

The Modular Specifications Security Framework

Hemi Leibowitz¹, Amir Herzberg², Ewa Syta³, and Sara Wrótniak²

¹ Dept. of Computer Science, Bar-Ilan University, Ramat Gan, Israel

² Dept. of Computer Science and Engineering, University of Connecticut, Storrs, CT

³ Dept. of Computer Science, Trinity College, Hartford, CT**

Abstract. We present the *Modular Specifications Security (MoSS) framework*, where security specifications are defined with respect to a specific *model predicate* \mathcal{M} . This allows analysis of even complex schemes and protocols, e.g., PKI schemes, under well-defined adversary, communication and synchronization models, in a modular and flexible way, and allows to analyze such schemes in both simplified and realistic models. The framework facilitates reuse of definitions, and, indeed, several of the model predicates and security specifications we define, seem ‘generic’ and reusable in analysis of other practical protocols.

1 Introduction

Cryptographic protocols are traditionally designed and analyzed under different assumptions, where the expected outcome is to achieve provable security. Most commonly, cryptographic protocols are specified and analyzed under specific adversary models. Other assumptions, related to communication and synchronization models are not always relevant, but when they are, highly simplified models are often applied, e.g., perfect synchronization (‘rounds model’). Understandably, since often there are no (relatively) simple tools that can be used to achieve stronger evaluation against more realistic models, such simplifications are used. However, this should serve as a motivation towards developing better evaluation frameworks, and should not be considered as ‘good enough’.

In particular, it is an established challenge to apply provable security on to applied cryptographic protocols, especially due to the realistic communication and synchronization models which such protocols operate under. A motivational example of such schemes are PKI schemes, which these aspects have crucial implications on their security. When considering the use of existing rigorous frameworks for analysis of distributed algorithms, e.g., the I/O automata [3] or reactive systems [5], the results are usually overly complex, even when focusing on relatively simple algorithms and requirements; applying them to complex tasks, such as PKI schemes and to cryptographic protocols, appears prohibitively challenging.

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

Instead, we present the *Modular Specifications Security (MoSS) framework*, which is a significant (positive) deviation from the ‘classical’ approach of defining security using games (experiments), with a given protocol and adversary. Namely, while in the ‘classical’ game-based approach, the model is an integral part of the game/experiment, in our framework, the model \mathcal{M} is a well-defined ‘standalone’ PPT algorithm, *separated* from the actual experiment. This approach taken by the MoSS framework, yields an important result. The MoSS framework allows a clear and concise method to take any well-defined security specification ξ and formally prove whether a given protocol \mathcal{P} *satisfies* specification ξ (with negligible advantage or with advantage ϵ) *under model \mathcal{M}* .

The *Modular Specifications Security (MoSS) framework* cleanly separates the specification and analysis of security into three components: the *execution process*, *model* and *specifications*. The first component is a well-defined, and relatively straightforward, *adversary-driven execution process* (Algorithm 1). This execution process defines precisely the process of *executing a protocol \mathcal{P} under adversary \mathcal{A}* , giving the adversary extensive control over the operation of the environment, including communication, local-clock values, inputs from the application, and faults. (The execution process only imposes some basic limitations, e.g., events are ordered, rather than potentially only partially-ordered, as allowed in [3, 5].)

The second component of the framework is the definition of *models*, where a model \mathcal{M} is a pair consisting of a *model predicate* $\mathcal{M}.\pi$ and a *base function* $\mathcal{M}.\beta$. A model predicate $\mathcal{M}.\pi$ classifies the execution of a protocol \mathcal{P} under adversary \mathcal{A} as ‘valid’ (\top) or ‘invalid’ (\perp), effectively enforcing one or multiple restrictions on the adversarial control of the execution, including initialization assumptions, limitations on number and type of faults, maximal delay and/or maximal clock drift. Furthermore, a model predicate may enforce restrictions on the use of resources by the adversary, such as numbers of queries or running time. The base function specifies the ‘base’ probability that a model predicate $\mathcal{M}.\pi$ outputs \perp . The sum of the base function $\mathcal{M}.\beta$ and the *advantage function* ϵ (or in the asymptotic definitions, $\mathcal{M}.\beta$ and a polynomial) limits the maximum probability that the predicate outputs \perp for any given parameters p .

The third component of the framework is the definition of *specifications*, which are similar to models except that they are used to define what is provided *by* a protocol or scheme, given some model. A specification ξ is a pair, $\xi = (\xi.\pi, \xi.\beta)$, where the specification predicate $\xi.\pi$ and the base function $\xi.\beta$ are used in the same way as model predicates and base functions.

We believe that this approach is not only likely to be beneficial for many cryptographic protocols and distributed systems, but would also make it easier to compare models and specifications between different works. Furthermore, we expect the separation between the model aspects and the problem-specific aspects of security definitions to allow cleaner definitions, as well as to allow the reuse of well-defined models \mathcal{M} .

Organization. Section 2 presents the basic integrity and availability (IA) execution model of the framework and asymptotic definitions for models and

specifications. Section 3 discusses the CIA execution process, which is an extension of the IA execution process which additionally allows confidentiality (and other indistinguishability-based) models and specifications. Section 4 extends the framework to allow concrete security (rather than asymptotic) definitions for models and specifications. Section 5 provides further model and specification predicates. We conclude and discuss future work in Section 6.

2 The MoSS Framework Building Blocks

Traditionally, research onto cryptographic schemes and protocols tend to focus on the adversary’s capabilities and whether some properties can be achieved under a specific adversary model or not. This is acceptable, when the analyzed system is not concerned with other assumptions. For example, when defining the security of encryption schemes or zero-knowledge proofs, there is no significance to the properties of the communication between the parties, or to the synchronization of their clocks. However, communication and synchronization assumptions are crucial factors in the operation and security of distributed systems. For instance, in PKI schemes, some of the properties are *time related*, and as such, the scheme’s security strongly depends on the communication and synchronization assumptions. Therefore, there is a clear need to bootstrap the communication and synchronization assumptions into the analysis of such systems; otherwise, analysis that overlooks them completely misses its purpose, and it is therefore, incomplete.

One way to address this challenge would be to adopt a specific, simple model for communication and synchronization, e.g., the synchronous ‘rounds model’, where all parties operate in lockstep, round by round, and messages sent at round i are delivered in round $i + 1$. Using such a fixed model has the benefit of making it easy to focus on the cryptographic aspects, e.g., present rigorous definitions and prove security (by reductions). However, real systems are more complex; clocks are not fully synchronized, and communication is rarely, if ever, perfectly synchronous. This creates a key dilemma in the design and analysis of cryptographic protocols; should the design and analysis assume a simplified model and focus on the cryptographic aspects of the protocol, or should it use a more realistic model, one that takes into account communication and clock synchronization and allows a *genuine* analysis of practical systems?

The reader might consider this dilemma as a ‘no-brainer’; obviously, an analysis should always be with respect to a genuine (as possible) model. Unfortunately, in many cases, it is extremely non-trivial, and as a result, many designs and analyses do choose to focus on the cryptographic, ignoring some key-aspects along the way. We now details how the MoSS framework solves that by splitting the design into three components: a generic *adversary-driven execution process*, *models*, and security *specifications*. We begin with how the framework’s basic execution process works.

2.1 Exec-IA: An Adversary-Driven Execution Process

We now present the adversary-driven execution process $\mathbf{Exec-IA}_{\mathcal{A},\mathcal{P}}(1^\kappa)$ as defined by the pseudo-code in Algorithm 1. The execution process specifies the details of running a given protocol \mathcal{P} with a given adversary \mathcal{A} , both modeled as efficient (PPT) functions. The execution process does not enforce any assumptions or model on the communication, synchronization, or inputs to the protocol; all of that is controlled entirely by the adversary \mathcal{A} . Furthermore, the execution allows the adversary to set and retrieve the state of any entity as well as messages sent and received, allowing for different failure models.

Notation. To allow the process to apply to protocols with multiple functions and operations, we define the entire protocol \mathcal{P} as a *single* PPT algorithm and use parameters to specify the exact operations and their inputs. Specifically, execution of an operation defined by \mathcal{P} over entity i , is expressed using the following notation: $\mathcal{P}(s, opr, inp, clk)$, where s is the *local state* of entity i , opr identifies the specific ‘operation’ or ‘function’ to be invoked, inp is the set of inputs to that operation/function, and clk is the current value of the local clock of entity i . The output of such execution is a tuple $(s, out, sec-out)$, where s is the state of entity i *after* the operation’s execution, out is the regular output of the executed function, and $sec-out$ is ‘secure output’ - namely, $sec-out$ contains values to be shared with entities through a secure channel, which is not accessible to \mathcal{A} .

The execution process reserves three specific strings, ‘Set-state’, ‘Set-output’ and ‘Get-state’, to define *generic operations*, which are *not* part of \mathcal{P} . We use these reserved operations to express specific adversarial capabilities which are allowed by some adversarial models. Namely, an adversary can expose the local state (‘Get-state’) of an entity, e.g., corresponding to a ‘honest but curious’ adversary; an adversary can control the state (‘Set-state’) of an entity; and an adversary can control the output (‘Set-output’), including the secure output, of an entity, e.g., corresponding to a ‘malicious’ adversary (‘byzantine’ fault model).

Additionally, we assume that \mathcal{P} includes an operation ‘Sec-in’. The execution process uses this operation to *securely* give values to an entity. Specifically, an adversary can set the operation to ‘Sec-in’ and choose some previous secure output (without seeing it - i.e., \mathcal{A} can only choose the event *during which* it was output), which will be given as input to \mathcal{P} . This provides a way to assume that certain information is securely shared, which may be used, e.g., if anonymous or confidential channels must be assumed without providing appropriate mechanisms at the protocol level.

We use the *index notation* to refer to cells of ‘arrays’. For example, $out[e]$ refers to the value of the e^{th} entry of the array out . Specifically, e represents the index execution events within the execution process. Note, however, that e does not represent any clock value and is independent of any clock synchronization model, and specifically, is not controlled by the adversary. Rather, e allows us to index a sequence of operations performed within one execution loop, and precisely refer to each event of the execution. The local clock clk and real time clock τ are used to represent and model different clock assumptions. Thus, to refer

to the local clock value and the real time clock value of event e , the execution process uses $clk[e]$ and $\tau[e]$, respectively. The execution process does not place any restrictions on these values and enforcing any constraints is left to the appropriate clock synchronization model. The specific clock synchronization model provided in this work ensures that $\tau[e]$ is monotonously increasing and uses this property to enforce other communication and clock properties (see Sections 5.1.3 and 5.1.4).

Construction. We now discuss the three main components of the execution process (Algorithm 1), that is, the initialization, main loop execution and termination.

Initialization. In line 1, we allow the adversary to set its state and to choose the set of entities \mathbf{N} . In line 2, we set the initial state for each entity i by invoking the ‘Init’ operation of the protocol \mathcal{P} with inputs (i, p) ; in other words, each entity receives its identifier i and the security parameters p and performs its initialization operation. In line 3, we initialize e , which we use to index the events of the execution, i.e., e is incremented by one (line 5) each time we complete one ‘execution loop’ (lines 4-14).

Main execution loop (lines 4-14). The design allows the adversary \mathcal{A} to have a generous control over the execution. Specifically, in each event e , \mathcal{A} determines (line 6) an operation $opr[e]$ to be applied to an entity $ent[e] \in \mathbf{N}$, with input $inp[e]$ and its local clock value $clk[e]$. Additionally, the adversary selects the global, real time clock value set to $\tau[e]$.

After the adversary defines the specific operation and inputs (line 6), the event is executed (lines 7-11). There are five options for each event specified as an operation $opr[e]$. More concretely, if \mathcal{A} sets $opr[e] = \text{‘Set-state’}$ (line 7), then the state $s_{ent[e]}$ of entity $ent[e]$ is set to $inp[e]$; if \mathcal{A} sets $opr[e] = \text{‘Set-output’}$ (line 8) then the outputs $out[e], sec-out[e][\cdot]$ are set to values chosen by \mathcal{A} as $inp[e]$; and if \mathcal{A} sets $opr[e] = \text{‘Get-state’}$ (line 9), then the output of the state $s_{ent[e]}$ of entity $ent[e]$, including any private state, e.g., private keys, is exposed to \mathcal{A} and saved as $out[e]$.

Alternatively, if \mathcal{A} sets $opr[e] = \text{‘Sec-in’}$ (line 10), then entity $ent[e]$ can receive values securely from another entity. This is done through the *sec-out* mechanism. Specifically, \mathcal{A} chooses some event $\hat{e} = inp[e]$, and the protocol receives $ent[\hat{e}], sec-out[\hat{e}][ent[e]]$ as input, which is the identifier of the ‘sender’ ($ent[\hat{e}]$) along with the value(s) that were *securely* sent in event \hat{e} to entity $ent[e]$. Note that \mathcal{A} is never given the contents of *sec-out* during the execution process.

Otherwise, if $opr[e]$ is set to any other operation (line 11), then the process invokes the protocol \mathcal{P} specific function over the state $s_{ent[e]}$ of entity $ent[e]$, with inputs $opr[e], inp[e], clk[e]$. This results in a new state $s_{ent[e]}$ and output $out[e]$ for entity $ent[e]$. The execution process allows the adversary to set and get the state of any entity; however, a specific execution model \mathcal{M} may forbid such operations, e.g., return \perp for executions where the adversary performs them. In line 13, the adversary processes the output $out[\hat{e}]$ of the operation $opr[\hat{e}]$. The adversary may modify its state $s_{\mathcal{A}}$, and outputs a value $out_{\mathcal{A}}$; when $out_{\mathcal{A}} \neq \perp$, the execution moves to the termination phase; otherwise the loop continues.

Algorithm 1 Adversary-Driven Execution Process $\text{Exec-IA}_{\mathcal{A}, \mathcal{P}}(p)$

1: $(s_{\mathcal{A}}, \mathbf{N}) \leftarrow \mathcal{A}(p)$ ▷ Initialize adversary

2: $\forall i \in \mathbf{N} : s_i \leftarrow \mathcal{P}(\perp, \text{'Init'}, (i, p), \perp)$ ▷ Initialize entities' local state

3: $e \leftarrow 0$ ▷ Initialize loop's counter

4: **repeat**

5: $e \leftarrow e + 1$ ▷ Advance loop's counter

6: $(ent[e], opr[e], inp[e], clk[e], \tau[e]) \leftarrow \mathcal{A}(s_{\mathcal{A}})$ ▷ \mathcal{A} selects entity $ent[e]$, operation $opr[e]$, input $inp[e]$, clock $clk[e]$, and real time $\tau[e]$ for event e

7: **if** $opr[e] = \text{'Set-state'}$ **then** ▷ When \mathcal{A} wants to change the local state of entity $ent[e]$
 $s_{ent[e]}, out[e], sec-out[e][\cdot] \leftarrow inp[e], \perp, \perp$

8: **else if** $opr[e] = \text{'Set-output'}$ **then** ▷ When \mathcal{A} wants to change the output of entity $ent[e]$
 $out[e], sec-out[e][\cdot] \leftarrow inp[e]$

9: **else if** $opr[e] = \text{'Get-state'}$ **then** ▷ When \mathcal{A} wants to get the current local state of entity $ent[e]$
 $out[e], sec-out[e][\cdot] \leftarrow s_{ent[e]}, \perp$

10: **else if** $opr[e] = \text{'Sec-in'}$ **then** ▷ When \mathcal{A} wants entity $ent[e]$ to receive secure output
 \mathcal{A} specifies a previous event in $inp[e]$ and then the process uses $(ent[inp[e]], sec-out[inp[e]][ent[e]])$ as input to \mathcal{P}
 $(s_{ent[e]}, out[e], sec-out[e][\cdot]) \leftarrow \mathcal{P}(s_{ent[e]}, \text{'Sec-in'}, (ent[inp[e]], sec-out[inp[e]][ent[e]]), clk[e])$
 ▷ \mathcal{P} returns the state of entity $ent[e]$, the output, and the secure output

11: **else** ▷ Otherwise, \mathcal{A} wants to execute operation $opr[e]$ of \mathcal{P} with input $inp[e]$ over entity $ent[e]$ with local clock of $clk[e]$
 $(s_{ent[e]}, out[e], sec-out[e][\cdot]) \leftarrow \mathcal{P}(s_{ent[e]}, opr[e], inp[e], clk[e])$

12: **end if** ▷ Inform \mathcal{A} of the value of $out[e]$ and allow \mathcal{A} to decide whether to continue ($out_{\mathcal{A}} = \perp$), or to terminate the loop ($out_{\mathcal{A}} \neq \perp$)

13: $(s_{\mathcal{A}}, out_{\mathcal{A}}, \mathbf{F}) \leftarrow \mathcal{A}(s_{\mathcal{A}}, out[e])$ ▷ to continue ($out_{\mathcal{A}} = \perp$), or to terminate the loop ($out_{\mathcal{A}} \neq \perp$)

14: **until** $out_{\mathcal{A}} \neq \perp$

15: $T \leftarrow (out_{\mathcal{A}}, e, \mathbf{N}, \mathbf{F}, ent[\cdot], opr[\cdot], inp[\cdot], clk[\cdot], \tau[\cdot], out[\cdot], sec-out[\cdot][\cdot])$ ▷ Output

16: **Return** T

Termination (lines 15-16). Upon termination, the process returns the *execution transcript* T (line 16), containing the relevant values from the execution. Namely, T contains the values of $ent[\hat{e}]$, $opr[\hat{e}]$, $inp[\hat{e}]$, $clk[\hat{e}]$, $\tau[\hat{e}]$, $out[\hat{e}]$ and $sec-out[\hat{e}][\cdot]$ for each of the events $\hat{e} \leq e$, where e is the index of the last event, as well as the adversary’s output $out_{\mathcal{A}}$, the index of the last event e , the set of entities \mathbf{N} , and the set of faulty entities \mathbf{F} (produced in line 13). Private values such as entities’ private keys are not part of the execution transcript unless provided as $inp[\hat{e}]$ for $opr[\hat{e}] = \text{‘Set-state’}$ or properly extracted using $opr[\hat{e}] = \text{‘Get’}$. We allow \mathcal{A} to output \mathbf{F} , so an adversary model can specify which entities are included in \mathbf{F} (i.e., which entities are considered ‘faulty’).

Limitations. The execution process supports a large variety of models. For example, the adversary may control the state (‘Set-state’), learn the state (‘Get-state’), and directly control the output (‘Set-output’) of every party $ent[\hat{e}] \in \mathbf{N}$, but this may be restricted (or fully prohibited) by specific model \mathcal{M} , allowing different fault models (honest-but curious, threshold, adaptive, proactive, etc.). However, **Exec. IA** allows the adversary to control all inputs events, while typical definitions of confidentiality and indistinguishability specifications provide randomized inputs which are not directly observable by the adversary. To allow such definitions, we extend the **Exec. IA** process to the **Exec. CIA** process in section 3.

2.2 Specifications

To describe both the model *assumed by* a protocol or scheme (i.e., different adversary, communication, clock-synchronization assumptions, etc.) and the property/ies (relating to security, communication, etc.) *provided by* the protocol or scheme, we define *specifications* and *advantage functions*. Specifications are used to describe both the assumptions (referred to as the models) and the goals for protocols or schemes. We define a specification as a pair containing the *specification predicate* and the *base function*. Namely, a specification ξ is:

$$\xi = (\xi.\pi, \xi.\beta) \tag{1}$$

where $\xi.\pi$ is the specification predicate and $\xi.\beta$ is the base function.

A specification predicate takes an execution transcript T and a parameter p as input, and it returns \top if T satisfies the predicate π for the given value of p and \perp otherwise. (In the asymptotic definitions, the parameter p is a single unary value which we refer to as the *security parameter*.)

The base function gives the ‘base’ probability of success for an adversary; for example, for collision- or forgery-preventing specifications, the base function may be 2^{-l} , where l is the output block size of a hash function or a key length; for indistinguishability-based specifications (see section 3), the base function may be $\frac{1}{2}$ (a constant).

An advantage function ϵ is the maximum *advantage* probability that the adversary may have (under the specifications), *over* the ‘base’ probability. Namely, intuitively, a specification is *satisfied with advantage* ϵ as long as the adversary’s

success probability is below $\xi.\beta(p) + \epsilon(p)$. The sum $\xi.\beta(p) + \epsilon(p)$ for any given p should be in the range $[0, 1]$. However, for the asymptotic definitions, we do not use specific functions ϵ but rather say that a specification is *satisfied with negligible advantage* (definitions are given in the next subsections).

As mentioned above, specifications are used to describe both the assumptions and the goals of protocols or schemes - namely, there are two ‘types’ of specifications. The first are model specifications (referred to as *models*), which describe assumptions about adversary capabilities, clocks, communication channels, etc. The second are specifications which describe the qualities provided by the protocol or scheme. Note that specifications can easily be combined or defined in terms of each other; e.g., multiple sub-model predicates may be combined into a model predicate which is the conjunction of the sub-model predicates.

In the following subsections, we further describe specifications. First, we focus on models. We define what it means for an adversary to IA-satisfy a model with negligible advantage, and we give an example of a model predicate, namely, the predicate $\mathcal{M}_{\Delta_{com}}^{Com}$. Later, we define what it means for a protocol to IA-satisfy a specification with negligible advantage under some model, and we give an example of a specification predicate, namely, the predicate `AuthComRcv`.

2.3 Models

The execution process described in subsection 2.1 specifies the details of running a protocol \mathcal{P} against an adversary \mathcal{A} ; however, it does not, on its own, *restrict* the adversary, nor it considers the *adversary model*, i.e., the exact capabilities of the adversary. Furthermore, other assumptions, i.e., communication and synchronization assumptions are also decoupled from the execution process. The rationale behind this design is that the security of distributed systems and cryptographic protocols should be stated and analyzed with respect to a specific *model*. Therefore, by separating the execution process from the model, this separation allows to take a transcript T of any execution of protocol \mathcal{P} with adversary \mathcal{A} , and examine whether transcript T satisfies the predicate $\mathcal{M}.\pi$ of a specific model \mathcal{M} . By the term *model*, we mean a combination of *adversary capabilities*, *communication assumptions* and *clock-synchronization assumptions*:

Adversary capabilities: The adversary capabilities (often referred to as the adversary model), define the computational resources of the adversary, e.g., probabilistic polynomial time (PPT), as well as other capabilities, e.g., from cipher-text only (CTO) to chosen ciphertext attacks (CCA) (for encryption schemes).

Communication assumptions: The communication assumptions define the properties of the underlying communication mechanism, such as reliable/unreliable communication, FIFO or non-FIFO, authenticated or not, bounded/fixed delay or asynchronous, and so on.

Clock-synchronization assumptions: The clock-synchronization assumptions define the availability and properties of per-entity clocks. Common models include purely asynchronous clocks (no synchronization), bounded-drift clocks, and synchronized clocks.

This section is structured as follows. We begin with the definitions of what it means to *IA-satisfy a model \mathcal{M} with negligible advantage* in §2.3.1. Then, in §5.1.3, we give the $\mathcal{M}_{\Delta_{com}}^{Com}$ predicate as an example. In §5.1.1, §5.1.2, §5.1.3, and §5.1.4, we present additional examples of model predicates representing various assumptions.

2.3.1 Model Satisfying Adversary

Models are used to restrict the capabilities of the adversary as well as the events that can happen in the execution process. This includes limiting of the possible faults, defining initialization assumptions, and defining the communication and synchronization models. Hence, we ensure that a given adversary \mathcal{A} followed a given model \mathcal{M} 's restrictions in a given execution transcript T by applying $\mathcal{M}.\pi$ to T and p , where p is a unary value $p = 1^\kappa$ for some $\kappa \in \mathbb{N}$. We say that \mathcal{A} *IA-satisfies \mathcal{M} with negligible advantage* if for every constant c , there is some $\kappa' \in \mathbb{N}$ such that the probability that $\mathcal{M}.\pi(T, 1^\kappa)$ returns \perp is less than $\mathcal{M}.\beta(1^\kappa) + \kappa^c$ for all $\kappa > \kappa'$, as defined next.

Definition 1 (Adversary \mathcal{A} IA-satisfies model \mathcal{M} with negligible advantage). *Let \mathcal{A} be a PPT algorithm, and let $\mathcal{M} = (\mathcal{M}.\pi, \mathcal{M}.\beta)$ be a model, where $\mathcal{M}.\pi$ is a predicate (i.e., outputs \top or \perp) and $\mathcal{M}.\beta$ is a function $\mathcal{M}.\beta(1^\kappa) : \{1\}^* \rightarrow [0, 1]$. We say that adversary \mathcal{A} IA-satisfies model \mathcal{M} with negligible advantage, denoted as $\mathcal{A} \models_{poly} \mathcal{M}$, if for every protocol \mathcal{P} and every constant c , there is some $\kappa' \in \mathbb{N}$ such that for all $\kappa > \kappa'$, the probability that the model predicate $\mathcal{M}.\pi$, applied to a random resulting **Exec IA** execution transcript T and security parameter 1^κ , is not satisfied, is less than $\mathcal{M}.\beta(1^\kappa) + \kappa^c$,*

$$\mathcal{A} \models_{poly} \mathcal{M} \stackrel{def}{=} \forall \mathcal{P}, c, \exists \kappa' \in \mathbb{N} \text{ s.t. } \forall \kappa > \kappa' : \Pr \left[\begin{array}{l} \mathcal{M}.\pi(T, 1^\kappa) = \perp : \\ T \leftarrow \mathbf{Exec IA}_{\mathcal{A}, \mathcal{P}}(1^\kappa) \end{array} \right] < \mathcal{M}.\beta(1^\kappa) + \kappa^c \quad (2)$$

2.3.2 $\mathcal{M}_{\Delta_{com}}^{Com}$: bounded-delay communication model predicate

As an example of a model predicate, we present the $\mathcal{M}_{\Delta_{com}}^{Com}$ model predicate, which ensures reliable, bounded-delay delivery of messages. (Further model predicates can be found in section 5.) The base function used with the $\mathcal{M}_{\Delta_{com}}^{Com}$ model predicate would usually be $\beta(1^\kappa) = 0$, since an adversary should have no significant advantage to prevent the timely delivery of messages (regardless of the value of the security parameter).

The adversary decides on the function $opr[\hat{e}_R]$ to be invoked at every event \hat{e}_R as well as the input $inp[\hat{e}_R]$. We assume a convention for *send* and *receive* events as follows. The adversary causes a message receipt event by setting $opr[\hat{e}_R]$ to ‘Incoming’ and $inp[\hat{e}_R]$ to (m, i_S) (where m is the message and $i_S \in \mathbb{N}$ is the purported sender). We use dot notation to refer to the message ($inp[\hat{e}_R].m$) and to the sender ($inp[\hat{e}_R].i_S$). Also, we allow the sender $ent[\hat{e}_S]$ to specify, as part

of its output $out[\hat{e}_S]$, one or more triplets of the form $(\text{'send'}, m, i_R)$, indicating the sending of message m to $i_R \in \mathbf{N}$.

The $\mathcal{M}_{\Delta_{com}}^{Com}$ model predicate ensures reliable, bounded-delay delivery of messages sent. Assume that at event \hat{e}_S of the execution, the output $out[\hat{e}_S]$ generated by $ent[\hat{e}_S]$, includes a $(\text{'send'}, m, j)$ triplet, i.e., $ent[\hat{e}_S]$ sends message m to $j \in \mathbf{N}$. If the $\mathcal{M}_{\Delta_{com}}^{Com}$ model predicate is true for this execution, then after at most Δ_{com} , if the execution did not terminate already, then entity j would receive m from $ent[\hat{e}_S]$. The $\mathcal{M}_{\Delta_{com}}^{Com}$ model predicate is shown in Algorithm 11.

Algorithm 2 $\mathcal{M}_{\Delta_{com}}^{Com}(T, p)$ Predicate

```

1: return (
2:    $\forall \hat{e}_S$  s.t.  $1 \leq \hat{e}_S < T.e$ :
3:     if  $(\exists (\text{'send'}, m, i_R) \in T.out[\hat{e}_S])$ 
4:       and  $T.\tau[T.e] \geq T.\tau[\hat{e}_S] + \Delta_{com}$ 
5:       and  $T.ent[\hat{e}_S] \in T.N - T.F$  )
6:     then  $\exists \hat{e}_R > \hat{e}_S$ 
7:       s.t.  $T.\tau[\hat{e}_S] + \Delta_{com} \geq T.\tau[\hat{e}_R]$ 
8:       and  $T.ent[\hat{e}_R] = i_R$ 
9:       and  $T.opr[\hat{e}_R] = \text{'Receive'}$ 
10:      and  $T.inp[\hat{e}_R] = (m, T.ent[\hat{e}_S])$ 
)

```

\triangleright If the output includes a send triple
 \triangleright And execution did not terminate yet after Δ_{com} real time
 \triangleright And the entity is honest
 \triangleright Then there is a later event
 \triangleright Within Δ_{com} real time
 \triangleright Where the entity is the intended recipient in the send triple
 \triangleright And which is a receive event
 \triangleright And in which the entity receives the message from the sender

We remark that $\mathcal{M}_{\Delta_{com}}^{Com}$ only applies when both sender and recipient are honest (i.e., in $\mathbf{N} - \mathbf{F}$), and $\mathcal{M}_{\Delta_{com}}^{Com}$ only ensures delivery and bounded delay. This still allows receipt of duplicate messages, which may involve unbounded delay, and receipt of messages that were never sent. To simplify $\mathcal{M}_{\Delta_{com}}^{Com}$, we use the adversary-controlled $\tau[\cdot]$ values (line 6 of Algorithm 1).

2.4 Security Specifications

The execution process and the model facilitate precise definition of security specifications, using predicates (‘experiments’) over the results of the execution. This is separate from the validation of the model, allowing modular specification of security specifications. Namely, different works may reuse the same security specifications (and execution process) but use other, possibly more realistic (and more complex) models, expressing different adversary capabilities, restrictions on usage, and assumptions on communication and synchronization. Similarly, different works may reuse the same models to study additional security specifications. The separation between the definition of the model and of the specifications also allows definition of *generic specification predicates*. Generic specification predicates are applicable to different protocols and problems. We identify four generic specification predicates, that appear relevant to many security protocols. These specification predicates focus on attributes of messages, i.e., non-repudiation,

and on detection of misbehaving entities (see §5.1.7). This approach is quite different from the current way of defining security for cryptographic schemes and protocols; it takes some time and effort to get used to the separate model and security definitions. However, we found that with a little use, the advantages become clear and the approach becomes natural and convenient, facilitating modularity and reuse of specification predicates and model predicates, and allowing for proper comparison of security guarantees between different schemes.

2.4.1 Model-Secure Specifications

To complete the presentation of the execution process, we now discuss how it is used to define specific security specifications and properties and to analyze whether these properties are ensured by a given protocol \mathcal{P} , under given model \mathcal{M} , interacting with any PPT adversary \mathcal{A} . While we present *game/experiment*-based definitions, future work may consider other forms of definitions, such as simulation-based.

A protocol \mathcal{P} would typically have multiple security properties, i.e., satisfy multiple security specifications. We define a *security specification* ξ as a pair $(\xi.\pi, \xi.\beta)$, where $\xi.\pi$ a predicate and $\xi.\beta$ is the base function. Again, the parameter p is a unary security parameter, $p = 1^\kappa$. Let b be the outcome of $\xi.\pi$ applied to $(T, 1^\kappa)$, where T is the transcript of the execution process ($T = \mathbf{Exec.IA}_{\mathcal{A}, \mathcal{P}}(1^\kappa)$) and p are the parameters, i.e., $b \leftarrow \xi.\pi(T, 1^\kappa)$; if $b = \perp$ then we say that *specification predicate* $\xi.\pi$ *was not satisfied* in this execution of \mathcal{P} , or that the *adversary won* in this execution; and if $b = \top$, then we say that *specification predicate* $\xi.\pi$ *was satisfied* in this execution, or that the *adversary lost*. We say that \mathcal{P} *IA-satisfies* ξ *with negligible advantage under model* \mathcal{M} if for every PPT adversary \mathcal{A} which IA-satisfies \mathcal{M} with negligible advantage, for every constant c , there is some $\kappa' \in \mathbb{N}$ such that the probability that $\xi.\pi(T, 1^\kappa)$ returns \perp , where $T \leftarrow \mathbf{Exec.IA}_{\mathcal{A}, \mathcal{P}}(1^\kappa)$, is less than $\xi.\beta(1^\kappa) + \kappa^c$ for all $\kappa > \kappa'$, as given in the next definition.

Definition 2 (Protocol \mathcal{P} IA-satisfies specification ξ with negligible advantage under model \mathcal{M}). *Let $\xi = (\xi.\pi, \xi.\beta)$ be a specification, where $\xi.\pi$ is a predicate and $\xi.\beta$ is a function $\xi.\beta(1^\kappa) : \{1\}^* \rightarrow [0, 1]$. We say that protocol \mathcal{P} IA-satisfies specification ξ under model \mathcal{M} , denoted $\mathcal{P} \stackrel{\mathcal{M}}{\underset{\text{poly}}{\models}} \xi$, if for every PPT adversary \mathcal{A} that IA-satisfies \mathcal{M} with negligible advantage, for every constant c , there is some $\kappa' \in \mathbb{N}$ such that for all $\kappa > \kappa'$, the probability that the specification predicate $\xi.\pi$, applied to a random resulting $\mathbf{Exec.IA}$ execution transcript T and security parameter 1^κ , is not satisfied, is less than $\xi.\beta(1^\kappa) + \kappa^c$, i.e.:*

$$\mathcal{P} \stackrel{\mathcal{M}}{\underset{\text{poly}}{\models}} \xi \stackrel{\text{def}}{=} \forall \mathcal{A} \in \text{PPT}, c, \text{ if } \mathcal{A} \underset{\text{poly}}{\models} \mathcal{M} \text{ then } \exists \kappa' \in \mathbb{N} \text{ s.t. } \forall \kappa > \kappa' : \quad (3)$$

$$\Pr \left[\begin{array}{l} \xi.\pi(T, 1^\kappa) = \perp : \\ T \leftarrow \mathbf{Exec.IA}_{\mathcal{A}, \mathcal{P}}(1^\kappa) \end{array} \right] < \xi.\beta(1^\kappa) + \kappa^c$$

2.4.2 AuthComRcv : authenticated-sender communication specification predicate

As an example of a specification predicate, we give the AuthComRcv predicate. For more specification predicates, see section 5.

The AuthComRcv model predicate verifies the authentic-sender property for all ‘Incoming’ events, which means that $out[\hat{e}_R].i_S$ indeed sent this message to $ent[\hat{e}_R]$, during some previous event $\hat{e}_S < \hat{e}_R$. The AuthComRcv model predicate is shown in Algorithm 3.

Algorithm 3 AuthComRcv (T, p) Predicate

```

1: return (
2:    $\forall \hat{e}_R$  s.t.  $1 \leq \hat{e}_R \leq T.e$ :
3:     if  $T.opr[\hat{e}_R] = \text{‘Receive’}$ :
4:       and  $T.out[\hat{e}_R] \neq \perp$ 
5:       and  $T.ent[\hat{e}_R], T.out[\hat{e}_R].i_S \in T.N - T.F$ 
6:     then  $\exists \hat{e}_S < \hat{e}_R$ 
7:       s.t.  $T.opr[\hat{e}_S] = \text{‘Send’}$ :
8:       and  $T.inp[\hat{e}_S].m = T.out[\hat{e}_R].m$ 
9:       and  $T.inp[\hat{e}_S].i_R = T.ent[\hat{e}_R]$ 
10:      and  $T.ent[\hat{e}_S] = T.out[\hat{e}_R].i_S$ 
)

```

\triangleright For each ‘Receive’ event
 \triangleright If the ‘Receive’ event was successful
 \triangleright And both receiver and sender are honest
 \triangleright Then there is a previous event
 \triangleright Which was a ‘Send’ event
 \triangleright Where the input message was $T.out[\hat{e}_R].m$
 \triangleright And where the intended recipient was $T.ent[\hat{e}_R]$
 \triangleright And where the entity was the sender output in event \hat{e}_R

Usually, the base function used with the AuthComRcv model predicate would be $\beta = 2^{-l}$, where l is the length of some tags or signatures used by the protocol or scheme for authentication. In other words, we would usually allow the adversary to have probability 2^{-l} to cause an entity to receive a forged message, but not (significantly) more than 2^{-l} .

3 Confidentiality, Integrity, and Availability Models and Specifications

The execution process $\mathbf{Exec.IA}_{\mathcal{A}, \mathcal{P}}(p)$ allows the adversary to choose all of the inputs to \mathcal{P} during the execution process. However, definitions of properties such as confidentiality often include randomized inputs (for instance, to allow ‘indistinguishability game’-based definitions). Thus, we define an extended execution process, $\mathbf{Exec.CIA}_{\mathcal{A}, \mathcal{P}}(p, b)$, which allows such indistinguishability-based definitions and specifications. Note that the only difference between $\mathbf{Exec.IA}_{\mathcal{A}, \mathcal{P}}(p)$ and $\mathbf{Exec.CIA}_{\mathcal{A}, \mathcal{P}}(p, b)$ is the additional input bit and an additional special operation ‘Challenge’ which uses this bit. Consequently, we can use $\mathbf{Exec.CIA}_{\mathcal{A}, \mathcal{P}}(p, b)$ also to run executions equivalent to those of $\mathbf{Exec.IA}_{\mathcal{A}, \mathcal{P}}(p)$ (i.e., simulate

$\mathbf{Exec_IA}_{\mathcal{A},\mathcal{P}}(p)$) simply by using an adversary model in which \mathcal{A} does not invoke the ‘Challenge’ operation. Therefore, $\mathbf{Exec_CIA}_{\mathcal{A},\mathcal{P}}(p, b)$ is an extension of $\mathbf{Exec_IA}_{\mathcal{A},\mathcal{P}}(p)$.

3.1 The Exec.CIA Process

$\mathbf{Exec_CIA}_{\mathcal{A},\mathcal{P}}(p, b)$, shown in Algorithm 4, takes an additional input bit b that is not controlled or visible to \mathcal{A} . Secondly, as part of $\mathbf{Exec_CIA}$, we define an additional special operation ‘Challenge’. \mathcal{A} can choose the ‘Challenge’ operation and specify two values in $inp[\hat{e}]$ (line 11 of Algorithm 4). The first value, $inp[\hat{e}][0]$, should be one of the operations of \mathcal{P} . The second input, $inp[\hat{e}][1]$, should be a tuple of two lists, where each list contains possible inputs to the operation specified in $inp[\hat{e}][0]$. (As usual, further assumptions about the inputs can be specified using a model.) Then $\mathbf{Exec_CIA}_{\mathcal{A},\mathcal{P}}(p, b)$ invokes \mathcal{P} with $inp[\hat{e}][0]$ as the operation, the b^{th} component of $inp[\hat{e}][1]$ (indexed from 0) as the input, and the state and clock values specified by \mathcal{A} as before. Note that \mathcal{A} does not know which element of $inp[\hat{e}][1]$ was sent to \mathcal{P} as the input.

\mathcal{A} can use the output of this operation, as well as the outputs of other operations executed during the process, to guess the value of b (0 or 1). (For example, \mathcal{P} might output the encryption of $inp[\hat{e}][1][b]$, and \mathcal{A} might try to guess b based on this encryption.) As before, \mathcal{A} sets the value of $out_{\mathcal{A}}$ at the end of each event, and the final value of $out_{\mathcal{A}}$ is returned as part of the execution transcript T . Consequently, we can define specifications in terms of the probability that $T.out_{\mathcal{A}}$ equals b . (Usually, we do not want $\Pr[T.out_{\mathcal{A}} = b]$ to be significantly better than for a random guess of the value of b .)

3.2 Specifications

Specifications are the same as defined in subsection 2.2, except that the specification predicates take *three* inputs (including a bit b). We give an example of such a specification predicate in subsection 3.4.

3.3 Models

In definition 3, we define *adversary \mathcal{A} CIA-satisfies model \mathcal{M} with negligible advantage*. The differences between the asymptotic definitions of ‘IA-satisfying’ and ‘CIA-satisfying’ a model are that in the ‘CIA-satisfying’ definition, the model predicate takes the additional input bit b , and the probability is taken over a transcript returned by $\mathbf{Exec_CIA}_{\mathcal{A},\mathcal{P}}(p, b)$ (instead of $\mathbf{Exec_IA}_{\mathcal{A},\mathcal{P}}(p)$), for a random bit b .

Definition 3 (Adversary \mathcal{A} CIA-satisfies model \mathcal{M} with negligible advantage). *Let \mathcal{A} be a PPT algorithm, and let $\mathcal{M} = (\mathcal{M}.\pi, \mathcal{M}.\beta)$ be a model, where $\mathcal{M}.\pi$ is a predicate over (T, p, b) and $\mathcal{M}.\beta$ is a function $\mathcal{M}.\beta(1^\kappa) : \{1\}^* \rightarrow [0, 1]$. We say that adversary \mathcal{A} CIA-satisfies model \mathcal{M} with negligible advantage, denoted as $\mathcal{A} \stackrel{\text{poly}}{=} \mathcal{M}$ (note that we reuse the IA notation for simplicity), if for*

Algorithm 4 Exec_CIA_{A, P}(p, b)
Adversary-Driven Execution Process with Challenge Bit

1: $(s_A, N) \leftarrow \mathcal{A}(p)$ ▷ Initialize adversary

2: $\forall i \in N : s_i \leftarrow \mathcal{P}(\perp, \text{'Init'}, (i, p), \perp)$ ▷ Initialize entities' local state

3: $e \leftarrow 0$ ▷ Initialize loop's counter

4: **repeat**

5: $e \leftarrow e + 1$ ▷ Advance loop's counter

6: $(ent[e], opr[e], inp[e], clk[e], \tau[e]) \leftarrow \mathcal{A}(s_A)$ ▷ \mathcal{A} selects entity $ent[e]$, operation $opr[e]$, input $inp[e]$, clock $clk[e]$, and real time $\tau[e]$ for event e

7: **if** $opr[e] = \text{'Set-state'}$ **then** ▷ When \mathcal{A} wants to change the local state of entity $ent[e]$
 $s_{ent[e]}, out[e], sec-out[e][\cdot] \leftarrow inp[e], \perp, \perp$

8: **else if** $opr[e] = \text{'Set-output'}$ **then** ▷ When \mathcal{A} wants to change the output of entity $ent[e]$
 $out[e], sec-out[e][\cdot] \leftarrow inp[e]$

9: **else if** $opr[e] = \text{'Get-state'}$ **then** ▷ When \mathcal{A} wants to get the current local state of entity $ent[e]$
 $out[e], sec-out[e][\cdot] \leftarrow s_{ent[e]}, \perp$

10: **else if** $opr[e] = \text{'Sec-in'}$ **then** ▷ When \mathcal{A} wants entity $ent[e]$ to receive secure output
 $(s_{ent[e]}, out[e], sec-out[e][\cdot]) \leftarrow \mathcal{P}(s_{ent[e]}, \text{'Sec-in'}, (ent[inp[e]], sec-out[inp[e]][ent[e]]), clk[e])$
 ▷ \mathcal{A} specifies a previous event in $inp[e]$ and then the process uses $(ent[inp[e]], sec-out[inp[e]][ent[e]])$ as input to \mathcal{P}
 ▷ \mathcal{P} returns the state of entity $ent[e]$, the output, and the secure output
 ▷ When \mathcal{A} wants a challenge, \mathcal{A} specifies the operation in $inp[e][0]$ and two inputs in $inp[e][1]$, but only the b^{th} component of $inp[e][1]$ (indexed from 0) is used as input to \mathcal{P}

11: **else if** $opr[e] = \text{'Challenge'}$ **then** ▷ Otherwise, \mathcal{A} wants to execute operation $opr[e]$ of \mathcal{P} with input $inp[e]$ over entity $ent[e]$ with local clock of $clk[e]$
 $(s_{ent[e]}, out[e], sec-out[e][\cdot]) \leftarrow \mathcal{P}(s_{ent[e]}, inp[e][0], inp[e][1][b], clk[e])$

12: **else** ▷ \mathcal{P} returns the state of entity $ent[e]$, the output, and the secure output
 $(s_{ent[e]}, out[e], sec-out[e][\cdot]) \leftarrow \mathcal{P}(s_{ent[e]}, opr[e], inp[e], clk[e])$

13: **end if**

14: $(s_A, out_A, F) \leftarrow \mathcal{A}(s_A, out[e])$ ▷ Inform \mathcal{A} of the value of $out[e]$ and allow \mathcal{A} to decide whether to continue ($out_A = \perp$), or to terminate the loop ($out_A \neq \perp$)

15: **until** $out_A \neq \perp$

16: $T \leftarrow (out_A, e, N, F, ent[\cdot], opr[\cdot], inp[\cdot], clk[\cdot], \tau[\cdot], out[\cdot], sec-out[\cdot][\cdot])$ ▷ Output

17: **Return** T

every protocol \mathcal{P} , every constant c , and $b \stackrel{\$}{\leftarrow} \{0, 1\}$, there is some $\kappa' \in \mathbb{N}$ such that for all $\kappa > \kappa'$, the probability that the model predicate $\mathcal{M}.\pi$, applied to a random resulting **Exec_CIA** execution transcript T , security parameter 1^κ , and $b \in \{0, 1\}$, is not satisfied, is less than $\mathcal{M}.\beta(1^\kappa) + \kappa^c$,

$$\mathcal{A} \models_{\text{poly}} \mathcal{M} \stackrel{\text{def}}{=} \forall \mathcal{P}, c, b \stackrel{\$}{\leftarrow} \{0, 1\}, \exists \kappa' \in \mathbb{N} \text{ s.t. } \forall \kappa > \kappa' : \Pr \left[\begin{array}{l} \mathcal{M}.\pi(T, 1^\kappa, b) = \perp : \\ T \leftarrow \mathbf{Exec_CIA}_{\mathcal{A}, \mathcal{P}}(1^\kappa, b) \end{array} \right] < \mathcal{M}.\beta(1^\kappa) + \kappa^c \quad (4)$$

3.4 Security Specifications

In definition 4, we define *protocol \mathcal{P} CIA-satisfies specification ξ with negligible advantage under model \mathcal{M}* . Similarly to CIA models, the specification predicate takes the additional input bit b , and the probability is taken over a transcript returned by **Exec_CIA** $_{\mathcal{A}, \mathcal{P}}(p, b)$, for a random bit b .

Definition 4 (Protocol \mathcal{P} CIA-satisfies specification ξ with negligible advantage under model \mathcal{M}). Let $\xi = (\xi.\pi, \xi.\beta)$ be a specification, where $\xi.\pi$ is a predicate over (T, p, b) and $\xi.\beta$ is a function $\xi.\beta(1^\kappa) : \{1\}^* \rightarrow [0, 1]$. We say that *protocol \mathcal{P} CIA-satisfies specification ξ with negligible advantage under model \mathcal{M}* , denoted $\mathcal{P} \models_{\text{poly}}^{\mathcal{M}} \xi$ (note that we reuse the IA notation for simplicity), if for every PPT adversary \mathcal{A} that CIA-satisfies \mathcal{M} and $b \stackrel{\$}{\leftarrow} \{0, 1\}$, for every constant c , there is some $\kappa' \in \mathbb{N}$ such that for all $\kappa > \kappa'$, the probability that the specification predicate $\xi.\pi$, applied to a random resulting **Exec_CIA** execution transcript T , security parameter 1^κ , and $b \in \{0, 1\}$, is not satisfied, is less than $\xi.\beta(1^\kappa) + \kappa^c$, i.e.:

$$\mathcal{P} \models_{\text{poly}}^{\mathcal{M}} \xi \stackrel{\text{def}}{=} \forall \mathcal{A} \in \text{PPT}, b \stackrel{\$}{\leftarrow} \{0, 1\}, c, \text{ if } \mathcal{A} \models_{\text{poly}} \mathcal{M} \text{ then } \exists \kappa' \in \mathbb{N} \text{ s.t. } \forall \kappa > \kappa' : \Pr \left[\begin{array}{l} \xi.\pi(T, 1^\kappa, b) = \perp : \\ T \leftarrow \mathbf{Exec_CIA}_{\mathcal{A}, \mathcal{P}}(1^\kappa, b) \end{array} \right] < \xi.\beta(1^\kappa) + \kappa^c \quad (5)$$

As mentioned previously, we can now define specifications in terms of the probability that $T.out_{\mathcal{A}}$ equals b . In this way, we define the Indistinguishability Specification Predicate IND^π in Algorithm 5.

Algorithm 5 $\text{IND}^\pi(T, p, b)$ Predicate

```

1: return  $\neg$ (
2:    $T.out_{\mathcal{A}} = b$ 
3:   and  $\pi(T, p, b)$ 
)

```

▷ \mathcal{A} guessed the value of b correctly

By using specific predicates π , we can use the IND^π specification predicate to define more specific specification predicates; for example, we use the **MsgConf**

predicate, shown below, to define $\text{IND}^{\text{MsgConf}}$, which is used to define message confidentiality for the SecCom protocol.

The MsgConf Predicate

The MsgConf model predicate is given in Algorithm 6. The model ensures that:

- \mathcal{A} only asks for ‘Send’ challenges (since we are only concerned with whether or not \mathcal{A} can distinguish outputs of ‘Send’). We assume that \mathcal{P} includes a ‘Send’ operation, which takes as input a message and the identifier of the recipient.
- During all ‘Send’ challenges, messages are only sent from one specific entity i_S to one specific entity i_R
- During each ‘Send’ challenge, \mathcal{A} specifies two messages of equal length and the same recipient in the two possible inputs. That is, \mathcal{A} specifies $\text{inp}[\hat{e}][1] = ((m_0, i_R), (m_1, i_R))$, where $|m_0| = |m_1|$. This ensures that \mathcal{A} does not distinguish the messages based on their lengths.
- \mathcal{A} does not use the ‘Receive’ operation to decrypt any output of a ‘Send’ challenge.

Algorithm 6 $\text{MsgConf}(T, p, b)$ Predicate

```

1: return (
2:    $\exists i_S, i_R$  s.t. : ▷ There is one specific sender  $i_S$ 
3:      $\forall \hat{e}$  s.t.  $1 \leq \hat{e} \leq T.e$ : ▷ and one specific receiver  $i_R$ 
4:       if  $T.\text{opr}[\hat{e}] = \text{‘Challenge’}$  then ▷ S.t. in any ‘Challenge’ event
5:          $T.\text{inp}[\hat{e}][0] = \text{‘Send’}$  ▷ It is a ‘Send’ challenge
6:         and  $T.\text{ent}[\hat{e}] = i_S$  ▷ The entity is  $i_S$ 
7:         and  $|T.\text{inp}[\hat{e}][1][0].m| = |T.\text{inp}[\hat{e}][1][0].m|$  ▷  $\mathcal{A}$  chose two messages of equal
▷ length
8:         and  $T.\text{inp}[\hat{e}][1][0].i_R = T.\text{inp}[\hat{e}][1][0].i_R = i_R$  ▷ And the recipient for both mes-
▷ sages is  $i_R$ 
▷ And  $\mathcal{A}$  did not use a ‘Receive’
9:         and  $\nexists \hat{e}'$  s.t.  $T.\text{opr}[\hat{e}'] = \text{‘Receive’}$  ▷ event to decrypt the output of
▷ the challenge
10:        and  $T.\text{inp}[\hat{e}'].c = T.\text{out}[\hat{e}].c$ 
11:        and  $T.\text{ent}[\hat{e}'] = i_R$ 
12:        and  $T.\text{inp}[\hat{e}'].i_S = i_S$ 
)

```

4 Concrete Security Models and Specifications

We extend the framework to allow concrete security definitions.

4.1 The Exec Process

The execution transcripts returned by **Exec.IA** and **Exec.CIA** already allow for specifications that refer to the concrete security parameters such as block

length, key length, and number of queries; however, they do not provide information about the time (number of steps) taken by the adversary when it is invoked during the execution. For this reason, we define an execution process **Exec**, which runs **Exec.CIA** but additionally ensures that the total number of steps taken by \mathcal{A} is returned as part of the execution transcript.

To count the number of steps of \mathcal{A} , we use **StepCount**, which runs \mathcal{A} and counts the number of steps taken by \mathcal{A} . Namely:

$$\mathbf{Exec}_{\mathcal{A},\mathcal{P}}(p, b) = \mathbf{Exec.CIA}_{\mathbf{StepCount}(\mathcal{A}),\mathcal{P}}(1^\kappa, b) \quad (6)$$

Whenever **Exec.CIA** invokes the adversary with some input, e.g. $s_{\mathcal{A}}$, then $\mathbf{StepCount}(\mathcal{A})(s_{\mathcal{A}})$ is called. Then, **StepCount** gives $s_{\mathcal{A}}$ as input to \mathcal{A} , runs \mathcal{A} , and adds the number of steps taken by \mathcal{A} to a counter that keeps track of the total number of steps taken by \mathcal{A} . If \mathcal{A} outputs \perp , $\mathbf{StepCount}(\mathcal{A})$ also outputs \perp . When \mathcal{A} finally outputs a value $\neq \perp$ to end the execution, then $\mathbf{StepCount}(\mathcal{A})$ returns this value *and* the total number of steps taken by \mathcal{A} . Consequently, the resulting execution transcript T is the same as if **Exec.CIA** $_{\mathcal{A},\mathcal{P}}(1^\kappa, b)$ had been run, except that $T.out_{\mathcal{A}}$ additionally contains the total number of steps taken by \mathcal{A} .

4.2 Specifications

Specifications are defined as before, but now we finally use the advantage function as well, and we no longer assume that the parameter p is a unary security parameter. Instead, depending on the particular protocol or scheme under consideration, p may contain, e.g., numbers of queries, block length, and key length. Notice that this allows specifications to include restrictions on various parameters, including numbers of queries and the number of steps that \mathcal{A} takes. We use *Params* to indicate the set of possible (allowed) values for p . This is also the domain of both the base function and the advantage function.

We say that (a protocol or an adversary) *satisfies ξ with advantage ϵ* if, for a random resulting execution transcript T , $b \xleftarrow{\$} \{0, 1\}$, and $p \in Params$, the probability that $\xi.\pi(T, p, b)$ returns \perp is bounded by $\xi.\beta(p) + \epsilon(p)$.

4.3 Models

Definition 5 (Adversary \mathcal{A} satisfies model \mathcal{M} with advantage ϵ). *Let \mathcal{A} be a PPT algorithm, and let $\mathcal{M} = (\mathcal{M}.\pi, \mathcal{M}.\beta)$ be a model, where $\mathcal{M}.\pi$ is a predicate (i.e., outputs \top or \perp) and $\mathcal{M}.\beta$ is a function $\mathcal{M}.\beta(p) : Params \rightarrow [0, 1]$, where *Params* is the set of possible values of the parameters. Furthermore, let ϵ also be a function $\epsilon(p) : Params \rightarrow [0, 1]$, such that, for any given p , holds $\mathcal{M}.\beta(p) + \epsilon(p) \in [0, 1]$. We say that adversary \mathcal{A} satisfies model \mathcal{M} with advantage ϵ , denoted as $\mathcal{A} \models_{\epsilon} \mathcal{M}$, if for every protocol \mathcal{P} , every $p \in Params$, and $b \xleftarrow{\$} \{0, 1\}$ the probability that the model predicate $\mathcal{M}.\pi$, applied to a random*

resulting execution transcript T , $p \in Params$, and $b \stackrel{\$}{\leftarrow} \{0, 1\}$ is not satisfied, is bounded by $\xi.\beta(p) + \epsilon(p)$, i.e.:

$$\mathcal{A} \models_{\epsilon} \mathcal{M} \stackrel{def}{=} \forall \mathcal{P}, p \in Params, b \stackrel{\$}{\leftarrow} \{0, 1\} : \Pr \left[\begin{array}{l} \mathcal{M}.\pi(T, p, b) = \perp \\ T \leftarrow \mathbf{Exec}_{\mathcal{A}, \mathcal{P}}(p, b) \end{array} \right] \leq \mathcal{M}.\beta(p) + \epsilon(p) \quad (7)$$

4.4 Security Specifications

Definition 6 (Protocol \mathcal{P} satisfies specification ξ with advantage ϵ under model \mathcal{M}). Let $\xi = (\xi.\pi, \xi.\beta)$ be a specification, where $\xi.\pi$ is a predicate over (T, p, b) and $\xi.\beta$ is a function $\xi.\beta(p) : Params \rightarrow [0, 1]$, where $Params$ is the set of possible values of the parameters. Furthermore, let $\xi.\epsilon$ also be a function $\xi.\epsilon(p) : Params \rightarrow [0, 1]$, such that, for any given p , holds $\xi.\beta(p) + \epsilon(p) \in [0, 1]$, and let \mathcal{M} be a model. We say that protocol \mathcal{P} satisfies specification ξ with advantage ϵ under model \mathcal{M} , denoted $\mathcal{P} \models_{\epsilon}^{\mathcal{M}} \xi$, if for every PPT adversary \mathcal{A} that CIA-satisfies \mathcal{M} , every $p \in Params$, and $b \stackrel{\$}{\leftarrow} \{0, 1\}$, the probability that the specification predicate $\xi.\pi$, applied to a random resulting execution transcript T , $p \in Params$, and $b \in \{0, 1\}$, is not satisfied, is bounded by $\xi.\beta(p) + \epsilon(p)$, i.e.:

$$\mathcal{P} \models_{\epsilon}^{\mathcal{M}} \xi \stackrel{def}{=} \forall \mathcal{A} \in PPT, p \in Params, b \stackrel{\$}{\leftarrow} \{0, 1\}, \text{ if } \mathcal{A} \models \mathcal{M} \text{ then :} \Pr \left[\begin{array}{l} \xi.\pi(T, p, b) = \perp \\ T \leftarrow \mathbf{Exec}_{\mathcal{A}, \mathcal{P}}(p, b) \end{array} \right] \leq \xi.\beta(p) + \epsilon(p) \quad (8)$$

5 Further Specification Predicates

5.1 Model Predicates

5.1.1 $\mathcal{M}_{\text{SecInit}}^{\hat{r}\text{-rounds}}$: the \hat{r} -rounds secure initialization model predicate

Cryptographic protocols are often designed assuming a secure initialization process, e.g., assuming shared secret keys. However, in the execution process (Algorithm 1), entities can only communicate via the adversary. As a result, we cannot simply *assume* shared secret keys, but the entities can use their local randomness to generate secret keys, and they can communicate, using cryptography, to securely establish shared secret values. We next define a simple secure initialization model predicate, $\mathcal{M}_{\text{SecInit}}^{\hat{r}\text{-rounds}}$. This model predicate ensures \hat{r} secure ‘rounds’ of $|\mathbf{N}|$ events each, where in event \hat{e} (where $1 \leq \hat{e} \leq \hat{r} \cdot |\mathbf{N}|$) holds:

- Entities are invoked with the special operation ‘Init’, i.e., $opr[\hat{e}] = \text{‘Init’}$, and in ‘round robin’, i.e., $ent[\hat{e}] = \hat{e} \bmod |\mathbf{N}|$ (where $\mathbf{N} = \{1, 2, \dots\}$). Note, in particular, that this prevents the adversary, during the initialization, from invoking the special ‘Set-state’, ‘Set-output’, and ‘Get-state’ operations to control the state or output of an entity (‘Set-state’ or ‘Set-output’) or to expose the state of an entity (‘Get-state’).

- Authenticated, reliable communication. Namely, every message received by entity i_R from entity i_S at round $2 \leq r \leq \hat{r}$, was indeed sent by i_S in the previous round to i_R ; and vice versa, i.e., every message sent by i_S to i_R at round $1 \leq r \leq (\hat{r} - 1)$, is correctly received by i_R , from sender i_S , in the next round.

It is convenient to capture each of these two aspects by a separate model predicate, i.e.:

$$\mathcal{M}_{\text{SecInit}}^{\hat{r}\text{-rounds}}(T, 1^\kappa) = \mathcal{M}_{\text{InitOps}}^{\hat{r}\text{-rounds}}(T, 1^\kappa) \wedge \mathcal{M}_{\text{InitCom}}^{\hat{r}\text{-rounds}}(T, 1^\kappa) \quad (9)$$

where $\mathcal{M}_{\text{InitOps}}^{\hat{r}\text{-rounds}}$ captures the first aspect ('operations') and $\mathcal{M}_{\text{InitCom}}^{\hat{r}\text{-rounds}}$ captures the second aspect ('communications'). We now define each of these more precisely; for convenience, let $\mathbf{N} = \{1, 2, \dots\}$. The $\mathcal{M}_{\text{InitOps}}^{\hat{r}\text{-rounds}}$ predicate is shown in Algorithm 7.

Algorithm 7 $\mathcal{M}_{\text{InitOps}}^{\hat{r}\text{-rounds}}(T, p)$ Predicate

```

1: return (
2:    $\forall i \in T.\mathbf{N}, r \in \{0, \dots, \hat{r} - 1\}$  :  $\triangleright$  For each entity  $i$  and each round  $r$ 
3:      $(T.\text{ent}[i + r \cdot |T.\mathbf{N}|] = i)$   $\triangleright$  'Round robin'
4:     and  $(T.\text{opr}[i + r \cdot |T.\mathbf{N}|] = \text{'Init'})$   $\triangleright$  Operation is 'Init'
)

```

To define the $\mathcal{M}_{\text{InitCom}}^{\hat{r}\text{-rounds}}$ (sub)model predicate, we assume the following conventions. To cause an entity to receive a message, the adversary includes the triple ('auth_init_recv', m, i_S) (where m is the message and $i_S \in \mathbf{N}$ is the purported sender) as part of the input $\text{inp}[\hat{e}_R]$. The sender indicates the sending of message m to $i_R \in \mathbf{N}$ by specifying, as part of its output $\text{out}[\hat{e}_S]$, a triplet ('auth_init_send', m, i_R). The $\mathcal{M}_{\text{InitCom}}^{\hat{r}\text{-rounds}}$ predicate is shown in Algorithm 8.

Algorithm 8 $\mathcal{M}_{\text{InitCom}}^{\hat{r}\text{-rounds}}(T, p)$ Predicate

```

1: return (
2:    $\forall i_S, i_R \in T.\mathbf{N}, r \in \{0, \dots, \hat{r} - 2\}, m \in \{0, 1\}^*$  :  $\triangleright$  For pairs of entities in  $T.\mathbf{N}$ , for each round except the last one, and for any message  $m$ 
3:      $(\text{'auth\_init\_send'}, m, i_R) \in T.\text{out}[i_S + r' \cdot |T.\mathbf{N}|]$   $\triangleright$   $i_S$  sends  $m$  to  $i_R$ 
4:     if and only if
5:      $(\text{'auth\_init\_recv'}, m, i_S) \in T.\text{inp}[i_R + (r' + 1) \cdot |T.\mathbf{N}|]$   $\triangleright$   $i_R$  receives  $m$  from  $i_S$  in the next round
)

```

5.1.2 $\mathcal{M}^{|\mathbf{F}| \leq f}$: up to f Byzantine (malicious) faults model predicate

We next define $\mathcal{M}^{|\mathbf{F}| \leq f}$, a specific adversary model predicate allowing the adversary to choose, and completely control, up to f of the entities in \mathbf{N} . We refer

to such failures, where the adversary is allowed to completely control the entity, as *malicious* or *Byzantine* faults. We use f as a function applied to the total number of entities $|\mathbb{N}|$. We refer to this particular faults model as $\mathcal{M}^{|\mathbb{F}| \leq f}$, where $f : \mathbb{N} \rightarrow \mathbb{N}$ bounds the number of faulty entities as a function of the total number of entities. Specifically, the adversary may corrupt entities by performing the ‘Get-state’, ‘Set-state’, and ‘Set-output’ operations. To enforce the model predicate, we simply ensure that the ‘Get-state’, ‘Set-state’, and ‘Set-output’ operations can be applied only to entities in \mathbb{F} , and that $|\mathbb{F}| \leq f(|\mathbb{N}|)$. We define the $\mathcal{M}^{|\mathbb{F}| \leq f}$ predicate in Algorithm 9.

Algorithm 9 $\mathcal{M}^{|\mathbb{F}| \leq f}(T, p)$ Predicate

```

1: return (
2:   ( $|T.\mathbb{F}| \leq f(|T.\mathbb{N}|)$ )                                      $\triangleright$  Max size of  $T.\mathbb{F}$  is not exceeded
3:   and  $\forall \hat{e}$  s.t.  $1 \leq \hat{e} \leq T.e$ :                              $\triangleright$  For each event
4:     if  $T.opr[\hat{e}] \in \{\text{‘Get-state’}, \text{‘Set-state’}, \text{‘Set-output’}\}$   $\triangleright$  If operation is ‘Get-state’, ‘Set-
5:     then  $T.ent[\hat{e}] \in T.\mathbb{F}$                                         $\triangleright$  state’, or ‘Set-output’
                                                                     $\triangleright$  Then entity is in  $T.\mathbb{F}$ 
)

```

PoC-PKI, our proof-of-concept PKI, uses $\mathcal{M}^{|\mathbb{F}| \leq f}$ with f defined as: $f(n) = \lfloor (n/3) \rfloor$. Namely, the faulty entities can arbitrarily misbehave so long the adversary controls at most a third of all entities. Note that $\mathcal{M}^{|\mathbb{F}| \leq f}$ only restricts the set \mathbb{F} of faulty entities, which does not yet restrict the adversary’s ability to interfere with the communication, clock, and local inputs of any entity - including non-faulty entities (i.e., in $\mathbb{N} - \mathbb{F}$). By using $\mathcal{M}^{|\mathbb{F}| \leq f}$ in conjunction with additional model predicates, e.g., $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom}}$ and $\mathcal{M}_{\Delta_{clk}}^{\text{CLK}}$ (described next), we can also restrict the adversary in such additional ways, as desired for a particular analysis. Notice, in particular, that $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom}}$, which defines restrictions on the communication events, completely excludes events where the sender or recipient are faulty (i.e., in \mathbb{F}). This allows the adversary to completely control all of faulty entities.

5.1.3 $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom}}$: authentic-sender, bounded-delay communication model predicate

We next present $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom}}$, an authentic-sender, bounded-delay communication model predicate. It is convenient to define $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom}}$ as a conjunction of two simpler predicates: $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom-rcv}}$, ensuring authentic-sender for message-receive events, and $\mathcal{M}_{\Delta_{com}}^{\text{Com}}$, ensuring reliable, bounded-delay for message-send events. Namely:

$$\mathcal{M}_{\Delta_{com}}^{\text{AuthCom}}(T, 1^\kappa) = \mathcal{M}_{\Delta_{com}}^{\text{Com}}(T, 1^\kappa) \wedge \mathcal{M}_{\Delta_{com}}^{\text{AuthCom-rcv}}(T, 1^\kappa) \quad (10)$$

We first present $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom-rcv}}$, which ensures authentic-sender for message-receive events. The adversary decides on the function $opr[\hat{e}_R]$ to be invoked at

every event \hat{e}_R as well as the input $inp[\hat{e}_R]$. We assume a convention for *send* and *receive* events as follows. The adversary causes a message receipt event by setting $opr[\hat{e}_R]$ to ‘Receive’ and $inp[\hat{e}_R]$ to (m, i_S) (where m is the message and $i_S \in \mathbf{N}$ is the purported sender). We use dot notation to refer to the message ($inp[\hat{e}_R].m$) and to the sender ($inp[\hat{e}_R].i_S$). Also, we allow the sender $ent[\hat{e}_S]$ to specify, as part of its output $out[\hat{e}_S]$, one or more triplets of the form $(\text{‘send’}, m, i_R)$, indicating the sending of message m to $i_R \in \mathbf{N}$.

The authentic-sender property ($\mathcal{M}_{\Delta_{com}}^{\text{AuthCom-rcv}}$ model predicate) implies that $inp[\hat{e}_R].i_S$ indeed sent this message to $ent[\hat{e}_R]$, during some previous event $\hat{e}_S < \hat{e}_R$. The $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom-rcv}}$ model predicate is shown in Algorithm 10.

Algorithm 10 $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom-rcv}}$ (T, p) Predicate

```

1: return (
2:    $\forall \hat{e}_R$  s.t.  $1 \leq \hat{e}_R \leq T.e$ :
3:     if  $T.opr[\hat{e}_R] = \text{‘Receive’}$ :
4:       and  $T.ent[\hat{e}_R], T.inp[\hat{e}_R].i_S \in T.N - T.F$ 
5:       then  $\exists \hat{e}_S < \hat{e}_R$ 
6:         s.t.  $(\text{‘send’}, T.inp[\hat{e}_R].m, T.ent[\hat{e}_R]) \in T.out[\hat{e}_S]$ 
7:         and  $T.ent[\hat{e}_S] = T.inp[\hat{e}_R].i_S$ 
)

```

\triangleright For each message-receive event
 \triangleright If both receiver and purported sender are honest
 \triangleright Then there is a previous event
 \triangleright In which an entity sent the message to the receiver
 \triangleright And that entity was the purported sender

The $\mathcal{M}_{\Delta_{com}}^{\text{Com}}$ model predicate ensures reliable, bounded-delay delivery of messages sent. Assume that at event \hat{e}_S of the execution, the output $out[\hat{e}_S]$ generated by $ent[\hat{e}_S]$, includes a $(\text{‘send’}, m, j)$ triplet, i.e., $ent[\hat{e}_S]$ sends message m to $j \in \mathbf{N}$. If the $\mathcal{M}_{\Delta_{com}}^{\text{Com}}$ model predicate is true for this execution, then after at most Δ_{com} , if the execution did not terminate already, then entity j would receive m from $ent[\hat{e}_S]$. The $\mathcal{M}_{\Delta_{com}}^{\text{Com}}$ model predicate is shown in Algorithm 11.

Algorithm 11 $\mathcal{M}_{\Delta_{com}}^{\text{Com}}$ (T, p) Predicate

```

1: return (
2:    $\forall \hat{e}_S$  s.t.  $1 \leq \hat{e}_S < T.e$ :
3:     if  $(\exists (\text{‘send’}, m, i_R) \in T.out[\hat{e}_S])$ 
4:       and  $T.\tau[T.e] \geq T.\tau[\hat{e}_S] + \Delta_{com}$ 
5:       and  $T.ent[\hat{e}_S] \in T.N - T.F$ 
6:       then  $\exists \hat{e}_R > \hat{e}_S$ 
7:         s.t.  $T.\tau[\hat{e}_S] + \Delta_{com} \geq T.\tau[\hat{e}_R]$ 
8:         and  $T.ent[\hat{e}_R] = i_R$ 
9:         and  $T.opr[\hat{e}_R] = \text{‘Receive’}$ 
10:        and  $T.inp[\hat{e}_R] = (m, T.ent[\hat{e}_S])$ 
)

```

\triangleright If the output includes a send triple
 \triangleright And execution did not terminate yet after Δ_{com} real time
 \triangleright And the entity is honest
 \triangleright Then there is a later event
 \triangleright Within Δ_{com} real time
 \triangleright Where the entity is the intended recipient in the send triple
 \triangleright And which is a receive event
 \triangleright And in which the entity receives the message from the sender

We remark that: $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom}}$ only applies when both sender and recipient are honest (i.e., in $\mathbf{N} - \mathbf{F}$); $\mathcal{M}_{\Delta_{com}}^{\text{AuthCom}}$ only ensures delivery, sender authentication and bounded delay. This still allows receipt of duplicate messages, which may involve unbounded delay. To simplify $\mathcal{M}_{\Delta_{com}}^{\text{Com}}$, we use the adversary-controlled $\tau[\cdot]$ values (line 6 of Algorithm 1). For this to be meaningful, we depend on the synchronization properties of the $\mathcal{M}_{\Delta_{clk}}^{\text{CLK}}$ model predicate, discussed next.

5.1.4 $\mathcal{M}_{\Delta_{clk}}^{\text{CLK}}$: bounded-drift clock synchronization assumptions

Finally, We present $\mathcal{M}_{\Delta_{clk}}^{\text{CLK}}$, which models the bounded-drift clock synchronization assumptions. We split this into two predicates: $\mathcal{M}_{\Delta_{clk}}^{\text{Drift}}$, which limits the drift between the clock values $clk[\hat{e}]$ (provided by the adversary as input to the protocol) and the real time values $\tau[\hat{e}]$; and $\mathcal{M}_{\Delta_{clk}}^{\text{Wake-up}}$, which provides a ‘wake-up service’ to the protocol. Namely:

$$\mathcal{M}_{\Delta_{clk}}^{\text{CLK}}(T, 1^\kappa) = \mathcal{M}_{\Delta_{clk}}^{\text{Drift}}(T, 1^\kappa) \wedge \mathcal{M}_{\Delta_{clk}}^{\text{Wake-up}}(T, 1^\kappa) \quad (11)$$

We begin with $\mathcal{M}_{\Delta_{clk}}^{\text{Drift}}$, which bounds the clock drift. It enforces two requirements on the execution: each local-clock value ($clk[\hat{e}]$) must be within Δ_{clk} drift from the real time $\tau[\hat{e}]$, and the real time values should be monotonously increasing. See Algorithm 12.

Algorithm 12 $\mathcal{M}_{\Delta_{clk}}^{\text{Drift}}(T, p)$ Predicate

```

1: return (
2:    $\forall \hat{e} \text{ s.t. } 1 \leq \hat{e} \leq T.e:$                                  $\triangleright$  For each event
3:      $|T.clk[\hat{e}] - T.\tau[\hat{e}]| \leq \Delta_{clk}$                         $\triangleright$  Local clock is within  $\Delta_{clk}$  drift
4:   and                                                          $\triangleright$  from real time
5:     if  $\hat{e} \geq 2$                                                 $\triangleright$  And if  $\hat{e}$  is not the first event
6:     then  $T.\tau[\hat{e}] \geq T.\tau[\hat{e} - 1]$                         $\triangleright$  Then the real time is  $\geq$  the real
                                                                    time at the previous event
)

```

As a special case, when $\Delta_{clk} = 0$, this function defines a model where the local clocks are fully synchronized, i.e., there is no difference between entities’ clocks. Finally, $\mathcal{M}_{\Delta_{clk}}^{\text{Wake-up}}$ provides a ‘wake-up service’ allowing the protocol to perform time-driven activities and ensuring that appropriate functions are invoked properly. This is ensured by requiring that if (*Sleep*, x) was part of the output $out[\hat{e}]$ (indicating that entity $ent[\hat{e}]$ was ‘put to sleep’ for x time) and execution did not terminate by ‘real’ time $\tau[\hat{e}] + x + \Delta_{clk}$, then at some event $\hat{e}' > \hat{e}$ (where $\tau[\hat{e}']$ was within Δ_{clk} from $\tau[\hat{e}] + x$), the same entity ($ent[\hat{e}']$) was indeed ‘Woken up’. The $\mathcal{M}_{\Delta_{clk}}^{\text{Wake-up}}$ predicate appears in Algorithm 13.

Algorithm 13 $\mathcal{M}_{\Delta_{clk}}^{\text{Wake-up}}(T, p)$ Predicate

```
1: return (
2:    $\forall \hat{e}$  s.t.  $1 \leq \hat{e} \leq T.e$ :
3:     if ( ('Sleep',  $x$ )  $\in T.out[\hat{e}]$ 
4:         and  $T.\tau[T.e] \geq T.\tau[\hat{e}] + x + \Delta_{clk}$  )
5:       then  $\exists \hat{e}' > \hat{e}$ 
6:         s.t.  $|T.\tau[\hat{e}'] - T.\tau[\hat{e}] - x| \leq \Delta_{clk}$ 
7:         and  $T.ent[\hat{e}'] = T.ent[\hat{e}]$ 
8:         and  $T.opr[\hat{e}'] = \text{'Wake-up'}$ 
   )
```

▷ For each event \hat{e}
▷ If the output includes a
▷ ('Sleep', x) tuple
▷ And execution did not terminate yet after $x + \Delta_{clk}$ real time
▷ Then there is a later event
▷ With real time x greater than at \hat{e} (within Δ_{clk})
▷ In which the entity is the same as in \hat{e}
▷ And the operation is 'Wake-up'

5.1.5 Securely shared keys ($\mathcal{M}_{\text{SecKeyShare}}$) specification predicate

Sometimes, especially when analyzing shared-key protocols, it may be useful to *assume* that the entities securely share keys before communicating with each other - i.e., to assume that the keys are correctly shared without the interference of the adversary. For this purpose, we define the model $\mathcal{M}_{\text{SecKeyShare}}$, shown in Algorithm 14, which ensures that values (keys) are securely shared (using the *sec-out* mechanism in the execution process) before entities send or receive messages. Specifically, before any 'Send' or 'Receive' event, the sender of the message must have securely shared values with the receiver, and the receiver of the message must have securely shared values with the sender. We assume that the operation 'Sec-channel-setup' is used to share these values.

5.1.6 Verifiable Attribution of Statements (VAS) specification predicate

The output of many protocols may include *attributable statements*. An *attributable statement* is a tuple (m, σ, i) , where m is a string, $i \in \mathbb{N}$ is the *purported origin* of the statement, and σ provides *evidence* (typically, a signature), allowing attribution of statement m to entity i . We next explain the *validation* process, which uses the evidence σ to establish if i has, in fact, originated m .

We focus on the typical case, where attribution is based on the use of a *digital signature scheme* \mathcal{S} , applied by the protocol \mathcal{P} . Namely, σ is the result of applying the signing algorithm $\mathcal{S}.\text{Sign}$ to the message m , using some (private) signing key sk belonging to the origin i . Therefore, we say that the attributable statement (m, σ, i) is *valid*, i.e., that σ really 'proves' that i is the origin of m , if $\mathcal{S}.\text{Ver}(pk, m, \sigma) = \top$, where pk is the public signature-verification key of i , i.e., the public key that validates signatures computed using sk . This *attributes* the message m to the 'owner' of the public key pk (and the corresponding signing key sk). To attribute m to i , it remains to establish the association between i and the public key pk , i.e., to attribute pk , and messages verified by it, to i . We focus on protocols where this association is known and secure ('off-band'), e.g., CA public keys in PKI schemes.

Algorithm 14 $\mathcal{M}_{\text{SecKeyShare}}(T, p)$ Predicate

```

1: return  $\top$  if (
2:    $\forall \hat{e}''$  s.t.  $T.\text{opr}[\hat{e}''] = \text{'Send'}$  or  $T.\text{opr}[\hat{e}''] = \text{'Receive'}$ :
3:     if  $T.\text{opr}[\hat{e}''] = \text{'Send'}$ 
4:       then  $i_S, i_R \leftarrow T.\text{ent}[\hat{e}''], T.\text{inp}[\hat{e}''].i_R$ 
5:       else if  $T.\text{opr}[\hat{e}''] = \text{'Receive'}$ 
6:         then  $i_S, i_R \leftarrow T.\text{inp}[\hat{e}''].i_S, T.\text{ent}[\hat{e}'']$ 
7:      $\exists \hat{e}, \hat{e}'$ 
8:       s.t.  $\hat{e} < \hat{e}' < \hat{e}''$ 
9:       and  $T.\text{opr}[\hat{e}] = \text{'Sec-channel-setup'}$ 
10:      and  $T.\text{ent}[\hat{e}] = i_S$ 
11:      and  $T.\text{sec-out}[\hat{e}][i_R] \neq \perp$ 
12:      and  $T.\text{opr}[\hat{e}'] = \text{'Sec-in'}$ 
13:      and  $T.\text{ent}[\hat{e}'] = i_R$ 
14:      and  $T.\text{inp}[\hat{e}'] = \hat{e}$ 
15:   and  $\exists \hat{e}, \hat{e}'$ 
16:     s.t.  $\hat{e} < \hat{e}' < \hat{e}''$ 
17:     and  $T.\text{opr}[\hat{e}] = \text{'Sec-channel-setup'}$ 
18:     and  $T.\text{ent}[\hat{e}] = i_R$ 
19:     and  $T.\text{sec-out}[\hat{e}][i_S] \neq \perp$ 
20:     and  $T.\text{opr}[\hat{e}'] = \text{'Sec-in'}$ 
21:     and  $T.\text{ent}[\hat{e}'] = i_S$ 
22:     and  $T.\text{inp}[\hat{e}'] = \hat{e}$ 
23: )

```

\triangleright For each message 'Send' event
 \triangleright and each message 'Receive' event
 \triangleright Determine the sender and receiver of the message

\triangleright Previously (before the 'Send' or 'Receive' event)
 \triangleright Keys were securely sent
 \triangleright From the sender of the message
 \triangleright To the receiver of the message
 \triangleright And they were securely received
 \triangleright By the receiver of the message

\triangleright And also previously (before the 'Send' or 'Receive' event)
 \triangleright Keys were securely sent
 \triangleright From the receiver of the message
 \triangleright To the sender of the message
 \triangleright And securely received
 \triangleright By the sender of the message

We formalize this by assuming that each entity $i \in \mathbf{N}$ *identifies* its public key pk by outputting the pair ('public key', pk) $\in out[\hat{e}]$, in some event \hat{e} ; namely, we use 'public key' as a 'label', to identify output of the public key. Typically, entities output the public key when they generate the key, i.e., $ent[\hat{e}] = i$, possibly as an initialization operation, i.e. $opr[\hat{e}] = 'Init'$. Notice that entities may often also send their public keys to each other using the ('send', m, i_R) output convention described in § 5.1.3; however, we prefer to keep the two conventions separate, since we believe that not every protocol that uses verification of attribution would necessarily send public keys in precisely the same way.

More precisely, the following *Key Attribution Predicate* V_{ka} outputs \top if entity i has identified pk as its public key in a given transcript T output by an execution of the protocol \mathcal{P} (Algorithm 1):

$$V_{ka}(i, pk, T) = \{\exists \hat{e} \text{ s.t. } T.ent[\hat{e}] = i \wedge ('public\ key', pk) \in T.out[\hat{e}]\} \quad (12)$$

We now define the Verifiable Attribution of Statements specification predicate. The adversary \mathcal{A} 'wins' in the experiment if its output $out_{\mathcal{A}}$ includes both a valid attributable statement (m, σ, i) for non-faulty entity $i \in \mathbf{N} - \mathbf{F}$ and a verification key pk associated with i , yet i did *not* originate m . To allow us to identify events \hat{e} in which an entity $i = ent[\hat{e}]$ intentionally signed message m , we adopt the following convention: whenever signing a message m , the party adds the pair ('signed', m) as part of its output, i.e., ('signed', m) $\in out[\hat{e}]$. Since this is *always* done, whenever the protocol signs a message, we will *not* explicitly include the ('signed', m) pairs as part of the output, which would make the pseudo-code cumbersome. Note that often the entity will also send the signed message, however, different protocols may send in different ways, hence this convention makes it easier to define the specification predicate.

The specification predicate is defined with respect to specific signature scheme \mathcal{S} , and the V_{ka} predicate defined above (Eq. 12). For simplicity, and since \mathcal{S} is typically obvious (as part of \mathcal{P}), we do not explicitly specify \mathcal{S} as a parameter of the specification predicate. The VAS predicate is shown in Algorithm 15.

Algorithm 15 Verifiable Attribution of Statements Predicate $VAS(T, p)$

```

1:  $(m, \sigma, i, pk) \leftarrow T.out_{\mathcal{A}}$ 
2: return  $\neg$ (
3:    $i \in T.N - T.F$   $\triangleright i$  is an honest entity
4:   and  $\mathcal{S}.Ver(pk, m, \sigma) = \top$   $\triangleright m$  was signed by the owner of  $pk$ 
5:   and  $V_{ka}(i, pk, T) = \top$   $\triangleright i$  identified  $pk$  as its public key
6:   and  $\nexists \hat{e}$  s.t.:
        $T.ent[\hat{e}] = i$ 
       and ('signed',  $m$ )  $\in T.out[\hat{e}]$   $\triangleright$  Yet,  $i$  never indicated that it signed  $m$ 
)

```

5.1.7 Generic misbehavior detection specification predicates

Many security protocols are required to be *resilient to misbehaviors*, i.e., to achieve their goals even if some of the entities, say entities in $F \subset N$, are faulty, and may misbehave (arbitrarily or in some specified manner). This resiliency to faulty, misbehaving entities is often based on *detection of misbehavior*; furthermore, often, many security protocols are required only to *detect misbehaviors*, which would be followed by taking some additional measures to deter and/or neutralize an attack. While misbehavior can be detected in different ways, detection is typically based either on some *evidence* that a certain entity is dishonest, where the evidence should be *verifiable by any third party*, or based on an *accusation*, where one entity (the *accuser*) accuses another entity (the *suspect*) of some misbehavior. Such an accusation may not be true, and therefore, it is harder to use this approach to deter and/or neutralize the attack; however, many misbehaviors do not leave any *evidence* verifiable by a third party, in which case, accusations may provide some security benefits, e.g., *detection* of the attack. A typical example of such misbehavior that does not leave any evidence is when a party *fails to act* in a required way, e.g., to send a required message or response; such failure may be plausibly blamed on communication issues, or on failure of the intended recipient. Often, a party, say Alice, detects such failure, say of Mal, to send a required message, after Alice waits for some *maximum delay*, and then Alice issues an ‘accusation’ against Mal, to alert others; for example, see [4]. An honest entity would only accuse a misbehaving party; however, because an accusation cannot be verified, a misbehaving entity could falsely accuse anyone, even an honest entity.

To formalize these concepts, we define two specification predicates: one to ensure that honest entities cannot be ‘framed’ as misbehaving, i.e., evidences are always verifiable with correct outcome, and another one to express that honest entities never accuse other honest entities, i.e., only accuse misbehaving entities.

The Non-frameability specification and Proof of Misbehavior. The first security specification predicate is called *non-frameability* (of honest entities), and ensures that a specific protocol would not allow any entity to produce a *valid Proof of Misbehavior* of a non-faulty entity. The specification predicate is therefore defined with respect to a given *Proof of Misbehavior Validation Predicate* V_{PoM} , which receives two inputs: a Proof-Validation Key pk and a purported-proof ζ . The output of $V_{PoM}(pk, \zeta)$ is \top if and only if ζ is a valid Proof of Misbehavior, as indicated by pk ; i.e., a misbehavior by an entity who knows the corresponding private key, typically, the ‘owner’ of pk , which can be validated using the Key Attribution Predicate V_{ka} . The natural way is to define the Proof of Misbehavior Validation Predicate V_{PoM} to be *protocol specific*, as the notions of misbehavior, and valid proof of misbehavior, depend on the specific protocol specifications. We specify for \mathcal{P} a special *stateless* operation $opr = V_{PoM}$, which does not modify the state or depend on it, or on the local clock. Abusing notation, we denote this operation simply as $\mathcal{P}.V_{PoM}(pk, \zeta)$. The use of a protocol-defined $\mathcal{P}.V_{PoM}$ allows us to define, below, the Non-frameability specification predicate.

Let $V_{PoM} : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{\top, \perp\}$ be a predicate. The Non-frameability predicate, shown in Algorithm 16, returns \perp if the adversary was able to output a Proof of Misbehavior for an honest entity, and \top otherwise.

Algorithm 16 Non-frameability Predicate $NF(T, p)$

```

1:  $(i, \zeta, pk) \leftarrow T.out_A$ 
2: return  $\neg$ (
3:    $V_{ka}(i, pk, T)$   $\triangleright$   $i$  identified  $pk$  as its public key
4:   and  $\mathcal{P}.V_{PoM}(pk, \zeta)$   $\triangleright$   $\zeta$  is a valid Proof of Misbehavior by the owner of  $pk$ 
5:   and  $i \in T.N - T.F$   $\triangleright$   $i$  is an honest entity
)

```

5.1.8 Accusations and the No False Accusations specification predicate

Recall that in the execution process, the adversary can use the ‘Set-output’, ‘Set-state’, and ‘Get-state’ operations to set the output and the state of a party and to learn the state of a party; we refer to such party as *faulty*, and denote by F the set of faulty parties in an execution. In many protocols, one party, say Alice, may *detect* that another party, say Mal, is faulty, typically, by receiving an invalid message from Mal - or simply by *not* receiving a message expected from Mal by a specific ‘deadline’ (for bounded-delay communication models).

Intuitively, the *No False Accusations (NFA)* specification predicate states that a non-faulty entity $a \notin F$ (Alice), would *never (falsely) accuse* of a fault, another non-faulty entity, $b \notin F$ (Bob). To properly define this specification predicate, we first define a convention for one party, say $a \in N$ (for ‘Alice’), to output an Indicator of Accusation, i.e., ‘accuse’ another party, say $i_M \in N$ (for ‘Mal’), of a fault. Specifically, we say that at event \hat{e}_A of the the execution, entity $ent[\hat{e}_A]$ *accuses* entity i_M (Mal), if $out[\hat{e}_A]$ is a triplet of the form (IA, i_M, x) . The last value in this triplet, x , should contain the clock value at the *first* time that Alice accused Mal; we discuss this in section 5 as the value x is not relevant for the specification predicate, and is just used as a convenient convention for some protocols.

The No False Accusations (NFA) predicate NFA checks whether the adversary was able to cause one honest entity, say Alice, to accuse another honest entity, say Bob (i.e., both Alice and Bob are in $N - F$). Namely, $NFA(T, 1^\kappa)$ returns \perp only if $T.out[e] = (IA, j, x)$, for some $j \in T.N$, and both j and $T.ent[e]$ are honest (i.e., $j, T.ent[e] \in T.N - T.F$). See Algorithm 17.

Algorithm 17 No False Accusations Predicate $NFA(T, p)$

```
1: return ¬(  
2:    $T.ent[T.e] \in T.N - T.F$  ▷  $T.ent[T.e]$  is an honest entity  
3:   and  $\exists j \in T.N - T.F, x$  s.t.  $(IA, j, x) \in T.out[T.e]$  ▷  $T.ent[T.e]$  accused an honest  
entity  
)
```

As noted above, in an accusation, the output $out[\hat{e}_A]$ contains a triplet of the form (IA, i_M, x) , where x is a clock value and *should* be the clock value at the *first* time that Alice accused Mal. We found this convenient in the definition of protocol-specific specifications where a party may accuse another party multiple times, and the specification is related to the time of the *first* accuse event. To allow the use of this convention, we define the following ‘technical’ specification predicate which merely confirms that honest entities always indicate, in any accuse event, the time of the first time they accused the same entity.

To simplify the predicate, let $fc(i, i_M, T)$ be the value of $T.clk[\hat{e}]$, where \hat{e} is the first event in T in which entity i accused entity $i_M \in T.N$ (or \perp if no such event exists). The Use First-Accuse Time (UFAT) predicate UFAT is defined in Algorithm 18.

Algorithm 18 Use First-Accuse Time Predicate $UFAT(T, p)$

```
1: return ¬(  
2:    $T.ent[T.e] \in T.N - T.F$  ▷  $T.ent[T.e]$  is an honest entity  
3:   and  $\exists i_M \in T.N$  s.t.  $(IA, i_M, x) \in T.out[T.e]$  ▷  $T.ent[T.e]$  did not indicate the  
   and  $x \neq fc(T.ent[T.e], i_M, T)$  ▷ first time of accusation in an  
   accusation  
)
```

6 Conclusions and Future Work

In this work, we discussed the lack of appropriate framework that allows rigorous analysis of applied, stateful cryptographic protocols, with different communication, synchronization and adversary models. We addressed this challenge by presenting the *Modular Specifications Security framework*, where security is defined with respect to a specific *model predicate* \mathcal{M} . The framework appears useful for various tasks; it allows comparison of protocols based on the specifications they satisfy and the models they assume. Definitions of models and specifications may be reused across different types of protocols and schemes; we identified several ‘generic’ specifications, which appear to be applicable to many different tasks.

We see this work as only the stepping stone to further research into the Modular Specifications Security framework. Specifically, into applying the framework onto various schemes, e.g., PKI schemes, in a way that will meet all desired specifications, minimizing or avoiding simplifications.

Another important challenge is to extend the Modular Specifications Security framework to support secure compositions, possibly following UC [1]. Specifically, it would be interesting to extend [2], which presents a UC definition for a simplified basic PKI (a reduced functionality of the X.509 PKI).

References

1. Canetti, R.: Universally Composable Security: A New Paradigm for Cryptographic Protocols. In: Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on. pp. 136–145. IEEE (2001), online at <https://eprint.iacr.org/2000/067.pdf>, last updated Dec. 2018.
2. Canetti, R., Shahaf, D., Vald, M.: Universally Composable Authentication and Key-exchange with Global PKI. Cryptology ePrint Archive, Report 2014/432 (2014), <https://eprint.iacr.org/2014/432>
3. Kaynar, D.K., Lynch, N., Segala, R., Vaandrager, F.: The theory of timed i/o automata. Synthesis Lectures on Distributed Computing Theory **1**(1), 1–137 (2010)
4. Leibowitz, H., Piotrowska, A.M., Danezis, G., Herzberg, A.: No Right to Remain Silent: Isolating Malicious Mixes. In: 28th USENIX Security Symposium (USENIX Security 19). pp. 1841–1858 (2019)
5. Pfitzmann, B., Waidner, M.: A model for asynchronous reactive systems and its application to secure message transmission. In: Security and Privacy, 2001. S&P 2001. Proceedings. 2001 IEEE Symposium on. pp. 184–200. IEEE (2001)