

# Provable Security for PKI Schemes

Hemi Leibowitz\*  
School of Computer Science, The  
College of Management Academic  
Studies  
Rishon Lezion, Israel

Amir Herzberg  
Dept. of Computer Science and  
Engineering, University of  
Connecticut  
Storrs, CT

Ewa Syta†  
Dept. of Computer Science, Trinity  
College  
Hartford, CT

## ABSTRACT

PKI provides a critical foundation to applied cryptographic protocols. However, there are no rigorous security specifications for PKI, and therefore, no PKI schemes were proven secure. This is problematic considering the extensive reliance on PKI, the multiple failures of PKI systems, and the fact that some proposed and deployed PKI schemes have complex design and advanced goals. The lack of specifications and proofs for PKI schemes means that applied cryptographic systems that use PKI are analyzed by adopting overly simplified models of the PKI, often, simply assuming secure public keys.

We present game-based security specifications for PKI schemes, and prove the security of the two most important and widely deployed schemes: PKIX and Certificate Transparency (CT), both based on version 3 of the X.509 standard, and using the (standard) CRL revocation mechanism. The proof shows a reduction from an adversary that ‘wins’ the PKI-specifications game to an adversary that ‘wins’ against the underlying signature scheme or hash function. This is the first reduction-based definition and proof of security for a realistic PKI scheme.

## 1 INTRODUCTION

*Public Key Infrastructure (PKI)* provides an essential foundation for applications that rely on public key cryptography, which is crucial to ensure security in open networks and systems. Early PKI ideas were proposed in 1978 [21], and the first version of the X.509 standard [8] was published in 1988. Since then, the deployment of PKI has been dominated by X.509, specifically, by the IETF PKIX standard, which adopts version 3 of X.509 (X.509v3) for Internet protocols, most notably, TLS/SSL [35]. *Certificate Transparency (CT)* [34, 24] is a recent, widely-deployed extension to PKIX, motivated by multiple PKI failures, mainly, rogue certificates issued by corrupt or negligent CAs. A significant number of other PKI schemes were proposed recently, with different goals and properties, and different, non-trivial designs, including [38, 10, 30, 20, 42, 3, 45, 40, 41, 28, 12, 2, 43, 22].

Considering the importance, variety and complexity of (some) PKI schemes, it is essential to ensure their security. However, this work is the first to define the security of a (non-trivial<sup>1</sup>) PKI scheme. This situation stands in sharp contrast to the accepted norms in

(applied and theoretical) modern cryptography, which require well-defined security requirements and reduction-based proofs of security. These norms began in the 1980s with the seminal papers defining secure encryption [14] and secure signature schemes [15].

The lack of rigorous specifications and analysis for PKI schemes is especially alarming, considering that PKI provides a critical infrastructure to applied cryptography, i.e., security of many applied cryptographic systems depend on the security of the underlying PKI. The extensive efforts to prove the security of cryptographic protocols may be moot when these protocols depend on an insecure PKI scheme. The concerns are even greater, considering that attacks against PKI are not only a theoretical threat, but are a major concern in practice. To address these concerns, we present the *rigorous (game-based) definitions of security requirements for PKI schemes*, and then the *first reduction-based proofs of security* for practical PKI schemes; see Table 1.

Defining and proving security for PKI schemes is challenging, especially for post-X.509 schemes, whose requirements (goals) and designs are more advanced and complex. Different PKI schemes may assume different communication, synchronization and adversary models; even the set of entities may differ, e.g., CT also introduces *loggers* and *monitors* in addition to *Certificate Authorities (CAs)*. Existing works often use informal security requirements and models for PKI schemes, often tailored to a specific construction. In fact, as we now discuss, even intuitively straightforward definitions have many subtleties, which may be overlooked or ignored.

Consider *accountability*, the most basic security property of PKI schemes. Intuitively, accountability is the ability to *identify* the CA that is *responsible* for the issuing of an unauthorized certificate  $\psi$ . However, *who is the responsible CA* for  $\psi$ ? Instinctively, we may expect this to be the CA identified as the *issuer* of  $\psi$ , denoted as  $\psi.issuer$ . That said, surely  $\psi.issuer$  should not be held accountable for  $\psi$ , if there is no guarantee that  $\psi.issuer$  indeed issued  $\psi$ . Namely, if the public verification key  $pk$  used to validate  $\psi$  is *not* a correct public-key of  $\psi.issuer$ . Even more clearly,  $\psi.issuer$  can only be considered accountable if it is a real, supposedly trustworthy entity. To illustrate, consider a scenario where a rogue CA issues a certificate  $\psi'$  which fraudulently specifies  $pk$  as a public verification key of  $\psi.issuer$ . According to the aforementioned instinct,  $\psi'$  would allow to ‘frame’  $\psi.issuer$  to be held accountable for  $\psi$ , even though  $\psi.issuer$  never issued  $\psi$ . Instead, the CA responsible (accountable) in this case for  $\psi$  should be the CA responsible for the fraudulent certificate  $\psi'$ . The definition of accountability, which we present later, supports a wide range of PKI schemes, and, in particular, the widely-used *certificate chains* mechanism, introduced in X.509 version 3, which allows intermediate CAs to issue certificates.

*Revocation*, supported by most applied PKIs, is also non-trivial to define. In fact, we found it necessary to define two variants: Full and

\*The work was partially completed during the author’s PhD studies at the Dept. of Computer Science, Bar-Ilan University, Israel

†The work was partially completed during a visiting position at the Dept. of Computer Science and Engineering, University of Connecticut, Storrs, CT

<sup>1</sup>A naive PKI ideal-functionality is presented in [7, 4, 13]; this functionality does not include basic functionalities such as *certificate chains*, *revocation* and *transparency*.

**Table 1: PKI requirements defined in this work, and properties we prove for prominent PKI systems.**

PKI scheme	Requirements				
	Existential Unforgeability (Definition 6)	Accountability (Definition 7)	$\Delta$ -Revocation (Definition 8)		$\Delta$ -Transparency (Definition 9)
			Weak	Full	
PKIX (X.509 version 3 with CRL)	✓ (Lemma 2)	✓ (Lemma 10)	✓ (Lemma 12)	✗ (Claim 3)	✗ (n/a)
PKIX with CT	✓ (Lemma 2)	✓ (Theorem 5)	✓ (Theorem 5)	✗ (Claim 3)	✗ (Claim 7)
PKIX with CT assuming honest loggers					✓ (Theorem 5)

Weak. The Full variant is more intuitive; basically, a PKI scheme ensures *Full  $\Delta$ -Revocation* if a certificate  $\psi$  revoked by a benign CA at some time  $t$ , will not be considered as valid by any benign party after time  $t + \Delta$ . However, we realized that PKIX fails to ensure Full  $\Delta$ -Revocation (Claim 3), and only ensures a weaker notion, which we call *Weak  $\Delta$ -Revocation*.

*Transparency* is a relatively recently identified goal for PKIs, which aims to ensure that certificates are available for scrutiny, in the form of a *public log* where certificates must appear (within bounded time after being issued). This requirement was identified when, after PKI failures, it was realized that the accountability requirement may not suffice, if the attacker’s use of a rogue certificate may go undetected. We show that PKIX, complemented by the *Certificate Transparency (CT)* specification [23, 25], ensures transparency if (and only if) we can trust the integrity of the CT *Logger* entities.

We also define a PKI *existential unforgeability* requirement, which ensures that if  $\psi$  is a valid certificate, as validated by the public key generated by a benign  $\psi$ .*issuer*, then  $\psi$ .*issuer* has indeed certified  $\psi$ . For X.509-based PKIs, unforgeability follows easily from the corresponding property of the underlying signature scheme; the property may not follow so simply for some other PKIs.

A challenge in defining security requirements for PKI schemes is that a scheme may satisfy a requirement only under specific models (assumptions), including *adversary*, *communication*, *synchronization* and other models. In particular, both PKIX and CT require some level of clock synchronization, to ensure consistent interpretation of validity periods. To ensure transparency, CT also requires *benign loggers* as well as reliable communication with bounded delays<sup>2</sup>.

To meet this challenge, we use the *Modular Security Specifications (MoSS) framework* [16], which allows the PKI security requirements to be applied under diverse, rigorously-defined adversary, synchronization, network and other models (assumptions).

We present pseudocode to rigorously define two PKI schemes, which are (minimally simplified<sup>3</sup> versions of) the most well-known

<sup>2</sup>CT includes some, but insufficient defenses against rogue loggers. Version 1 of the CT specifications [34] incorrectly states that the public logs can be *untrusted*; the CT 2.0 specifications [36] states that loggers should be *trusted* third-parties.

<sup>3</sup>Obviously, we could not include all aspects of the PKIX and CT specifications, but we believe that we included all or most of the significant aspects. Our most significant simplification is that we focus on the standardized CRL revocation mechanism, while current browsers mostly use proprietary revocation mechanisms. It appears feasible to extend our analysis to cover these and other aspects which we simplified, although, realistically, this may require automation. Note also that PKI standards are often ambiguous and not fully defined; for example, both CT specifications [34, 36] mention a gossip mechanism, but do not specify it. Hence, specifications can be implemented in different ways, and possibly some implementations will be insecure while others may be secure.

and widely-deployed PKI schemes: PKIX (X.509 version 3 with CRL) and CT (PKIX extended to support certificate transparency). We then rigorously prove security requirements satisfied by these schemes (see Table 1).

*Out of scope: selection of trust anchors and certificate-chain constraints.* We do not address how relying parties select their *trust anchors*, i.e., the identities of the ‘root CAs’; and we mostly ignore constraints on the allowed certificate-chains, such as the *name*, *length* and *policy constraints* (defined in X.509v3 and PKIX). A model of such *trust decisions* for PKI systems was proposed by Mauer [29], subsequently extended by [27, 6], and others [17, 26, 46, 39, 5, 19]. These solutions are complementary and orthogonal to our results.

**Contributions.** This work:

- (1) Presents the first definition of a (non-trivial) PKI scheme, e.g., supporting certificate chains.
- (2) Presents the first rigorous security requirements for PKI schemes, including *accountability*, *revocation*, *transparency* and *unforgeability*.
- (3) Presents the first precise and complete specification, analysis and proof of security for the two most widely-deployed PKI standards, PKIX and CT.
- (4) Introduces and constructs a *certificate scheme* (Section 2), an abstraction for applying signatures to structured information. Certificate schemes simplify definition and analysis of PKI schemes, and may have additional applications.

**Organization.** In Section 2 we define *Certificate schemes* and their security. In Section 3 we define a PKI scheme and in Section 4 we define security requirements for PKI schemes. In Section 5 we present the specifications of PKIX and CT, and in Section 6 we analyze their security. Finally, Section 7 concludes and discusses future work.

## 2 CERTIFICATE SCHEMES

PKI schemes define how to issue, manage and use *certificates*. Usually, e.g., in X.509, a certificate is a signed object, containing some *certified information*. Different PKIs may certify different information (fields), use different encodings and different signature algorithms, and, in principle, may even use a different design (i.e., not a signature over an object). For example, in X.509 and PKIs based on it, certain certified information is encoded as a ‘field’, while other information is encoded as an ‘extension’.

In this section, we define *certificate schemes*, which provide an abstract definition of a certificate. Our definition allows a uniform treatment of different types of certificates and certified information.

This abstraction allows us to define *PKI schemes* and their security requirements by having each PKI scheme use a specific certificate scheme.

## 2.1 Certificate Scheme

DEFINITION 1 (Certificate scheme). A certificate scheme  $C$  is a set of four PPT algorithms:

$$C = (\text{KG}, \text{Certify}, \text{Verify}, \text{Extract})$$

where:

- $\text{KG}(1^n) \rightarrow (\{0, 1\}^*, \{0, 1\}^*)$ : key generation algorithm. Returns a pair  $(sk, pk)$  of keys, where  $sk$  is a private signing (certification) key and  $pk$  is a public verification key.
- $\text{Certify}_{sk}(tbc) \rightarrow \psi$ : certification (signing) algorithm. Returns a binary string  $\psi$  which is called the certificate of the input,  $tbc$  (to be certified), signed using the private certifying (signing) key  $sk$ . The input  $tbc$  is a set of pairs,  $tbc = \{(f_1, v_1), (f_2, v_2), \dots\}$ , where the first element of each pair ( $f_i$ ) is called the field name and the second element ( $v_i$ ) is called the field value.
- $\text{Verify}_{pk}(\psi) \rightarrow \{\top, \perp\}$ : certificate verification algorithm. Returns  $\top$  if  $\psi$  was correctly signed according to the public verification key  $pk$ , otherwise, outputs  $\perp$ .
- $\text{Extract}(\psi, \text{FIELD}) \rightarrow \{0, 1\}^* \cup \{\perp\}$ : field value extraction algorithm. Returns the value of the inputted field name  $\text{FIELD}$  in  $\psi$ .

We say that  $C$  ensures correctness if for every  $(sk, pk) \leftarrow \text{KG}(1^n)$  and for every set of name-value pairs  $tbc$ , the following two correctness properties hold:

- (1) Verification correctness:

$$\text{Verify}_{pk}(\text{Certify}_{sk}(tbc)) = \top \quad (1)$$

- (2) Extraction correctness: Let  $\psi \leftarrow \text{Certify}_{sk}(tbc)$ . Then:

$$\text{Extract}(\psi, \text{field}) = \begin{cases} \text{val} & \text{if } (\text{field}, \text{val}) \in tbc \\ \perp & \text{otherwise} \end{cases} \quad (2)$$

We say that  $C$  ensures Existential Unforgeability if for every PPT  $\mathcal{A}$ , the probability of  $\mathcal{A}$  to win in the Existential Certificate Forgery game,  $\text{Pr}[\text{ECF}(n)]$ , is negligible, where

$$\text{ECF}(n) \equiv \left( \begin{array}{l} \text{Verify}_{pk}(\psi) = \top, \text{ where:} \\ (sk, pk) \leftarrow \text{KG}(1^n), \\ \psi \leftarrow \mathcal{A}^{\text{Certify}_{sk}(\cdot)}(pk), \text{ and} \\ \mathcal{A} \text{ didn't receive } \psi \text{ from } \text{Certify}_{sk}(\cdot) \end{array} \right) \quad (3)$$

We say that  $C$  is secure if it ensures both correctness and existential unforgeability.

## 2.2 Security of Simple Certificate Schemes

All existing PKIs use a simple certificate scheme which applies a public-key signature to an encoding of the input sequence of name-value pairs ( $tbc$ ). Let us define such schemes and prove their security.

DEFINITION 2. Let  $\Theta$  be an invertible function from sequences of name-value pairs to binary strings, and let  $\mathcal{S} = (\mathcal{S}.\text{KG}, \mathcal{S}.\text{Sign}, \mathcal{S}.\text{Verify})$  be an (existentially-unforgeable) signature scheme. The  $C^{\mathcal{S}, \Theta}$  certificate scheme is defined as:

$$\begin{aligned} C^{\mathcal{S}, \Theta}.\text{KG}(1^n) &\equiv \mathcal{S}.\text{KG}(1^n) \\ C^{\mathcal{S}, \Theta}.\text{Certify}_{sk}(tbc) &\equiv \Theta(\{(\text{'tbc'}, tbc), (\text{'\sigma'}, \mathcal{S}.\text{Sign}_{sk}(\Theta(tbc)))\}) \\ C^{\mathcal{S}, \Theta}.\text{Verify}_{pk}(\lambda) &\equiv \mathcal{S}.\text{Verify}_{pk}(\Theta^{-1}(\lambda)[\text{'tbc'}], \Theta^{-1}(\lambda)[\text{'\sigma'}]) \\ C^{\mathcal{S}, \Theta}.\text{Extract}(\lambda, f) &\equiv \begin{cases} \text{val} & \text{if } (f, \text{val}) \in \Theta^{-1}(\lambda)[\text{'tbc'}] \\ \perp & \text{otherwise} \end{cases} \end{aligned}$$

LEMMA 1. If  $\mathcal{S}$  is an existentially-unforgeable signature scheme, and  $\Theta$  is an invertible function, then  $C^{\mathcal{S}, \Theta}$ , defined in Definition 2, is a secure certificate scheme.

*Proof:* Direct reduction to the security of  $\mathcal{S}$ .  $\square$

*Dot notation for schemes and certificates.* We use dot notation for the functions of a scheme, above and elsewhere, to avoid ambiguity when referring to functions of different schemes. We also use dot notation for fields in a certificate, i.e., instead of referring to the value  $\text{val}$  of field  $f$  in certificate  $\psi$  by the cumbersome notation  $\text{Extract}(\psi, f)$ , we use the compact notation  $\psi.f$ . Namely,  $\psi.f = \text{Extract}(\psi, f)$ , and specifically,  $\psi.f = \perp$  if the certificate  $\psi$  does not include a name-value pair for a field name  $f$ .

## 2.3 Common Certificate Fields

PKI standards, most notably X.509 and PKIX, define and refer to specific named values in certificates. We refer to all of these named values as ‘fields’. We list below some of the important fields<sup>4</sup>, to which we also refer in our security requirements.

- $\psi.\text{issuer}$ : the entity that issued the certificate. In X.509, this is encoded as a *distinguished name* in the *Issuer* certificate field. X.509 certificates may also include an *Issuer Alternative Name (IAN)* extension that may include additional identifiers for the issuer, e.g., as a DNS name.
- $\psi.\text{from}$ : the date and time at which the certificate becomes valid. In X.509, this is encoded in the *notBefore* entry of the *Validity* certificate field.
- $\psi.\text{to}$ : the date and time at which the certificate should expire (become invalid). In X.509, this is encoded in the *notAfter* entry of the *Validity* certificate field.
- $\psi.\text{subject}$ : an identifier for the subject of  $\psi$ . In X.509, this is encoded as a *distinguished name* in the *Subject* certificate field. X.509 certificates may also include a *Subject Alternative Name (SAN)* extension, that may include additional identifiers for the subject, e.g., as a DNS name.
- $\psi.\text{pk}$ : a public key certified as the public key of the subject. In X.509, the subject public key is one component of the *Subject Public Key Info* certificate-field.
- $\psi.\text{is\_CA}$ : a Boolean which is true if the issuer authorizes the subject to be a CA, i.e., to issue public key certificates. In X.509, this is a part of the *Basic Constraints* extension.
- $\psi.\text{type}$ : the *certificate type*. This field is not explicitly defined in X.509; see discussion below.

<sup>4</sup>We use the term ‘field’ for all of these named values, and have abstracted away the encoding. However, specific PKIs, in particular, X.509, may use different terms in addition or instead of ‘field’, and may use different encodings for different fields. For example, X.509 uses the term ‘field’ for the *subject* and *issuer* fields (of the certificate) as well as for the *pathLenConstraint* field (which is a part of the *Basic Constraints* certificate extension). X.509 also uses other terms for named values, e.g., it refers to the *dnsName* component (in the *subjectAltName* extension) and to the *cA* boolean (in the *BasicConstraints* extension).

## 2.4 Types of Certificates

The aforementioned common certificate fields mostly correspond to certificate fields defined by X.509. One important exception is the *type* field, which is not one of the X.509 fields. We use the *type* field to distinguish between *different types of certificates*. X.509 explicitly refers to two types of certificates: *public key certificates* and *attribute certificates*, which we identify by  $\psi.type = PubKey$  (for public key certificates) or  $\psi.type = ATTR$  (for attribute certificates). Public key certificates are signed by a certificate authority (CA), and associate a public key with a particular subject ('owner' of the public key), whereas *attribute certificates* do not contain a public key (i.e., if  $\psi.type = ATTR$ , then  $\psi.pk = \perp$ ).

Note that our definition of a certificate can be applied also to other types of signed information consisting of multiple name-value pairs. Important, standardized examples include:

- *Certificate Revocation Lists (CRLs)*, used to identify revoked and non-revoked certificates. CRLs are certificates since they are signed (certified) objects with multiple fields (e.g., validity period and issuer).
- *Signed Tree Head (STH)*. Certificate Transparency (CT) uses *Signed Tree Head (STH)* certificates, signed (certified) by the CT loggers.
- *Route Origin Authorizations (ROAs)*, defined in the Resource Public Key Infrastructure (RPKI) [32]. ROAs are signed (certified) by the (certified) owner of a routing resource such as IP prefix or Autonomous System Number (ASN), and include multiple fields, e.g., validity period and issuer.

## 3 PKI SCHEMES AND TRANSPARENT PKI SCHEMES

We define *PKI schemes* and *transparent PKI schemes*, which are PKI schemes with additional entities and functions used to ensure transparency (§3.5). These sections provide necessary background: *basic PKI entities* (§3.1) and basic PKI functions, for evaluating the *validity* of a certificate (§3.2), for handling *certificate chains* (§3.3) and for certifying (issuing) and revoking certificates (§3.4).

### 3.1 PKI Entities

A PKI scheme  $\mathcal{P}$  is defined by a set of functions, some stateful and some stateless, and a set  $N$  of stateful entities. Entities in  $N$  can perform the stateful functions, e.g., issue certificates. We refer to the entities in  $N$  that issue public key certificates as *certificate authorities (CAs)*. There could be other entities in  $N$ , e.g., Certificate Transparency (CT) uses *loggers* and *monitors*. Entities in  $N$  may be *honest (benign)* or *corrupt*, i.e., controlled by an adversary. *Relying parties* are entities which use only stateless PKI functions, in particular, the certificate validation function, which allows the relying parties to decide whether to *rely on the certificate* (i.e., use the certified public key) or not.

The state of the entities in  $N$  can be initialized using a dedicated *initialization* operation, denoted as  $\mathcal{P}.Init$ . Typically,  $\mathcal{P}.Init$  outputs a *self-certified public key certificate*, i.e., a certificate  $\psi$  which certifies a key for its issuer ( $\psi.issuer = \psi.subject$ ), and is validated (successfully) using the certified key  $\psi.pk$ .

### 3.2 Certificate Validity

We determine the validity of a certificate  $\psi$  using the stateless *certificate validation predicate*  $\mathcal{P}.Valid(clk, \psi, store, \omega)$  (Definition 4). An entity evaluates validity with respect to its current clock value  $clk$ , the set  $store$  of self-certified *trust-anchor* certificates and auxiliary information  $\omega$ .

In existing PKIs, validation uses the  $C.Verify$  operation of an underlying certificate scheme, typically, by establishing trust in a public validation key  $pk$  of  $\psi.issuer$ . This usually uses *certificate chains* from  $\omega$ , introduced next. We use chains to establish trust in  $pk$  (from the trust-anchor certificates in  $store$ ), and validate that  $\psi$  and certificates in the chains were not revoked, e.g., using a CRL (§3.4).

### 3.3 Certificate Chains

We now define *certificate chains*<sup>5</sup>, which allow validation of certificates signed by a non-root CA, i.e., a CA which does not have a self-signed certificate in  $store$ .

**DEFINITION 3 (Certificate chain).** Let  $\xi = \{\psi_0, \psi_1, \dots\}$  be a sequence of certificates. We say that  $\xi$  is a certificate chain for certificate scheme  $C$  if for every issued certificate  $\psi_{i-1}$  and its issuer certificate  $\psi_i$ , for  $0 < i < |\xi|$ , holds:

- (1) *Correct certification:*  $C.Verify_{\psi_i.pk}(\psi_{i-1}) = \top$ .
- (2) *The issuer of  $\psi_{i-1}$  is the subject of the issuer certificate  $\psi_i$ , i.e.,  $\psi_{i-1}.issuer = \psi_i.subject$ .*
- (3) *The issuer certificate is a public-key certificate for a CA, i.e.:  $\psi_i.is\_CA = \top$  and  $\psi_i.type = PubKey$ .*

See pseudocode in Algorithm 1.

---

#### Algorithm 1 Certificate chain predicate

---

**Input:** certificate chain  $\xi$ .

**Output:**  $\top$  if  $\xi$  is a well structured chain or  $\perp$  otherwise.

```

1: procedure ISCHAINC( $\xi$ )
2:   return  $\forall i : 0 < i < |\xi| :$ 
3:      $C.Verify_{\xi[i].pk}(\xi[i-1]) \wedge$  ▷ All non-first certs in the chain
4:      $\xi[i-1].issuer = \xi[i].subject \wedge$  ▷ Are properly certified
5:      $\xi[i].type = PubKey \wedge$  ▷ Are for issuer of prev cert
6:      $\xi[i].is\_CA$  ▷ Are PubKey certs
7: end procedure ▷ Are for a CA

```

---

The stateless operation  $\mathcal{P}.Chains(clk, \psi, store, \omega) \rightarrow \{\xi_i\}$  identifies one or more chains used to validate a given certificate  $\psi$ . The certificate chains returned by  $\mathcal{P}.Chains$  are used to ensure *accountability*, as defined in §4.3.

### 3.4 Certifying and Revoking Certificates

A certificate  $\psi$  is issued using the private certification key of a CA, using the PKI's  $\mathcal{P}.Certify$  operation, namely,  $\mathcal{P}.Certify(st, clk, tbc) \rightarrow (st, \psi/\perp)$ , which takes as input the entity's local state  $st$ , local clock  $clk$  and the data to be certified  $tbc$ , and outputs an updated state  $st$ , and a signed certificate  $\psi$  if successful or  $\perp$  otherwise. Typically,  $\mathcal{P}.Certify$  uses the certify function  $C.Certify_{sk}(tbc)$  of an underlying certificate scheme  $C$  (see Definition 2) where  $tbc$  is

<sup>5</sup>The term 'chain of trust' is also sometimes used.

the information to be certified using the entity's private key  $sk$  extracted from its local state  $st$ .

Since certificates are typically issued for a specific time period, most PKI schemes provide a way to revoke certificates before their expiration date, for example, if a certificate is found to be fraudulently issued or the corresponding private key exposed. Revocation is done by the issuer using a dedicated revoke operation, denoted as  $\mathcal{P}.\text{Revoke}$ . The  $\mathcal{P}.\text{Revoke}$  operation takes as input a certificate  $\psi$  and outputs whether the revocation was successful or not, i.e.:  $\mathcal{P}.\text{Revoke}(st, clk, \psi) \rightarrow (st, \top/\perp)$ . For example,  $\mathcal{P}.\text{Revoke}$  would return  $\perp$  if attempting to revoke an already expired or revoked certificate, or a certificate not issued by this issuer.

Most PKIs use some form of a non-revocation certificate to allow a relying party to verify that a certificate was not revoked (at a given time). X.509 defines two non-revocation mechanisms, *certificate revocation lists (CRLs)* and the *online certificate status protocol (OCSP)*. Relying parties can verify that a certificate was not revoked at time  $clk$  by obtaining a CRL or OCSP response, which are valid at time  $clk$ .

### 3.5 Definitions of PKI Schemes

**DEFINITION 4** (PKI scheme). *A PKI scheme  $\mathcal{P}$  is a set containing (at least) the following PPT algorithms:*

- $\mathcal{P}.\text{Init}(st, clk, params) \rightarrow (st, \psi)$ : Takes as input the state<sup>6</sup>  $st$ , local clock  $clk$ , and parameters  $params$ , and returns the initialized local state  $st$  after performing initialization based on the input parameters  $params$  on time  $clk$  and a self-certified certificate  $\psi$ .
- $\mathcal{P}.\text{Certify}(st, clk, tbc) \rightarrow (st, \psi/\perp)$ : Takes as input the state  $st$ , local clock  $clk$  and information to be certified  $tbc$ , and outputs an updated state  $st$  and either a certificate  $\psi$  or a failure indicator  $\perp$ .
- $\mathcal{P}.\text{Revoke}(st, clk, \psi) \rightarrow (st, \top/\perp)$ : Takes as input the state  $st$ , local clock  $clk$ , and a certificate  $\psi$ , and outputs an updated state  $st$  and  $\top$  if  $\psi$  was revoked successfully or  $\perp$  if the revocation failed.
- $\mathcal{P}.\text{Valid}(clk, \psi, store, \omega) \rightarrow (\top/\perp)$ : This (stateless) algorithm takes as input local clock  $clk$ , a certificate  $\psi$ , a root store  $store$ , and auxiliary information  $\omega$ , and outputs either  $\top$  or  $\perp$ .
- $\mathcal{P}.\text{Chains}(clk, \psi, store, \omega) \rightarrow \{\xi_i\}$ : This (stateless) algorithm receives the same inputs as  $\mathcal{P}.\text{Valid}$ , and, if  $\psi$  is valid, returns one or more certificate chains  $\xi_i$ .

PKI schemes might include additional inputs or operations. In particular, we next define *transparent* PKI schemes; such schemes have additional operations used to ensure that valid certificates will be publicly available (i.e., “transparent”), allowing to detect discrepancies and suspect certificates.

**DEFINITION 5** (Transparent PKI scheme). *A transparent PKI scheme  $\mathcal{P}$  is a PKI scheme with the following additional PPT algorithms:*

- $\mathcal{P}.\text{Monitor}(st, clk, \iota) \rightarrow st$ : Takes as input state  $st$ , local clock  $clk$  and entity identifier  $\iota \in \mathbb{N}$ , and outputs an updated state  $st$ , and it starts to monitor (certificates logged by)  $\iota$ .

<sup>6</sup>The state  $st$  is given in the input for a technical reason: the execution process treats  $\text{Init}$  like other operations, and therefore, also passes the state  $st$  to it. Typically, the input state is ignored.

- $\mathcal{P}.\text{Lookup}(st, clk, subject) \rightarrow (st, \Psi)$ : Takes as input state  $st$ , local clock  $clk$ , and an identifier  $subject$ , and outputs an updated state  $st$  and a set  $\Psi = \{(\psi_1, \omega_1), (\psi_2, \omega_2), \dots\}$  of all pairs  $(\psi_i, \omega_i)$  known to the entity s.t.  $\psi_i$  is the certificate of the given subject, i.e.,  $\psi_i.subject = subject$ , and  $\omega_i$  is the corresponding auxiliary information for  $\psi_i$ .

## 4 PKI REQUIREMENTS

In this section we define the requirements that a PKI scheme should satisfy. Formally-defined requirements are necessary to (formally) prove whether or not a given PKI implementation satisfies specific requirements, and if so, under what assumptions.

### 4.1 Modular Security Specifications

PKI schemes, both deployed and proposed, vary greatly in their designs, operate under different models (assumptions) and aim to satisfy different requirements (goals). To define requirements which apply to different PKI schemes, even if they assume different models, we use the *Modular Security Specifications (MoSS) Framework* [16] for our PKI specifications. MoSS separates the definition of *requirements* from the *models*, which simplifies the definition of requirements and allows evaluating the requirements satisfied by PKI schemes under different models. For example, we show that CT provides transparency under a model that assumes honest loggers, but not under a model that allows corrupt loggers (Table 1).

To illustrate the importance of separating models from requirements, consider the following *simplified  $\Delta$ -Revocation* requirement: a certificate revoked at time  $clk$  by an honest CA would be considered invalid at any time after  $clk + \Delta$ . The delay  $\Delta$  can be a function of network delays, clock bias and design decisions (e.g., periodicity of issuing CRLs). Satisfaction of the  $\Delta$ -Revocation requirement depends on the clock synchronization and delay models.

For completeness, we recall here and in Appendix A.1 concepts of MoSS [16] used in this work.

MoSS defines an execution process  $\text{Exec}_{\mathcal{A}, \mathcal{P}}(params)$  in which an adversary  $\mathcal{A}$  has complete control over the environment, the execution of a protocol  $\mathcal{P}$  and all external inputs (parameters) of the protocol ( $params$ ). The output of the execution of protocol  $\mathcal{P}$  with parameters  $params$  is called a *transcript*, denoted as  $T \leftarrow \text{Exec}_{\mathcal{A}, \mathcal{P}}(params)$ . A transcript contains the set of entities  $T.N$  and the subset of faulty entities  $T.F \subseteq T.N$ . It also contains, for each step in the execution, the operation invoked, the inputs and the outputs of the protocol, the output of the adversary  $T.out_{\mathcal{A}}$ , and other relevant values of the execution. To illustrate, given a transcript  $\text{logm}$  and certificate  $\psi$ , we can check if there is an event  $e$  in  $T$  in which the adversary invoked the  $\mathcal{P}.\text{Certify}$  operation on entity  $\psi.issuer$  with input  $\psi.tbc$  and received  $\psi$  as output by using the following predicate:

$$\exists e \text{ s.t. } T.opr[e] = \text{'Certify'} \wedge T.ent[e] = \psi.issuer \wedge T.inp[e] = \psi.tbc \wedge T.out[e] = \psi$$

A MoSS model  $\mathcal{M}$  is defined using an associated predicate  $\pi_{\mathcal{M}}$ , and similarly, a MoSS requirement  $\mathcal{R}$  is defined using an associated predicate  $\pi_{\mathcal{R}}$ . The predicates are applied to the transcript of an execution of a protocol and define if the execution satisfied the requirement or model. For each model and predicate, we can define

a base function  $\beta(params)$ , applied to the parameters  $params$ , to allow some probability for an execution to not satisfy the model or requirement. In this work we omit  $\beta$  and write  $(\pi)$  as a shorthand for  $(\pi, \beta)$ , since our models and requirements have  $\beta(params) = 0$ .

A PPT adversary  $\mathcal{A}$  satisfies model  $\mathcal{M} = (\pi_{\mathcal{M}})$ , if the probability that  $\pi_{\mathcal{M}}(T) = \perp$  is negligible, for a random transcript of any protocol interacting with  $\mathcal{A}$ . A protocol  $\mathcal{P}$  satisfies requirement  $\mathcal{R} = (\pi_{\mathcal{R}})$  given model  $\mathcal{M}$ , if the probability that  $\pi_{\mathcal{R}}(T) = \perp$  is negligible for a random transcript of  $\mathcal{P}$  interacting with any PPT adversary  $\mathcal{A}$  that satisfies model  $\mathcal{M}$ .

The rest of this section defines several PKI requirements: the existential unforgeability requirement in §4.2, the accountability requirement in §4.3, two versions (Weak and Full) of the  $\Delta$ -Revocation requirement in §4.4, and the  $\Delta$ -Transparency requirement in §4.5.

## 4.2 The Existential Unforgeability Requirement

The existential unforgeability requirement of PKI schemes is similar to the corresponding property for signature and certificate schemes (Definition 1). It ensures that every certificate  $\psi$  issued by an honest entity  $\iota$  and successfully validated using  $\iota$ 's public key, was issued following a request to do so, i.e.,  $\iota$  invoked  $\mathcal{P}.\text{Certify}$  to produce  $\psi$ . The only exception are *self-signed certificates*, i.e., certificates whose *issuer* and *subject* fields are the same and certify the same public key used to validate them. CAs *self-certify* their public keys after generating them (typically, at initialization). The unforgeability requirement ensures that honest entities will not be considered as issuers for certificates that they did not issue. Unforgeability typically follows easily from the fact that certificates are issued by a secure certificate scheme; see Lemma 2.

**DEFINITION 6** (PKI Existential Unforgeability). *The PKI existential unforgeability requirement  $\mathcal{R}_{\text{EUF}}$  is defined as  $\mathcal{R}_{\text{EUF}} = (\pi_{\text{EUF}})$ , where  $\pi_{\text{EUF}}$  is defined in Algorithm 2.*

The following lemma shows that both PKIX and CT ensure existential unforgeability under the  $\mathcal{M}^{\text{F}}$  model. This model, defined in [16], ensures that all entities follow the protocol correctly, except for entities in T.F (the set of faulty entities outputted by the adversary).

**LEMMA 2.** *The PKIX and CT PKI schemes, presented in §5, ensure existential unforgeability under the  $\mathcal{M}^{\text{F}}$  model, if they use a secure (existentially unforgeable) certificate scheme  $C$ .*

*Proof:* In both PKIX and CT, non-faulty entities (1) output a self-signed certificate only using a public key generated by the entity using  $C.\text{KG}$  (in their  $\text{Init}$  operations), (2) never output the corresponding private key, and (3) only apply the corresponding private key, using  $C.\text{Certify}$ , to certify  $\psi.tbc$ . The claim follows by reduction to the unforgeability of  $C$ .  $\square$

## 4.3 The Accountability Requirement

Intuitively, we<sup>7</sup> use the term *accountability* for the ability to *identify*, for every certificate  $\psi_0$ , an *accountable* (responsible) *trust anchor* for  $\psi_0$ . We identify the accountable trust anchor by a chain<sup>8</sup> of valid certificates leading to a self-signed certificate in *store*, which

<sup>7</sup>We do not use the term accountability in a legal sense.

<sup>8</sup>Stronger accountability requirements are possible, e.g., requiring at least two 'independently-accountable' chains to different root CAs, or allowing only limited-length of certificate chains.

## Algorithm 2 $\pi_{\text{EUF}}$ : Existential Unforgeability Requirement

---

**Input:** transcript  $T$ .  
**Output:**  $\top$  if for the certificate  $\psi$  outputted by the adversary in the execution transcript  $T$ , where  $\psi.\text{issuer}$  is an identifier of a benign entity and  $\psi$  is valid with respect to a public key self-signed by  $\psi.\text{issuer}$ , holds that  $\psi.\text{issuer}$  was requested to certify  $\psi.tbc$ . Otherwise,  $\perp$ .

---

```

1: procedure  $\pi_{\text{EUF}}(T)$ 
2:    $clk, \psi, store, \omega \leftarrow T.out_{\mathcal{A}}$  ▷ Extract adversary's output
3:   if CERTIFYREQUESTED( $T, \psi$ ) then return  $\top$  ▷ Requested, i.e., not forged!
4:   if  $\left[ \begin{array}{l} \mathcal{P}.\text{Valid}(clk, \psi, store, \omega) \wedge \\ \psi.\text{issuer} \in T.N - T.F \wedge \\ (\forall \xi \in \omega \text{ s.t. } \xi[0] = \psi : \\ \text{SELF CERTIFIED}(T, \psi.\text{issuer}, \xi[1].pk)) \wedge \\ (\forall \psi' \in store \text{ s.t. } \psi'.subject = \psi.\text{issuer} : \\ \text{SELF CERTIFIED}(T, \psi'.subject, \psi'.pk)) \end{array} \right]$  ▷  $\psi$  is valid
▷  $\psi$ 's issuer is a benign entity
▷ All chains by honest CAs in  $\omega$ 
▷ to validate  $\psi$  are correct
▷ All honest root CAs' keys
▷ in store are real
▷ Forgery! Adversary wins
5:   then return  $\perp$ 
6:   end if
7:   return  $\top$  ▷ Adversary failed
8: end procedure
9: procedure CERTIFYREQUESTED( $T, \psi$ )
10:  return  $\exists e \text{ s.t. } T.opr[e] = \text{'Certify'}$  ▷ Certify operation invoked
11:   $T.ent[e] = \psi.\text{issuer}$  ▷ on  $\psi.\text{issuer}$ 
12:   $T.inp[e] = \psi.tbc$  ▷ for certificate  $\psi$ 
13: end procedure
14: procedure SELF CERTIFIED( $T, \iota, pk$ )
15:  return  $\exists e, \psi' \text{ s.t. } T.ent[e] = \iota \wedge$  ▷ Entity  $\iota$ 
16:   $T.out[e] = \psi' \wedge$  ▷ outputted certificate  $\psi'$ 
17:   $\psi'.pk = pk \wedge$  ▷ certifying given  $pk$ ,
18:   $C.\text{Verify}_{pk}(\psi') \wedge$  ▷ signed by same  $pk$ , and
19:   $\psi'.\text{issuer} = \psi'.subject = \iota$  ▷ using the same identity
20: end procedure
```

---

provides *non-repudiation*. Accountability is an important PKI requirement; it provides a (reactive) defense by identifying a root CA that is accountable for rogue yet valid certificates, regardless of whether the CA acted intentionally or negligently. A root CA which is found accountable for a rogue certificate, once or repeatedly, may be removed from root stores or otherwise penalized.

To identify a CA accountable for a valid  $\psi$ , we use the  $\mathcal{P}.\text{Chains}$  function to extract the set of auxiliary information  $\omega$  used to validate  $\psi$ . In an accountable PKI,  $\omega$  contains at least one *valid certificate chain*  $\xi$  that begins with  $\psi$  (i.e.,  $\xi[0] = \psi$ ). We say that a  $\xi$  is a valid certificate chain, if it contains a (self-signed) trust anchor  $\xi[i] \in store$ , and the certs from  $\xi[0]$  to  $\xi[i-1]$  are valid and form a chain (Definition 3); see the  $\text{ISVALIDCHAIN}$  function in Algorithm 3. The issuer (which is also the subject) of the trust anchor<sup>9</sup>  $\xi[i]$  is *accountable* for  $\psi$ .

We express accountability as a predicate  $\pi_{\text{ACC}}$  (Algorithm 3) and use it to define the accountability requirement in Definition 7.

**DEFINITION 7** (Accountability). *The accountability requirement  $\mathcal{R}_{\text{ACC}}$  is defined as  $\mathcal{R}_{\text{ACC}} = (\pi_{\text{ACC}})$ ; see Alg. 3.<sup>10</sup>*

## 4.4 The Weak and Full $\Delta$ -Revocation Requirements

Most PKIs allow a certificate to be revoked by its issuer, i.e., invalidated prior to its specified expiration date. Revocation may be used

<sup>9</sup>Typically, but not necessarily, the trust anchor  $\xi[i]$  is the last certificate in  $\xi$ .  
<sup>10</sup>Line 5 of  $\pi_{\text{ACC}}$  checks  $\mathcal{P}.\text{Valid}$  after checking  $\text{ISVALIDCHAIN}(\xi)$ . Why? See App. D.

**Algorithm 3**  $\pi_{\text{ACC}}$ : Accountability requirement

---

**Input:** transcript  $T$ .  
**Output:**  $\top$  if the adversary loses, i.e., does not output a valid certificate, or outputs a cert  $\psi$  for which the PKI identifies a valid certificate chain which contains an accountable CA (in *store*).  
 Otherwise, i.e., if adversary wins, output  $\perp$ .

1: **procedure**  $\pi_{\text{ACC}}(T)$   
 2:    $(clk, \psi, store, \omega) \leftarrow T.out_{\mathcal{A}}$  ▷ Adversary's output  
 3:   **return**  $\neg \mathcal{P}.\text{Valid}(clk, \psi, store, \omega) \vee$  ▷ Certificate is not valid  
 4:    $(\exists \xi \in \mathcal{P}.\text{Chains}(clk, \psi, store, \omega), j < |\xi|) \text{ s.t. :}$  ▷ Or, for some chain  $\xi$  and  $j < |\xi|$ :  
 5:    $\left( \begin{array}{l} \text{ISCHAIN}(\xi) \wedge \xi[j] = \psi \wedge \\ \mathcal{P}.\text{Valid}(clk, \xi[j+1], store, \omega) \end{array} \right)$  ▷ which contains  $\psi$  and whose next cert,  $\xi[j+1]$ , is valid  
 6: **end procedure**

---

for a variety of reasons, including loss or compromise of the private key corresponding to a certified public key, or when the issuer determines that the certificate contains incorrect information. We focus on two variants of the basic requirement for revocation which we call the *Weak* and *Full*  $\Delta$ -Revocation requirements. We first discuss the stronger and simpler Full  $\Delta$ -Revocation requirement.

*Full  $\Delta$ -Revocation.* The requirement ensures that a certificate  $\psi$  issued by a benign CA  $\psi.\text{issuer}$ , and revoked by  $\psi.\text{issuer}$  at time  $clk$ , will not be considered as valid at any time following  $clk + \Delta$ . The ‘grace period’  $\Delta$  is typically required for three reasons: (1) the bias between the clock of  $\psi.\text{issuer}$  and the clock of the relying party validating  $\psi$ , (2) the communication delay from the  $\psi.\text{issuer}$  to the relying parties, and (3) the time for any proof that  $\psi$  is non-revoked that was issued prior to its revocation becomes stale, i.e., expire. Full  $\Delta$ -Revocation ensures that the revocation operation by a benign CA is always effective (after at most  $\Delta$ ); a relying party cannot be misled to rely on a revoked certificate (for more than  $\Delta$ ).

A natural (but unfortunately not sufficient) approach to ensure Full  $\Delta$ -Revocation is to validate  $\psi$  using a *non-revocation certificate*  $\psi_{NR} \in \omega$  that certifies that  $\psi$  was not revoked until a given time  $\psi_{NR}.\tau$ ; as a result, the certificate will be considered by relying parties as valid until time  $\psi_{NR}.\tau + \Delta$ . In X.509 and PKIs based on it, including PKIX and CT, the non-revocation certificate<sup>11</sup>  $\psi_{NR}$  must indicate the same issuer as of  $\psi$ , i.e.,  $\psi_{NR}.\text{issuer} = \psi.\text{issuer}$ .

However, this design fails to ensure the Full  $\Delta$ -Revocation requirement. We present the *Zombie certificate attack*, where an attacker is able to cause  $\psi$  to be valid at time  $t \in [\psi.\text{from}, \psi.\text{to}]$  although it was revoked by a benign issuer  $\psi.\text{issuer}$ .

**CLAIM 3** (The Zombie certificate attack). *PKIX and CT, using CRLs or OCSP for revocation, fail to ensure the Full  $\Delta$ -Revocation requirement.*

*Proof (informal).* We present the proof for PKIX using CRLs, but it applies with trivial changes also to CT and when using OCSP for revocation. The proof is by a counter example: the *Zombie-certificate attack*.

Let  $\psi$  be a certificate issued by a benign CA  $\psi.\text{issuer}$ , with validity period  $[\psi.\text{from}, \psi.\text{to}]$ , which is later revoked by  $\psi.\text{issuer}$ , at time  $t_R$  s.t.  $\psi.\text{from} < t_R < \psi.\text{to} - \Delta$ . We show how an attacker is able to certify a fake-yet-valid CRL which indicates that  $\psi$  was not revoked and is still valid during  $[t_R, \psi.\text{to}]$ , i.e., after  $\psi$  was revoked by the benign issuer  $\psi.\text{issuer}$ .

<sup>11</sup>X.509 defines two types of non-revocation certificates: the *Certificate Revocation List (CRL)* [9] and the *Online Certificate Status Protocol (OCSP)* [33].

The attacker controls a rogue CA  $\iota_R \in T.F$  which is able to issue valid CA certificates. This may be the root CA that is the trust anchor of  $\psi$ 's certificate chain, or an intermediate CA that has a certificate chain rooted by the same trust anchor of  $\psi$ 's<sup>12</sup>. The attacker generates a keypair  $(sk_R, pk_R) \leftarrow C.\text{KG}(1^\beta)$  and then certifies a fake CA-certificate  $\psi_R$  for  $\psi.\text{issuer}$ , i.e.,  $\psi_R.pk = pk_R$  (which, of course, is not the public key used by the benign  $\psi.\text{issuer}$ , e.g.,  $C.\text{Verify}_{\psi_R.pk}(\psi) = \perp$ ). Let  $\omega_R$  denote the auxiliary information required to validate  $\psi_R$ , including a valid certificate chain  $\xi_R$  from a root CA. If  $\iota_R$  is a root CA, then  $\xi_R$  will contain only  $\psi_R$  and the self-signed certificate of  $\iota_R$ .

The attacker then uses  $sk_R$  to certify  $CRL_R$ , a fake-yet-valid CRL, where  $\psi$  is *not* marked as revoked. Hence,  $\psi$  would be considered valid during  $[t_R, \psi.\text{to}]$  when provided together with the (fake yet valid) CRL  $CRL_R$  and with  $\omega_R$  (containing the (misleading yet valid) certificate chain  $\xi_R$ , validating the fake CA-certificate  $C_R$ ). Namely, the scheme fails to satisfy the Full  $\Delta$ -Revocation requirement.  $\square$

*Notes:* (1) A naive ‘fix’ is to require the non-revocation certificate (CRL/OCSP) to be signed by the same key used to certify  $\psi$ . However, this would make it impossible for a benign issuer  $\psi.\text{issuer}$  to change its key if necessary, e.g., due to exposure. (2) Short-lived certificates are a simple way to ensure Full, using  $\psi.\text{to} - \psi.\text{from} \leq \Delta$ .

*Weak  $\Delta$ -Revocation.* This requirement *allows*  $\psi$  to be considered valid even after being revoked by its benign CA  $\psi.\text{issuer}$ , provided that the PKI can identify a valid-yet-rogue CA certificate  $\psi_R$  that is responsible for  $\psi$  being valid after it was revoked. If the PKI can identify the rogue CA certificate  $C_R$ , and also ensures *accountability*, then we can identify a root CA which is accountable for the  $\psi_R$ . We identify  $\psi_R$  as a rogue CA certificate, since it certifies public key  $\psi_R.pk$  for  $\psi.\text{issuer}$ , although the benign  $\psi.\text{issuer}$  did not output a corresponding self-signed certificate, i.e., it didn't generate  $\psi_R.pk$ .

We define both Full and Weak  $\Delta$ -Revocation requirements in Definition 8, using the function  $f_{\Delta\text{Rev}}$  presented in algorithm 25. In §6 we show that PKIX and CT ensure Weak  $\Delta$ -Revocation.

**DEFINITION 8** ( $\Delta$ -Revocation requirements). *We define the Full and Weak  $\Delta$ -revocation predicates as:*

$$\begin{aligned} \pi_{\text{Full}-\Delta\text{Rev}}(T) &= \begin{cases} \top & \text{if } f_{\Delta\text{Rev}}(T) = \top, \\ \perp & \text{otherwise} \end{cases} \\ \pi_{\text{Weak}-\Delta\text{Rev}}(T) &= \begin{cases} \perp & \text{if } f_{\Delta\text{Rev}}(T) = \perp, \\ \top & \text{otherwise} \end{cases} \end{aligned}$$

And define the Full and Weak  $\Delta$ -Revocation requirements as:

$$\begin{aligned} \mathcal{R}_{\text{Full}-\Delta\text{Rev}} &= (\pi_{\text{Full}-\Delta\text{Rev}}) \\ \mathcal{R}_{\text{Weak}-\Delta\text{Rev}} &= (\pi_{\text{Weak}-\Delta\text{Rev}}) \end{aligned}$$

## 4.5 The $\Delta$ -Transparency Requirement

Accountability, as described in §4.3, is a *deterrent* against rogue or negligent behavior by root CAs. For many years this reactive measure was viewed as a sufficient defense, under the assumption that root CAs, and CAs certified by them, were respectable and trustworthy entities who would not risk, intentionally or otherwise, being implicated in issuing rogue certificates. However, repeated cases of rogue certificates issued by compromised or dishonest CAs,

<sup>12</sup>RFC5280 states that “The trust anchor for the certification path MUST be the same as the trust anchor used to validate the target certificate”.



have proven this assumption to be overly optimistic. It turned out that punishing root CAs is non-trivial: beyond negative publicity, punishment was often ineffective [37, 1, 18].

Furthermore, ‘punishment’ requires that a rogue certificate is *discovered*. An attacker could reduce the risk of discovery by minimizing the exposure of the rogue certificate. Efforts such as Perspectives Project [44] and the EFF SSL Observatory [11] provide some assistance to the discovery of rogue certificates, but were not sufficient to address this concern.

This motivated *transparent* PKI designs, most notably, the standardized and deployed *Certificate Transparency (CT)* [23, 25]. Certificate transparency requires a certificate  $\psi$  to be certified by one or more *loggers*. Loggers are parties committed to include (transparent) certificates in a *public log* they maintain, and to make this log *available* to third parties called *monitors*. Each monitor keeps tabs on the certificates logged by (one or, usually, more) loggers; in addition, monitors may detect suspect certificates and inform interested parties, such as domain owners and relying parties.

In other words, transparency prevents a CA from ‘quietly’ generating a rogue certificate  $\psi$ , and using  $\psi$  to attack select victims. Transparency facilitates early detection of rogue certificates issued by a corrupt, compromised or negligent CA. By demanding that a valid certificate must be transparent, we ensure the detection of rogue and suspect certificates.

There is some unavoidable delay from the time that a certificate is submitted to the log and until the relevant monitors are aware of it. This delay is mainly due to significant allowed delay from receiving a certificate and until including it in a new signed log (STH).

Informally, a PKI scheme  $\mathcal{P}$  satisfies the  $\Delta$ -Transparency requirement if no adversary  $\mathcal{A}$  can produce, with non-negligible probability, a pair  $(\psi, \omega)$  where  $\psi$  is a certificate *valid* w.r.t. the auxiliary information  $\omega$  on time  $clk$ , and yet, there exists a benign monitor  $\iota_M$  which is unaware of  $\psi$ , even though  $\iota_M$  monitors a log in the set of logs  $\psi.logs$  in which  $\psi$  is (allegedly) published, prior to  $\psi.\tau$ .

We express this informal definition as a predicate  $\pi_{\Delta Tra}$  (Algorithm 26), and present the formal definition of  $\Delta$ -Transparency in Definition 9.

**DEFINITION 9 ( $\Delta$ -Transparency).** *The  $\Delta$ -Transparency requirement  $\mathcal{R}_{\Delta Tra}$  is defined as  $\mathcal{R}_{\Delta Tra} = (\pi_{\Delta Tra})$ , where  $\pi_{\Delta Tra}$  is defined in Algorithm 26.*

## 5 PROVABLY-SECURE PKI SCHEMES

In this section, we present the design of two important PKI schemes: PKIX and CT, where our design of CT is based on the CT 2.0 specification [36], using PKIX for aspects not covered in [36]. In both PKIX and CT, we use the CRL revocation mechanism. The full PKIX and CT specifications are far too complex to present and analyze in this paper; as such, we considered a significantly simplified design but we retained all aspects related to security considerations, e.g., by considering certificate chains (although we ignore the length, name and policy constraint extensions).

### 5.1 PKIX (with CRLs)

*Entities and state.* In PKIX, all entities (in  $N$ ) are certificate authorities (CAs), i.e., they issue and revoke public key certificates. Some

of these are *root CAs* (trust anchors), which relying parties trust directly (by maintaining their self-signed certificates in *store*), and others are *intermediate CAs*, which relying parties trust based on a certificate chain ending at a root CA. Each CA maintains a local state  $st$  which contains the following information:

- $st.i$  : the identifier of the CA.
- $st.sk$  : the CA’s secret (signing) key.
- $st.pk$  : the CA’s (public) verification key.
- $st.CRL$  : the list of all revoked certificates.
- $st.\Delta_r$  : CRL’s validity period.
- $st.certs$  : the set of certificates issued by the entity.

*Certificate fields and scheme.*  $PKIX_C$  uses a simple certification scheme  $C$  (Definition 2), defined in [31]. The  $PKIX_C$  implementation (Algorithms 4-8) uses the common certificate fields mentioned in §2.3, with the addition of the serial number of public key certificates  $\psi.serial$ .

*Auxiliary information.* In order for a certificate  $\psi$  to be considered valid in  $PKIX_C$ , it must have a valid chain-of-trust to a trusted root CA and a CRL that proves that  $\psi$  was not revoked. Specifically, the auxiliary information  $\omega$  should contain the following information:

- $\omega.\psi_{CA}$ , a certificate of the CA that issued  $\psi$  and  $\omega.\omega_{CA}$ , auxiliary information required to validate  $\omega.\psi_{CA}$ .
- $\omega.\psi_{CRL}$ , a CRL showing that  $\psi$  was not revoked, and  $\omega.\omega_{CRL}$ , auxiliary information required to validate  $\omega.\psi_{CRL}$ .

*Implementation.* We now present the implementation (Algorithms 4-8), according to Definition 4, with the addition of one more operation:  $PKIX_C.GetCRL$ , which produces the current CRL that is properly certified (signed). This corresponds to the CRL’s distribution point mechanism described in [31].

$PKIX_C.Init$  (Algorithm 4). The  $PKIX$  Init function generates the CA’s keypair (line 3) and self-certifies the public key (line 5). It also initializes the state  $st$  of the CA: the CA’s identity  $st.i$ , the  $st.\Delta_r$  parameter (the CRL validity period), both from the *params* input, and the set of certificates issued ( $st.certs$ ) and revoked ( $st.CRL$ ) by the CA, both initially empty.

$PKIX_C.Certify$  (Algorithm 5).  $PKIX$  certifies only public-key certificates. Therefore, the algorithm sets the certificate’s type to be of public key (line 2), and then signs  $tbc$  using  $C.Certify(st, clk, tbc)$  and the CA’s secret signing key (line 3). The data to be certified and the signature over it are bundled to form the certificate  $\psi$  (line 4). Then, a copy of  $\psi$  is stored locally (line 5) and the algorithm outputs  $\psi$  along with its auxiliary information  $\omega$  (line 6).

$PKIX_C.Revoke$  (Algorithm 6). The algorithm verifies that the inputted parameters are legitimate, i.e., the certificate to be revoked was indeed issued by the CA (line 1). If so, the algorithm adds the certificate to the list of revoked certificates (line 2).

$PKIX_C.Valid$  (Algorithm 8). The (stateless)  $PKIX.Valid$  algorithm starts by checking if the inputted certificate  $\psi$  is one of the trusted certificates in *store*. If so, the only additional check required is to ensure that  $\psi$  did not expire (line 1). Next, the algorithm ensures that the inputted certificate is certified correctly, has a valid certificate chain  $\xi$  and both  $\psi$  and  $\xi$  were not revoked. To that end, the algorithm ensures that (line 2):



- (1) The auxiliary information  $\omega$  contains a certificate chain  $\xi$  which starts with  $\psi$ .
- (2)  $\xi$  is well structured and all the certificates are: correctly certified, did not expire and they were not revoked.
- (3)  $\xi$  contains a trusted certificate, i.e., in *store*.

To check that the certificates in  $\xi$  were not revoked, the algorithm ensures that each certificate  $\psi'$  has a *valid* CRL in  $\omega$  which does include  $\psi'$  (line 11). Namely, a valid CRL is:

- (1) A valid certificate of type 'CRL' (line 9).
- (2) Was issued by the same issuer of the certificate (line 10).
- (3) The CRL certificate is valid. Since the CRL certificate is a certificate like any other certificate, this is checked using a recursive call to  $\text{PKIX}_C.\text{Valid}$  (line 12).

If any of the aforementioned checks fail, the algorithm outputs  $\perp$ , otherwise, the certificate is considered valid and therefore, the algorithm outputs  $\top$ .

$\text{PKIX}_C.\text{Chains}$  (Algorithm 7). In PKIX, all the certificate chains, as well as CRL, are given in the  $\omega$  input. Therefore the Chains function simply returns the  $\omega$  input.

$\text{PKIX}_C.\text{GetCRL}$  (Algorithm 9). First, the algorithm removes all the certificates that expired from the current CRL (line 1). Then, it generates an updated CRL, i.e., defines a data to be certified  $tbc$  with CRL type (line 2), sets the issuer and validity period (lines 3-5) and the list of revoked certificates (line 6). Finally, the algorithm signs  $tbc$  using  $C.\text{Certify}(st, clk, tbc)$  and the CA's secret signing key (line 7) and outputs the CRL certificate along with the local state (lines 8-9).

## 5.2 Certificate Transparency (CT)

We now describe Certificate Transparency (CT), an extension to PKIX defined in [34, 36], as a transparent PKI scheme (Definition 5). Specifically, in Algorithms 11-24 we present the CT algorithms, utilizing, as subroutines, the PKIX algorithms (Algorithms 4-8).

*Convention:* CT uses a few timing parameters ( $\Delta_r$ ,  $\Delta_{MMD}$ ,  $\Delta_{clk}$  and  $\Delta_{com}$ ), known to all parties. Formally, these should be part of the scheme name, but we omit them to avoid clutter.

*Entities and state.* In CT, we consider, in addition to certificate authorities (with the same local state information as defined for PKIX), also two additional types of entities (in  $N$ ): *loggers* and *monitors*. Each logger maintains a log of certificates (*st.mylog*). Each monitor keeps the list of loggers it oversees (in *st.loggers*), and for each of these loggers, the set of certificates they logged (*st.log*).

*SCTs and pre-certificates.* CT certificates use the same certification schemes and certificate fields as in PKIX, and adds a new type of certified-objects, the *signed certificate timestamp (SCT)*. In CT, a certificate is valid only together with two<sup>13</sup> valid SCTs, from different loggers. A CA wishing to issue a CT certificate, typically first certifies and submits to loggers, using the  $\text{CT}.\text{AddPreChain-Req}$

operation, a so-called *precertificate*, which is a PKIX-valid<sup>14</sup> certificate. After validating the (pre)certificate, the loggers add it to the log, certify and return the corresponding SCT. We consider the SCT as a field of the CT certificate; the  $\text{CT}.\text{Valid}$  operation requires a valid SCT field<sup>15</sup>.

*Implementation.* Algorithms 11-24 present the implementation of CT. The operations defined include these required from any PKI (Definition 4) and from any transparent PKI (Definition 5), and operations<sup>16</sup> defined in CT specifications [34]:  $\text{GetSTH-Req}$ ,  $\text{GetSTH-Resp}$ ,  $\text{AddPreChain-Req}$ ,  $\text{GetEntries-Req}$  and  $\text{GetEntries-Resp}$ . One more operation, not explicitly defined in the CT specifications [34], is  $\text{CT}.\text{Wakeup}$ . The Wake-up operation is required to perform scheduled operations defined in [34], such as updating the log's root signature and periodically retrieving the updated logs by monitors.

To schedule a Wakeup operation at time  $t$  with input  $data$ , the protocol outputs ('Wake-up',  $t$ ,  $data$ ). The  $\pi_{\Delta_{clk}}^{\text{Wake-up}}$  model predicate ensures that a  $\text{CT}.\text{Wakeup}$  operation is invoked, with input  $data$ , on the entity requesting Wake-up, at real time within window of  $\Delta_{clk}$  before and after time  $t$ , i.e., in  $[t - \Delta_{clk}, t + \Delta_{clk}]$ .

We describe the implementation, focusing on the differences from PKIX. The revocation mechanism used in PKIX and CT is identical, hence, the  $\text{CT}.\text{Revoke}$  (Algorithm 13) and  $\text{CT}.\text{GetCRL}$  (Algorithm 14) algorithms simply invoke their PKIX counterparts. Similarly,  $\text{CT}.\text{Chains}$  (Algorithm 15) does not introduce any new functionality, and simply invokes its PKIX counterpart.

$\text{CT}.\text{Init}$  (Algorithm 11). The algorithm adds the initialization details for loggers and monitors. First, the algorithm calls  $\text{PKIX}.\text{Init}$  to handle any of PKIX related details (line 1). Then, the algorithm performs additional operations for loggers and monitors. First, the algorithm requests a Wake-up within  $\Delta_{MMD}$  (line 3). Finally, for monitors, the algorithm initializes an empty set of monitored logs and stores the set of trusted loggers' certificates, while for loggers, the algorithm initializes an empty maintained log (lines 4-9).

$\text{CT}.\text{Certify}$  (Algorithm 12). In addition to the 'regular' certificates issued in PKIX, CT also allows to issue pre-certificates. To that end,  $\text{CT}.\text{Certify}$  checks if the inputted data to be certified contains the 'pre-certificate' type (line 1), and if so, adds the 'poison' extension as described in RFC6962 [34] (line 2). In any case, the certificate is generated using the  $\text{PKIX}.\text{Certify}$  algorithm (line 4).

$\text{CT}.\text{Valid}$  (Algorithm 16). First, use  $\text{PKIX}.\text{Valid}$  to validate that  $\psi$  is a PKIX-valid certificate. From line 3, we verify that  $\psi$  is submitted with two valid SCTs, issued for  $\psi$ , from different trusted loggers (i.e., with certificates from *store.logger*).

$\text{CT}.\text{Monitor}$  (Algorithm 17). The algorithm receives a logger identifier  $\iota$  and adds it to the list of monitored logs (line 1). As a result, starting from the next time the  $\text{CT}.\text{Wakeup}$  algorithm is invoked,  $\iota$  will be monitored along with the other loggers.

<sup>13</sup>The decision of which logger-SCTs are sufficient can differ among relying parties; for simplicity, and since it appears irrelevant to our analysis, we assume two SCTs from different loggers are required.

<sup>14</sup>Precertificates may include a critical 'poison' extension that makes them PKIX-invalid, to prevent their use with relying parties that accept PKIX certificates. For simplicity, we ignore this, and consider them PKIX-valid.

<sup>15</sup>The RFC allows the SCT to be encoded as a certificate extension or to be used as a separate data structure in certificate validation.

<sup>16</sup>We keep the operation names from [34].

**Algorithm 4** PKIX<sub>C</sub>.Init: initialization

**Input:** local state  $st$ , local clock  $clk$  and parameters  $params$ , including identifier  $params.i$ , security parameter  $params.k$ , and acceptable revocation delay  $params.\Delta_r$ .

**Output:** updated state and self-signed certificate  $\psi$ .

**procedure** PKIX<sub>C</sub>.Init( $st, clk, params$ )

- 1:  $st.\Delta_r \leftarrow params.\Delta_r$  ▷ Save  $\Delta_r$  (validity interval for CRLs)
- 2:  $st.i \leftarrow params.i$  ▷ Save identity ( $i$ )
- 3:  $(st.sk, st.pk) \leftarrow C.KG(params.k)$  ▷ Generate signature key pair
- 4:  $tbc \leftarrow ('pk', st.pk) + params.tbc$  ▷ Set public verification key
- 5:  $\psi \leftarrow \{PKIX_C.Certify(tbc)\}$  ▷ Issue self-signed certificate  
Initialize empty sets of certificates  
certified and revoked by the current entity
- 6:  $(st.certs, st.CRL) \leftarrow \perp$
- 7: **return** ( $st, \psi$ )

**end procedure**

**Algorithm 5** PKIX<sub>C</sub>.Certify: certificate issuance

**Input:** local state  $st$ , local clock  $clk$  and data to be certified  $tbc$ .

**Output:** updated state and certificate  $\psi$ .

**procedure** PKIX<sub>C</sub>.Certify( $st, clk, tbc$ )

- 1: **if**  $tbc.subject = st.i$  **then return** ( $st, \perp$ ) ▷ Self-signed cert only from Init!
- 2:  $tbc.type \leftarrow 'public-key'$  ▷ Certify only public-key certs
- 3:  $\sigma \leftarrow C.Certify_{st.sk}(tbc)$  ▷ Sign certificate
- 4:  $\psi \leftarrow (tbc, \sigma)$  ▷ Certified data
- 5:  $st.certs += \psi$  ▷ Store locally
- 6: **return** ( $st, \psi$ )

**end procedure**

**Algorithm 6** PKIX<sub>C</sub>.Revoke: certificate revocation

**Input:** local state  $st$ , local clock  $clk$  and certificate  $\psi$ .

**Output:** updated state,  $\top$  if revocation completed successfully, and  $\perp$  otherwise.

**procedure** PKIX<sub>C</sub>.Revoke( $st, clk, \psi$ )

- 1: **if**  $\psi \in st.certs$  **then** ▷  $\psi$  was issued by this CA
- 2:  $st.CRL += \psi.serial$  ▷ Add to local CRL
- 3: **return** ( $st, \top$ ) ▷  $\psi$  revoked successfully
- 4: **end if**
- 5: **return** ( $st, \perp$ ) ▷ Failure

**end procedure**

**Algorithm 7** PKIX<sub>C</sub>.Chains:

**Input:** local clock  $clk$ , certificate  $\psi$ , root store  $store$  and auxiliary information  $\omega$ .

**Output:** certificate chains for  $\psi$ .

**procedure** PKIX<sub>C</sub>.Chains( $clk, \psi, store, \omega$ )

- 1: **return**  $\omega$  ▷  $\omega$  contains all the relevant data

**end procedure**

**Algorithm 8** PKIX<sub>C</sub>.Valid: certificate validation

**Input:** local clock  $clk$ , certificate  $\psi$ , root store  $store$  and auxiliary information  $\omega$ .

**Output:**  $\top$  if  $\psi$  is a valid certificate, and  $\perp$  otherwise.

**procedure** PKIX<sub>C</sub>.Valid( $clk, \psi, store, \omega$ )

- 1: **if**  $\psi \in store \wedge clk \in [\psi.from, \psi.to]$  **then return**  $\top$  ▷ Root certs are valid
- 2: **if** 
$$\left[ \begin{array}{l} \exists \xi \in \omega \text{ s.t. } \xi[0] = \psi \wedge \\ \text{IsCHAIN}^C(\xi) \wedge \\ (\exists i)(\xi[i] \in store) \wedge \\ (\forall \psi' \in \xi)(clk \in [\psi'.from, \psi'.to]) \wedge \\ \text{CHAINISNOTREVOKED}(\xi) \end{array} \right]$$
 ▷  $\omega$  has a chain  $\xi$  for  $\psi$   
▷  $\xi$  is correctly structured  
▷  $\xi$  ends in a trusted CA  
▷  $\xi$ 's certs did not expire  
▷  $\xi$ 's certs are non-revoked
- 3: **then return**  $\top$
- 4: **end if**
- 5: **return**  $\perp$  ▷ Otherwise,  $\psi$  is invalid

**end procedure**

- 6: **procedure** CHAINISNOTREVOKED( $\xi$ )
- 7: **return**  $\forall \psi' \in \xi$  ▷ For every certificate in  $\xi$
- 8:  $\exists \xi_{CRL} \in \omega$  s.t.: ▷ Exists a CRL cert chain
- 9:  $\xi_{CRL}[0].type = CRL \wedge$  ▷  $\xi_{CRL}[0]$  is a CRL
- 10:  $\xi_{CRL}[0].issuer = \psi'.issuer \wedge$  ▷ Issued by  $\psi'$ .issuer
- 11:  $\psi'.serial \notin \xi_{CRL}[0].CRL \wedge$  ▷  $\psi'$  is not in the CRL
- 12:  $PKIX_C.Valid(clk, \xi_{CRL}[0], store, \omega - \{\xi\})$  ▷ And  $\xi_{CRL}[0]$  is valid
- 13: **end procedure**

**Algorithm 9** PKIX<sub>C</sub>.GetCRL: retrieve CRL

**Input:** local state  $st$  and local clock  $clk$ .

**Output:** updated state and current CRL.

**procedure** PKIX<sub>C</sub>.GetCRL( $st, clk$ )

- 1:  $st.CRL -= \left\{ \begin{array}{l} \psi.serial \in st.CRL \\ \text{s.t. } \psi.to < clk \end{array} \right\}$  ▷ Remove expired certificates from CRL
- 2:  $tbc \leftarrow (type, 'CRL')$  ▷ Certificate of type CRL
- 3:  $tbc.issuer \leftarrow st.i$  ▷ The issuer of the CRL
- 4:  $tbc.from \leftarrow clk$  ▷ When was issued
- 5:  $tbc.to \leftarrow clk + st.\Delta_r$  ▷ Valid until
- 6:  $tbc.CRL \leftarrow st.CRL$  ▷ CRL information
- 7:  $\sigma \leftarrow C.Certify_{st.sk}(tbc)$  ▷ Sign certificate
- 8:  $\psi_{CRL} \leftarrow (tbc, \sigma)$  ▷ Certified data
- 9: **return** ( $st, \psi_{CRL}$ )

**end procedure**

**Figure 1: PKIX (X.509 version 3 with CRL) implementation**

CT.Lookup (Algorithm 18). The algorithm returns the set of certificates that were issued to the given *subject*, based on the local copies of the logs that the monitor maintains (lines 2-4).

CT.Wakeup (Algorithm 19). First, the algorithm asks for another Wake-up after another  $\Delta_{MMD}$  interval (line 1). In a logger, the algorithm updates the *signed tree hash* over the local log maintained by the logger (lines 3-9), and stores the STH locally (line 10) so it can be retrieved by monitors in CT.GetSTH-Req (Algorithm 21). In a monitor, CT.Wakeup outputs a request to get the latest STH for each of the loggers overseen by this monitor (line 12).

CT.AddPreChain-Req (Algorithm 20). Add a (pre)certificate to the log. The algorithm ensures that the (pre)certificate is a valid PKIX certificate, not already logged, and issued before at most the maximal delay and clock drift (line 1). Then, the algorithm adds the certificate to the local log (line 2), generates a Signed Certificate Timestamp (SCT; lines 3-8) and outputs the SCT (line 9).

CT.GetSTH-Req (Algorithm 21). The algorithm outputs the latest STH (line 1), that was generated in CT.Wakeup (Algorithm 19).

CT.GetSTH-Resp (Algorithm 22). The algorithm uses the logger's certificate, which is stored in the monitor's local state (line 3), to check that the given STH was indeed signed by the logger, and that

new certificates, in the STH, were added to the log (line 4). If so, the monitor stores the new STH in a temporary location (line 5), until the newly logged certificates will be retrieved and the STH can be verified as a valid root that reflects the addition of the newly logged certificates. To that end, the algorithm outputs a ‘GetEntries’ request, to be sent to the logger, to which the logger should respond by sending the entries (of the newly logged certificates) (line 6).

CT.GetEntries-Req (*Algorithm 23*). The algorithm is asked to output a subset of the logged certificates, starting from index *start* until index *end*. Thus, the algorithm takes the relevant subset (line 4), generates a signature over it (lines 1-6) and outputs it (line 7).

CT.GetEntries-Resp (*Algorithm 24*). The algorithm uses the logger’s certificate, which is stored in the monitor’s local state (line 2) to check that the input was indeed signed by the logger (line 3). Then, the algorithm calculates the updated tree hash by adding the newly logged certificates to the local copy of the log (lines 4-5). If the STH that was stored temporarily in CT.GetSTH-Resp (*Algorithm 22*) indeed matches the STH that was calculated in line 5, the monitor inspects the newly logged certificates according to its implementation (line 9). Finally, the STH and new certificates are added to the local state (lines 7-8).

## 6 SECURITY ANALYSIS

In this section, we analyze the security of PKIX with and without the CT extension against the security requirements defined in Section 4. First, we define the model specifications (assumptions) in §6.1, then, we show which requirements each PKI satisfies under these model assumptions in §6.2.

### 6.1 Models

Different protocols assume different models; even the same protocol may assume different models in order to ensure different requirements. The  $\text{PKIX}_C$  and  $\text{CT}_C^{M\mathcal{T}^h}$  protocols assume models based on the following predicates, defined in [16]; these model predicates formalize standard adversary, communication and synchronization assumptions.

- The  $\pi_{\Delta_{clk}}^{\text{Drift}}$  predicate ensures clock drifts from real time are bounded by  $\Delta_{clk}$ .
- The  $\pi_{\Delta_{clk}}^{\text{Wake-up}}$  predicate ensures *Wake-up* operations are invoked within  $\Delta_{clk}$  of the requested (real) time.
- The  $\pi_{\Delta_{com}}^{\text{Com}}$  predicate ensures reliable communication between non-faulty parties, with delays bounded by  $\Delta_{com}$ . Namely:
  - (1) Whenever an entity  $\iota$  includes in its output a triple of the form  $(\alpha\text{-req}, \iota', x)$ , where  $\alpha\text{-req}$  is one of the request operations of CT, then, within  $\Delta_{com}$ , there is a  $\alpha\text{-req}$  event in  $\iota'$  with input  $x$ .
  - (2) Whenever an entity  $\iota'$  outputs a pair of the form  $(\alpha\text{-resp}, y)$  in the output of an  $\alpha\text{-req}$  operation, and this  $\alpha\text{-req}$  was invoked by some entity  $\iota$ , then, within  $\Delta_{com}$ , there is a  $\alpha\text{-resp}$  event in  $\iota$  with input  $x$ .
- The  $\pi^F$  predicate ensures that all entities follow the protocol correctly, except for entities in T.F (the set of faulty entities outputted by the adversary).

- The  $\pi^{\text{F,HL}}$  predicate extends  $\pi^F$  to also ensure *Honest Loggers (HL)*. Namely, if any benign monitor has a ‘Monitor’ operation for logger  $\iota$ , then  $\iota$  has to be benign.

To illustrate, we include the  $\pi_{\Delta_{clk}}^{\text{Drift}}$  model predicate from [16] in Algorithm 10.

---

#### Algorithm 10 The $\pi_{\Delta_{clk}}^{\text{Drift}}$ model predicate

---

1:	<b>return</b> $\forall \hat{e} \in \{1, \dots, T.e\}$ :	▷ For each event
2:	$ T.\text{clk}[\hat{e}] - T.\tau[\hat{e}]  \leq \Delta_{clk} \wedge$	▷ Local clock within $\Delta_{clk}$ drift from real time
3:	<b>if</b> $\hat{e} \geq 2$ <b>then</b> $T.\tau[\hat{e}] \geq T.\tau[\hat{e} - 1]$	▷ The real time difference is monotonically increasing

---

### 6.2 Security Analysis

**THEOREM 4.** *If  $C$  is a secure certificate scheme, then  $\text{PKIX}_C$  satisfies the following requirements:*

- *Existential unforgeability, under model  $(\pi^F)$ .*
- *Accountability, under the trivial (always true) model.*
- *Weak  $\Delta_{\text{Rev}}$ -Revocation, where  $\Delta_{\text{Rev}} = \Delta_r + \Delta_{clk}$ , under model  $(\pi_{\Delta_r}^F \wedge \pi_{\Delta_{clk}}^{\text{Drift}})$ .*

**PROOF.** We proved existential unforgeability in Lemma 2. We prove accountability and weak  $\Delta_{\text{PKIX}}$ -revocation in Lemmas 10 and 12, respectively. The proofs are presented in appendices B and C, respectively; we present their sketches below.

*Accountability.* Assume, to the contrary, that  $\text{PKIX}_C$  fails to ensure accountability. Namely, with non-negligible probability, the accountability predicate returns  $\perp$ , i.e., for a random transcript  $T$ :

$$\pi_{\text{ACC}}(T) = \perp \quad (4)$$

Following the implementation of  $\pi_{\text{ACC}}$ , let:

$$(clk, \psi, store, \omega) \leftarrow T.out_{\mathcal{A}} \quad [\text{Alg. 3 line 2}]$$

then, following Equation (4):

$$\mathcal{P}.\text{Valid}(clk, \psi, store, \omega) = \top \quad [\text{Alg. 3 line 3}]$$

and yet:

$$\nexists \xi \in \mathcal{P}.\text{Chains}(clk, \psi, store, \omega), j < |\xi| \text{ s.t.} \quad [\text{Alg. 3 line 4}]$$

$$\xi[j] = \psi \wedge \text{ISCHAIN}(\xi) \wedge \mathcal{P}.\text{Valid}(clk, \xi[j+1], store, \omega) \quad [\text{Alg. 3 line 5}]$$

However, according to the implementation of  $\text{PKIX}_C.\text{Valid}$  (*Algorithm 8*), the algorithm only outputs  $\top$  if  $\text{PKIX}_C.\text{Chains}$  outputs a valid certificate chain for the same inputs, showing a contradiction. Hence,  $\text{PKIX}_C$  satisfies accountability.

*Weak  $\Delta_{\text{PKIX}}$ -Revocation.* We prove in the following steps:

- (1) We first define the  $\overline{\text{Exp}}_{\mathcal{A}, \mathcal{P}}^{\text{Forge}}$  (*params*) experiment, based on the  $\pi_{\text{Weak}-\Delta_{\text{Rev}}}$  predicate (*Definition 8*). The  $\overline{\text{Exp}}_{\mathcal{A}, \mathcal{P}}^{\text{Forge}}$  (*params*) experiment executes an adversary  $\mathcal{A}$  with the PKI scheme  $\mathcal{P}$ . Adversary  $\mathcal{A}$  is given public validation key  $pk$  but not the corresponding signing/certifying key  $sk$ , and can use an oracle to the  $C.\text{Certify}_{sk}$  operation. Adversary  $\mathcal{A}$  ‘wins’ if it outputs a certificate  $\psi$  which is valid (according to  $pk$ ), without  $\mathcal{A}$  asking the oracle to certify  $\psi$ . See *Definition 18*.

- (2) After that, we show that the existence of an adversary  $\mathcal{A}$  that 'wins' the  $\overline{\text{Exp}}_{\mathcal{A}, \text{PKIX}_C}^{\text{Forge}}(\text{params})$  game with non-negligible probability contradicts the security of the certification scheme  $C$ . See Claim 11.
- (3) We conclude by showing that if  $\text{PKIX}_C$  does not satisfy the weak  $\Delta_{\text{Rev}}$ -Revocation requirement, we can construct an adversary  $\mathcal{A}_{\text{Rev}}$  that wins the  $\overline{\text{Exp}}_{\mathcal{A}_{\text{Rev}}, \text{PKIX}_C}^{\text{Forge}}(\text{params})$  game with non-negligible probability. See Lemma 12.  $\square$

**THEOREM 5.** *Let  $C$  be a secure certificate scheme and let  $h$  be an ACR (any collision resistant) keyed hash function. Then,  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$  satisfies the following requirements:*

- Existential unforgeability, under model  $(\pi^F)$ .
- Accountability, under the trivial (always true) model.
- Weak  $\Delta_{\text{Rev}}$ -Revocation, where  $\Delta_{\text{Rev}} = \Delta_r + \Delta_{\text{clk}}$ , under model  $(\pi^F \wedge \pi_{\Delta_{\text{clk}}}^{\text{Drift}})$ .

**PROOF.** For existential unforgeability, see Lemma 2. For accountability and weak  $\Delta_{\text{Rev}}$ -Revocation, the implementation related to certificates issuance and revocation in  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$  is identical to  $\text{PKIX}_C$ , since the implementation in Algorithms 11-24 directly invokes the corresponding  $\text{PKIX}_C$  algorithms. The additional implementation in  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$  only regards transparency, and in particular, can only issue SCTs, STHs and signatures over newly added certificates to the log, but cannot issue public-key certificates or CRL certificates. Therefore, the proofs that  $\text{PKIX}_C$  satisfies these properties also hold for  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$ ; otherwise, the same adversary that breaks these properties in  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$  also breaks them in  $\text{PKIX}_C$ .  $\square$

Next, Theorem 6 proves that CT ensures also the  $\Delta_{\text{Tra}}$ -Transparency requirement under the *CT honest-logger model*<sup>17</sup>, which we define as:

$$\mathcal{M}^{\text{CT-HL}} \equiv \left( \pi^{\text{F,HL}} \wedge \pi_{\Delta_{\text{clk}}}^{\text{Drift}} \wedge \pi_{\Delta_{\text{com}}}^{\text{Com}} \wedge \pi_{\Delta_{\text{clk}}}^{\text{Wake-up}} \right) \quad (5)$$

**THEOREM 6.** *Let  $C$  be a secure certificate scheme,  $h$  be an ACR keyed hash,  $\mathbb{N}$  be a set of entities, and  $\Delta_{\text{Tra}} = 2 \cdot \Delta_{\text{MMD}} + 3 \cdot \Delta_{\text{clk}} + 5 \cdot \Delta_{\text{com}}$ . Then,  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$  satisfies the  $\Delta_{\text{Tra}}$ -Transparency requirement, under model  $\mathcal{M}^{\text{CT-HL}}$ .*

**PROOF.** Suppose  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$  does not satisfy the  $\Delta_{\text{Tra}}$ -Transparency requirement under model  $\mathcal{M}^{\text{CT-HL}}$ . Namely, there exists an adversary  $\mathcal{A}$  that produces, with non-negligible probability, a certificate  $\psi$  that is valid at time  $\tau \geq \psi.\text{from} + \Delta_{\text{Tra}}$ , such that a benign monitor  $\iota_M$  who is monitoring a logger of  $\psi$  since  $\tau - \Delta_{\text{Tra}}$  is, at time  $\tau$ , unaware of  $\psi$ . Let us denote this logger by  $\ell$ ; from the honest-logger assumption of  $\mathcal{M}^{\text{CT-HL}}$ , we know that  $\ell$  is benign.

Since  $\psi$  is valid at time  $\tau$ , then, from lines 2-5 of  $\text{CT.Valid}$  (Alg. 16), it follows that  $\psi$  must contain an SCT, denoted  $\psi_{\text{sct}}$ , which is properly certified by the benign logger  $\ell$ . Furthermore,  $\psi_{\text{sct}}.\text{timestamp} - \Delta_{\text{clk}} \leq \tau$ .

From Theorem 5,  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$  satisfies existential unforgeability. Hence,  $\psi_{\text{sct}}$  must have been issued by logger  $\ell$ , using its self-signed

key.  $\text{CT.AddPreChain-Req}$  (Algorithm 20) is the only operation where  $\ell$  issues an SCT certificate, and the value of  $\psi_{\text{sct}}.\text{timestamp}$  would be the value of  $\ell$ 's clock at the time. Let  $\tau_{\text{SCT}}$  denote the real time when  $\ell$  issued the SCT; from the  $\pi_{\Delta_{\text{clk}}}^{\text{Drift}}$  model predicate,  $\tau_{\text{SCT}} \leq \psi_{\text{sct}}.\text{timestamp} + \Delta_{\text{clk}}$ . On the other hand,  $\ell$  would issue the SCT only if its clock upon receiving  $\psi$  is at most  $\psi_{\text{sct}}.\text{timestamp} \leq \psi.\text{from} + \Delta_{\text{clk}} + \Delta_{\text{com}}$ .

Since  $\ell$  is benign, it asks for Wake-up every  $\Delta_{\text{MMD}}$  (see the  $\text{CT.Init}$  and  $\text{CT.Wakeup}$  operations); by the  $\pi_{\Delta_{\text{clk}}}^{\text{Wake-up}}$  model predicate, the operation will occur after (at most)  $\Delta_{\text{MMD}} + \Delta_{\text{clk}}$ , i.e., at or before  $\tau_{\text{SCT}} + \Delta_{\text{MMD}} + \Delta_{\text{clk}}$ . At this time,  $\ell$  updates its STH to reflect the logging of  $\psi$  (and possibly other certificates). The STH is signed by  $\ell$ , and includes the Merkle-tree (accumulator) digest of the log of all certificates logged by  $\ell$  until this time.

The benign monitor  $\iota_M$  that monitors  $\ell$  also invokes  $\text{CT.Wakeup}$  every  $\Delta_{\text{MMD}}$ . In every invocation of  $\text{CT.Wakeup}$ ,  $\iota_M$  requests a  $\text{CT.GetSTH-Req}$  operation from every logger that  $\iota_M$  is monitoring, which includes  $\ell$ . From the  $\pi_{\Delta_{\text{com}}}^{\text{Com}}$  model, the  $\text{CT.GetSTH-Req}$  operation in  $\ell$  occurs at most  $\Delta_{\text{com}}$  afterwards. Namely, a  $\text{CT.GetSTH-Req}$  operation, invoked by  $\iota_M$ , must happen in  $\ell$  when  $\ell$ 's clock is at or before  $\tau_{\text{SCT}} + \Delta_{\text{MMD}} + \Delta_{\text{clk}} + \Delta_{\text{com}}$  (including the possible impact of the clock bias and of the network delay).

From the  $\pi_{\Delta_{\text{com}}}^{\text{Com}}$  predicate, the corresponding  $\text{CT.GetSTH-Resp}$  event in  $\iota_M$  occurs within  $\Delta_{\text{com}}$ . As shown in  $\text{CT.GetSTH-Resp}$  (Algorithm 22), the monitor  $\iota_M$  would identify that new certificates were logged and ask for them, i.e., invoke  $\text{CT.GetEntries-Req}$  at  $\ell$ ; and the logger  $\ell$  immediately responds by sending the certificates, including  $\psi$ . This round-trip of communication will take at most  $2 \cdot \Delta_{\text{com}}$ ; hence,  $\iota_M$  receives  $\psi$  at or before time  $\tau_\psi$ , defined as:

$$\begin{aligned} \tau_\psi &\equiv \tau_{\text{SCT}} + 2 \cdot \Delta_{\text{MMD}} + \Delta_{\text{clk}} + 4 \cdot \Delta_{\text{com}} \\ &\leq \psi_{\text{sct}}.\text{timestamp} + 2 \cdot \Delta_{\text{MMD}} + 2 \cdot \Delta_{\text{clk}} + 4 \cdot \Delta_{\text{com}} \\ &\leq \psi.\text{from} + 2 \cdot \Delta_{\text{MMD}} + 3 \cdot \Delta_{\text{clk}} + 5 \cdot \Delta_{\text{com}} \\ &= \psi.\text{from} + \Delta_{\text{Tra}} \end{aligned} \quad (6)$$

Which contradicts the assumption that  $\mathcal{A}$  produces, with non-negligible probability, a certificate  $\psi$  that is valid at time  $\tau \geq \psi.\text{from} + \Delta_{\text{Tra}}$  but not known to  $\iota_M$ .  $\square$

Finally, we show that CT fails to satisfy the  $\Delta_{\text{Tra}}$ -Transparency requirement, if we allow rogue loggers, i.e., under model:

$$\mathcal{M}_{\Delta_{\text{clk}}, \Delta_{\text{com}}} \equiv \left( \pi^{\text{F}} \wedge \pi_{\Delta_{\text{clk}}}^{\text{Drift}} \wedge \pi_{\Delta_{\text{com}}}^{\text{Com}} \wedge \pi_{\Delta_{\text{clk}}}^{\text{Wake-up}} \right) \quad (7)$$

**CLAIM 7.** *Let  $C$  be a secure certificate scheme,  $h$  be an ACR keyed hash,  $\mathbb{N}$  be a set of entities, and  $\Delta_{\text{Tra}} = 2 \cdot \Delta_{\text{MMD}} + 3 \cdot \Delta_{\text{clk}} + 5 \cdot \Delta_{\text{com}}$ . Then,  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$  does not satisfy the  $\Delta_{\text{Tra}}$ -Transparency correctness requirement under model  $\mathcal{M}_{\Delta_{\text{clk}}, \Delta_{\text{com}}}$ .*

*Argument.* Let  $\mathcal{A}$  be an adversary that controls a logger  $\ell \in T.F$  for a transcript  $T$ .  $\mathcal{A}$  instructs  $\ell$  to execute  $\text{CT}_C^{\mathcal{M}^{\mathcal{T}^h}}$  as described in §5.2, with the single change of eliminating line 2 from Algorithm 20. As a result,  $\ell$  generates an SCT to every valid certificate, yet  $\ell$  never logs these certificates, and therefore, never informs monitors about them. This means that despite the fact that there is a valid SCT from  $\ell$  for each of the certificates, none of these certificates will be known to monitors. Hence, the existence of such  $\mathcal{A}$  means that

<sup>17</sup>In Claim 7, below, we show that the honest-logger assumption is necessary.

the  $\Delta_{\text{Tra}}$ -Transparency requirement is not satisfied under model  $\mathcal{M}_{\Delta_{\text{clk}}, \Delta_{\text{com}}}$ .  $\square$

## 7 CONCLUSIONS AND FUTURE WORK

We presented specifications for secure PKI schemes, and applied them to prove the security of two applied PKI schemes: PKIX and Certificate Transparency (CT). Future work may apply our specifications to other PKI schemes, present additional specifications for PKI schemes (capturing additional properties, e.g., privacy properties), and apply a similar approach to other practical cryptographic protocols, e.g., blockchain protocols.

## REFERENCES

- [1] Hadi Asghari, Michel Van Eeten, Axel Armbak, and Nico ANM van Eijk. 2013. Security Economics in the HTTPS Value Chain. In *Twelfth Workshop on the Economics of Information Security (WEIS 2013)*, Washington, DC.
- [2] Louise Axon and Michael Goldsmith. 2017. PB-PKI: A Privacy-aware Blockchain-based PKI. In *SECRYPT*.
- [3] David Basin, Cas Cremers, Tiffany Hyun-Jin Kim, Adrian Perrig, Ralf Sasse, and Pawel Szalachowski. 2014. ARPki: Attack Resilient Public-Key Infrastructure. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 382–393.
- [4] Alexandra Boldyreva, Marc Fischlin, Adriana Palacio, and Bogdan Warinschi. 2007. A Closer Look at PKI: Security and Efficiency. In *International Workshop on Public Key Cryptography*. Springer, 458–475.
- [5] Johannes Braun. 2015. *Maintaining Security and Trust in Large Scale Public Key Infrastructures*. Ph.D. Dissertation. Technische Universität.
- [6] Johannes Braun, Franziskus Kiefer, and Andreas Hülsing. 2013. Revocation & Non-Repudiation: When the first destroys the latter. In *European Public Key Infrastructure Workshop*. Springer, 31–46.
- [7] Ran Canetti, Daniel Shahaf, and Margarita Vald. 2016. Universally Composable Authentication and Key-exchange with Global PKI. In *Public-Key Cryptography–PKC 2016*. Springer, 265–296.
- [8] BLUE BOOK CCITT. 1988. Recommendations X. 509 and ISO 9594-8. *Information Processing Systems-OSI-The Directory Authentication Framework (Geneva: CCITT)*.
- [9] David Cooper, Stefan Santesson, Stephen Farrell, Sharon Boeyen, Russell Housley, and William Polk. 2008. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. Tech. rep.
- [10] Peter Eckersley. 2012. Sovereign Key Cryptography for Internet Domains. <https://git.ietf.org/?p=sovereign-keys.git;a=blob;f=sovereign-key-design.txt;hb=HEAD>. (2012).
- [11] Electronic Frontier Foundation (EFF). [n. d.] The EFF SSL Observatory. Retrieved May 30, 2019 from <https://www.eff.org/observatory>.
- [12] Conner Fromknecht, Dragos Velicanu, and Sophia Yakoubov. 2014. A Decentralized Public Key Infrastructure with Identity Retention. *IACR Cryptology ePrint Archive*, 2014, 803.
- [13] Sebastian Gajek, Mark Manulis, Olivier Pereira, Ahmad-Reza Sadeghi, and Jörg Schwenk. 2008. Universally Composable Security Analysis of TLS. In *International Conference on Provable Security*. Springer, 313–327.
- [14] Shafi Goldwasser and Silvio Micali. 1984. Probabilistic Encryption. *Journal of Computer and System Sciences*, 28, 2, 270–299.
- [15] Shafi Goldwasser, Silvio Micali, and Ronald L Rivest. 1988. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on computing*, 17, 2, 281–308.
- [16] Amir Herzberg, Hemi Leibowitz, Ewa Syta, and Sara Wrótniak. 2021. MoSS: Modular Security Specifications framework. In *CRYPTO' 2021*. <https://eprint.iacr.org/2020/1040>, 33–63.
- [17] Amir Herzberg, Yosi Mass, Joris Mihaeli, Dalit Naor, and Yiftach Ravid. 2000. Access Control Meets Public Key Infrastructure, Or: Assigning Roles to Strangers. In *Proceeding 2000 IEEE Symposium on Security and Privacy*. S&P 2000. IEEE, 2–14.
- [18] Joel Hruska. 2015. Apple, Microsoft buck trend, refuse to block unauthorized Chinese root certificates. ExtremeTech. (Apr. 2015).
- [19] Jingwei Huang and David M Nicol. 2017. An anatomy of trust in public key infrastructure. *International Journal of Critical Infrastructures*, 13, 2-3, 238–258.
- [20] Tiffany Hyun-Jin Kim, Lin-Shung Huang, Adrian Perrig, Collin Jackson, and Virgil Gligor. 2013. Accountable Key Infrastructure (AKI): A Proposal for a Public-Key Validation Infrastructure. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 679–690.
- [21] Loren M Kohnfelder. 1978. *Towards a practical public-key cryptosystem*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [22] Murat Yasin Kubilay, Mehmet Sabir Kiraz, and Haci Ali Mantar. 2018. CertLedger: A new PKI model with Certificate Transparency based on blockchain. *arXiv preprint arXiv:1806.03914*.
- [23] Ben Laurie. 2014. Certificate transparency. *Communications of the ACM*, 57, 10, 40–46.
- [24] Ben Laurie and Emilia Kasper. 2012. Revocation Transparency. *Google Research*, September.
- [25] Ben Laurie, Adam Langley, Emilia Kasper, Eran Messeri, and Rob Stradling. 2019. Certificate transparency version 2.0. *IETF TRANS (Public Notary Transparency) WG, Internet-Draft*, (Nov. 2019).
- [26] Dimitrios Lekkas. 2003. Establishing and managing trust within the Public Key Infrastructure. *Computer Communications*, 26, 16, 1815–1825.
- [27] John Marchesini and Sean Smith. 2005. Modeling Public Key Infrastructure in the Real World. In *European Public Key Infrastructure Workshop*. Springer, 118–134.
- [28] Stephanos Matsumoto and Raphael M Reischuk. 2017. IKP: Turning a PKI Around with Decentralized Automated Incentives. In *Security and Privacy (SP), 2017 IEEE Symposium on*. IEEE, 410–426.
- [29] Ueli Maurer. 1996. Modelling a Public-Key Infrastructure. In *European Symposium on Research in Computer Security*. Springer, 325–350.
- [30] Marcela S Melara, Aaron Blankstein, Joseph Bonneau, Edward W Felten, and Michael J Freedman. 2015. CONIKS: Bringing Key Transparency to End Users. In *USENIX Security Symposium*, 383–398.
- [31] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk. 2008. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280 (Proposed Standard). RFC. Updated by RFCs 6818, 8398, 8399. Fremont, CA, USA: RFC Editor, (May 2008). doi: 10.17487/RFC5280.
- [32] M. Lepinski and S. Kent. 2012. An Infrastructure to Support Secure Internet Routing. RFC 6480 (Informational). RFC. Fremont, CA, USA: RFC Editor, (Feb. 2012). doi: 10.17487/RFC6480.
- [33] S. Santesson, M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams. 2013. X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP. RFC 6960 (Proposed Standard). RFC. Updated by RFC 8954. Fremont, CA, USA: RFC Editor, (June 2013). doi: 10.17487/RFC6960.
- [34] B. Laurie, A. Langley, and E. Kasper. 2013. Certificate Transparency. RFC 6962 (Experimental). RFC. Obsoleted by RFC 9162. Fremont, CA, USA: RFC Editor, (June 2013). doi: 10.17487/RFC6962.
- [35] E. Rescorla. 2018. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Proposed Standard). RFC. Fremont, CA, USA: RFC Editor, (Aug. 2018). doi: 10.17487/RFC8446.
- [36] B. Laurie, E. Messeri, and R. Stradling. 2021. Certificate Transparency Version 2.0. RFC 9162 (Experimental). RFC. Fremont, CA, USA: RFC Editor, (Dec. 2021). doi: 10.17487/RFC9162.
- [37] Steven B Roosa and Stephen Schultze. 2010. The "Certificate Authority" Trust Model for SSL: A Defective Foundation for Encrypted Web Traffic and a Legal Quagmire. *Intellectual property & technology law journal*, 22, 11, 3.
- [38] Mark Dermot Ryan. 2014. Enhanced certificate transparency and end-to-end encrypted mail. In *NDSS*.
- [39] Wazan Ahmad Samer, Laborde Romain, Barrere Francois, and Benzekri AbdelMalek. 2011. A formal model of trust for calculating the quality of X. 509 certificate. *Security and Communication Networks*, 4, 6, 651–665.
- [40] Ewa Syta, Iulia Tamas, Dylan Visher, David Isaac Wolinsky, and Bryan Ford. 2015. Certificate Cothority: Towards Trustworthy Collective CAs. *Hot Topics in Privacy Enhancing Technologies (HotPETs)*, 7.
- [41] Ewa Syta, Iulia Tamas, Dylan Visher, David Isaac Wolinsky, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Ismail Khoffi, and Bryan Ford. 2016. Keeping Authorities "Honest or Bust" with Decentralized Witness Cosigning. In *Security and Privacy (SP), 2016 IEEE Symposium on*. Ieee, 526–545.
- [42] Pawel Szalachowski, Stephanos Matsumoto, and Adrian Perrig. 2014. PoliCert: Secure and Flexible TLS Certificate Management. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 406–417.
- [43] Alin Tomescu and Srinivas Devadas. 2017. Catena: Efficient Non-equivocation via Bitcoin. In *2017 38th IEEE Symposium on Security and Privacy (SP)*. IEEE, 393–409.
- [44] Dan Wendlandt, David G Andersen, and Adrian Perrig. 2008. Perspectives: Improving SSH-style Host Authentication with Multi-Path Probing. In *USENIX Annual Technical Conference*. Vol. 8, 321–334.
- [45] Jiangshan Yu, Vincent Cheval, and Mark Ryan. 2016. DTKI: A New Formalized PKI with Verifiable Trusted Parties. *The Computer Journal*, 59, 11, 1695–1713.
- [46] Michelle Zhou, Prithvi Bisht, and VN Venkatakrishnan. 2011. Strengthening XSRF Defenses for Legacy Web Applications Using Whitebox Analysis and Transformation. In *Information Systems Security*. Springer, 96–110.

## A CRYPTOGRAPHIC PRIMITIVES

### A.1 MoSS

First, Definition 10 defines the advantage of adversary  $\mathcal{A}$  against protocol  $\mathcal{P}$ , for given some specification predicate  $\pi$ .

DEFINITION 10 (Advantage of adversary  $\mathcal{A}$  against protocol  $\mathcal{P}$  for specification predicate  $\pi$ ). *Let  $\mathcal{A}, \mathcal{P}$  be algorithms and let  $\pi$  be a specification predicate. The advantage of adversary  $\mathcal{A}$  against protocol  $\mathcal{P}$  for specification predicate  $\pi$  is defined as:*

$$\epsilon_{\mathcal{A}, \mathcal{P}}^{\pi}(params) \stackrel{\text{def}}{=} \Pr \left[ \begin{array}{l} \pi(T) = \perp, \text{ where} \\ T \leftarrow \text{Exec}_{\mathcal{A}, \mathcal{P}}(params) \end{array} \right] \quad (8)$$

Now, Definition 11 defines an adversary that satisfies a given model  $\mathcal{M}$ .

DEFINITION 11 (Adversary  $\mathcal{A}$  satisfies model  $\mathcal{M}$ ). *Let  $\mathcal{A} \in PPT$ , and let  $\mathcal{M}$  be a model specification, i.e.,  $\mathcal{M} = (\pi, \beta)$ . We say that adversary  $\mathcal{A}$  satisfies model  $\mathcal{M}$ , denoted  $\mathcal{A} \models_{\text{poly}} \mathcal{M}$ , if for every protocol  $\mathcal{P} \in PPT$  and  $params \in \{0, 1\}^*$ , the advantage of  $\mathcal{A}$  against  $\mathcal{P}$  for  $\pi$  is at most negligibly greater than  $\beta(params)$ , i.e.:*

$$\mathcal{A} \models_{\text{poly}} \mathcal{M} \stackrel{\text{def}}{=} \left[ \begin{array}{l} (\forall \mathcal{P} \in PPT, params \in \{0, 1\}^*) : \\ \epsilon_{\mathcal{A}, \mathcal{P}}^{\pi}(params) \leq \beta(params) + \text{Negl}(params.1^k) \end{array} \right] \quad (9)$$

Finally, Definition 12 defines what it means for a protocol  $\mathcal{P}$  to satisfy a specific requirement under a specific model.

DEFINITION 12 (Protocol  $\mathcal{P}$  satisfies requirement  $\mathcal{R}$  under model  $\mathcal{M}$ ). *Let  $\mathcal{P} \in PPT$ , and let  $\mathcal{R}$  be a requirement specification, i.e.,  $\mathcal{R} = (\pi, \beta)$ . We say that protocol  $\mathcal{P}$  satisfies requirement  $\mathcal{R}$  under model  $\mathcal{M}$ , denoted  $\mathcal{P} \models_{\text{poly}}^{\mathcal{M}} \mathcal{R}$ , if for every PPT adversary  $\mathcal{A}$  that satisfies  $\mathcal{M}$  and every parameters  $params \in \{0, 1\}^*$ , the advantage of  $\mathcal{A}$  against  $\mathcal{P}$  for  $\pi$  is at most negligibly greater than  $\beta(params)$ , i.e.:*

$$\mathcal{P} \models_{\text{poly}}^{\mathcal{M}} \mathcal{R} \stackrel{\text{def}}{=} \left[ \begin{array}{l} (\forall \mathcal{A} \in PPT \text{ s.t. } \mathcal{A} \models_{\text{poly}} \mathcal{M}, params \in \{0, 1\}^*) : \\ \epsilon_{\mathcal{A}, \mathcal{P}}^{\pi}(params) \leq \beta(params) + \text{Negl}(params.1^k) \end{array} \right] \quad (10)$$

### A.2 Accumulator scheme and the Merkle Tree construction

We define an accumulator scheme and its security properties.

An accumulator receives as input an ordered set of elements  $X = \{x_i\}_{i=1}^k \in \mathcal{X}^k$ , and outputs their *digest*  $\Delta$ , as well as a *proof of inclusion (PoI)*  $\pi_i$  for each element  $accEl_i$  in the input. We will later define and discuss extensions, most notably, *dynamic/additive accumulators*, that allow to add elements to the accumulator in multiple operations.

DEFINITION 13 (An accumulator scheme). *An accumulator scheme  $\mathcal{ACC}$  is a tuple of algorithms  $(\mathcal{ACC}.Setup, \mathcal{ACC}.Add, \mathcal{ACC}.Verify)$  with element space  $\mathcal{ACC}.X$ , where the algorithms are defined as:*

- $\mathcal{ACC}.Setup(1^\lambda) \rightarrow pk$ : input security parameter  $\lambda \in \mathbb{N}$  outputs a public key  $pk$ .
- $\mathcal{ACC}.Add_{pk}(X) \rightarrow (\Delta, \Pi)$ : add (accumulate) a sequence of elements  $X = \{x_i\}_{i=1}^k \in \mathcal{ACC}.X^k$ . Output the digest  $\Delta$  of  $X$  and a set  $\Pi = \{\pi_i\}_{i=1}^k$  of proofs-of-inclusions, where  $\pi_i$  is the proof of inclusion for input element  $x_i$ .

- $\mathcal{ACC}.Verify_{pk}(\Delta, x, \pi) \rightarrow \{0, 1\}$ : output 1 (true) if  $\pi$  is a valid proof of inclusion for  $x$  and 0 (false) otherwise. As we define below, a valid proof of inclusion shows that there is a set  $X$  such that  $x \in X$  and  $\text{Add}_{pk}(X)$  outputs  $\Delta$ .

We next define three basic requirements from accumulator schemes: *correctness* and *unforgeability* of the proofs-of-inclusion, and *collision resistance* for the digest. We begin with the most basic property, *correctness*.

DEFINITION 14 (Correctness of Accumulator Scheme). *An accumulator  $\mathcal{ACC}$  is said to satisfy correctness, if the following holds for every security parameter  $\lambda \in \mathbb{N}$ , input sequence  $X = \{x_i\}_{i=1}^k \in \mathcal{X}^k$  and  $i \in [1, \dots, k]$ . Let  $pk \leftarrow \mathcal{ACC}.Setup(\lambda)$  and  $(\Delta, \Pi) \leftarrow \mathcal{ACC}.Add(X)$ . Then, for every  $x \in X$ , there exists  $\pi \in \Pi$  s.t.  $\mathcal{ACC}.Verify_{pk}(\Delta, x, \pi) = 1$ .*

Note: typically, each PoI is for the corresponding input element, i.e.,  $\mathcal{ACC}.Verify_{pk}(\Delta, x_i, \pi_i) = 1$ . However, this is not required by the definition above.

The next property we define is *collision resistance*. The collision-resistance of accumulators is the natural extension of the collision-resistance property of hash functions. We present only the keyed ‘any collision resistance’ variant; keyless variants and ‘target collision resistance’ can be defined similarly.

DEFINITION 15 (Collision-Resistant Accumulator). *We say that an accumulator scheme  $\mathcal{ACC}$  is Collision-Resistant if for any PPT adversary  $\mathcal{A}$  and security parameter  $\lambda \in \mathbb{N}$  holds:*

$$\Pr \left[ \mathcal{ACC}.Add(X) = \mathcal{ACC}.Add(X') \mid \left\{ \begin{array}{l} pk \leftarrow \mathcal{ACC}.Setup(\lambda) \\ (X, X') \leftarrow \mathcal{A}(pk) \end{array} \right\} \right] \in \text{Negl}(\lambda)$$

The final security property which we define in this section is *existential unforgeability*. Unforgeability extends collision-resistance, to ensure the integrity of specific elements in the accumulated sequence. The existential unforgeability definition is similar to the existential unforgeability of signature and MAC schemes.

DEFINITION 16 (Existentially Unforgeable Accumulator). *We say that an accumulator scheme  $\mathcal{ACC}$  is Existentially Unforgeable if for any PPT adversary  $(\mathcal{A}_1, \mathcal{A}_2)$  and security parameter  $\lambda \in \mathbb{N}$  holds:*

$$\Pr [\mathcal{ACC}.Verify(\Delta, x, \pi_{\mathcal{A}}) = 1 \wedge x \notin X] \in \text{Negl}(\lambda)$$

Where  $pk, \Delta, x$  and  $\pi_{\mathcal{A}}$  are defined by the following process:

$$\begin{aligned} pk &\leftarrow \mathcal{ACC}.Setup(\lambda); (s, X) \leftarrow \mathcal{A}_1(pk); \\ (\Delta, \Pi) &\leftarrow \mathcal{ACC}.Add(X); (x, \pi_{\mathcal{A}}) \leftarrow \mathcal{A}_2(s, \Delta, \Pi) \end{aligned}$$

### A.3 The CT Merkle-Tree Accumulator

In this subsection, we present and prove the security of the CT Merkle-Tree Accumulator. This construction is a keyed-version of the construction used by the Certificate Transparency 2.0 standard, defined in Section 2.1 of RFC 9162 [36].

DEFINITION 17 (The CT Merkle-Tree Accumulator). *Let  $h$  be a keyed (any-collision-resistant) hash function. The  $h$ -CT Merkle-Tree Accumulator, denoted  $\mathcal{MT}^h$ , is defined by the following algorithms:*

- $\mathcal{MT}^h.Setup(1^\lambda)$  outputs a public key  $pk \leftarrow \{0, 1\}^\lambda$ .
- $\mathcal{MT}^h.Add_{pk}(X)$  outputs  $(\Delta_{pk}(X), \pi_{pk}(X))$ , where  $\Delta, \pi$  are computed as follows.  
First define  $\Delta_{pk}(X)$ :  
– If  $X$  is an empty list then  $\Delta_{pk}(X) = h_{pk}(\epsilon)$ , where  $\epsilon$  denotes the empty string.

- If  $|X| = 1$ , i.e.,  $X$  contains a single entry  $x_1$ , then  $\Delta_{pk}(X) = h_{pk}(0x00||x_1)$ . This computation is sometimes called a leaf hash.
- If  $|X| > 1$ , then we compute  $\Delta_{pk}(X)$  as follows. Let  $n = |X|$ , i.e.,  $X = \{x_1, \dots, x_n\}$ , and let  $k = \max_{i \in \mathbb{N}} \{2^i < n\}$ , i.e., the largest power of two smaller than  $n$  (i.e.,  $k < n \leq 2k$ ). We compute  $\Delta_{pk}(X)$  recursively using  $\Delta_{pk}(X) = h(0x01||\Delta_{pk}(\{x_1, \dots, x_k\})||\Delta_{pk}(\{x_{k+1}, \dots, x_n\}))$ .
- We omit the (similar) definition of  $M\mathcal{T}^h.\text{Verify}_{pk}(\Delta, x, \pi)$  and  $\pi_{pk}(X)$ .

Finally, we state security properties of the CT Merkle-Tree Accumulator.

LEMMA 8. *The CT Merkle-Tree Accumulator is correct, collision-resistant and unforgeable.*

## B PKIX ENSURES ACCOUNTABILITY

PKIX ensures accountability by ensuring that for a certificate  $\psi$  to be valid, it must come with a valid certificate chain. Let us first state and prove this claim, which we use both to prove accountability (next) and to prove weak  $\Delta$  revocation (later).

CLAIM 9. *(A valid certificate must come with a valid chain) If, for some values  $clk, \xi, store, \omega$  holds  $\text{PKIX}_C.\text{Valid}(clk, \psi, store, \omega) = \top$ , then*

$$\begin{aligned} \exists \xi \in \text{PKIX}_C.\text{Chains}(clk, \psi, store, \omega) \text{ s.t. :} \\ \xi[0] = \psi \wedge \text{ISCHAIN}^C(\xi) \wedge (\exists i)(\xi[i] \in \text{store}) \wedge \\ (\forall \psi' \in \xi)(clk \in [\psi'.\text{from}, \psi'.\text{to}]) \wedge \text{CHAINSNOTREVOKED}(\xi) \end{aligned}$$

PROOF. The claim follows immediately from the implementation of  $\text{PKIX}.\text{Valid}$  (Algorithm 8 line 2), which checks exactly for this condition.  $\square$

Let us now prove that PKIX ensures accountability.

LEMMA 10. *Let  $C$  be a secure certificate scheme. Then,  $\text{PKIX}_C$  satisfies the accountability requirement.*

PROOF. Assume to the contrary that  $\text{PKIX}_C$  does not satisfy the accountability requirement. Hence, there exists a PPT adversary  $\mathcal{A}_{\text{ACC}}$  such that:

$$\epsilon_{\mathcal{A}_{\text{ACC}}, \text{PKIX}}^{\pi_{\text{ACC}}}(params) \notin \text{Negl}(params.1^\kappa) \quad (11)$$

Where  $\pi_{\text{ACC}}$  is defined in Algorithm 3.

In other words, from Definition 7,  $\mathcal{A}_{\text{ACC}}$  satisfies

$$\Pr \left[ T \leftarrow \text{Exec}_{\mathcal{A}_{\text{ACC}}, \text{PKIX}}(params) \right] \notin \text{Negl}(params.1^\kappa) \quad (12)$$

Note that for  $\pi_{\text{ACC}}(T) = \perp$ , the adversary must have produced a certificate  $\psi$  with auxiliary information  $\omega$  that is valid, i.e.:

$$\text{PKIX}.\text{Valid}(clk, \psi, store, \omega) = \top$$

and yet, no valid chain exists. However, Claim 9 shows this cannot hold. Therefore, there cannot exist such adversary  $\mathcal{A}_{\text{ACC}}$  that produces an input which  $\text{PKIX}.\text{Valid}$  classifies as valid without  $\text{PKIX}.\text{Chains}$  able to produce a valid certificate chain, in contrast to the assumption that such adversary exists.  $\square$

## C PKIX SATISFIES THE WEAK REVOCATION REQUIREMENT

First, we define the  $\overline{\text{Exp}}_{\mathcal{A}, \mathcal{P}}^{\text{Forge}}(params)$  experiment, a modification of the  $\pi_{\text{Weak}-\Delta_{\text{Rev}}}$  predicate (Definition 8). In the  $\overline{\text{Exp}}_{\mathcal{A}, \mathcal{P}}^{\text{Forge}}(params)$  experiment, an adversary  $\mathcal{A}_{pk}^{C.\text{Certify}_{sk}(\cdot)}$  is given a verification key  $pk$  and has an oracle access to a corresponding signing oracle  $C.\text{Certify}_{sk}(\cdot)$ . The adversary wins if after the execution of PKI scheme  $\mathcal{P}$ , the adversary produces a certificate  $\psi$  that validates successfully using the public verification key  $pk$ , without the adversary asking the oracle to certify  $\psi$ .

DEFINITION 18 (The  $\overline{\text{Exp}}_{\mathcal{A}, \mathcal{P}}^{\text{Forge}}(params)$  experiment). *Let  $\mathcal{P}$  be a PKI scheme and  $\mathcal{A}$  be a PPT algorithm. We define  $\overline{\text{Exp}}_{\mathcal{A}, \mathcal{P}}^{\text{Forge}}(params)$  as the following random process:*

- (1) Generate key pair  $(sk, pk) \leftarrow C.\text{KG}(params.1^\kappa)$ .
- (2) Let  $\mathcal{A}_{pk}^{C.\text{Certify}_{sk}(\cdot)}$  denote algorithm  $\mathcal{A}$  with input  $pk$  and oracle access to the function  $C.\text{Certify}_{sk}(\cdot)$ .
- (3) Execute  $\mathcal{A}_{pk}^{C.\text{Certify}_{sk}(\cdot)}$  with  $\mathcal{P}$ , except for a random entity  $\iota \xleftarrow{R} T.N$ , where entity  $\iota$  also executes  $\mathcal{P}$ , except:
  - (a) In Init,  $\iota$  sets  $(st.sk, st.pk) \leftarrow (\perp, pk)$ , and
  - (b)  $\iota$  replaces calls to  $C.\text{Certify}_{st.sk}(\cdot)$  with calls to  $C.\text{Certify}_{sk}(\cdot)$ . Let  $T$  denote the resulting transcript.
- (4) Let  $\psi \leftarrow T.\text{out}_{\mathcal{A}}$ . Return 1 ( $\mathcal{A}$  ‘won’) if:
  - (a)  $C.\text{Verify}_{pk}(\psi) = \top$
  - (b)  $\mathcal{A}$  did not use the oracle access on  $\psi.tbc$ .
  - (c)  $\iota$  is an honest entity, i.e.,  $\iota \in T.N - T.F$ . Otherwise, return 0 ( $\mathcal{A}$  ‘lost’).

We next observe that that if  $C$  is a secure certification scheme, then no PPT adversary  $\mathcal{A}$  ‘wins’ the  $\overline{\text{Exp}}_{\mathcal{A}, \text{PKIX}_C}^{\text{Forge}}$  experiment with non-negligible probability.

CLAIM 11. *If there is a PPT adversary  $\mathcal{A}$  s.t.:*

$$\Pr \left[ \overline{\text{Exp}}_{\mathcal{A}, \text{PKIX}_C}^{\text{Forge}}(params) = 1 \right] \notin \text{Negl}(1^\kappa) \quad (13)$$

*then  $C$  is not a secure certification scheme.*

PROOF. Follows by reduction to the existential unforgeability of  $C$ . Namely, if there is such an adversary  $\mathcal{A}$ , then there is also an adversary  $\hat{\mathcal{A}}$  that wins the certificate forgery game against  $C$ . Adversary  $\hat{\mathcal{A}}$  runs the  $\overline{\text{Exp}}_{\mathcal{A}, \text{PKIX}_C}^{\text{Forge}}$  experiment, except that  $\hat{\mathcal{A}}$  outputs the certificate  $\psi$  that  $\mathcal{A}$  produces to win the  $\overline{\text{Exp}}_{\mathcal{A}, \text{PKIX}_C}^{\text{Forge}}$  experiment (i.e.,  $\psi \leftarrow T.\text{out}_{\mathcal{A}}$ ). Since  $\mathcal{A}$  ‘wins’ with non-negligible probability, it follows that  $\hat{\mathcal{A}}$  will also succeed in forgery with non-negligible probability, contradicting the assumed security of  $C$ , as defined in Definition 1.  $\square$

LEMMA 12. *Let  $C$  be a secure certificate scheme, let  $N$  be a set of entities and let  $\Delta_{\text{Rev}} = \Delta_r + \Delta_{clk}$ , where  $\Delta_r$  is the acceptable revocation delay that all entities in  $N$  are initialized with (or smaller), i.e.,  $params.\Delta_r \leq \Delta_r$ . Then,  $\text{PKIX}_C$  satisfies the weak  $\Delta_{\text{Rev}}$ -Revocation requirement under model  $(\pi_{\Delta_r}^F \wedge \pi_{\Delta_{clk}}^{\text{Drift}})$ .*

PROOF. Assume to the contrary, that  $\text{PKIX}_C$  fails to satisfy the weak  $\Delta_{\text{Rev}}$ -Revocation under model  $(\pi_{\Delta_r}^F \wedge \pi_{\Delta_{clk}}^{\text{Drift}})$ . Namely, there



is a PPT adversary  $\mathcal{A}$  that satisfies model  $(\pi^F \wedge \pi_{\Delta_{clk}}^{\text{Drift}})$  and yet, has non-negligible advantage against the  $\pi_{\text{Weak}-\Delta_{\text{Rev}}}$  requirement (Definition 8), i.e.:

$$\epsilon_{\mathcal{A}, \text{PKIX}_C}^{\pi_{\text{Weak}-\Delta_{\text{Rev}}}}(\text{params}) \notin \text{Negl}(\text{params}.1^\kappa) \quad (14)$$

From Definitions 8 and 10, this implies that:

$$\Pr \left[ \begin{array}{l} f_{\Delta_{\text{Rev}}}(T) = \perp, \text{ where} \\ T \leftarrow \text{Exec}_{\mathcal{A}, \text{PKIX}_C}(\text{params}) \end{array} \right] \notin \text{Negl}(\text{params}.1^\kappa) \quad (15)$$

For a given  $T$ , let  $(\text{clk}, \psi, \text{store}, \omega) \leftarrow T.\text{out}_{\mathcal{A}}$ . From Algorithm 25,  $f_{\Delta_{\text{Rev}}}(T) = \perp$  only if:

$$\begin{aligned} & \psi.\text{issuer} \in T.\text{N} - T.\text{F} \wedge \\ & \text{PKIX}_C.\text{Valid}(\text{clk}, \psi, \text{store}, \omega) \wedge \\ & \text{REVOCATIONREQUESTED}(\text{clk}, \psi, T, \Delta_{\text{Rev}}) \wedge \\ & (\text{CANFINDFRAUDULENTCERT}(T, \text{clk}, \psi, \text{store}, \omega) = \perp) \end{aligned} \quad (16)$$

Since  $\text{REVOCATIONREQUESTED}(\tau, \psi, T, \Delta_{\text{Rev}})$  returns  $\top$ , then  $\psi$  must have been previously issued and then revoked at  $\psi.\text{issuer}$ , both before  $\tau - \Delta_{\text{Rev}}$ ; let  $\tau_R \leq \tau - \Delta_{\text{Rev}}$  denote the time of revocation. Being non-faulty,  $\psi.\text{issuer}$  have executed the  $\text{PKIX}_C.\text{Certify}$  and  $\text{PKIX}_C.\text{Revoke}$  operations correctly (Algorithms 5 and 6), and also includes  $\psi$  in every CRL certified in the  $\text{PKIX}_C.\text{GetCRL}$  operation after  $\tau_R$ . As can be confirmed in the pseudocode of  $\text{PKIX}$  (Figure 1),  $\text{PKIX}_C.\text{GetCRL}$  is the only operation in which  $\text{PKIX}_C$  certifies a CRL (i.e., certificate with type CRL).

Since  $\text{PKIX}_C.\text{Valid}(\text{clk}, \psi, \text{store}, \omega)$  return  $\top$ , i.e.,  $\psi$  is valid at time  $\tau$  with auxiliary information  $\omega$ , it follows from the implementation of  $\text{PKIX}_C.\text{Valid}$  that  $\omega$  must contain a chain  $\xi_{\text{CRL}}$  such that  $\xi_{\text{CRL}}[0]$  is a valid CRL at time  $\tau$  that does *not* list  $\psi$  as revoked. From line 10 of section 8, we know that  $\xi_{\text{CRL}}[0].\text{issuer} = \psi.\text{issuer}$ .

However, we next show that  $\psi.\text{issuer}$  cannot have issued  $\xi_{\text{CRL}}[0]$ . Assume otherwise, i.e., that  $\xi_{\text{CRL}}[0]$  was issued by  $\psi.\text{issuer}$  (and does *not* list  $\psi$  as revoked). Since  $\psi.\text{issuer}$  is benign, then it must have issued  $\xi_{\text{CRL}}[0]$  before revoking  $\psi$ , i.e., at some time  $\tau'_R < \tau_R$ . Also,  $\xi_{\text{CRL}}[0].\text{from}$  is the clock value of  $\psi.\text{issuer}$  at  $\tau'_R$ . Since the clock drift is bounded by  $\Delta_{\text{clk}}$ , except with negligible probability ( $\pi_{\Delta_{\text{clk}}}^{\text{Drift}}$ ), we have  $\xi_{\text{CRL}}[0].\text{from} \leq \tau'_R + \Delta_{\text{clk}} < \tau_R + \Delta_{\text{clk}}$ . And, again since  $\psi.\text{issuer}$  is benign, we know that  $\xi_{\text{CRL}}[0].\text{to} = \xi_{\text{CRL}}[0].\text{from} + \Delta_r$ . Since  $\tau_R \leq \tau - \Delta_{\text{Rev}}$ , we have:

$$\begin{aligned} \xi_{\text{CRL}}[0].\text{to} &= \psi_R.\text{from} + \Delta_r \\ &< \tau_R + (\Delta_{\text{clk}} + \Delta_r) \\ &= (\tau - \Delta_{\text{Rev}}) + \Delta_{\text{Rev}} \\ &= \tau \end{aligned} \quad (17)$$

Namely,  $\xi_{\text{CRL}}[0].\text{to} < \tau$ , which contradicts the assumption that  $\xi_{\text{CRL}}[0]$  is valid at time  $\tau$ . This shows that the (benign)  $\psi.\text{issuer}$  cannot have issued  $\xi_{\text{CRL}}[0]$ .

Still, the attacker obtained a valid CRL  $\xi_{\text{CRL}}[0]$  in some other way, i.e., not issued by  $\psi.\text{issuer}$ , but with  $\xi_{\text{CRL}}[0].\text{issuer} = \psi.\text{issuer}$ . Suppose this is the case; recall that  $\xi_{\text{CRL}}[0]$  must be valid, specifically,  $\text{PKIX}.\text{Valid}(\tau, \xi_{\text{CRL}}[0], \text{store}, \omega - \{\xi\})$  (line 12 of Algorithm 8). Hence, from Algorithm 8, there is some chain<sup>18</sup>  $\xi' \in \omega$  which includes a trust-anchor certificate  $\xi'[i] \in \text{store}$ , such that  $\xi_{\text{CRL}}[0] = \xi'[0]$  and  $\text{IsCHAIN}^C(\xi')$  returns  $\top$ .

<sup>18</sup>Intuitively, we expect  $\xi'$  to be the same as  $\xi_{\text{CRL}}$ , but this is not required.

Hence, from Algorithm 1,  $C.\text{Verify}_{\xi'[1].pk}(\xi_{\text{CRL}}[0])$  holds and  $\xi_{\text{CRL}}[0].\text{issuer} = \xi'[1].\text{subject}$ . Recall that  $\xi_{\text{CRL}}[0].\text{issuer} = \psi.\text{issuer}$  and that  $\psi.\text{issuer}$  is benign. Let  $\hat{\psi}$  denote the self-signed certificate outputted by  $\psi.\text{issuer}$  upon its initialization. We consider two cases: (1)  $\xi'[1].pk = \hat{\psi}.pk$  and (2)  $\xi'[1].pk \neq \hat{\psi}.pk$ .

**Case (1):**  $\xi'[1].pk = \hat{\psi}.pk$ . Since  $\psi.\text{issuer}$  is benign, then it never exposes the corresponding private signing key (generated upon initialization and used to sign  $\hat{\psi}$  and other CRLs and public-key certificates issued by  $\psi.\text{issuer}$ ). In fact,  $\psi.\text{issuer}$  only uses this signing key to certify CRLs when doing the  $\text{PKIX}_C.\text{GetCRL}$  operation. Recall that  $\xi_{\text{CRL}}[0]$  is a valid CRL not issued by  $\psi.\text{issuer}$  and that  $C.\text{Verify}_{\xi'[1].pk}(\xi_{\text{CRL}}[0])$  holds. Since the corresponding private key was never exposed or used to certify  $\xi_{\text{CRL}}[0]$ , such an adversary contradicts the existential-unforgeability property of the certificate scheme  $C$ .

**Case (2):**  $\xi'[1].pk \neq \hat{\psi}.pk$ . From Claim 9,  $\xi_{\text{CRL}}[0]$  has also a valid chain; but then surely  $\text{CANFINDFRAUDULENTCERT}(T, \text{clk}, \psi, \text{store}, \omega)$  would return  $\top$ , which contradicts our assumption.  $\square$

## D EXPLANATION OF A SUBTLE POINT REGARDING THE ACCOUNTABILITY REQUIREMENT

Just before submission, a careful reviewer, which we will maintain anonymous for the submission, asked us a very natural question regarding the accountability requirement predicate (Algorithm 3). Rephrased, he asked why does the predicate, in line 5, check first that  $\xi$  is a chain ( $\text{IsCHAIN}(\xi)$ ) and then also that  $\xi[j+1]$  is valid. Isn't the fact that  $\xi[j+1]$  is valid follows from the fact that  $\xi$  is a chain?

The answer is: both checks are necessary. The  $\text{IsCHAIN}(\xi)$  function, Algorithm 1, checks (only) that  $\xi$  is a certificate chain, but does not check that each certificate  $\xi[j]$  along the chain is  $\mathcal{P}.\text{Valid}$ . Namely,  $\text{IsCHAIN}(\xi)$  (Algorithm 1) checks that for every 'chain link', i.e., for every  $j : 0 < j < |\xi|$ , holds:

- Line 3: Each certificate  $\xi[j-1]$  in the chain (except the last one) is certified by the next certificate on the chain  $\xi[j]$ .
- Line 4: The issuer of  $\xi[j-1]$  is the subject of  $\xi[j]$ .
- Line 5: Each cert  $\xi[j]$  in the chain, except the first one, is a public key certificate. Notice that it is important to allow the first certificate on the chain,  $\xi[0]$ , to be a non-public-key certificate, in particular, a CRL.
- Line 5: Each certificate  $\xi[j]$  in the chain, except the first one, indicates that the certified entity  $\xi[j].\text{subject}$  is a CA.

All of these checks are necessary for the accountability predicate - but not sufficient. We also need to confirm that every non-first certificate along the chain,  $\xi[j]$ , is  $\mathcal{P}.\text{Valid}$  (at the current time, with the current  $\omega$  and  $\text{store}$ ). This is checked by the (recursive) call to  $\mathcal{P}.\text{Valid}$  at the end of line 5 of the accountability requirement predicate (Definition 7).

A related question may be, why do we allow  $\psi$  to appear in an arbitrary location ( $j$ ) along the chain  $\xi$ , and do not require it to be the very first certificate along the chain. In fact, we could have required  $\psi$  to be the first certificate along the chain  $\xi$ ; however, that would have made the use of the  $\mathcal{P}.\text{Valid}$  quite inconvenient. This is since

**Algorithm 11** CT.Init: initialization

**Input:** local state  $st$ , local clock  $clk$  and parameters  $params$  which has the same parameters as in Algorithm 4 with the addition of the maximum merge delay  $\Delta_{MMD}$  and a trusted root store  $store$ .

**Output:** updated state, the output of Algorithm 4 and the next wakeup time.

```

procedure CT.Init ( $st, clk, params$ )
1:  ( $st, out$ )  $\leftarrow$  PKIX.Init( $st, clk, params$ )            $\triangleright$  See Algorithm 4
2:  if  $st.role \in \{\text{'Logger'}, \text{'Monitor'}\}$  then          $\triangleright$  Only monitors and loggers
3:     $out \ +=$  ('Wakeup',  $clk + st.\Delta_{MMD}$ )            $\triangleright$  Add next wakeup time
4:    if  $st.role = \text{'Monitor'}$  then                      $\triangleright$  If monitor
5:       $st.loggers \leftarrow \perp$                         $\triangleright$  Init empty list of monitored loggers
6:       $st.log \leftarrow \perp$                             $\triangleright$  Init empty list of monitored logs
7:    else                                              $\triangleright$  If logger
8:       $st.mylog \leftarrow \perp$                           $\triangleright$  Initialize local log
9:    end if
10: end if
11: return ( $st, out$ )                                  $\triangleright$  Output
end procedure

```

**Algorithm 12** CT.Certify: certificate issuance

**Input:** local state  $st$ , local clock  $clk$  and data to be certified  $tbc$ .

**Output:** updated state and certificate  $\psi$ .

```

procedure CT.Certify ( $st, clk, tbc$ )
1:  if  $tbc.type = \text{'pre-certificate'}$  then              $\triangleright$  If pre-certificate
2:    Add critical 'poison' extension to  $tbc$           $\triangleright$  'Poison' extension
3:  end if
4:  return ( $st, \text{PKIX.Certify}(tbc)$ )                  $\triangleright$  See Algorithm 5
end procedure

```

**Algorithm 13** CT.Revoke: certificate revocation

**Input:** local state  $st$ , local clock  $clk$  and certificate  $\psi$ .

**Output:** the output of PKIX.Revoke.

```

procedure CT.Revoke ( $st, clk, \psi$ )
1:  return PKIX.Revoke( $\psi$ )                              $\triangleright$  See Algorithm 6
end procedure

```

**Algorithm 14** CT.GetCRL: retrieve CRL

**Input:** local state  $st$  and local clock  $clk$ .

**Output:** the output of PKIX.GetCRL.

```

procedure CT.GetCRL ( $st, clk$ )
1:  return PKIX.GetCRL( $st, clk$ )                          $\triangleright$  See Algorithm 9
end procedure

```

**Algorithm 15** CT.Chains:

**Input:** local clock  $clk$ , certificate  $\psi$ , root store  $store$  and auxiliary information  $\omega$ .

**Output:** the output of PKIX.Chains.

```

procedure CT.Chains ( $clk, \psi, store, \omega$ )
1:  return PKIX.Chains( $clk, \psi, store, \omega$ )            $\triangleright$  See Algorithm 7
end procedure

```

**Algorithm 16** CT.Valid: certificate validation

**Input:** local clock  $clk$ , certificate  $\psi$ , root store  $store$  and auxiliary information  $\omega$ .

**Output:**  $\top$  if  $\psi$  is a valid certificate, and  $\perp$  otherwise.

```

procedure CT.Valid ( $clk, \psi, store, \omega$ )
1:  return PKIX.Valid( $clk, \psi, store.CAs, \omega$ )  $\wedge$             $\triangleright$  See Algorithm 8
2:     $(\forall i \in \{1, 2\}) \left( \begin{array}{l} \exists \psi_i^{sct} \in \psi.SCT \wedge \\ \exists \psi_i^f \in store.logger \end{array} \right) \text{ s.t.:$             $\triangleright \psi$  contains two SCTs
3:     $\psi_i^{sct}.type = \text{'SCT'} \wedge$                               $\triangleright \psi_i^{sct}$  is an SCT
4:     $C.Verify_{\psi_i^f.pk}(\psi_i^{sct}) \wedge$                           $\triangleright$  Logger  $i$  signed  $\psi_i^{sct}$ 
5:     $\psi_i^{sct}.tbc.cert = \psi.tbc - \psi.SCT \wedge$               $\triangleright$  Both SCTs are for  $\psi$ 
6:     $\psi_1^{sct}.issuer \neq \psi_2^{sct}.issuer$                     $\triangleright$  SCTs from different loggers
end procedure

```

**Algorithm 17** CT.Monitor: add to monitored logs

**Input:** local state  $st$ , local clock  $clk$  and log's identifier  $i$ .

**Output:** updated state.

```

procedure CT.Monitor ( $st, clk, i$ )
1:  Add  $i$  to  $st.loggers$                                   $\triangleright$  Add  $i$  to monitored logs
2:  return ( $st, \text{'Monitor'}, st.loggers$ )                  $\triangleright$  Output
end procedure

```

**Algorithm 18** CT.Lookup: subject's certificates lookup

**Input:** local state  $st$ , local clock  $clk$  and subject's identifier  $subject$ .

**Output:** updated state and all the certificates issued to  $subject$  that are known to the entity.

```

procedure CT.Lookup ( $st, clk, subject$ )
1:   $certs \leftarrow \perp$                                       $\triangleright$  Initialize set to  $\perp$ 
2:  if  $st.role = \text{'Monitor'}$  then
3:     $certs \leftarrow \left\{ \psi \mid \begin{array}{l} \exists i \in st.loggers, \psi \in st.log[i] \text{ s.t.} \\ \psi.subject = subject \end{array} \right\}$             $\triangleright$   $subject$ 's certificates
4:  end if
5:  return ( $st, certs$ )                                    $\triangleright$  Output
end procedure

```

**Algorithm 19** CT.Wakeup: scheduled operations

**Input:** local state  $st$ , local clock  $clk$  and wake up related information  $data$ .

**Output:** updated state and next wakeup time.

```

procedure CT.Wakeup ( $st, clk, data$ )
1:   $out \leftarrow \text{'Wakeup'}, clk + st.\Delta_{MMD}$             $\triangleright$  Set next wakeup
2:  if  $st.role = \text{'Logger'}$  then
3:     $tbc \leftarrow (type, \text{'STH'})$                           $\triangleright$  Certificate of type STH
4:     $tbc.issuer \leftarrow st.i$                             $\triangleright$  The issuer of the STH
5:     $tbc.from \leftarrow clk$                                 $\triangleright$  When was issued
6:     $tbc.to \leftarrow clk + st.\Delta_{MMD}$                   $\triangleright$  Time for next STH
7:     $tbc.STH \leftarrow \mathcal{MT}^H(st.log)$                     $\triangleright$  See CT Merkle-Tree Accumulator (Definition 17)
8:     $tbc.size \leftarrow |st.log|$                           $\triangleright$  Log size
9:     $\sigma \leftarrow C.Certify_{st.sk}(tbc)$                 $\triangleright$  Sign certificate
10:    $st.\psi_{STH} \leftarrow (tbc, \sigma)$                   $\triangleright$  Certified data
11: else if  $st.role = \text{'Monitor'}$  then
12:    $out \ += \{(\text{'GetSTH'}, i) \mid \forall i \in st.loggers\}$     $\triangleright$  Ask loggers for latest STH
13: end if
14: return ( $st, out$ )                                      $\triangleright$  Output
end procedure

```

Figure 2: CT (PKIX with certificate transparency) implementation

**Algorithm 20** CT.AddPreChain-Req: add certificate to log

**Input:** local state  $st$ , local clock  $clk$ , (pre)certificate  $\psi$  and auxiliary information  $\omega$ .  
**Output:** updated state and SCT  $\psi_{SCT}$ .

**procedure** CT.AddPreChain-Req ( $st, clk, \psi, \omega$ )

```

1:  if ( PKIX.Valid( $\psi, st.store, clk, \omega$ )  $\wedge$ 
      ( $\psi \notin st.mylog \wedge clk \leq \psi.from + \Delta_{clk} + \Delta_{com}$ ) ) then
2:     $st.mylog += \psi$  ▷ Store locally
3:     $tbc.type \leftarrow 'SCT'$  ▷ Certificate of type SCT
4:     $tbc.issuer \leftarrow st.i$  ▷ The issuer of the SCT
5:     $tbc.timestamp \leftarrow clk$  ▷ Timestamp (when issued)
6:     $tbc.cert \leftarrow \psi$  ▷ Certificate information
7:     $\sigma \leftarrow C.Certify_{st.sk}(tbc)$  ▷ Sign certificate
8:     $\psi_{SCT} \leftarrow (tbc, \sigma)$  ▷ Certified data
9:    return ( $st, ('AddPreChain-Resp', \psi_{SCT})$ ) ▷ Respond with SCT
10:  end if
11:  return ( $st, \perp$ ) ▷ Invalid request
end procedure

```

**Algorithm 21** CT.GetSTH-Req: retrieve current STH

**Input:** local state  $st$ , local clock  $clk$ .  
**Output:** updated state and STH  $\psi_{STH}$ .

**procedure** CT.GetSTH-Req ( $st, clk$ )

```

1:  return ( $st, ('GetSTH-Resp', st.\psi_{STH})$ ) ▷ Output
end procedure

```

**Algorithm 22** CT.GetSTH-Resp: store latest STH

**Input:** local state  $st$ , local clock  $clk$ , root store  $store$  and STH certificate  $\psi_{STH}$ .  
**Output:** updated state and request for newly added log entries.

**procedure** CT.GetSTH-Resp ( $st, clk, store, \psi$ )

```

1:   $out \leftarrow \perp$ 
2:   $log \leftarrow st.log[\psi_{STH}.issuer]$  ▷ Logger's local data
3:   $\psi_\ell \leftarrow store.loggers[\psi_{STH}.issuer]$  ▷ Logger's certificate
4:  if  $C.Verify_{\psi_\ell.pk}(\psi_{STH}) \wedge |log.entries| < \psi_{STH}.size$  then ▷ Valid STH
5:     $log.newSTH \leftarrow \psi_{STH}$  ▷ Store new STH
6:     $out \leftarrow ('GetEntries-Req', \psi_{STH}.issuer,$ 
7:       $start = |log.entries|, end = \psi_{STH}.size)$  ▷ Ask for new certs
8:  end if
9:  return ( $st, out$ ) ▷ Output
end procedure

```

**Algorithm 23** CT.GetEntries-Req: retrieve log entries

**Input:** local state  $st$ , local clock  $clk$ , the number of the first entry requested  $start$  and the number of the last entry requested  $end$ .  
**Output:** updated state and signed list of entries  $\psi_{entries}$ .

**procedure** CT.GetEntries-Req ( $st, clk, start, end$ )

```

1:   $tbc \leftarrow (type, 'get-entries')$  ▷ Certificate type
2:   $tbc.issuer \leftarrow st.i$  ▷ The issuer
3:   $tbc.from \leftarrow clk$  ▷ When was issued
4:   $tbc.entries \leftarrow st.mylog[start : end]$  ▷ Desired certificates
5:   $\sigma \leftarrow C.Certify_{st.sk}(tbc)$  ▷ Sign certificate
6:   $\psi_{entries} \leftarrow (tbc, \sigma)$  ▷ Certified data
7:  return ( $st, ('GetEntries-Resp', \psi_{entries})$ ) ▷ Output
end procedure

```

**Algorithm 24** CT.GetEntries-Resp: store log entries

**Input:** local state  $st$ , local clock  $clk$ , root store  $store$  and certificate  $\psi$ .  
**Output:** updated state.

**procedure** CT.GetEntries-Resp ( $st, clk, store, \psi$ )

```

2:   $\psi_\ell \leftarrow store.loggers[\psi.issuer]$  ▷ Logger's certificate
3:  if  $C.Verify_{\psi_\ell.pk}(\psi)$  then ▷ Valid certification
4:     $log \leftarrow st.log[\psi.issuer]$  ▷ Logger's local data
5:     $root \leftarrow \mathcal{M}^H(log.entries + \psi.entries)$  ▷ Calc new root
6:    if  $root = log.newSTH$  then ▷ Verify Merkle root
7:       $log.entries += \psi.entries$  ▷ Update local certs
8:       $log.STH \leftarrow log.newSTH$  ▷ Update local root
9:      Optionally: examine new certificates for problems, e.g., potential phishing
10:   end if
11:  end if
12:  return  $st$  ▷ Output
end procedure

```

**Figure 3: CT implementation contd.**

$\mathcal{P}$ . Chains would need to return a separate chain for every certificate along  $\xi$ , the chains getting shorter and shorter until the trivial chain that only contains a trust anchor certificate. For example, if  $\psi$  is provided with the (short) chain  $\xi[0] = \psi$ ,  $\xi[1] = \psi'$  and  $\xi[2] \in store$ , then, to ensure that  $\mathcal{P}.Valid(clk, \psi', store, \omega)$  would hold, we would need  $\omega$  to also contain the chain  $\xi'$  s.t.  $\xi'[0] = \xi[1] = \psi'$  and

$\xi'[1] = \xi[2] \in store$ . And this is even without considering the CRL chains required to validate  $\psi$  and  $\psi'$ . So, such an accountability predicate would be inconvenient, unnatural and inelegant, and definitely would require us to define a very strange implementations for the PKIs (and different from the real implementations).

**Algorithm 25**  $f_{\Delta_{\text{Rev}}}$ : Revocation function

**Input:** transcript  $T$ .

**Output:**  $\top$  if the adversary fails to present a 'revoked-yet-valid' certificate  $\psi$  (which is valid on local clock  $clk$  although its benign issuer  $\psi.issuer$  revoked it before  $\tau - \Delta_{\text{Rev}}$ ); 'Weak' if the adversary presents a 'revoked-yet-valid' certificate  $\psi$ , and the PKI identifies a fraudulent yet valid certificate which is valid with respect to a key not declared by its honest issuer; and  $\perp$  if the adversary presents a 'revoked-yet-valid' certificate *and* the PKI fails to identify such a certificate.

```

1: procedure  $f_{\Delta_{\text{Rev}}}(T)$ 
2:    $(clk, \psi, store, \omega) \leftarrow T.out_{\mathcal{A}}$  ▷ Extract adversary's output
3:   if  $\left[ \begin{array}{l} \psi.issuer \in T.N - T.F \wedge \\ \mathcal{P}.Valid(clk, \psi, store, \omega) \wedge \\ REVOCATIONREQUESTED(clk, \psi, T, \Delta_{\text{Rev}}) \end{array} \right]$  then
4:     if  $CANFINDFRAUDULENTCERT(T, clk, \psi, store, \omega)$  then
5:       return 'Weak'
6:     else
7:       return  $\perp$ 
8:     end if
9:   end if
10:  return  $\top$ 
11: end procedure
12: procedure  $CANFINDFRAUDULENTCERT(T, clk, \psi, store, \omega)$ 
13:  Output:  $\top$  if  $\mathcal{P}.Chains$  returns a chain  $\xi$  which contains some 'valid yet fraudulent' public key certificate  $\xi[i]$ . By saying that  $\xi[i]$  is fraudulent we mean that the subject  $\xi[i].subject$  is a benign entity, but the key certified is not the key generated (and self-certified) by  $\xi[i].subject$ .
14:  for  $\xi \in \mathcal{P}.Chains(\psi, store, \tau, \omega)$  do
15:    if  $\left[ \begin{array}{l} \exists i \text{ s.t. :} \\ \mathcal{P}.Valid(clk, \xi[i], store, \omega) \wedge \\ \xi[i].type = PubKey \wedge \\ \xi[i].subject \in T.N - T.F \wedge \\ SELF CERTIFIED(T, \xi[i].subject, \xi[i].pk) = \perp \end{array} \right]$  then return  $\top$ 
16:  end for
17:  return  $\perp$ 
18: end procedure
19: procedure  $REVOCATIONREQUESTED(clk, \psi, T, \Delta_{\text{Rev}})$ 
20:  return  $\exists e, e_R \text{ s.t. } e < e_R \wedge$  ▷  $\psi$  issued, then revoked
21:     $T.opr[e] = \text{'Certify'} \wedge$  ▷  $e$  was a Certify operation
22:     $T.opr[e_R] = \text{'Revoke'} \wedge$  ▷  $e_R$  was a Revoke operation
23:     $T.inp[e] = \psi \wedge$  ▷  $\psi$  issued by  $\psi.issuer$ 
24:     $T.inp[e_R] = \psi \wedge$  ▷  $\psi$  revoked by  $\psi.issuer$ 
25:     $T.\tau[e] \leq T.\tau[e_R] \leq clk - \Delta_{\text{Rev}}$  ▷ both at least  $T.\tau[e]$  before local clock  $clk$ 
26: end procedure

```

**Algorithm 26**  $\pi_{\Delta_{\text{Tra}}}$ : the  $\Delta_{\text{Tra}}$ -Transparency requirement predicate

**Input:** transcript  $T$ .

**Output:**  $\perp$  if the adversary outputs a certificate  $\psi$  and auxiliary information  $\omega$ , s.t.  $\psi$  is valid on local clock  $clk$  given  $\omega$ , although  $\psi.from \leq \tau - \Delta_{\text{Tra}}$  and there exists a benign monitor  $t_M$ , unaware of  $\psi$  at local clock  $clk$ , and monitoring a log in  $\psi.logs$  since  $\tau - \Delta_{\text{Tra}}$  or earlier. Otherwise, the adversary failed, and the output is  $\top$ .

```

1: procedure  $\pi_{\Delta_{\text{Tra}}}(T)$ 
2:    $(clk, \psi, store, \omega, t_M) \leftarrow T.out_{\mathcal{A}}$  ▷ Extract adversary's output
3:   return  $\perp$  if  $\mathcal{P}.Valid(clk, \psi, store, \omega) \wedge$  ▷  $\psi$  is valid certificate at  $\tau$ 
4:      $\psi.from \leq \tau - \Delta_{\text{Tra}}$  ▷  $\psi$  issued at least  $\Delta_{\text{Tra}}$  before  $\tau$ , and
5:      $ISBENIGNMONITOR(T, \psi, t_M, \tau - \Delta_{\text{Tra}}) \wedge$  ▷  $t_M$  monitors a log in  $\psi.logs$  since  $\tau - \Delta_{\text{Tra}}$ 
6:      $MONITORISUNAWARE(T, \psi, t_M, \tau)$  ▷ Yet  $t_M$  is unaware of  $\psi$  at  $\tau$ 
7:   return  $\top$ 
8: end procedure
9: procedure  $ISBENIGNMONITOR(T, \psi, t_M, \tau')$ 
10:  return  $t_M \in T.N - T.F \wedge$  ▷  $t_M$  is benign, and
11:     $\exists e, \Lambda \text{ s.t.}$  ▷ for some event  $e$  and set of logs  $\Lambda$ ,
12:     $T.ent[e] = t_M \wedge$  ▷  $t_M$  was invoked
13:     $T.out[e] = (\text{'Monitor'}, \Lambda) \wedge$  ▷ to monitor set of logs  $\Lambda$ 
14:     $\psi.logs \cap \Lambda \neq \emptyset \wedge$  ▷ which includes a log in  $\psi.logs$ 
15:     $T.\tau[e] \leq \tau'$  ▷ before  $\psi.\tau$ 
16: end procedure
17: procedure  $MONITORISUNAWARE(T, \psi, t_M, \tau)$ 
18:  return  $\exists e \text{ s.t. } T.opr[e] = \text{'Lookup'} \wedge$  ▷ Lookup operation was invoked
19:     $T.ent[e] = t_M \wedge$  ▷ on  $t_M$ 
20:     $T.inp[e] = \psi.subject \wedge$  ▷ w.r.t  $\psi.subject$ 
21:     $T.\tau[e] > \tau \wedge$  ▷ sometime after  $\psi.\tau + \Delta$ 
22:     $\psi \notin T.out[e]$  ▷ but  $t_M$  was unaware of  $\psi$ 
23: end procedure

```

**Figure 4: Remaining PKI security requirements**