If  $\mathcal{A}$  did not abort, output result}

Include the missing static code from the AUC transformation  $\mathcal{T}_1(\cdot)$  here, i.e., include additional code from Figure 2 and 3 here.

<sup>a</sup>Here:  $\text{Roles}_{\mathcal{F}} := \{\text{node}\}$ .

Fig. 44: The asynchronous MPC functionality  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$ .

Description of  $\mathcal{F}_{\text{judgeParams}}^{\text{async-MPC}} = (\text{judgeParams})$ :

**Participating role(s):**  $\{\text{judgeParams}\}$   
**Corruption model:** *incorruptible*

Description of  $M_{\text{judgeParams}}$ :

**Implemented role(s):**  $\{\text{judgeParams}\}$

**CheckID**( $pid, sid, role$ ):

Accept all messages with the same  $sid$ .

**Main:**

**recv** (BreakAccProp, verdict, toBreak, internalState) **from** I/O:

Let  $ids \subset \{0, 1\}^*$  be the IDs part from *toBreak*.

**if**  $toBreak \subset \{\text{correctness} \times ids\} \wedge \nexists p \in \mathcal{P}$  ( $\mathcal{P}$  extracted from *internalState*) with  $\text{state}[p] = \text{finishedOutput} \wedge$

$[\forall id_1, id_2, \text{ such that } id_1 \neq id_2 \text{ where } \text{verdict}[id_1] \neq \varepsilon, \text{ resp. } \text{verdict}[id_2] \neq \varepsilon \wedge$

it holds true that  $\text{verdict}[id_1] = \text{verdict}[id_2] \wedge$

all verdicts in *verdict* are individual verdicts:

$\{\mathcal{A} \text{ aborts computation, only allowed if restrictions are met}$

**reply** (BreakAccProp, true,  $\varepsilon$ )

$\{\text{Accept abort}$

**else:**

**reply** (BreakAccProp, false,  $\varepsilon$ )

$\{\text{Decline abort}$

**recv** (GetJudicialReport, msg, internalState) **from** I/O:

$\{\text{Generate judicial report}$

**reply** (GetJudicialReport,  $\varepsilon$ )

$\{\text{Return state variable as report}$

**recv** (Corrupt, ( $id, sid, \text{judge}$ ), internalState) **from** I/O:

**if**  $id = (\text{local}, pid, role) \wedge (pid, sid, role) \in \text{CorruptionSet}$ :

**reply** true

**else:**

**reply** false

$\left\{ \begin{array}{l} \mathcal{F}_{\text{judgeParams}}^{\text{KE}} \text{ allows to corrupt a local judge iff the} \\ \text{accompanied initiator or responder is corrupted} \end{array} \right.$

**recv** (CorruptionStatus?, ( $id, sid, \text{judge}$ ), internalState) **from** I/O:

**if**  $(id, sid, \text{judge}) \in \text{CorruptionSet}$ :

**reply** true

**else if**  $id = (\text{local}, pid, role) \wedge (pid, sid, role) \in \text{CorruptionSet}$  :

**reply** true

**else:**

**reply** false

$\left\{ \begin{array}{l} \mathcal{F}_{\text{judgeParams}}^{\text{KE}} \text{ interprets a local judge as corrupted iff} \\ \text{the accompanied initiator or responder is corrupted} \end{array} \right.$

Fig. 45: The judge parameter functionality  $\mathcal{F}_{\text{judgeParams}}^{\text{async-MPC}}$  for  $\mathcal{F}_{\text{async-MPC}}^{\text{acc}}$ .