

# Modeling Computational Security in Long-Lived Systems

Ran Canetti<sup>1,2</sup>, Ling Cheung<sup>2</sup>, Dilsun Kaynar<sup>3</sup>, Nancy Lynch<sup>2</sup>, and Olivier Pereira<sup>4</sup>

<sup>1</sup> IBM T. J. Watson Research Center

<sup>2</sup> Massachusetts Institute of Technology

<sup>3</sup> Carnegie Mellon University

<sup>4</sup> Université catholique de Louvain,

**Abstract.** For many cryptographic protocols, security relies on the assumption that adversarial entities have limited computational power. This type of security degrades progressively over the lifetime of a protocol. However, some cryptographic services, such as timestamping services or digital archives, are *long-lived* in nature; they are expected to be secure and operational for a very long time (i.e., super-polynomial). In such cases, security cannot be guaranteed in the traditional sense: a computationally secure protocol may become insecure if the attacker has a super-polynomial number of interactions with the protocol.

This paper proposes a new paradigm for the analysis of long-lived security protocols. We allow entities to be active for a potentially unbounded amount of real time, provided they perform only a polynomial amount of work *per unit of real time*. Moreover, the space used by these entities is allocated dynamically and must be polynomially bounded. We propose a new notion of *long-term implementation*, which is an adaptation of computational indistinguishability to the long-lived setting. We show that long-term implementation is preserved under polynomial parallel composition and exponential sequential composition. We illustrate the use of this new paradigm by analyzing some security properties of the long-lived timestamping protocol of Haber and Kamat.

## 1 Introduction

*Computational security in Long-lived systems:* Security properties of cryptographic protocols typically hold only against resource-bounded adversaries. Consequently, mathematical models for representing and analyzing security of such protocols usually represent all participants as resource-bounded computational entities. The predominant way of formalizing such bounds is by representing all entities as time-bounded machines, specifically, polynomial-time machines (a very partial list of works representative of this direction includes [11, 5, 19, 7, 10, 3, 2, 8]).

This modeling approach has been successful in capturing the security of protocols for many cryptographic tasks. However, it has a fundamental limitation: it assumes that the analyzed system runs for only a relatively “short” time. In particular, since all entities are polynomially-bounded (in the security parameter), the system’s execution must end after a polynomial amount of time. This type of modeling is inadequate for analyzing security properties of protocols that are supposed to run for a “long” time, that is, an amount of time that is not bounded by a polynomial.

There are a number of natural tasks for which one would indeed be interested in the behavior of systems that run for a long time. Furthermore, a number of protocols have been developed for such tasks. However, none of the existing models for analyzing security against computationally bounded adversaries is adequate for asserting and proving security properties of protocols for such “long-lived” tasks.

One such task is *proactive security* [18]. Here, some secret information is distributed among several parties, in a way that allows the parties to jointly reconstruct the information, while preventing an adversary that breaks into any small subset of the parties from reconstructing the information. Furthermore, the parties periodically engage in a protocol for “refreshing” their shares in a way that guarantees secrecy of the information even if all parties are broken into multiple times, as long as not too many parties are broken into *between two refreshes*. The overall intention is to provide *long-lived* security of the system. Another such task is *forward secure signatures* [1, 6], where the system runs for a “long” time, and the signer periodically refreshes its secret key so that an adversary that corrupts the signer cannot forge signatures that bear time prior to the time of corruption. *Forward secure encryption* [1, 9] is defined analogously. Yet another task of the same flavor is timestamping [4, 12, 13]. Although the literature contains protocols for these long-lived tasks, we do not currently have the analytical tools to formulate and prove interesting assertions about their security.

*Related works* A first suggestion for an approach might be to use existing models, such as the PPT calculus [16], the Reactive Simulatability [3], or the Universally Composable security frameworks [7], with a sufficiently large value of

the security parameter. However, this would be too limited for our purpose in that it would force protocols to protect against an overly powerful adversary *even in the short run*, while not providing any useful information in the long run. Similarly, turning to information theoretic security notions is not appropriate in our case because unbounded adversaries would be able to break computationally secure schemes instantaneously. We are interested in a notion of security that can protect protocols against an adversary that runs for a long time, but is only “reasonably powerful” at any point in time. Recently, Müller-Quade and Unruh proposed a notion of *long-term security* for cryptographic protocols [17]. However, they consider adversaries that try to derive information from the protocol transcript *after* protocol conclusion. This work does not consider long-lived protocol execution and, in particular, the adversary of [17] has polynomially bounded interactions with the protocol parties, which is not suitable for the analysis of long-lived tasks such as those we described above.

*Our approach:* In this paper, we propose a new mathematical model for analyzing the security of such *long-lived systems*. Our understanding of a long-lived system is that some protocol parties, including adversaries, may be active for an unbounded amount of real time, subject to the condition that only a polynomial amount of work can be done per unit of real time. Other parties may be active for only a short time, as in traditional settings. Thus, the adversary’s interaction with the system is unbounded, and the adversary may perform an unbounded number of computation steps during the entire protocol execution. This renders traditional security notions insufficient: computationally and even statistically secure protocols may fail if the adversary has unbounded interactions with the protocol.

Modeling long-lived systems requires significant departures from standard cryptographic modeling. First and foremost, unbounded entities cannot be modeled as *probabilistic polynomial time (PPT)* Turing machines. In search of a suitable alternative, we see the need to distinguish between two types of unbounded computation: steps performed steadily over a long period of time, versus those performed very rapidly in a short amount of time. The former conforms with our understanding of boundedness, while the latter does not. Guided by this intuition, we introduce real time explicitly into a basic probabilistic automata model, the Task-PIOA model [8], and impose computational restrictions in terms of *rates*, i.e., number of computation steps per unit of real time.

Another interesting challenge is the restriction on space, which traditionally is not an issue because PPT Turing machines can, by their nature, access only a polynomially bounded amount of space. In the long-lived setting, space restriction warrants explicit consideration. For instance, we would like to model dynamic allocation of space, as new entities are invoked and old entities die off. We achieve this by restricting the use of state variables. In particular, all state variables of a dormant entity (either not yet invoked or already dead) are set to a special null value  $\perp$ . A system is regarded as bounded only if, at any point in its execution, only a bounded amount of space is needed to maintain all variables with non- $\perp$  values. For example, a sequential composition (in the temporal sense) of an unbounded number of entities is bounded if each entity uses a bounded amount of space.

Having appropriate restrictions on space and computation rates, we then define a new *long-term implementation relation*,  $\leq_{\text{neg,pt}}$ , for long-lived systems. This is intended to extend the familiar notion of *computational indistinguishability*, where two systems (*real* and *ideal*) are deemed equivalent if their behaviors are indistinguishable from the point of view of a computationally bounded environment. However, notice that, in the long-lived setting, an environment with super-polynomial run time can typically distinguish the two systems trivially, e.g., by launching brute force attacks. This is true even if the environment has bounded computation rate. Therefore, our definition cannot rule out significant degradation of security in the overall lifetime of a system. Instead, we require that the *rate* of degradation is small at any point in time; in other words, the probability of a *new* successful attack during any polynomial-bounded window of time remains bounded during the lifetime of the system.

To capture this intuition, we extend the ideal systems traditionally used in cryptography by allowing them to take some designated *failure* steps, which allow an ideal system to take actions that could only occur in the real world, e.g., accepting forgeries as valid signatures, or producing ciphertexts that could allow recovering the corresponding plaintext. However, if failure steps are prevented from some time  $t$ , then the ideal system starts following the specified ideal behavior.

Our long-term implementation relation  $\leq_{\text{neg,pt}}$  requires that the real system approximates the ideal’s system’s handling of failures. More precisely, we quantify over all real time points  $t$  and require that the real and ideal systems are computationally indistinguishable up to time  $t + q$  (where  $q$  is polynomial in the security parameter), even if no failure steps are taken by the ideal system in the interval  $[t, t + q]$ . Notice that we do allow failure steps before time  $t$ . This expresses the idea that, despite any security breaches that may have occurred before time  $t$ , the success probability of a *fresh* attack in the interval  $[t, t + q]$  is small. Our formal definition of  $\leq_{\text{neg,pt}}$  includes one more generalization:

it considers failure steps in the real system as well as the ideal system, in both cases before the same real time  $t$ . This natural extension is intended to allow repeated use of  $\leq_{\text{neg,pt}}$ , in verifying protocols using several levels of abstraction.

We show that  $\leq_{\text{neg,pt}}$  is transitive, and is preserved under the operations of polynomial parallel composition and exponential sequential composition. The sequential composition result highlights the power of our model to formulate and prove properties of an exponential number of entities in a meaningful way.

*Example: Digital Timestamping:* As a proof of concept, we analyze some security properties of the digital timestamping protocol of Haber et al. [4, 12, 13], which was designed to address the problem of content integrity in long-term digital archives. In a nutshell, a digital timestamping scheme takes as input a document  $d$  at a specific time  $t_0$ , and produces a certificate  $c$  that can be used later to verify the existence of  $d$  at time  $t_0$ . The security requirement is that timestamp certificates are difficult to forge. Haber et al. note that it is inadvisable to use a single digital signature scheme to generate all timestamp certificates, even if signing keys are refreshed periodically. This is because, over time, any single signature scheme may be weakened due to advances in algorithmic research and/or discovery of vulnerabilities. Haber et al. propose a solution in which timestamps must be renewed periodically by generating a new certificate for the pair  $\langle d, c \rangle$  using a new signature scheme. Thus, even if the signature scheme used to generate  $c$  is broken in the future, the new certificate  $c'$  still provides evidence that  $d$  existed at the time  $t_0$  stated in the original certificate  $c$ .

We model the protocol of Haber et al. as the composition of a dispatcher component and a sequence of signature services. Each signature service “wakes up” at a certain time and is active for a specified amount of time before becoming dormant again. This can be viewed as a regular update of the signature service, which may entail a simple refresh of the signing key, or the adoption of a new signing algorithm. The dispatcher component accepts various timestamp requests and forwards them to the appropriate signature service. We show that the composition of the dispatcher and the signature services is indistinguishable from an ideal system, consisting of the same dispatcher composed with ideal signature functionalities. Specifically, this guarantees that the probability of a new forgery is small at any given point in time, regardless of any forgeries that may have happened in the past.

## 2 Task-PIOAs

We build our new framework using task-PIOAs [8], which are a version of Probabilistic Automata [20], augmented with an oblivious scheduling mechanism based on tasks. A task is a set of related actions (e.g., actions representing the same activity but with different parameters). We view tasks as basic groupings of events, both for real time scheduling and for imposing computational bounds (cf. Sections 3 and 4). In this section, we review basic notations related to task-PIOAs.

*Notations:* Given a set  $S$ , let  $\text{Disc}(S)$  denote the set of discrete probability measures on  $S$ . For  $s \in S$ , let  $\delta(s)$  denote the *Dirac* measure on  $s$ , i.e.,  $\delta(s)(s) = 1$ .

Let  $V$  be a set of variables. Each  $v \in V$  is associated with a (*static*) *type*  $\text{type}(v)$ , which is the set of all possible values of  $v$ . We assume that  $\text{type}(v)$  is countable and contains the special symbol  $\perp$ . A *valuation*  $s$  for  $V$  is a function mapping every  $v \in V$  to a value in  $\text{type}(v)$ . The set of all valuations for  $V$  is denoted  $\text{val}(V)$ . Given  $V' \subseteq V$ , a valuation  $s'$  for  $V'$  is sometimes referred to as a *partial valuation* for  $V$ . Observe that  $s'$  induces a (full) valuation  $\iota_V(s')$  for  $V$ , by assigning  $\perp$  to every  $v \notin V'$ . Finally, for any set  $S$  with  $\perp \notin S$ , we write  $S_\perp := S \cup \{\perp\}$ .

*PIOA:* We define a *probabilistic input/output automaton (PIOA)* to be a tuple  $\mathcal{A} = \langle V, S, s^{\text{init}}, I, O, H, \Delta \rangle$ , where:

- (i)  $V$  is a set of *state variables* and  $S \subseteq \text{val}(V)$  is a set of *states*;
- (ii)  $s^{\text{init}} \in S$  is the *initial state*;
- (iii)  $I, O$  and  $H$  are countable and pairwise disjoint sets of actions, referred to as *input, output and hidden actions*, respectively;
- (iv)  $\Delta \subseteq S \times (I \cup O \cup H) \times \text{Disc}(S)$  is a *transition relation*.

The set  $\text{Act} := I \cup O \cup H$  is the *action alphabet* of  $\mathcal{A}$ . If  $I = \emptyset$ , then  $\mathcal{A}$  is said to be *closed*. The set of *external* actions of  $\mathcal{A}$  is  $I \cup O$  and the set of *locally controlled* actions is  $O \cup H$ . An *execution* is a sequence  $\alpha = q_0 a_1 q_1 a_2 \dots$  of alternating states and actions where  $q_0 = s^{\text{init}}$  and, for each  $\langle q_i, a_{i+1}, q_{i+1} \rangle$ , there is a transition  $\langle q_i, a_{i+1}, \mu \rangle \in \Delta$  with  $q_{i+1} \in \text{Supp}(\mu)$ . A sequence obtained by restricting an execution of  $\mathcal{A}$  to external actions is called a *trace*. We write  $s.v$  for the value of variable  $v$  in state  $s$ . An action  $a$  is *enabled* in a state  $s$  if  $\langle s, a, \mu \rangle \in \Delta$  for some  $\mu$ . We require that  $\mathcal{A}$  satisfies the following conditions.

- **Input Enabling:** For every  $s \in S$  and  $a \in I$ ,  $a$  is enabled in  $s$ .
- **Transition Determinism:** For every  $s \in S$  and  $a \in Act$ , there is at most one  $\mu \in \text{Disc}(S)$  with  $\langle s, a, \mu \rangle \in \Delta$ . We write  $\Delta(s, a)$  for such  $\mu$ , if it exists.

Parallel composition for PIOAs is based on synchronization of shared actions. PIOAs  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are said to be *compatible* if  $V_i \cap V_j = Act_i \cap Act_j = O_i \cap O_j = \emptyset$  whenever  $i \neq j$ . In that case, we define their *composition*  $\mathcal{A}_1 \parallel \mathcal{A}_2$  to be  $\langle V_1 \cup V_2, S_1 \times S_2, \langle s_1^{\text{init}}, s_2^{\text{init}} \rangle, (I_1 \cup I_2) \setminus (O_1 \cup O_2), O_1 \cup O_2, H_1 \cup H_2, \Delta \rangle$ , where  $\Delta$  is the set of triples  $\langle \langle s_1, s_2 \rangle, a, \mu_1 \times \mu_2 \rangle$  satisfying: (i)  $a$  is enabled in some  $s_i$ , and (ii) for every  $i$ , if  $a \in Act_i$ , then  $\langle s_i, a, \mu_i \rangle \in \Delta_i$ , otherwise  $\mu_i = \delta(s_i)$ . It is easy to check that input enabling and transition determinism are preserved under composition. Moreover, the definition of composition can be generalized to any finite number of components.

*Task-PIOAs:* To resolve nondeterminism, we make use of the notion of tasks introduced in [14, 8]. Formally, a *task-PIOA* is a pair  $\langle \mathcal{A}, \mathcal{R} \rangle$  where  $\mathcal{A}$  is a PIOA and  $\mathcal{R}$  is a partition of the locally-controlled actions of  $\mathcal{A}$ . The equivalence classes in  $\mathcal{R}$  are called *tasks*. For notational simplicity, we often omit  $\mathcal{R}$  and refer to the task-PIOA  $\mathcal{A}$ . The following additional axiom is assumed.

- **Action Determinism:** For every state  $s$  and every task  $T$ , at most one action  $a \in T$  is enabled in  $s$ .

Unless otherwise stated, terminologies are inherited from the PIOA setting. For instance, if some  $a \in T$  is enabled in a state  $s$ , then  $T$  is said to be *enabled* in  $s$ .

*Example 1 (Clock automaton).* Figure 1 describes a simple task-PIOA  $\text{Clock}(\mathbb{T})$ , which has a tick( $t$ ) output action for every  $t$  in some discrete time domain  $\mathbb{T}$ . For concreteness, we assume that  $\mathbb{T} = \mathbb{N}$ , and write simply  $\text{Clock}$ .  $\text{Clock}$  has a single task tick, consisting of all tick( $t$ ) actions. These clock ticks are produced in order, for  $t = 1, 2, \dots$ . In Section 3, we will define a mechanism that will ensure that each tick( $t$ ) occurs exactly at real time  $t$ .

Clock( $\mathbb{T}$ )	
<b>Signature</b>	<b>Tasks</b>
Input:	tick = {tick(*)}
none	
Output:	<b>States</b>
tick( $t : \mathbb{T}$ ), $t > 0$	count $\in \mathbb{T}$ , initially 0
<b>Transitions</b>	
tick( $t$ )	
Precondition:	
count = $t - 1$	
Effect:	
count := $t$	

**Fig. 1.** Task-PIOA Code for  $\text{Clock}(\mathbb{T})$

*Operations:* Given compatible task-PIOAs  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , we define their *composition* to be  $\langle \mathcal{A}_1 \parallel \mathcal{A}_2, \mathcal{R}_1 \cup \mathcal{R}_2 \rangle$ . Note that  $\mathcal{R}_1 \cup \mathcal{R}_2$  is an equivalence relation because compatibility requires disjoint sets of locally controlled actions. Moreover, it is easy to check that action determinism is preserved under composition.

We also define a *hiding* operator: given  $\mathcal{A} = \langle V, S, s^{\text{init}}, I, O, H, \Delta \rangle$  and  $S \subseteq O$ ,  $\text{hide}(\mathcal{A}, S)$  is the task-PIOA given by  $\langle V, S, s^{\text{init}}, I, O', H', \Delta \rangle$ , where  $O' = O \setminus S$  and  $H' = H \cup S$ . This prevents other PIOAs from synchronizing with  $\mathcal{A}$  via actions in  $S$ : any PIOA with an action in  $S$  in its signature is no longer compatible with  $\mathcal{A}$ .

*Executions and traces:* A *task schedule* for a closed task-PIOA  $\langle \mathcal{A}, \mathcal{R} \rangle$  is a finite or infinite sequence  $\rho = T_1, T_2, \dots$  of tasks in  $\mathcal{R}$ . This induces a well-defined run of  $\mathcal{A}$  as follows.

- From the start state  $s^{\text{init}}$ , we *apply* the first task  $T_1$ : due to action- and transition-determinism,  $T_1$  specifies at most one transition from  $s^{\text{init}}$ ; if such a transition exists, it is taken, otherwise nothing happens.
- Repeat with remaining  $T_i$ 's.

Such a run gives rise to a unique *probabilistic execution*, which is a probability distribution over executions in  $\mathcal{A}$ . For finite  $\rho$ , let  $\text{lstate}(\mathcal{A}, \rho)$  denote the state distribution of  $\mathcal{A}$  after executing according to  $\rho$ . A state  $s$  is said to be *reachable* under  $\rho$  if  $\text{lstate}(\mathcal{A}, \rho)(s) > 0$ . Moreover, the probabilistic execution induces a unique *trace distribution*  $\text{tdist}(\mathcal{A}, \rho)$ , which is a probability distribution over the set of traces of  $\mathcal{A}$ . We refer to [8] for more details on these constructions.

### 3 Real Time Scheduling Constraints

In this section, we describe how to model entities with unbounded lifetime but bounded processing rates. A natural approach is to introduce real time, so that computational restrictions can be stated in terms of the number of steps performed per unit real time. Thus, we define a *timed* task schedule  $\tau$  for a closed task-PIOA  $\langle \mathcal{A}, \mathcal{R} \rangle$  to be a finite or infinite sequence  $\langle T_1, t_1 \rangle, \langle T_2, t_2 \rangle, \dots$  such that:  $T_i \in \mathcal{R}$  and  $t_i \in \mathbb{R}_{\geq 0}$  for every  $i$ , and  $t_1, t_2, \dots$  is non-decreasing.

The *limit time*, denoted  $\text{ltime}(\tau)$ , is defined as follows.

- If  $\tau$  is empty, then  $\text{ltime}(\tau) := 0$ .
- If  $t_1, t_2, \dots$  is bounded, then  $\text{ltime}(\tau) := \lim_{i \rightarrow \infty} t_i$ , otherwise  $\text{ltime}(\tau) := \infty$ .

Following [15], we associate lower and upper real time bounds to each task. If  $l$  and  $u$  are, respectively, the lower bound and upper bound for a task  $T$ , then the amount of time between consecutive occurrences of  $T$  is at least  $l$  and at most  $u$ . To limit computational power, we impose a rate bound on the number of occurrences of  $T$  within an interval  $I$ , based on the length of  $I$ . A burst bound is also included for modeling flexibility.

Formally, a *bound map* for a task-PIOA  $\langle \mathcal{A}, \mathcal{R} \rangle$  is a tuple  $\langle \text{rate}, \text{burst}, \text{lb}, \text{ub} \rangle$  such that: (i)  $\text{rate}, \text{burst}, \text{lb} : \mathcal{R} \rightarrow \mathbb{R}_{\geq 0}$ , (ii)  $\text{ub} : \mathcal{R} \rightarrow \mathbb{R}_{> 0}^{\infty}$ , and (iii) for all  $T \in \mathcal{R}$ ,  $\text{lb}(T) \leq \text{ub}(T)$ . To ensure that  $\text{rate}$  and  $\text{ub}$  can be satisfied simultaneously, we require  $\text{rate}(T) \geq 1/\text{ub}(T)$  whenever  $\text{rate}(T) \neq 0$  and  $\text{ub}(T) \neq \infty$ . From this point on, we assume that every task-PIOA is associated with a particular bound map.

Given a timed schedule  $\tau$  and a task  $T$ , let  $\text{proj}_T(\tau)$  denote the result of removing all pairs  $\langle T_i, t_i \rangle$  with  $T_i \neq T$ . Let  $I$  be any left-closed interval with left endpoint 0. We say that  $\tau$  is *valid* for the interval  $I$  (under a bound map  $\langle \text{rate}, \text{burst}, \text{lb}, \text{ub} \rangle$ ) if the following hold for every task  $T$ .

- (i) If the pair  $\langle T, t \rangle$  appears in  $\tau$ , then  $t \in I$ .
- (ii) If  $\text{lb}(T) > 0$ , then: (a) if  $\langle T, t \rangle$  is the first element of  $\text{proj}_T(\tau)$ , then  $t \geq \text{lb}(T)$ ; (b) for every interval  $I'$  of a non-negative real length less than  $\text{lb}(T)$ ,  $\text{proj}_T(\tau)$  contains at most one element  $\langle T, t \rangle$  with  $t \in I'$ .
- (iii) If  $\text{ub}(T) \neq \infty$ , then, for every interval  $I' \subseteq I$  of a non-negative real length greater than  $\text{ub}(T)$ ,  $\text{proj}_T(\tau)$  contains at least one element  $\langle T, t \rangle$  with  $t \in I'$ .
- (iv) For any  $d \in \mathbb{R}_{\geq 0}$  and any interval  $I'$  of length  $d$ ,  $\text{proj}_T(\tau)$  contains at most  $\text{rate}(T) \cdot d + \text{burst}(T)$  elements  $\langle T, t \rangle$  with  $t \in I'$ .

We sometimes say that a task schedule  $\tau$  is valid, without specifying an interval, to mean that it is valid for the interval  $[0, \text{ltime}(\tau)]$ .

Note that every timed schedule  $\tau$  projects to an untimed schedule  $\rho$  by removing all real time information  $t_i$ , thereby inducing a trace distribution  $\text{tdist}(\mathcal{A}, \tau) := \text{tdist}(\mathcal{A}, \rho)$ . The set of trace distributions induced by all valid timed schedules for  $\mathcal{A}$  and  $\langle \text{rate}, \text{burst}, \text{lb}, \text{ub} \rangle$  is denoted  $\text{TrDists}(\mathcal{A}, \langle \text{rate}, \text{burst}, \text{lb}, \text{ub} \rangle)$ . Since the bound map is typically fixed, we often omit it and write  $\text{TrDists}(\mathcal{A})$ .

In a parallel composition  $\mathcal{A}_1 \parallel \mathcal{A}_2$ , the composite bound map is the union of component bound maps:

$$\langle \text{rate}_1 \cup \text{rate}_2, \text{burst}_1 \cup \text{burst}_2, \text{lb}_1 \cup \text{lb}_2, \text{ub}_1 \cup \text{ub}_2 \rangle.$$

This is well defined since the task partition of  $\mathcal{A}_1 \parallel \mathcal{A}_2$  is  $\mathcal{R}_1 \cup \mathcal{R}_2$ .

*Example 2 (Bound map for Clock).* We use upper and lower bounds to ensure that Clock's internal counter evolves at the same rate as real time. Namely, we set  $\text{lb}(\text{tick}) = \text{ub}(\text{tick}) = 1$ . The rate and burst bounds are also set to 1. It is not hard to see that, regardless of the system of automata with which Clock is composed, we always obtain the unique sequence  $\langle \text{tick}, 1 \rangle, \langle \text{tick}, 2 \rangle, \dots$  when we project a valid schedule to the task tick.

Note that we use real time solely to express constraints on task schedules. We do not allow computationally-bounded system components to maintain real-time information in their states, nor to communicate real-time information to each other. System components that require knowledge of time will maintain discrete approximations to time in their states, based on inputs from Clock.

## 4 Complexity Bounds

We are interested in modeling systems that run for an unbounded amount of real time. During this long life, we expect that a very large number of components will be active at various points in time, while only a small proportion of them will be active simultaneously. During the life time of a long-lived system, especially for systems such as those that use short-lived cryptographic primitives, it is natural to expect that many components will become obsolete or die, and will be replaced with other components. Defining complexity bounds in terms of the total number of components would then introduce unrealistic security constraints. Therefore, we find it more reasonable to define complexity bounds in terms of the characteristics of the components that are simultaneously active at any point in time.

To capture these intuitions, we define a notion of *step bound*, which limits the amount of computation a task-PIOA can perform, and the amount of space it can use, in executing a single step. By combining the step bound with the rate and burst bounds of Section 3, we obtain an *overall bound*, encompassing both bounded memory and bounded computation rates.

Note that we do not model situations where the rates of computation, or the computational power of machines, increases over time. This is an interesting direction in which the current research could be extended.

*Step Bound:* We assume some standard bit string encoding for Turing machines and for the names of variables, actions, and tasks. We also assume that variable valuations are encoded in the obvious way, as a list of name/value pairs. Let  $\mathcal{A}$  be a task-PIOA with variable set  $V$ . Given state  $s$ , let  $\hat{s}$  denote the partial valuation obtained from  $s$  by removing all pairs of the form  $\langle v, \perp \rangle$ . We have  $\iota_V(\hat{s}) = s$ , therefore no information is lost by reducing  $s$  to  $\hat{s}$ . This key observation allows us to represent a “large” valuation  $s$  with a “condensed” partial valuation  $\hat{s}$ .

Let  $p \in \mathbb{N}$  be given. We say that a state  $s$  is  $p$ -bounded if the encoding of  $\hat{s}$  is at most  $p$  bits long. The task-PIOA  $\mathcal{A}$  is said to have *step bound*  $p$  if the following hold.

- (i) For every variable  $v \in V$ ,  $\text{type}(v) \subseteq \{0, 1\}^p$ .
- (ii) The name of every action, task, and variable of  $\mathcal{A}$  has length at most  $p$ .
- (iii) The initial state  $s^{\text{init}}$  is  $p$ -bounded.
- (iv) There exists a deterministic Turing machine  $M_{\text{enable}}$  satisfying: for every  $p$ -bounded state  $s$ ,  $M_{\text{enable}}$  on input  $\hat{s}$  outputs the list of tasks enabled in  $s$ .
- (v) There exists a probabilistic Turing machine  $M_{\mathcal{R}}$  satisfying: for every  $p$ -bounded state  $s$  and task  $T$ ,  $M_{\mathcal{R}}$  on input  $\langle \hat{s}, T \rangle$  decides whether  $T$  is enabled in  $s$ . If so,  $M_{\mathcal{R}}$  computes and outputs a new partial valuation  $\hat{s}'$ , along with the unique  $a \in T$  that is enabled in  $s$ . The distribution on  $\iota_V(\hat{s}')$  coincides with  $\Delta(s, a)$ .
- (vi) There exists a probabilistic Turing machine  $M_I$  satisfying: for every  $p$ -bounded state  $s$  and action  $a$ ,  $M_I$  on input  $\langle \hat{s}, a \rangle$  decides whether  $a$  is an input action of  $\mathcal{A}$ . If so,  $M_I$  computes a new partial valuation  $\hat{s}'$ . The distribution on  $\iota_V(\hat{s}')$  coincides with  $\Delta(s, a)$ .
- (vii) The encoding of  $M_{\text{enable}}$  is at most  $p$  bits long, and  $M_{\text{enable}}$  terminates after at most  $p$  steps on every input. The same hold for  $M_{\mathcal{R}}$  and  $M_I$ .

Thus, step bound  $p$  limits the size of action names, which often represent protocol messages. It also limits the number of tasks enabled from any  $p$ -bounded state (Condition (iv)) and the complexity of individual transitions (Conditions (v) and (vi)). Finally, Condition (vii) requires all of the Turing machines to have description bounded by  $p$ .

Lemma 1 guarantees that a task-PIOA with step bound  $p$  will never reach a state in which more than  $p$  variables have non- $\perp$  values. The proof is a simple inductive argument.

**Lemma 1.** *Let  $\mathcal{A}$  be a task-PIOA with step bound  $p$ . For every valid timed task schedule  $\tau$  and every state  $s$  reachable under  $\tau$ , there are at most  $p$  variables  $v$  such that  $s.v \neq \perp$ .*

*Proof.* By the definition of step bounds, we have  $s^{\text{init}}$  is  $p$ -bounded. For a state  $s'$  reachable under schedule  $\tau'$ , let  $s$  be a state immediately preceding  $s'$  in the probabilistic execution induced by  $\tau'$ . Thus  $s$  is reachable under some prefix of  $\tau$ . If the transition from  $s$  to  $s'$  is locally controlled, we use the fact that  $M_{\mathcal{R}}$  always terminates after at most  $p$  steps, therefore every possible output, including  $\hat{s}'$ , has length at most  $p$ . This implies  $\hat{s}'$  is a partial valuation on at most  $p$  variables. If the transition from  $s$  to  $s'$  is an input, we follow the same argument with  $M_I$ .  $\square$

Given a closed (i.e., no input actions) task-PIOA  $\mathcal{A}$  with step bound  $p$ , one can easily define a Turing machine  $M_{\mathcal{A}}$  with a combination of nondeterministic and probabilistic branching that simulates the execution of  $\mathcal{A}$ . Lemma 1 can be used to show that the amount of work tape needed by  $M_{\mathcal{A}}$  is polynomial in  $p$ . This is reminiscent of the PSPACE complexity class, except that our setting introduces bounds on the computation rate, and allows probabilistic choices.

Lemma 2 says that, when we compose task-PIOAs in parallel, the complexity of the composite is proportional to the sum of the component complexities. The proof is similar to that of the full version of [8, Lemma 4.2]. We also note that the hiding operator introduced in Section 2 preserves step bounds.

**Lemma 2.** *Suppose  $\{\mathcal{A}_i \mid 1 \leq i \leq b\}$  is a compatible set of task-PIOAs, where each  $\mathcal{A}_i$  has step bound  $p_i \in \mathbb{N}$ . The composition  $\parallel_{i=1}^b \mathcal{A}_i$  has step bound  $c_{\text{comp}} \cdot \sum_{i=1}^b p_i$ , where  $c_{\text{comp}}$  is a fixed constant.*

*Overall Bound:* We now combine real time bounds and step bounds. To do so, we represent global time using the clock automaton Clock (Figure 1). Let  $p \in \mathbb{N}$  be given and let  $\mathcal{A}$  be a task-PIOA compatible with Clock. We say that  $\mathcal{A}$  is  $p$ -bounded if the following hold:

- (i)  $\mathcal{A}$  has step bound  $p$ .
- (ii) For every task  $T$  of  $\mathcal{A}$ ,  $\text{rate}(T)$  and  $\text{burst}(T)$  are both at most  $p$ .
- (iii) Let  $\pi$  be any execution that has nonzero probability under some valid schedule, and  $t \in \mathbb{N}$ . Then the total number of tasks enabled in states in  $\pi$  in which  $\text{count} = t$  is at most  $p$ .

Conditions (i) and (ii) are self-explanatory. Condition (iii) is a technical condition that ensures that the enabling of tasks does not change too rapidly. Without such a restriction,  $\mathcal{A}$  could cycle through a large number of tasks between two clock ticks, without violating the rate bound of any individual task.

*Task-PIOA Families:* We now extend our definitions to task-PIOA families, indexed by a *security parameter*  $k$ . More precisely, a *task-PIOA family*  $\bar{\mathcal{A}}$  is an indexed set  $\{\mathcal{A}_k\}_{k \in \mathbb{N}}$  of task-PIOAs. Given  $p : \mathbb{N} \rightarrow \mathbb{N}$ , we say that  $\bar{\mathcal{A}}$  is  $p$ -bounded just in case: for all  $k$ ,  $\mathcal{A}_k$  is  $p(k)$ -bounded. If  $p$  is a polynomial, then we say that  $\bar{\mathcal{A}}$  is *polynomially bounded*. The notions of compatibility and parallel composition for task-PIOA families are defined pointwise. We now present an example of a polynomially bounded family of task-PIOAs—a signature service that we use in our digital timestamping example in Section 8.

*Example: Signature Service:* A signature scheme Sig consists of three algorithms: KeyGen, Sign and Verify. KeyGen is a probabilistic algorithm that outputs a signing-verification key pair  $\langle sk, vk \rangle$ . Sign is a probabilistic algorithm that produces a signature  $\sigma$  from a message  $m$  and the key  $sk$ . Finally, Verify is a deterministic algorithm that maps  $\langle m, \sigma, vk \rangle$  to a boolean. The signature  $\sigma$  is said to be *valid* for  $m$  and  $vk$  if  $\text{Verify}(m, \sigma, vk) = 1$ .

Let  $SID$  be a domain of service identifiers. For each  $j \in SID$ , we build a signature service as a family of task-PIOAs indexed by security parameter  $k$ . Specifically, we define three task-PIOAs,  $\text{KeyGen}(k, j)$ ,  $\text{Signer}(k, j)$ , and  $\text{Verifier}(k, j)$  for every pair  $\langle k, j \rangle$ , representing the key generator, signer, and verifier, respectively. We assume a function  $\text{alive} : \mathbb{T} \rightarrow 2^{SID}$  such that, for every  $t$ ,  $\text{alive}(t)$  is the set of services alive at time  $t$ . The lifetime of each service  $j$  is then given by  $\text{aliveTimes}(j) := \{t \in \mathbb{T} \mid j \in \text{alive}(t)\}$ ; we assume this to be a finite set of consecutive numbers.

For every security parameter  $k$ , we assume the following finite domains:  $RID_k$  (request identifiers),  $M_k$  (messages to be signed) and  $\Sigma_k$  (signatures). The representations of elements in these domains are bounded by  $p(k)$ , for some polynomial  $p$ . Similarly, the domain  $\mathbb{T}_k$  consists of natural numbers representable using  $p(k)$  bits. Each of the components  $\text{KeyGen}(k, j)$ ,  $\text{Signer}(k, j)$ , and  $\text{Verifier}(k, j)$  has a set of input actions  $\text{tick}(t)$ ,  $t \in \mathbb{T}_k$ , which are intended to match with corresponding outputs from the clock automaton Clock (Figure 1). These inputs allow each component to record discrete time information in its state variable *clock*.

$\text{KeyGen}(k, j)$  chooses a signing key  $mySK$  and a corresponding verification key  $myVK$ . It does this exactly once, at any time when service  $j$  is alive. It outputs the two keys separately, via actions  $\text{signKey}(sk)_j$  and  $\text{verKey}(vk)_j$ . The signing key goes to  $\text{Signer}(k, j)$ , while the verification key goes to  $\text{Verifier}(k, j)$ .

The code for  $\text{KeyGen}(k, j)$  is given in Figure 2. As we mentioned before, the  $\text{tick}(t)$  action brings in the current time. If  $j$  is alive at time  $t$ , then *clock* is set to the current time  $t$ . Also, if  $j$  has just become alive, as evidenced by the fact that the *awake* flag is currently  $\perp$ , the *awake* flag is set to *true*. On the other hand, if  $j$  is no longer alive at time  $t$ , all variables are set to  $\perp$ .

The *chooseKeys* action uses  $\text{KeyGen}_j$  to choose the key pair, and is enabled only when  $j$  is awake and the keys are currently  $\perp$ . Note that the  $\text{KeyGen}$  algorithm is indexed by  $j$ , because different services may use different algorithms. The same applies to  $\text{Sign}_j$  in  $\text{Signer}(k, j)$  and  $\text{Verify}_j$  in  $\text{Verifier}(k, j)$ . The *signKey* and *verKey* actions output the keys, and they are enabled only when  $j$  is awake and the keys have been chosen.

$\text{Signer}(k, j)$  receives the signing key from another component, e.g.,  $\text{KeyGen}(k, j)$ . It then responds to signing requests by running the  $\text{Sign}_j$  algorithm on the given message  $m$  and the received signing key  $sk$ . Figure 3 presents the code for  $\text{Signer}(k, j)$ , which is fairly self-explanatory.

KeyGen( $k : \mathbb{N}, j : SID$ )

### Signature

Input:

tick( $t : \mathbb{T}_k$ )

Output:

signKey( $sk : 2^k$ )<sub>j</sub>

verKey( $vk : 2^k$ )<sub>j</sub>

Internal:

chooseKeys<sub>j</sub>

### Transitions

tick( $t$ )

Effect:

if  $j \in \text{alive}(t)$  then

$clock := t$

  if  $awake = \perp$  then

$awake := true$

  else

$awake, clock, mySK,$

$myVK := \perp$

chooseKeys<sub>j</sub>

Precondition:

$awake = true$

$mySK = myVK = \perp$

Effect:

$\langle mySK, myVK \rangle$

$\leftarrow \text{KeyGen}_j(1^k)$

### Tasks

verKey<sub>j</sub> = {verKey(\*)<sub>j</sub>}

signKey<sub>j</sub> = {signKey(\*)<sub>j</sub>}

chooseKeys<sub>j</sub> = {chooseKeys<sub>j</sub>}

### States

$awake : \{true\}_\perp, \text{init } \perp$

$clock : (\mathbb{T}_k)_\perp, \text{init } \perp$

$mySK : (2^k)_\perp, \text{init } \perp$

$myVK : (2^k)_\perp, \text{init } \perp$

signKey( $sk$ )<sub>j</sub>

Precondition:

$awake = true$

$sk = mySK \neq \perp$

Effect:

  none

verKey( $vk$ )<sub>j</sub>

Precondition:

$awake = true$

$vk = myVK \neq \perp$

Effect:

  none

**Fig. 2.** Task-PIOA Code for KeyGen( $k, j$ )

The data type  $\text{que}_k$  represents queues with maximum length  $p(k)$ , where  $p$  is a polynomial. The enqueue operation automatically discards the new entry if the queue is already of length  $p(k)$ . This models the fact that  $\text{Signer}(k, j)$  has a bounded amount of memory. For concreteness, we assume here that  $p$  is the constant function  $\underline{1}$  for the queues *toSign* and *signed*.

$\text{Verifier}(k, j)$  accepts verification requests and simply runs the  $\text{Verify}_j$  algorithm. The code appears in Figure 4. Again, all queues have maximum length 1.

Assuming the algorithms  $\text{KeyGen}_j$ ,  $\text{Sign}_j$  and  $\text{Verify}_j$  are polynomial time, it not hard to check that the composite  $\text{KeyGen}(k, j) \parallel \text{Signer}(k, j) \parallel \text{Verifier}(k, j)$  has step bound  $p(k)$  for some polynomial  $p$ . If  $\text{rate}(T)$  and  $\text{burst}(T)$  are at most  $p(k)$  for every  $T$ , then the composite is  $p(k)$ -bounded. The family  $\{\text{KeyGen}(k, j) \parallel \text{Signer}(k, j) \parallel \text{Verifier}(k, j)\}_{k \in \mathbb{N}}$  is therefore polynomially bounded.

## 5 Long-Term Implementation Relation

Much of modern cryptography is based on the notion of computational indistinguishability. For instance, an encryption algorithm is (chosen-plaintext) secure if the ciphertexts of two distinct but equal-length messages are indistinguishable from each other, even if the plaintexts are generated by the distinguisher itself. The key assumption is that the distinguisher is computationally bounded, so that it cannot launch a brute force attack. In this section, we adapt this notion of indistinguishability to the long-lived setting.

We define an implementation relation based on closing environments and acceptance probabilities. Let  $\mathcal{A}$  be a closed task-PIOA with output action  $\text{acc}$  and task  $\{\text{acc}\}$ . Let  $\tau$  be a timed task schedule for  $\mathcal{A}$ . The *acceptance probability* of  $\mathcal{A}$  under  $\tau$  is:  $\mathbf{P}_{\text{acc}}(\mathcal{A}, \tau) := \Pr[\beta \text{ contains } \text{acc} : \beta \leftarrow_{\mathbb{R}} \text{tdist}(\mathcal{A}, \tau)]$ ; that is, the probability that a trace drawn from the distribution  $\text{tdist}(\mathcal{A}, \tau)$  contains the action  $\text{acc}$ . If  $\mathcal{A}$  is not necessarily closed, we include a closing



environment. A task-PIOA  $\text{Env}$  is an *environment* for  $\mathcal{A}$  if it is compatible with  $\mathcal{A}$  and  $\mathcal{A} \parallel \text{Env}$  is closed. From here on, we assume that every environment has output action  $\text{acc}$ .

In the short-lived setting, we say that a system  $\mathcal{A}_1$  implements another system  $\mathcal{A}_2$  if every run of  $\mathcal{A}_1$  can be “matched” by a run of  $\mathcal{A}_2$  such that no polynomial time environment can distinguish the two runs. As we discussed in the introduction, this type of definition is too strong for the long-lived setting, because we must allow environments with unbounded total run time (as long as they have bounded rate and space).

For example, consider the timestamping protocol of [12, 13] described in Section 1. After running for a long period of real time, a distinguisher environment may be able to forge a signature with non-negligible probability. As a result, it can distinguish the real system from an ideal timestamping system, in the traditional sense. However, the essence of the protocol is that such failures can in fact be tolerated, because they do not help the environment to forge *new* signatures, after a new, uncompromised signature service becomes active.

This timestamping example suggests that we need a new notion of long-term implementation that makes meaningful security guarantees in any polynomial-bounded window of time, in spite of past security failures. Our new implementation relation aims to capture this intuition.

First we define a comparability condition for task-PIOAs:  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are said to be *comparable* if they have the same external interface, that is,  $I_1 = I_2$  and  $O_1 = O_2$ . In this case, every environment  $E$  for  $\mathcal{A}_1$  is also an environment for  $\mathcal{A}_2$ , provided  $E$  is compatible with  $\mathcal{A}_2$ .

Let  $\mathcal{A}_1$  and  $\mathcal{A}_2$  be comparable task-PIOAs. To model security failure events in both automata, we let  $F_1$  be a set of designated *failure tasks* of  $\mathcal{A}_1$ , and let  $F_2$  be a set of *failure tasks* of  $\mathcal{A}_2$ . We assume that each task in  $F_1 \cup F_2$  has  $\infty$  as its upper bound.

Given  $t \in \mathbb{R}_{\geq 0}$  and an environment  $\text{Env}$  for both  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , we consider two experiments. In the first experiment,  $\text{Env}$  interacts with  $\mathcal{A}_1$  according to some valid task schedule  $\tau_1$  of  $\mathcal{A}_1 \parallel \text{Env}$ , where  $\tau_1$  does not contain any tasks from  $F_1$  from time  $t$  onwards. In the second experiment,  $\text{Env}$  interacts with  $\mathcal{A}_2$  according to some valid task schedule  $\tau_2$  of  $\mathcal{A}_2 \parallel \text{Env}$ , where  $\tau_2$  does not contain any tasks from  $F_2$  from time  $t$  onwards. Our definition requires that the first experiment “approximates” the second one, that is, if  $\mathcal{A}_1$  acts ideally (does not perform any of the failure tasks in  $F_1$ ) after time  $t$ , then it simulates  $\mathcal{A}_2$ , also acting ideally after time  $t$ .

More specifically, we require that, for any valid  $\tau_1$ , there exists a valid  $\tau_2$  as above such that the two executions are identical up to time  $t$  from the point of view of the environment. That is, the acceptance probabilities in these experiments are the same up to time  $t$  and  $\text{Env}$  has the same state distribution immediately before time  $t$ . Moreover, the two executions are overall *computationally indistinguishable*, namely, the difference in acceptance probabilities in these two experiments is negligible provided  $\text{Env}$  is computationally bounded.

Given a task schedule  $\tau = \langle T_1, t_1 \rangle, \langle T_2, t_2 \rangle, \dots$ , let  $\text{trunc}_{\geq t}(\tau)$  denote the result of removing all pairs  $\langle T_i, t_i \rangle$  with  $t_i \geq t$ . If  $\tau$  is a schedule of  $\mathcal{A} \parallel \mathcal{B}$ , then we define  $\text{proj}_{\mathcal{B}}(\tau)$  to be the result of removing all  $\langle T_i, t_i \rangle$  where  $T_i$  is *not* a task of  $\mathcal{B}$ . Moreover, let  $\text{lstate}_{\mathcal{B}}(\mathcal{A} \parallel \mathcal{B}, \tau)$  denote the final state distribution of  $\mathcal{B}$  after executing with  $\mathcal{A}$  under the schedule  $\tau$  (assuming  $\tau$  is finite).

**Definition 1.** Let  $\mathcal{A}_1$  and  $\mathcal{A}_2$  be comparable task-PIOAs that are both compatible with  $\text{Clock}$ . Let  $F_1$  and  $F_2$  be sets of tasks of, respectively,  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , such that for any  $T \in (F_1 \cup F_2)$ ,  $\text{ub}(T) = \infty$ . Let  $p, q \in \mathbb{N}$  and  $\epsilon \in \mathbb{R}_{\geq 0}$  be given. Then we say that  $(\mathcal{A}_1, F_1) \leq_{p,q,\epsilon} (\mathcal{A}_2, F_2)$  provided that the following is true:

For every  $t \in \mathbb{R}_{\geq 0}$ , every environment  $\text{Env}$  of the form  $\text{Env}' \parallel \text{Clock}$  with  $\text{Env}'$   $p$ -bounded, and every valid timed schedule  $\tau_1$  for  $\mathcal{A}_1 \parallel \text{Env}$  for the interval  $[0, t + q]$  that does not contain any pairs of the form  $\langle T_i, t_i \rangle$  where  $T_i \in F_1$  and  $t_i \geq t$ , there exists a valid timed schedule  $\tau_2$  for  $\mathcal{A}_2 \parallel \text{Env}$  for the interval  $[0, t + q]$  such that:

- (i)  $\mathbf{P}_{\text{acc}}(\mathcal{A}_1 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_1)) = \mathbf{P}_{\text{acc}}(\mathcal{A}_2 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_2))$ ;
- (ii)  $\text{lstate}_{\text{Env}}(\mathcal{A}_1 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_1)) = \text{lstate}_{\text{Env}}(\mathcal{A}_2 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_2))$ ;
- (iii)  $\text{proj}_{\text{Env}}(\tau_1) = \text{proj}_{\text{Env}}(\tau_2)$ ;
- (iv)  $\tau_2$  does not contain any pairs of the form  $\langle T_i, t_i \rangle$  where  $T_i \in F_2$  and  $t_i \geq t$ ; and
- (v)  $|\mathbf{P}_{\text{acc}}(\mathcal{A}_1 \parallel \text{Env}, \tau_1) - \mathbf{P}_{\text{acc}}(\mathcal{A}_2 \parallel \text{Env}, \tau_2)| \leq \epsilon$ .

The following lemma says that  $\leq_{p,q,\epsilon}$  (Definition 1) is transitive up to additive errors.

**Lemma 3.** Let  $\mathcal{A}_1, \mathcal{A}_2$ , and  $\mathcal{A}_3$  be comparable task-PIOAs, and let  $F_1$  be a set of task of  $\mathcal{A}_1$ ,  $F_2$  be a set of tasks for  $\mathcal{A}_2$ , and  $F_3$  be a set of tasks of  $\mathcal{A}_3$  such that for any  $T \in (F_1 \cup F_2 \cup F_3)$ ,  $\text{ub}(T) = \infty$ . Let  $p, q \in \mathbb{N}$  and  $\epsilon \in \mathbb{R}_{\geq 0}$  be given. Assume that  $(\mathcal{A}_1, F_1) \leq_{p,q,\epsilon_1} (\mathcal{A}_2, F_2)$  and  $(\mathcal{A}_2, F_2) \leq_{p,q,\epsilon_2} (\mathcal{A}_3, F_3)$ . Then  $(\mathcal{A}_1, F_1) \leq_{p,q,\epsilon_1+\epsilon_2} (\mathcal{A}_3, F_3)$ .

*Proof.* Let  $t \in \mathbb{R}_{\geq 0}$ , a  $p$ -bounded environment  $\text{Env}$  of the form  $\text{Env}' \parallel \text{Clock}$ , and a valid timed schedule  $\tau_1$  for  $\mathcal{A}_1 \parallel \text{Env}$  for the interval  $[0, t + q]$  be given, where  $\tau_1$  does not contain any pairs of the form  $\langle T_i, t_i \rangle$  where  $T_i \in F_1$  and  $t_i \geq t$ .

Choose  $\tau_2$  for  $\mathcal{A}_2 \parallel \text{Env}$  according to the assumption  $(\mathcal{A}_1, F_1) \leq_{p,q,\epsilon_1} (\mathcal{A}_2, F_2)$ . Using  $\tau_2$ , choose  $\tau_3$  for  $\mathcal{A}_3 \parallel \text{Env}$  according to the assumption  $(\mathcal{A}_2, F_2) \leq_{p,q,\epsilon_2} (\mathcal{A}_3, F_3)$ .

Clearly, we have

- $\mathbf{P}_{\text{acc}}(\mathcal{A}_1 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_1))$   
 $= \mathbf{P}_{\text{acc}}(\mathcal{A}_2 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_2))$   
 $= \mathbf{P}_{\text{acc}}(\mathcal{A}_3 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_3));$
- $\text{lstate}_{\text{Env}}(\mathcal{A}_1 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_1))$   
 $= \text{lstate}_{\text{Env}}(\mathcal{A}_2 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_2))$   
 $= \text{lstate}_{\text{Env}}(\mathcal{A}_3 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_3));$
- $\text{proj}_{\text{Env}}(\tau_1) = \text{proj}_{\text{Env}}(\tau_2) = \text{proj}_{\text{Env}}(\tau_3).$

It is also immediate that  $\tau_3$  does not contain any pairs of the form  $\langle T_i, t_i \rangle$  where  $T_i \in F_3$  and  $t_i \geq t$ . Finally,

$$\begin{aligned} & | \mathbf{P}_{\text{acc}}(\mathcal{A}_1 \parallel \text{Env}, \tau_1) - \mathbf{P}_{\text{acc}}(\mathcal{A}_3 \parallel \text{Env}, \tau_3) | \\ & \leq | \mathbf{P}_{\text{acc}}(\mathcal{A}_1 \parallel \text{Env}, \tau_1) - \mathbf{P}_{\text{acc}}(\mathcal{A}_2 \parallel \text{Env}, \tau_2) | \\ & \quad + | \mathbf{P}_{\text{acc}}(\mathcal{A}_2 \parallel \text{Env}, \tau_2) - \mathbf{P}_{\text{acc}}(\mathcal{A}_3 \parallel \text{Env}, \tau_3) | \\ & \leq \epsilon_1 + \epsilon_2. \end{aligned} \quad \square$$

The relation  $\leq_{p,q,\epsilon}$  can be extended to task-PIOA families as follows. Let  $\bar{\mathcal{A}}_1 = \{(\bar{\mathcal{A}}_1)_k\}_{k \in \mathbb{N}}$  and  $\bar{\mathcal{A}}_2 = \{(\bar{\mathcal{A}}_2)_k\}_{k \in \mathbb{N}}$  be pointwise comparable task-PIOA families. Let  $\bar{F}_1$  be a family of sets such that each  $(\bar{F}_1)_k$  is a set of tasks of  $(\bar{\mathcal{A}}_1)_k$  and let  $\bar{F}_2$  be a family of sets such that each  $(\bar{F}_2)_k$  is a set of tasks of  $(\bar{\mathcal{A}}_2)_k$ , satisfying the condition that each task of those sets has an infinite upperbound. Let  $\epsilon : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$  and  $p, q : \mathbb{N} \rightarrow \mathbb{N}$  be given. We say that  $(\bar{\mathcal{A}}_1, \bar{F}_1) \leq_{p,q,\epsilon} (\bar{\mathcal{A}}_2, \bar{F}_2)$  just in case  $((\bar{\mathcal{A}}_1)_k, (\bar{F}_1)_k) \leq_{p(k),q(k),\epsilon(k)} ((\bar{\mathcal{A}}_2)_k, (\bar{F}_2)_k)$  for every  $k$ .

Restricting our attention to negligible error and polynomial time bounds, we obtain the long-term implementation relation  $\leq_{\text{neg,pt}}$ . Formally, a function  $\epsilon : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$  is said to be *negligible* if, for every constant  $c \in \mathbb{N}$ , there exists  $k_0 \in \mathbb{N}$  such that  $\epsilon(k) < \frac{1}{k^c}$  for all  $k \geq k_0$ . (That is,  $\epsilon$  diminishes more quickly than the reciprocal of any polynomial.) Given task-PIOA families  $\bar{\mathcal{A}}_1$  and  $\bar{\mathcal{A}}_2$  and task set families  $\bar{F}_1$  and  $\bar{F}_2$ , respectively, of  $\bar{\mathcal{A}}_1$  and  $\bar{\mathcal{A}}_2$ , we say that  $(\bar{\mathcal{A}}_1, \bar{F}_1) \leq_{\text{neg,pt}} (\bar{\mathcal{A}}_2, \bar{F}_2)$  if  $\forall p, q \exists \epsilon : (\bar{\mathcal{A}}_1, \bar{F}_1) \leq_{p,q,\epsilon} (\bar{\mathcal{A}}_2, \bar{F}_2)$ , where  $p, q$  are polynomials and  $\epsilon$  is a negligible function.

**Lemma 4 (Transitivity of  $\leq_{\text{neg,pt}}$ ).** *Let  $\bar{\mathcal{A}}_1, \bar{\mathcal{A}}_2$ , and  $\bar{\mathcal{A}}_3$  be comparable task-PIOA families. Let  $\bar{F}_1$  be a task set family of  $\bar{\mathcal{A}}_1$ , Let  $\bar{F}_2$  be a task set family of  $\bar{\mathcal{A}}_2$ , and let  $\bar{F}_3$  be a task set family of  $\bar{\mathcal{A}}_3$  (satisfying the upperbound condition). Suppose  $(\bar{\mathcal{A}}_1, \bar{F}_1) \leq_{\text{neg,pt}} (\bar{\mathcal{A}}_2, \bar{F}_2)$  and  $(\bar{\mathcal{A}}_2, \bar{F}_2) \leq_{\text{neg,pt}} (\bar{\mathcal{A}}_3, \bar{F}_3)$ . Then  $(\bar{\mathcal{A}}_1, \bar{F}_1) \leq_{\text{neg,pt}} (\bar{\mathcal{A}}_3, \bar{F}_3)$ .*

*Proof.* Given polynomials  $p$  and  $q$ , choose negligible functions  $\epsilon_1$  and  $\epsilon_2$  according to the assumptions. Then  $\epsilon_1 + \epsilon_2$  is negligible. By Lemma 3, we have  $(\bar{\mathcal{A}}_1, \bar{F}_1) \leq_{p,q,\epsilon_1+\epsilon_2} (\bar{\mathcal{A}}_3, \bar{F}_3)$ .

## 6 Ideal Signature Functionality

In this section, we specify an *ideal signature functionality*  $\text{SigFunc}$ , and show that it is implemented, in the sense of our  $\leq_{\text{neg,pt}}$  definition, by the real signature service of Section 4.

As with  $\text{KeyGen}$ ,  $\text{Signer}$ , and  $\text{Verifier}$ , each instance of  $\text{SigFunc}$  is parameterized with a security parameter  $k$  and an identifier  $j$ . The code for  $\text{SigFunc}(k, j)$  appears in Figure 5. It is very similar to the composition of  $\text{Signer}(k, j)$  and  $\text{Verifier}(k, j)$ . The important difference is that  $\text{SigFunc}(k, j)$  maintains an additional variable *history*, which records the set of signed messages. In addition,  $\text{SigFunc}(k, j)$  has an internal action  $\text{fail}_j$ , which sets a boolean flag *failed*. If *failed* = false, then  $\text{SigFunc}(k, j)$  uses *history* to answer verification requests: a signature is rejected if the submitted message is not in *history*, even if  $\text{Verify}_j$  returns 1. If *failed* = true, then  $\text{SigFunc}(k, j)$  bypasses the check on *history*, so that its answers are identical to those from the real signature service.

Recall that, for every task  $T$  of the real signature service,  $\text{rate}(T)$  and  $\text{burst}(T)$  are bounded by  $p(k)$  for some polynomial  $p$ . We assume that the same bound applies to  $\text{SigFunc}(k, j)$ . Since  $\text{aliveTimes}(j)$  is a finite set of consecutive numbers, it represents essentially an interval whose length is constant in the security parameter  $k$ . Therefore,  $p(k)$  gives rise to a bound  $p'(k)$  on the maximum number of signatures generated by  $\text{SigFunc}(k, j)$ , where  $p'$  is also polynomial. We set the maximum length of the queue *history* to  $p'(k)$ . All other queues have maximum length 1.

We claim that the real signature service implements the ideal signature functionality. The proof relies on a reduction to standard properties of a signature scheme, namely, completeness and existential unforgeability, as defined below.

**Definition 2.** A signature scheme  $\text{Sig} = \langle \text{KeyGen}, \text{Sign}, \text{Verify} \rangle$  is complete if  $\text{Verify}(m, \sigma, vk) = 1$  whenever  $\langle sk, vk \rangle \leftarrow \text{KeyGen}(1^k)$  and  $\sigma \leftarrow \text{Sign}(sk, m)$ . We say that  $\text{Sig}$  is existentially unforgeable under adaptive chosen message attacks (or EUF-CMA secure) if no probabilistic polynomial-time forger has non-negligible success probability in the following game.

**Setup** The challenger runs  $\text{KeyGen}$  to obtain  $\langle sk, vk \rangle$  and gives the forger  $vk$ .

**Query** The forger submits message  $m$ . The challenger responds with signature  $\sigma \leftarrow \text{Sign}(m, sk)$ . This may be repeated adaptively.

**Output** The forger outputs a pair  $\langle m^*, \sigma^* \rangle$  and he wins if  $m^*$  is not among the messages submitted during the query phase and  $\text{Verify}(m^*, \sigma^*, vk) = 1$ .

For all  $k \in \mathbb{N}$  and  $j \in \text{SID}$ , we define  $\text{RealSig}(j)_k$  to be  $\text{hide}(\text{KeyGen}(k, j) \parallel \text{Signer}(k, j) \parallel \text{Verifier}(k, j), \text{signKey}_j)$  and  $\text{IdealSig}(j)_k$  to be  $\text{hide}(\text{KeyGen}(k, j) \parallel \text{SigFunc}(k, j), \text{signKey}_j)$ .

These automata are gathered into families in the obvious way:  $\overline{\text{RealSig}}(j) := \{\text{RealSig}(j)_k\}_{k \in \mathbb{N}}$  and  $\overline{\text{IdealSig}}(j) := \{\text{IdealSig}(j)_k\}_{k \in \mathbb{N}}$ . Note that the hiding operation prevents the environment from learning the signing key.

**Theorem 1.** Let  $j \in \text{SID}$  be given. Suppose that  $\langle \text{KeyGen}_j, \text{Sign}_j, \text{Verify}_j \rangle$  is a complete and EUF-CMA secure signature scheme. Then  $\overline{\text{RealSig}}(j, \{\}) \leq_{\text{neg,pt}} \overline{\text{IdealSig}}(j, \{\text{fail}_j\})$ .

To prove Theorem 1, we show that, for every time point  $t$ , the environment cannot distinguish  $\text{RealSig}(j)_k$  from  $\text{IdealSig}(j)_k$  with high probability between time  $t$  and  $t + q(k)$ , where  $q$  is a polynomial. This holds even when the task  $\{\text{fail}_j\}$  is not scheduled in the interval  $[t, t + q]$ . The interesting case is when  $j$  is awakened *after* time  $t$ . That implies the *failed* flag is never set and  $\text{SigFunc}(k, j)$  uses *history* to reject forgeries.

We use the the EUF-CMA assumption to obtain a bound on the distinguishing probability of any environment. Essentially, we build a forger that emulates the execution of our various task-PIOAs under some valid schedule. When the environment interacts with the Signer and Verifier automata, this forger uses the signature oracle and verification algorithm in the EUF-CMA game. Moreover, the success probability of this forger is maximized over all environments satisfying a particular polynomial bound. (Note that, given polynomial  $p$  and security parameter  $k$ , there are only a finite number of  $p(k)$ -bounded environments.) Applying the definition of EUF-CMA security, we obtain the desired negligible bound on distinguishing probability.

*Proof (Proof of Theorem 1).* Unwinding the definition of  $\leq_{\text{neg,pt}}$  using the given failure sets, we need to show the following: for every polynomials  $p$  and  $q$ , there is a negligible function  $\epsilon$  such that, for every  $k \in \mathbb{N}$ ,  $t \in \mathbb{R}_{\geq 0}$ ,  $p(k)$ -bounded environment  $\text{Env}$  for  $\text{RealSig}(j)_k$ , and valid schedule  $\tau_1$  for  $\text{RealSig}(j)_k \parallel \text{Env}$  for the interval  $[0, t + q(k)]$ , there is a valid schedule  $\tau_2$  for  $\text{IdealSig}(j)_k \parallel \text{Env}$  such that

- (i)  $\mathbf{P}_{\text{acc}}(\text{RealSig}(j)_k \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_1))$  is the same as  $\mathbf{P}_{\text{acc}}(\text{IdealSig}(j)_k \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_2))$ ;
- (ii)  $\text{lstate}_{\text{Env}}(\text{RealSig}(j)_k \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_1))$  is the same as  $\text{lstate}_{\text{Env}}(\text{IdealSig}(j)_k \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_2))$ ;
- (iii)  $\text{proj}_{\text{Env}}(\tau_1) = \text{proj}_{\text{Env}}(\tau_2)$ ;
- (iv)  $\tau_2$  does not contain any pairs of the form  $\langle \text{fail}_j, t_i \rangle$  where  $t_i \geq t$ ;
- (v)  $\mathbf{P}_{\text{acc}}(\text{RealSig}(j)_k \parallel \text{Env}, \tau_1)$  is at most  $\epsilon(k)$  away from  $\mathbf{P}_{\text{acc}}(\text{IdealSig}(j)_k \parallel \text{Env}, \tau_2)$ .

Let polynomial  $p$  and  $q$  be given. We need to obtain a negligible  $\epsilon$  bound that makes all the conditions above satisfied for every  $k, t, p(k)$ -bounded  $\text{Env}$ , valid  $\tau_1$ , and some corresponding  $\tau_2$ .

Fix  $t_l$  and  $t_r$  to be time points such that  $[t_l, t_r] = \{t \in \mathbb{T} \mid j \in \text{alive}(t)\}$ . So, we know that both  $\text{RealSig}(j)_k$  and  $\text{IdealSig}(j)_k$  are dormant outside the interval  $[t_l, t_r]$ .

First consider the cases in which  $t_l < t$ . We obtain  $\tau_2$  by inserting  $\langle \{\text{fail}_j\}, t_l \rangle$  immediately after  $\langle \text{tick}, t_l \rangle$ . This sets the *failed* flag in  $\text{SigFunc}(k, j)$  to true immediately after *awake* becomes true. Notice that, if *failed* = true, the verify transition bypasses the check  $m \in \text{history}$  (Figure 5). In other words,  $\text{SigFunc}(k, j)$  answers verify requests in exactly the same way as  $\text{Verifier}(k, j)$ , using the Verify algorithm only. Furthermore, it is easy to check that *failed* remains true as long as  $\text{SigFunc}(k, j)$  is alive. Therefore,  $\text{IdealSig}(j)_k$  has exactly the same visible behavior as  $\text{RealSig}(j)_k$  and Conditions (i) through (v) above are satisfied if we choose  $\epsilon(k) = 0$ , for every  $k, p(k)$ -bounded  $\text{Env}$  and valid  $\tau_1$ .

Now, consider the cases in which  $t \leq t_l$ . Set  $\tau_2 := \tau_1$ . Since both  $\text{RealSig}(j)_k$  and  $\text{IdealSig}(j)_k$  are dormant during  $[0, t]$ , Conditions (i) and (ii) must hold. Condition (iii) is immediate and Condition (iv) holds because  $\text{fail}_j$  is not a task of  $\text{RealSig}(j)_k$ . It remains to argue that there exists a negligible function  $\epsilon$  such that Condition (v) is satisfied.

To this purpose, we rely on the EUF-CMA security of  $\text{Sig}$ . We however do not need to bound the success probability of one specific forger, as in the EUF-CMA definition, but the success probability of all forgers that satisfy fixed polynomial  $p$  and  $q$  bounds, for every time  $t$  and schedule  $\tau_1$ .

For every  $k \in \mathbb{N}$ , we define a time  $(t_{max})_k \leq t_l$ , a  $p(k)$ -bounded environment  $(Env_{max})_k$  for  $RealSig_k$ , and a valid schedule  $(\tau_{1max})_k$  for  $RealSig_k \parallel (Env_{max})_k$  for the time interval  $[0, (t_{max})_k + q(k)]$ , with the following property: for every time  $t \leq t_l$ , every  $p(k)$ -bounded environment  $Env$  for  $RealSig_k$ , and every valid schedule  $\tau_1$  for  $RealSig_k \parallel Env$  for the interval  $[0, t + q(k)]$ , we have:

$$\begin{aligned} & | \mathbf{P}_{acc}(RealSig(j)_k \parallel Env, \tau_1) - \mathbf{P}_{acc}(IdealSig(j)_k \parallel Env, \tau_1) | \\ & \leq | \mathbf{P}_{acc}(RealSig(j)_k \parallel (Env_{max})_k, (\tau_{1max})_k) - \mathbf{P}_{acc}(IdealSig(j)_k \parallel (Env_{max})_k, (\tau_{1max})_k) |. \end{aligned}$$

To see that such a  $(t_{max})_k$ ,  $(Env_{max})_k$  and  $(\tau_{1max})_k$  exist, it is enough to observe that there are only a finite number of times, environments and schedules respecting the  $t_l$ ,  $p(k)$  and  $q(k)$  bounds (up to isomorphism).

This means that is enough to show the existence of a negligible function  $\epsilon$  such that, for every  $k \in \mathbb{N}$ , we have:

$$| \mathbf{P}_{acc}(RealSig(j)_k \parallel (Env_{max})_k, (\tau_{1max})_k) - \mathbf{P}_{acc}(IdealSig(j)_k \parallel (Env_{max})_k, (\tau_{1max})_k) | \leq \epsilon(k).$$

Since  $Sig$  is complete, we observe that, for every value of  $k$ , the difference of acceptance probabilities of the two automata compared in Condition (v) can only be non-zero if  $(Env_{max})_k$  succeeds in producing a forged signature (that is, a valid signature for a message that was not signed by the  $Sign$  or  $SigFunc$  automata before) and in having this signature rejected when the  $verify$  and  $respVer$  actions of  $SigFunc$  execute.

We now use each  $(Env_{max})_k$  and  $(\tau_{1max})_k$  to define a probabilistic polynomial-time (non-uniform) forger  $G = \{G_k\}_{k \in \mathbb{N}}$  for  $Sig$ , in such a way that  $G_k$  essentially emulates an execution of the automaton  $IdealSig(j)_k \parallel (Env_{max})_k$  with schedule  $(\tau_{1max})_k$ .

More precisely,  $G_k$  successively reads all the tasks in the schedule  $(\tau_{1max})_k$ , and uses them to internally emulate an execution of  $IdealSig(j)_k \parallel (Env_{max})_k$ , up to the following exceptions:

1. when the  $\{verKey(*)\}$  task has to be emulated,  $G_k$  replaces the verification algorithm obtained when emulating the  $\{chooseKeys\}$  task with the one provided by  $Sig$  in the EUF-CMA game, and
2. when the  $\{sign(*,*)\}$  task has to be emulated,  $G_k$  obtains signatures by using the signing oracle available in the EUF-CMA game.

Furthermore,  $G_k$  stores a list of all messages that the emulated  $(Env_{max})_k$  asked to sign, and checks whether  $(Env_{max})_k$  ever asks for the verification of a message with a valid signature that is not in the list. If such a signature is produced,  $G_k$  outputs it as a forgery.

We observe that this emulation process is polynomial time-bounded because all transitions of the emulated systems are polynomial time-bounded, the total running time of the system is bounded by  $t_l + q(k)$ , and Condition (iii) on the overall bound of automata guarantees that no more than a polynomial number of transitions are performed per time unit.

We also observe that the two proposed exceptions in the emulation of the execution of  $IdealSig(j)_k \parallel (Env_{max})_k$  do not change the distribution of the messages that  $(Env_{max})_k$  sees, since the verification algorithm used by  $G_k$  is generated in the same way as  $KeyGen$  generates it, and since the message signatures are also produced in a valid way. Therefore, it is with the same probability that the environment distinguishes the two systems it is interacting with (that is, by producing a forgery early enough) in a real execution of the different automata and in the version emulated by  $G$ .

Now, the assumption that  $Sig$  is EUF-CMA secure guarantees that there exists a negligible function  $\epsilon$  bounding the success probability of  $G$ . Selecting this function  $\epsilon$  completes our proof.  $\square$

## 7 Composition Theorems

In practice, cryptographic services are seldom used in isolation. Most likely, different types of services operate in conjunction, interacting with each other and with multiple protocol participants. For example, a participant may submit a document to an encryption service to obtain a ciphertext, which is later submitted to a timestamping service. In such situations, it is important that the services are provably secure even in the context of composition.

In this section, we consider two types of composition. The first, *parallel composition*, is a combination of services that are active at the same time and may interact with each other. Given a polynomially bounded collection of real services such that each real service implement some ideal service, the parallel composition of the real services is guaranteed to implement that of the ideal services.

The second type, *sequential composition*, is a combination of services that are active in succession. The interaction between two distinct services is much more limited in this setting, because the earlier one must have finished execution

before the later one comes online. An example of such a collection is the signature services in the timestamping protocol of [13, 12], where each service is replaced by the next at regular intervals.

As in the parallel case, we prove that the sequential composition of real services implements the sequential composition of ideal services. We are able to relax the restriction on the number of components from polynomial to exponential.<sup>5</sup> This highlights a unique aspect of our implementation relation: essentially, we walk down the real time line and, at every point  $t$ , we focus on a polynomial length interval starting from  $t$ .

*Parallel Composition:* Using a standard hybrid argument, we show that the relation  $\leq_{p,q,\epsilon}$  (cf. Definition 1) is preserved under polynomial parallel composition, with some appropriate adjustment to the environment complexity bound and to the error in acceptance probability.

**Theorem 2.** *Let  $b \in \mathbb{N}$  be given and, for each  $1 \leq i \leq b$ , let  $\mathcal{A}_i^1$  and  $\mathcal{A}_i^2$  be comparable task-PIOAs.  $F_i^1$  be a set of tasks of  $\mathcal{A}_i^1$  and let  $F_i^2$  be a set of tasks of  $\mathcal{A}_i^2$  satisfying the upperbound condition. Let  $\hat{F}_1$  denote  $\bigcup_{i=1}^b F_i^1$  and let  $\hat{F}_2$  denote  $\bigcup_{i=1}^b F_i^2$ . Suppose there exists a non-decreasing function  $r : \mathbb{N} \rightarrow \mathbb{N}$  such that, for all  $i$ , both  $\mathcal{A}_i^1$  and  $\mathcal{A}_i^2$  are  $r(i)$ -bounded. Suppose further that  $\mathcal{A}_1^{\alpha_1}, \dots, \mathcal{A}_b^{\alpha_b}$  are pairwise compatible for any combination of  $\alpha_i \in \{1, 2\}$ .*

*Let  $p, p', q \in \mathbb{N}$  and  $\epsilon, \epsilon' \in \mathbb{R}_{\geq 0}$  be given, and assume the following.*

(1)  $p = c_{\text{comp}} \cdot (b \cdot r(b) + p')$ , where  $c_{\text{comp}}$  is the constant factor for composing task-PIOAs in parallel.

(2)  $\epsilon' = b \cdot \epsilon$ .

(3) For all  $i$ ,  $(\mathcal{A}_i^1, F_i^1) \leq_{p,q,\epsilon} (\mathcal{A}_i^2, F_i^2)$ .

*Then we have  $(\|\bigcup_{i=1}^b \mathcal{A}_i^1, \hat{F}_1) \leq_{p',q,\epsilon'} (\|\bigcup_{i=1}^b \mathcal{A}_i^2, \hat{F}_2)$ .*

*Proof.* Let  $t \in \mathbb{R}_{\geq 0}$  be given. Let  $\text{Env} = \text{Env}' \parallel \text{Clock}$  be a  $p'$ -bounded environment and let  $\tau_0$  be a valid timed task schedule for  $\|\bigcup_{i=1}^b \mathcal{A}_i^1 \parallel \text{Env}$  for the interval  $[0, t + q]$  where  $\tau_0$  contains no actions from  $\hat{F}_1$  occurring at  $t$  or later.

For each  $0 \leq i \leq b$ , let  $H_i$  denote  $\mathcal{A}_1^2 \parallel \dots \parallel \mathcal{A}_i^2 \parallel \mathcal{A}_{i+1}^1 \parallel \dots \parallel \mathcal{A}_b^1$ . In particular,  $H_0 = \|\bigcup_{i=1}^b \mathcal{A}_i^1$  and  $H_b = \|\bigcup_{i=1}^b \mathcal{A}_i^2$ . Similarly, let

$$\text{Env}_i := \mathcal{A}_1^2 \parallel \dots \parallel \mathcal{A}_{i-1}^2 \parallel \mathcal{A}_{i+1}^1 \parallel \dots \parallel \mathcal{A}_b^1 \parallel \text{Env}$$

for each  $1 \leq i \leq b$ . Note that every  $\text{Env}_i$  is  $p$ -bounded and is an environment for  $\mathcal{A}_i^1$  and  $\mathcal{A}_i^2$ . In fact, we have  $H_{i-1} \parallel \text{Env} = \mathcal{A}_i^1 \parallel \text{Env}_i$  and  $H_i \parallel \text{Env} = \mathcal{A}_i^2 \parallel \text{Env}_i$ .

Since  $\tau_0$  does not contain any tasks from  $\hat{F}_1$  at time  $t$  or later, it does not contain any tasks from  $F_1^1$  from time  $t$  or later. Since  $(\mathcal{A}_1^1, F_1^1) \leq_{p,q,\epsilon} (\mathcal{A}_1^2, F_1^2)$  and  $\tau_0$  is a valid schedule for  $\mathcal{A}_1^1 \parallel \text{Env}_1$  in which no tasks from  $F_1^1$  occur from time  $t$  onwards, we may choose a valid schedule  $\tau_1$  for  $\mathcal{A}_1^2 \parallel \text{Env}_1$  for the interval  $[0, t + q]$  such that

- (i)  $\mathbf{P}_{\text{acc}}(\mathcal{A}_1^1 \parallel \text{Env}_1, \text{trunc}_{\geq t}(\tau_0)) = \mathbf{P}_{\text{acc}}(\mathcal{A}_1^2 \parallel \text{Env}_1, \text{trunc}_{\geq t}(\tau_1))$ ;
- (ii)  $\text{lstate}_{\text{Env}_1}(\mathcal{A}_1^1 \parallel \text{Env}_1, \text{trunc}_{\geq t}(\tau_0)) = \text{lstate}_{\text{Env}_1}(\mathcal{A}_1^2 \parallel \text{Env}_1, \text{trunc}_{\geq t}(\tau_1))$ ;
- (iii)  $\text{proj}_{\text{Env}_1}(\tau_0) = \text{proj}_{\text{Env}_1}(\tau_1)$ ;
- (iv)  $\tau_1$  does not contain any pairs of the form  $\langle T_i, t_i \rangle$  where  $T_i \in F_1^2$  and  $t_i \geq t$ ;
- (v)  $|\mathbf{P}_{\text{acc}}(\mathcal{A}_1^1 \parallel \text{Env}_1, \tau_0) - \mathbf{P}_{\text{acc}}(\mathcal{A}_1^2 \parallel \text{Env}_1, \tau_1)| \leq \epsilon$ .

Repeating this argument, we choose valid schedules  $\tau_2, \dots, \tau_b$  for  $H_2 \parallel \text{Env}, \dots, H_b \parallel \text{Env}$ , respectively, all satisfying the appropriate five conditions. By Condition (i), we have

$$\mathbf{P}_{\text{acc}}(H_0 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_0)) = \mathbf{P}_{\text{acc}}(H_1 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_1)) = \dots = \mathbf{P}_{\text{acc}}(H_b \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_b)).$$

Also, since  $\text{Env}$  is part of every  $\text{Env}_i$ , Condition (ii) guarantees that

$$\text{lstate}_{\text{Env}}(H_0 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_0)) = \text{lstate}_{\text{Env}}(H_b \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_b)).$$

Similarly, Condition (iii) guarantees that  $\text{proj}_{\text{Env}}(\tau_0) = \text{proj}_{\text{Env}}(\tau_b)$ .

Using both Conditions (iii) and (iv), we can infer that  $\tau_b$  does not contain any pairs of the form  $\langle T_i, t_i \rangle$  where  $T_i \in \hat{F}_2 = \bigcup_{i=1}^b F_i^2$  and  $t_i \geq t$ . Finally,

$$\begin{aligned} & |\mathbf{P}_{\text{acc}}(\|\bigcup_{i=1}^b \mathcal{A}_i^1 \parallel \text{Env}, \tau_0) - \mathbf{P}_{\text{acc}}(\|\bigcup_{i=1}^b \mathcal{A}_i^2 \parallel \text{Env}, \tau_b)| \\ & \leq |\mathbf{P}_{\text{acc}}(H_0 \parallel \text{Env}, \tau_0) - \mathbf{P}_{\text{acc}}(H_1 \parallel \text{Env}, \tau_1)| + \dots \\ & \quad + |\mathbf{P}_{\text{acc}}(H_i \parallel \text{Env}, \tau_i) - \mathbf{P}_{\text{acc}}(H_{i+1} \parallel \text{Env}, \tau_{i+1})| + \dots \\ & \quad + |\mathbf{P}_{\text{acc}}(H_{b-1} \parallel \text{Env}, \tau_{b-1}) - \mathbf{P}_{\text{acc}}(H_b \parallel \text{Env}, \tau_b)| \\ & \leq b \cdot \epsilon = \epsilon'. \end{aligned}$$

□

<sup>5</sup> In our model, it is not meaningful to exceed an exponential number of components, because the length of the description of each component is polynomially bounded.

Using Theorem 2, it is not hard to prove that  $\leq_{\text{neg,pt}}$  is preserved under polynomial composition.

**Theorem 3 (Parallel Composition Theorem for  $\leq_{\text{neg,pt}}$ ).** *Let two sequences of task-PIOA families  $\bar{\mathcal{A}}_1^1, \bar{\mathcal{A}}_2^1, \dots$  and  $\bar{\mathcal{A}}_1^2, \bar{\mathcal{A}}_2^2, \dots$  be given, with  $\bar{\mathcal{A}}_i^1$  comparable to  $\bar{\mathcal{A}}_i^2$  for all  $i$ . Assume that  $\bar{\mathcal{A}}_1^{\alpha_1}, \bar{\mathcal{A}}_2^{\alpha_2}, \dots$  are pairwise compatible for any combination of  $\alpha_i \in \{1, 2\}$ . For each  $i$ , let  $\bar{F}_i^1$  be a family of sets such that  $(\bar{F}_i^1)_k$  is a set of tasks of  $(\bar{\mathcal{A}}_i^1)_k$  for every  $k$  and let  $\bar{F}_i^2$  be a family of sets such that  $(\bar{F}_i^2)_k$  is a set of tasks of  $(\bar{\mathcal{A}}_i^2)_k$  for every  $k$ , satisfying the upperbound condition.*

*Suppose there exist polynomials  $r, s : \mathbb{N} \rightarrow \mathbb{N}$  such that, for all  $i, k$ , both  $(\bar{\mathcal{A}}_i^1)_k$  and  $(\bar{\mathcal{A}}_i^2)_k$  are bounded by  $s(k) \cdot r(i)$ . Assume that  $r$  is non-decreasing and*

$$\forall p, q \exists \epsilon \forall i (\bar{\mathcal{A}}_i^1, \bar{F}_i^1) \leq_{p,q,\epsilon} (\bar{\mathcal{A}}_i^2, \bar{F}_i^2), \quad (1)$$

*where  $p, q$  are polynomials and  $\epsilon$  is a negligible function. (This is a strengthening of the statement  $\forall i (\bar{\mathcal{A}}_i^1, \bar{F}_i^1) \leq_{\text{neg,pt}} (\bar{\mathcal{A}}_i^2, \bar{F}_i^2)$ .) Let  $b$  be any polynomial. For each  $k$ , let  $(\hat{\mathcal{A}}^1)_k$  denote  $(\bar{\mathcal{A}}_1^1)_k \parallel \dots \parallel (\bar{\mathcal{A}}_{b(k)}^1)_k$ . Similarly for  $(\hat{\mathcal{A}}^2)_k$ . Also, let  $(\hat{F}_1)_k$  denote  $\bigcup_{i=1}^{b(k)} (\bar{F}_i^1)_k$  and let  $(\hat{F}_2)_k$  denote  $\bigcup_{i=1}^{b(k)} (\bar{F}_i^2)_k$ . Then we have  $(\hat{\mathcal{A}}^1, \hat{F}_1) \leq_{\text{neg,pt}} (\hat{\mathcal{A}}^2, \hat{F}_2)$ .*

*Proof.* By the definition of  $\leq_{\text{neg,pt}}$ , we need to prove the following:  $\forall p', q \exists \epsilon' (\hat{\mathcal{A}}^1, \hat{F}_1) \leq_{p',q,\epsilon'} (\hat{\mathcal{A}}^2, \hat{F}_2)$ , where  $p', q$  are polynomials and  $\epsilon'$  is a negligible function. Let polynomials  $p'$  and  $q$  be given and define  $p := c_{\text{comp}} \cdot (b \cdot (r \circ b) + p')$ , where  $c_{\text{comp}}$  is the constant factor for composing task-PIOAs in parallel. Now choose  $\epsilon$  using  $p, q$ , and Assumption (1). Define  $\epsilon' := b \cdot \epsilon$ .

Let  $k \in \mathbb{N}$  be given. We need to prove  $((\hat{\mathcal{A}}^1)_k, (\hat{F}_1)_k) \leq_{p'(k),q(k),\epsilon'(k)} ((\hat{\mathcal{A}}^2)_k, (\hat{F}_2)_k)$ . That is,

$$((\bar{\mathcal{A}}_1^1)_k \parallel \dots \parallel (\bar{\mathcal{A}}_{b(k)}^1)_k, \bigcup_{i=1}^{b(k)} (\bar{F}_i^1)_k) \leq_{p'(k),q(k),\epsilon'(k)} ((\bar{\mathcal{A}}_1^2)_k \parallel \dots \parallel (\bar{\mathcal{A}}_{b(k)}^2)_k, \bigcup_{i=1}^{b(k)} (\bar{F}_i^2)_k).$$

For every  $i$ , we know that  $(\bar{\mathcal{A}}_i^1)_k$  and  $(\bar{\mathcal{A}}_i^2)_k$  are bounded by  $(s(k) \cdot r)(i)$ . Also, by the choice of  $\epsilon$ , we have  $((\bar{\mathcal{A}}_i^1)_k, (\bar{F}_i^1)_k) \leq_{p(k),q(k),\epsilon(k)} ((\bar{\mathcal{A}}_i^2)_k, (\bar{F}_i^2)_k)$  for all  $i$ . Therefore, we may apply Theorem 2 to conclude that  $((\hat{\mathcal{A}}^1)_k, \bigcup_{i=1}^{b(k)} (\bar{F}_i^1)_k) \leq_{p'(k),q(k),\epsilon'(k)} ((\hat{\mathcal{A}}^2)_k, \bigcup_{i=1}^{b(k)} (\bar{F}_i^2)_k)$ . This completes the proof.  $\square$

*Sequential Composition:* We now treat the more interesting case, namely, exponential sequential composition. The first challenge is to formalize the notion of sequentiality. On a syntactic level, all components in the collection are combined using the parallel composition operator. To capture the idea of successive invocation, we introduce some auxiliary notions. Intuitively, we distinguish between *active* and *dormant* entities. Active entities may perform actions and store information in memory. Dormant entities have no available memory and do not enable locally controlled actions.<sup>6</sup> In Definition 3, we formalize the idea of an entity  $\mathcal{A}$  being active during one specific time interval. Then we introduce sequentiality in Definition 4.

**Definition 3.** *Let  $\mathcal{A}$  be a task-PIOA and let reals  $t_1 \leq t_2$  be given. We say that  $\mathcal{A}$  is restricted to the interval  $[t_1, t_2]$  if:*

- for any  $t < t_1$ , environment  $\text{Env}$  for  $\mathcal{A}$  of the form  $\text{Env}' \parallel \text{Clock}$ , valid schedule  $\tau$  for  $\mathcal{A} \parallel \text{Env}$  for  $[0, t]$ , and state  $s$  reachable under  $\tau$ , no locally controlled actions of  $\mathcal{A}$  are enabled in  $s$ , and  $s.v = \perp$  for every variable  $v$  of  $\mathcal{A}$ .
- the same for all  $t > t_2$ .

**Definition 4 (Sequentiality).** *Let  $\mathcal{A}_1, \mathcal{A}_2, \dots$  be pairwise compatible task-PIOAs. We say that  $\mathcal{A}_1, \mathcal{A}_2, \dots$  are sequential if there exist reals  $0 \leq t_1 < t_2 < \dots$  such that: for all  $i$ ,  $\mathcal{A}_i$  is restricted to  $[t_i, t_{i+1}]$ .*

Note that each  $\mathcal{A}_i$  may overlap with  $\mathcal{A}_{i+1}$  at the boundary time  $t_{i+1}$ . Lemma 5 below states the intuitive fact that no environment can distinguish two entities during an interval in which both entities are dormant.

**Lemma 5.** *Suppose  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are comparable task-PIOAs that are both restricted to the interval  $[t_1, t_2]$ . Let  $\text{Env}$  be an environment for both  $\mathcal{A}_1$  and  $\mathcal{A}_2$  and of the form  $\text{Env}' \parallel \text{Clock}$ . Let  $t \in \mathbb{R}_{\geq 0}$  and  $q \in \mathbb{N}$  be given. Suppose we have valid schedule  $\tau_1$  for  $\mathcal{A}_1 \parallel \text{Env}$  for the interval  $[0, t + q]$  and valid schedule  $\tau_2$  for  $\mathcal{A}_2 \parallel \text{Env}$  for the interval  $[0, t + q]$ , satisfying:*

<sup>6</sup> For technical reasons, dormant entities must synchronize on input actions. Some inputs cause dormant entities to become active, while all others are trivial loops on the null state.

- $\mathbf{P}_{\text{acc}}(\mathcal{A}_1 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_1)) = \mathbf{P}_{\text{acc}}(\mathcal{A}_2 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_2));$
- $\text{lstate}_{\text{Env}}(\mathcal{A}_1 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_1)) = \text{lstate}_{\text{Env}}(\mathcal{A}_2 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_2));$
- $\text{proj}_{\text{Env}}(\tau_1) = \text{proj}_{\text{Env}}(\tau_2).$

Assume further that either  $t_2 < t$  or  $t_1 > t + q$ . Then  $\mathbf{P}_{\text{acc}}(\mathcal{A}_1 \parallel \text{Env}, \tau_1) = \mathbf{P}_{\text{acc}}(\mathcal{A}_2 \parallel \text{Env}, \tau_2)$ .

*Proof.* First we consider the case  $t_2 < t$ . Since  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are restricted to the interval  $[t_1, t_2]$ , neither of them enables any output actions during the interval  $[t, t + q]$ . By assumption,  $\tau_1$  and  $\tau_2$  agree on the tasks of Env and the state distributions of Env just before time  $t$  are identical in the two experiments. Therefore, the probability that Env outputs acc during  $[t, t + q]$  must be identical in the two experiments. We also have the assumption that Env outputs acc with the same probability during  $[0, t)$ , therefore the acceptance probabilities are the same for the entire interval  $[0, t]$ .

Similarly, if  $t_1 > t + q$ , then neither  $\mathcal{A}_1$  nor  $\mathcal{A}_2$  enables any output actions during the interval  $[t, t + q]$ . Then we follow the same argument as above.  $\square$

Now we are ready to state the sequential composition theorems.

**Theorem 4.** Let  $\mathcal{A}_1^1, \mathcal{A}_2^1, \dots$  and  $\mathcal{A}_1^2, \mathcal{A}_2^2, \dots$  be two sequences of task-PIOAs such that  $\mathcal{A}_i^1$  and  $\mathcal{A}_i^2$  are comparable for every  $i$ . Assume that  $\mathcal{A}_1^{\alpha_1}, \mathcal{A}_2^{\alpha_2}, \dots$  are pairwise compatible for any combination of  $\alpha_i \in \{1, 2\}$ . Also, let  $L, \hat{p} \in \mathbb{N}$  be given such that both  $\|\mathcal{A}_{i=1}^L \mathcal{A}_i^1$  and  $\|\mathcal{A}_{i=1}^L \mathcal{A}_i^2$  are  $\hat{p}$ -bounded. Assume that both  $\mathcal{A}_1^1, \dots, \mathcal{A}_L^1$  and  $\mathcal{A}_1^2, \dots, \mathcal{A}_L^2$  are sequential for the same sequence of reals  $t_1 < \dots < t_{L+1}$ .

Let  $p, q \in \mathbb{N}$  and  $\epsilon \in \mathbb{R}_{\geq 0}$  be given. Suppose there are sets of tasks  $F_i^1, F_i^2, 1 \leq i \leq L$ , such that  $(\mathcal{A}_i^1, F_i^1) \leq_{p, q, \epsilon} (\mathcal{A}_i^2, F_i^2)$  for all  $i$ . Let  $\hat{F}_1$  denote  $\bigcup_{i=1}^L F_i^1$  and  $\hat{F}_2$  denote  $\bigcup_{i=1}^L F_i^2$ . Let  $b$  denote the largest number such that  $b$  consecutive  $t_i$ 's fall into a single closed interval of length  $q$ . (Such  $b$  must exist and is between 1 and  $L$ ). Let  $p' \in \mathbb{N}$  and  $\epsilon' \in \mathbb{R}_{\geq 0}$  be given, with  $\epsilon' \geq (b + 2) \cdot \epsilon$  and  $p \geq c_{\text{comp}} \cdot (\hat{p} + p')$  (where  $c_{\text{comp}}$  is the constant factor for parallel composition). Then we have  $(\|\mathcal{A}_{i=1}^L \mathcal{A}_i^1, \hat{F}_1) \leq_{p', q, \epsilon'} (\|\mathcal{A}_{i=1}^L \mathcal{A}_i^2, \hat{F}_2)$ .

In the statement of Theorem 4, the error in acceptance probability increases by a factor of  $b + 2$ , where  $b$  is the largest number of components that may be active in a closed time interval of length  $q$ . For example, if the life time of each component is  $\frac{q}{3}$ , then  $b$  is 5.<sup>7</sup> This is the key difference between parallel composition and sequential composition: for the former, error increases with the total number of components (namely,  $L$ ), and hence no more than a polynomial number of components can be tolerated. In the sequential case,  $L$  may be exponential, as long as  $b$  remains small. The proof of Theorem 4 involves a standard hybrid argument for active components, while dormant components are replaced without affecting the difference in acceptance probabilities.

*Proof (Proof of Theorem 4).* Let  $t \in \mathbb{R}_{\geq 0}$  be given. Let  $\text{Env} = \text{Env}' \parallel \text{Clock}$  be a  $p'$ -bounded environment and let  $\tau_0$  be a valid timed task schedule for  $(\|\mathcal{A}_{i=1}^L \mathcal{A}_i^1) \parallel \text{Env}$  for the interval  $[0, t + q]$  where  $\tau_0$  has no tasks from  $\hat{F}_1$  occurring at  $t$  or later. We need to find  $\tau_L$  for  $(\|\mathcal{A}_{i=1}^L \mathcal{A}_i^2) \parallel \text{Env}$  such that

- (i)  $\mathbf{P}_{\text{acc}}(\|\mathcal{A}_{i=1}^L \mathcal{A}_i^1 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_0)) = \mathbf{P}_{\text{acc}}(\|\mathcal{A}_{i=1}^L \mathcal{A}_i^2 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_L));$
- (ii)  $\text{lstate}_{\text{Env}}(\|\mathcal{A}_{i=1}^L \mathcal{A}_i^1 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_0)) = \text{lstate}_{\text{Env}}(\|\mathcal{A}_{i=1}^L \mathcal{A}_i^2 \parallel \text{Env}, \text{trunc}_{\geq t}(\tau_L));$
- (iii)  $\text{proj}_{\text{Env}}(\tau_0) = \text{proj}_{\text{Env}}(\tau_L);$
- (iv)  $\tau_L$  does not contain any pairs of the form  $\langle T_i, t_i \rangle$  where  $T_i \in \hat{F}_2$  and  $t_i \geq t$ ;
- (v)  $|\mathbf{P}_{\text{acc}}(\|\mathcal{A}_{i=1}^L \mathcal{A}_i^1 \parallel \text{Env}, \tau_0) - \mathbf{P}_{\text{acc}}(\|\mathcal{A}_{i=1}^L \mathcal{A}_i^2 \parallel \text{Env}, \tau_L)| \leq \epsilon'.$

Without loss of generality, assume there is an index  $i$  such that  $[t_i, t_{i+1}]$  intersects with  $[t, t + q]$ . Let  $l$  be the smallest such index. Recall from the assumptions that at most  $b$  consecutive  $t_i$ 's fall into a closed interval of length  $q$ . Therefore, we know that  $t_{l-1} < t$  and  $t_{l+b} > t + q$ .

The rest of the proof proceeds as in the proof of Theorem 2. Namely, we define

$$\text{Env}_i := \mathcal{A}_1^2 \parallel \dots \parallel \mathcal{A}_{i-1}^2 \parallel \mathcal{A}_{i+1}^1 \parallel \dots \parallel \mathcal{A}_l^1 \parallel \text{Env}$$

for each  $1 \leq i \leq L$ . Note that  $\text{Env}_i$  is  $p$ -bounded, therefore we may choose  $\tau_{i+1}$  using  $\tau_i$  and the assumption that  $(\mathcal{A}_i^1, F_i^1) \leq_{p, q, \epsilon} (\mathcal{A}_i^2, F_i^2)$ . Since  $\text{Env}$  is part of  $\text{Env}_i$  for every  $i$ , Conditions (i) through (iii) are clearly satisfied at every replacement step. Condition (iv) is satisfied because the following hold at every step  $i$ .

- The new task schedule  $\tau_{i+1}$  does not contain tasks from  $F_{i+1}^2$ .

<sup>7</sup> Recall that two components may be active simultaneously at the boundary time.

- Condition (iii) guarantees that  $\tau_{i+1}$  does not contain tasks from  $\bigcup_{j=1}^i F_j^2$ .

Finally, we consider Condition (v). There are two cases. If  $i < l-1$  or  $i \geq l+b$ , then we can apply Lemma 5 to conclude that  $\mathbf{P}_{\text{acc}}(\mathcal{A}_i^1 \parallel \text{Env}_i, \tau_i)$  in fact equals  $\mathbf{P}_{\text{acc}}(\mathcal{A}_i^2 \parallel \text{Env}_i, \tau_{i+1})$ . Otherwise,  $\mathbf{P}_{\text{acc}}(\mathcal{A}_i^1 \parallel \text{Env}_i, \tau_i)$  and  $\mathbf{P}_{\text{acc}}(\mathcal{A}_i^2 \parallel \text{Env}_i, \tau_{i+1})$  differ by at most  $\epsilon$ . Summing over all indices  $i$ , we have  $|\mathbf{P}_{\text{acc}}(\|\mathcal{A}_{i=1}^L \parallel \text{Env}, \tau_0) - \mathbf{P}_{\text{acc}}(\|\mathcal{A}_{i=1}^L \parallel \text{Env}, \tau_L)| \leq (b+2) \cdot \epsilon = \epsilon'$ .  $\square$

Using Theorem 4, it is straightforward to prove the sequential composition theorem for  $\leq_{\text{neg,pt}}$ .

**Theorem 5 (Sequential Composition Theorem for  $\leq_{\text{neg,pt}}$ ).** *Let two sequences of task-PIOA families  $\bar{\mathcal{A}}_1^1, \bar{\mathcal{A}}_2^1, \dots$  and  $\bar{\mathcal{A}}_1^2, \bar{\mathcal{A}}_2^2, \dots$  be given, with  $\bar{\mathcal{A}}_i^1$  comparable to  $\bar{\mathcal{A}}_i^2$  for all  $i$ . Assume that  $\bar{\mathcal{A}}_1^{\alpha_1}, \bar{\mathcal{A}}_2^{\alpha_2}, \dots$  are pairwise compatible for any combination of  $\alpha_i \in \{1, 2\}$ . For each  $i$ , let  $\bar{F}_i^1$  be a family of sets such that  $(\bar{F}_i^1)_k$  is a set of tasks of  $(\bar{\mathcal{A}}^2 = 1_i)_k$  for every  $k$  and let  $\bar{F}_i^2$  be a family of sets such that  $(\bar{F}_i^2)_k$  is a set of tasks of  $(\bar{\mathcal{A}}_i^2)_k$  for every  $k$  satisfying the upperbound condition. Let  $L : \mathbb{N} \rightarrow \mathbb{N}$  be an exponential function and, for each  $k$ , let  $(\hat{\mathcal{A}}^1)_k$  denote  $(\bar{\mathcal{A}}_1^1)_k \parallel \dots \parallel (\bar{\mathcal{A}}_{L(k)}^1)_k$ . Similarly for  $(\hat{\mathcal{A}}^2)_k$ . Also, let  $(\hat{F}_1)_k$  denote  $\bigcup_{i=1}^{L(k)} (\bar{F}_i^1)_k$  and let  $(\hat{F}_2)_k$  denote  $\bigcup_{i=1}^{L(k)} (\bar{F}_i^2)_k$ .*

*Let  $\hat{p}$  be a polynomial such that both  $\hat{\mathcal{A}}^1$  and  $\hat{\mathcal{A}}^2$  are  $\hat{p}$ -bounded. Suppose there exist a sequence of positive reals  $t_1 < t_2 < \dots$  such that, for each  $k$ , both  $(\bar{\mathcal{A}}_1^1)_k, \dots, (\bar{\mathcal{A}}_{L(k)}^1)_k$  and  $(\bar{\mathcal{A}}_1^2)_k, \dots, (\bar{\mathcal{A}}_{L(k)}^2)_k$  are sequential for the sequence  $t_1 < \dots < t_{L(k)+1}$ . Assume there is a constant real number  $c$  such that consecutive  $t_i$ 's are at least  $c$  apart.*

*Suppose that, for every pair of polynomials  $\langle p, q \rangle$ , there exists negligible function  $\epsilon$  such that  $(\bar{\mathcal{A}}_i^1, \bar{F}_i^1) \leq_{p,q,\epsilon} (\bar{\mathcal{A}}_i^2, \bar{F}_i^2)$  for all  $i$ . Then we have  $(\hat{\mathcal{A}}^1, \hat{F}_1) \leq_{\text{neg,pt}} (\hat{\mathcal{A}}^2, \hat{F}_2)$ .*

*Proof.* Let polynomials  $p', q$  be given and define  $p := c_{\text{comp}} \cdot (\hat{p} + p')$ , where  $c_{\text{comp}}$  is the constant factor for composing task-PIOAs in parallel. Choose  $\epsilon$  from  $p, q$  according to the assumption of the theorem. For each  $k$ , let  $b(k)$  be the ceiling of  $\frac{q(k)}{c} + 1$ . (The choice of  $b(k)$  ensures that at most  $b(k)$  consecutive  $t_i$ 's fall within any interval of length at most  $q(k)$ . This is necessary in order to apply Theorem 4.) Since  $c$  is constant,  $b$  is a polynomial. Define  $\epsilon' := b \cdot \epsilon$ .

For every  $k \in \mathbb{N}$ , we apply Theorem 4 to conclude that

$$((\bar{\mathcal{A}}_1^1)_k \parallel \dots \parallel (\bar{\mathcal{A}}_{L(k)}^1)_k, (\hat{F}_1)_k) \leq_{p'(k), q(k), \epsilon'(k)} ((\bar{\mathcal{A}}_1^2)_k \parallel \dots \parallel (\bar{\mathcal{A}}_{L(k)}^2)_k, (\hat{F}_2)_k).$$

That is,  $((\hat{\mathcal{A}}^1)_k, (\hat{F}_1)_k) \leq_{p'(k), q(k), \epsilon'(k)} ((\hat{\mathcal{A}}^2)_k, (\hat{F}_2)_k)$ . This completes the proof.  $\square$

## 8 Application: Digital Timestamping

In this section, we present a formal model of the digital timestamping protocol of Haber et al. (cf. Section 1). Recall the real and ideal signature services from Section 6. The timestamping protocol consists of a dispatcher component and a collection of real signature services. Similarly, the ideal protocol consists of the same dispatcher with a collection of ideal signature services. Using the parallel and sequential composition theorems (Theorems 3 and 5), we prove that the real protocol implements the ideal protocol with respect to the long-term implementation relation  $\leq_{\text{neg,pt}}$ . This result implies that, no matter what security failures (forgeries, guessed keys, etc.) occur up to any particular time  $t$ , new certifications and verifications performed by services that awaken after time  $t$  will still be correct (with high probability) for a polynomial-length interval of time after  $t$ .

Note that this result does *not* imply that any particular document is reliably certified for super-polynomial time. In fact, the protocol does not guarantee this: even if a document certificate is refreshed frequently by new services, there is at any time a small probability that the environment guesses the current certificate, thus creating a forgery. That probability, over super-polynomial time, becomes large. Once the environment guesses a current certificate, it can continue to refresh the certificate forever, thus maintaining the forgery.

Let  $SID$ , the domain of service names, be  $\mathbb{N}$ . In addition to  $\text{alive}$  and  $\text{aliveTimes}$  (cf. Section 4), we assume the following.

- $\text{pref} : \mathbb{T} \rightarrow SID$ . For every  $t \in \mathbb{T}$ , the service  $\text{pref}(t)$  is the designated signer for time  $t$ , i.e., any signing request sent by the dispatcher at time  $t$  goes to service  $\text{pref}(t)$ .
- $\text{usable} : \mathbb{T} \rightarrow 2^{SID}$ . For every  $t \in \mathbb{T}$ ,  $\text{usable}(t)$  specifies the set of services that are accepting new verification requests.

Assume, for every  $t \in \mathbb{T}$ ,  $\text{pref}(t) \in \text{usable}(t) \subseteq \text{alive}(t)$ . If a service is preferred, it accepts both signing and verification requests. If it is alive but not usable, no new verification requests are accepted, but those already submitted will still be processed.



*Dispatcher*: We define  $\text{Dispatcher}_k$  for each security parameter  $k$ . If the environment sends a first-time certificate request  $\text{reqCert}(rid, x)$ ,  $\text{Dispatcher}_k$  requests a signature from service  $j = \text{pref}(t)$  via the action  $\text{reqSign}(rid, \langle x, t, \perp \rangle)_j$ , where  $t$  is the clock reading at the time of  $\text{reqSign}$ . In this communication, we instantiate the message space  $M_k$  as  $X_k \times \mathbb{T}_k \times (\Sigma_k)_\perp$ , where  $X_k$  is the domain of documents to which timestamps are associated. After service  $j$  returns with action  $\text{respSign}(rid, \sigma)_j$ ,  $\text{Dispatcher}_k$  issues a new certificate via  $\text{respCert}(rid, \sigma, j)$ .

If a renew request  $\text{reqCert}(rid, x, t, \sigma_1, \sigma_2, j)$  comes in,  $\text{Dispatcher}_k$  first checks to see if  $j$  is still usable. If not, it responds with  $\text{respCert}(rid, \text{false})$ . Otherwise, it sends  $\text{reqVer}(rid, \langle x, t, \sigma_1 \rangle, \sigma_2)_j$  to service  $j$ . If service  $j$  answers affirmatively,  $\text{Dispatcher}_k$  sends a signature request  $\text{reqSign}(rid, \langle x, t, \sigma_2 \rangle)_{j'}$ , where  $j'$  is the current preferred service. When service  $j'$  returns with action  $\text{respSign}_{j'}(rid, \sigma_3)$ ,  $\text{Dispatcher}_k$  issues a new certificate via  $\text{respCert}(rid, \sigma_3, j')$ .

The code for  $\text{Dispatcher}$  appears in The task-PIOA code for the component  $\text{Dispatcher}$  appears in Figure 7. As a convention, we use  $\sigma_1, \sigma_2$  and  $\sigma_3$  to denote previous, current, and new signatures, respectively.

*Concrete Time Scheme*: Let  $d$  be a positive natural number. Each service  $j$  is alive from time  $(j-1) \cdot d$  to  $(j+2) \cdot d$ . Thus, at any given point in time, there can be at most three services that are concurrently alive. Moreover, service  $j$  is preferred for signing from time  $(j-1) \cdot d$  to  $j \cdot d$ , and is usable from time  $(j-1) \cdot d$  to  $(j+1) \cdot d$ . Between  $(j+1) \cdot d$  and  $(j+2) \cdot d$ , services  $j$  continues to process requests already submitted, without receiving new requests.

*Protocol Correctness*: For every security parameter  $k$ , let  $SID_k \subseteq SID$  denote the set of  $p(k)$ -bit numbers, for some polynomial  $p$ . Recall from Section 6 that  $\text{RealSig}(j)_k = \text{hide}(\text{KeyGen}(k, j) \parallel \text{Signer}(k, j) \parallel \text{Verifier}(k, j), \text{signKey}_j)$  and  $\text{IdealSig}(j)_k = \text{hide}(\text{KeyGen}(k, j) \parallel \text{SigFunc}(k, j), \text{signKey}_j)$ . Here we define

$$\text{RealSigSys}_k := \text{Dispatcher}_k \parallel (\parallel_{j \in SID_k} \text{RealSig}(j)_k) \quad \text{and} \quad \text{IdealSigSys}_k := \text{Dispatcher}_k \parallel (\parallel_{j \in SID_k} \text{IdealSig}(j)_k).$$

Next, define  $\overline{\text{RealSigSys}} := \{\text{RealSigSys}_k\}_{k \in \mathbb{N}}$  and  $\overline{\text{IdealSigSys}} := \{\text{IdealSigSys}_k\}_{k \in \mathbb{N}}$ . Our goal is to show that

$$(\overline{\text{RealSigSys}}, \emptyset) \leq_{\text{neg.pt}} (\overline{\text{IdealSigSys}}, \bar{F}),$$

where we use  $\emptyset$  for a family of empty failure sets and  $\bar{F}_k := \bigcup_{j \in SID_k} \{\{\text{fail}_j\}\}$  for every  $k$  (Theorem 6).

First we make a key observation.

**Lemma 6.** *Suppose we have  $k \in \mathbb{N}$ ,  $j \in SID_k$ . Then  $\text{RealSig}(j)_k$  is restricted to  $[(j-1) \cdot d, (j+2) \cdot d]$ . Similarly for  $\text{IdealSig}(j)_k$ .*

*Proof.* Suppose we have  $t < (j-1) \cdot d$ , environment  $\text{Env}$  for  $\text{RealSig}(j)_k$  of the form  $\text{Env}' \parallel \text{Clock}$ , valid schedule  $\tau$  for  $\text{RealSig}(j)_k \parallel \text{Env}$  for  $[0, t]$ , and state  $s$  reachable under  $\tau$ . Recall from Section 3 that, for every  $t' \in \mathbb{T}$ , the action  $\text{tick}(t')$  must take place at time  $t'$ . Therefore,  $\tau$  does not trigger a  $\text{tick}(t')$  action with  $t' \in [(j-1) \cdot d, (j+2) \cdot d]$ . On the other hand, all variables of  $\text{RealSig}(j)_k$  remains  $\perp$  unless such a  $\text{tick}(t')$  action takes place, so we can conclude that  $s.v = \perp$  for every variable  $v$  of  $\text{RealSig}(j)_k$ .

For  $t > (j+2) \cdot d$ , we know that  $\tau$  must have triggered the action  $\text{tick}((j+2) \cdot d)$ , which sets all variables of  $\text{RealSig}(j)_k$  to  $\perp$ . Moreover, every subsequent  $\text{tick}(t')$  has  $t' > t$ , therefore the variables remain  $\perp$ .

Finally, by inspection of the code for  $\text{RealSig}(j)_k$ , we know that no locally controlled actions are enabled if all variables are  $\perp$ .

The proof for  $\text{IdealSig}(j)_k$  is similar. □

For each  $i \in \{0, 1, 2\}$ , define  $\text{Real}_{i,k}$  to be the parallel composition of all  $\text{RealSig}(j)_k$  with  $(j-1) \bmod 3 = i$ . Let  $\overline{\text{Real}}_i$  be  $\{\text{Real}_{i,k}\}_{k \in \mathbb{N}}$ . By Lemma 6, we know that  $\text{RealSig}(i)_k, \text{RealSig}(i+3)_k, \dots$  are sequential. Thus, we have partitioned the collection of real signature services into three classes,  $\overline{\text{Real}}_0, \overline{\text{Real}}_1$ , and  $\overline{\text{Real}}_2$ , such that the services within each  $\overline{\text{Real}}_i$  are sequential. For instance, the first class consists of services 1, 4,  $\dots$ , which are alive in intervals  $[0, 3d], [3d, 6d], \dots$  respectively.

Define  $\text{Ideal}_{i,k}$  and  $\overline{\text{Ideal}}_i$  similarly. We make the following observations.

**Lemma 7.** *The task-PIOA families  $\overline{\text{Real}}_0, \overline{\text{Real}}_1, \overline{\text{Real}}_2, \overline{\text{Ideal}}_0, \overline{\text{Ideal}}_1$ , and  $\overline{\text{Ideal}}_2$  are polynomially bounded.*

**Lemma 8.** *The following hold for every  $k$ .*

1.  $\text{RealSig}(1)_k, \text{RealSig}(4)_k, \dots$  in  $\overline{\text{Real}}_0$  and  $\text{IdealSig}(1)_k, \text{IdealSig}(4)_k, \dots$  in  $\overline{\text{Ideal}}_0$  are sequential for the sequence  $0 < 3d < 6d < \dots$

2.  $\text{RealSig}(2)_k, \text{RealSig}(5)_k, \dots$  in  $\text{Real}_1$  and  $\text{IdealSig}(2)_k, \text{IdealSig}(5)_k, \dots$  in  $\text{Ideal}_1$  are sequential for the sequence  $d < 4d < \dots$
3.  $\text{RealSig}(3)_k, \text{RealSig}(6)_k, \dots$  in  $\text{Real}_2$  and  $\text{IdealSig}(3)_k, \text{IdealSig}(6)_k, \dots$  in  $\text{Ideal}_2$  are sequential for the sequence  $2d < 5d < \dots$

*Proof.* Follows directly from Lemma 6.

Since each ideal service  $j$  has the same lifetime as the real service  $j$ , we can apply Theorem 5 to show that  $\overline{\text{Real}}_i$  implements  $\overline{\text{Ideal}}_i$ . This is the core step in the proof of the following correctness theorem. Then, the result follows from the application of Theorem 3 to the parallel composition of  $\overline{\text{Real}}_i$  and  $\overline{\text{Dispatcher}}$ , and the parallel composition of  $\overline{\text{Ideal}}_i$  and  $\overline{\text{Dispatcher}}$ .

**Theorem 6.** *Assume the concrete time scheme described above and that every signature scheme used in the timestamping protocol is complete and existentially unforgeable. By Theorem 1, this implies  $(\overline{\text{RealSig}}(j), \emptyset) \leq_{\text{neg,pt}} (\overline{\text{IdealSig}}(j), \{\text{fail}_j\})$  for every  $j \in \text{SID}$ . Assume further that, for every pair of polynomials  $\langle p, q \rangle$ , there exists a negligible function  $\epsilon$  such that  $(\overline{\text{RealSig}}(j), \emptyset) \leq_{p,q,\epsilon} (\overline{\text{IdealSig}}(j), \{\text{fail}_j\})$  for every  $j \in \text{SID}$ . Then  $(\overline{\text{RealSigSys}}, \emptyset) \leq_{\text{neg,pt}} (\overline{\text{IdealSigSys}}, \overline{F})$ , where  $\overline{F}_k := \bigcup_{j \in \text{SID}_k} \{\{\text{fail}_j\}\}$  for every  $k$ .*

*Proof.* First we apply Theorem 5 three times to show:

1.  $(\overline{\text{Real}}_0, \emptyset) \leq_{\text{neg,pt}} (\overline{\text{Ideal}}_0, \overline{F}_0)$  where  $(\overline{F}_0)_k := \bigcup_{j \in \{1,4,\dots\}} \{\{\text{fail}_j\}\}$  for every  $k$ .
2.  $(\overline{\text{Real}}_1, \emptyset) \leq_{\text{neg,pt}} (\overline{\text{Ideal}}_1, \overline{F}_1)$  where  $(\overline{F}_1)_k := \bigcup_{j \in \{2,5,\dots\}} \{\{\text{fail}_j\}\}$  for every  $k$ .
3.  $(\overline{\text{Real}}_2, \emptyset) \leq_{\text{neg,pt}} (\overline{\text{Ideal}}_2, \overline{F}_2)$  where  $(\overline{F}_2)_k := \bigcup_{j \in \{3,6,\dots\}} \{\{\text{fail}_j\}\}$  for every  $k$ .

It is easy to see that for each  $i \in \{0, 1, 2\}$  and  $j \in \text{SID}$ ,  $\text{RealSig}_j \in \text{Real}_i$  is comparable to  $\text{IdealSig}_j \in \text{Ideal}_i$ . Observe also that compatibility conditions are also satisfied. The number of components in  $\text{Real}_{i,k}$  is bounded by the cardinality of the set  $\text{SID}_k$ . Since  $\text{SID}_k$  is the set of  $p(k)$ -bit numbers for some polynomial  $p$ , the size of  $\text{SID}_k$  is bounded by some exponential in  $k$ . We use this exponential for the  $L$  bound in Theorem 5. By Lemma 7 we know that conditions on the complexity bounds are met. By Lemma 8 we exhibit the needed sequence of positive reals for sequentiality. By Theorem 1, we have for every pair of polynomials  $p$  and  $q$ , there exists a negligible function such that  $(\text{RealSig}_j, \emptyset) \leq_{p,q,\epsilon} (\text{IdealSig}_j, \{\text{fail}_j\})$ . Hence, we can apply Theorem 5 to get 1–3 above.

Then, we apply Theorem 3 to  $\overline{\text{Dispatcher}} \parallel \overline{\text{Real}}_0 \parallel \overline{\text{Real}}_1 \parallel \overline{\text{Real}}_2$  and  $\overline{\text{Dispatcher}} \parallel \overline{\text{Ideal}}_0 \parallel \overline{\text{Ideal}}_1 \parallel \overline{\text{Ideal}}_2$ . In order to apply this theorem we first observe that  $\overline{\text{Dispatcher}}$  is comparable to  $\overline{\text{Dispatcher}}$ , and for each  $i \in \{0, 1, 2\}$  and  $j \in \text{SID}$ ,  $\overline{\text{RealSig}}_j \in \overline{\text{Real}}_i$  is comparable to  $\overline{\text{IdealSig}}_j \in \overline{\text{Ideal}}_i$ . Observe also that compatibility conditions are also satisfied.

It is also obvious that for every pair of polynomials  $p$  and  $q$ ,  $(\overline{\text{Dispatcher}}, \emptyset) \leq_{p,q,0} (\overline{\text{Dispatcher}}, \emptyset)$ , and we just showed that there are negligible functions  $\epsilon_i$  such that  $(\overline{\text{Real}}_i, \emptyset) \leq_{p,q,\epsilon_i} (\overline{\text{Ideal}}_i, \overline{F}_i)$  for each  $i \in \{0, 1, 2\}$ . The fact that each of the composed families is polynomially bounded, and that we are only considering the composition of a constant number of them (that is, 4) provides the  $r, s$  bounds and guarantees the uniformity condition (1) required for Theorem 3 (we can simply select the largest of the bounds of each individual families). Those observations are sufficient to apply Theorem 3, which yields the result we need.

## 9 Conclusion

We have introduced a new model for long-lived security protocols, based on task-PIOAs augmented with real-time task schedules. We express computational restrictions in terms of processing rates with respect to real time. The heart of our model is a long-term implementation relation,  $\leq_{\text{neg,pt}}$ , which expresses security in any polynomial-length interval of time, despite of prior security violations. We have proved polynomial parallel composition and exponential sequential composition theorems for  $\leq_{\text{neg,pt}}$ . Finally, we have applied the new theory to show security properties for a long-lived timestamping protocol.

This work suggests several directions for future work. First, for our particular timestamping case study, it would be interesting to define a higher-level “centralized” functionality specification for a long-lived timestamp service, and to use  $\leq_{\text{neg,pt}}$  to show that our ideal system, and hence, the real protocol, implements that specification.

We would also like to know whether or not it is possible to achieve stronger properties for long-lived timestamp services, such as reliably certifying a document for super-polynomial time.

It remains to use these definitions to study additional long-lived protocols and their security properties. The use of real time in the model should enable quantitative analysis of the rate of security degradation. Finally, it would be interesting to generalize the framework to allow the computational power of the various system components to increase with time.

## References

1. Ross Anderson. Two remarks on public key cryptology. Technical Report UCAM-CL-TR-549, University of Cambridge, dec 2002.
2. M. Backes, B. Pfitzmann, and M. Waidner. A general composition theorem for secure reactive systems. In M. Naor, editor, *First Theory of Cryptography Conference, TCC 2004*, volume 2951 of *LNCS*, pages 336–354, Cambridge, MA, USA, February 2004. Springer-Verlag.
3. Michael Backes, Birgit Pfitzmann, and Michael Waidner. Secure asynchronous reactive systems. Cryptology ePrint Archive, Report 2004/082, 2004. <http://eprint.iacr.org/>.
4. D. Bayer, S. Haber, and S. W. Stornetta. Improving the efficiency and reliability of digital time-stamping. In R. M. Capocalli, A. De Santis, , and U. Vaccaro, editors, *Sequences II: Methods in Communication, Security, and Computer Science*, pages 329–334. Springer-Verlag, 1993. (Proceedings of the Sequences Workshop, 1991).
5. M. Bellare and P. Rogaway. Entity authentication and key distribution. In *Proceedings of Advances in Cryptology: Crypto'93*, pages 232–249, Santa Barbara, USA, 1994. Springer-Verlag - LNCS Vol. 773.
6. Mihir Bellare and Sara K. Miner. A forward-secure digital signature scheme. In Michael J. Wiener, editor, *Advances in Cryptology - CRYPTO '99*, volume 1666 of *Lecture Notes in Computer Science*, pages 431–448. Springer, 1999.
7. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In Moni Naor, editor, *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, pages 136–145. IEEE Computer Society, 2001. Full version available on <http://eprint.iacr.org/2000/067>.
8. Ran Canetti, Ling Cheung, Dilsun Kaynar, Moses Liskov, Nancy Lynch, Olivier Pereira, and Roberto Segala. Analyzing security protocols using time-bounded Task-PIOAs. *Discrete Event Dynamic Systems*, 18(1), 2008. 49 p., to appear. (Full version available on <http://eprint.iacr.org/2005/452>).
9. Ran Canetti, Shai Halevi, and Jonathan Katz. A forward-secure public-key encryption scheme. In E. Biham, editor, *Advances in Cryptology — EUROCRYPT 2003*, number 2656 in *Lecture Notes in Computer Science*, pages 255–271. Springer, 2003.
10. O. Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, 2001 (reprint of 2003).
11. S. Goldwasser and L. Levin. Fair computation of general functions in presence of immoral majority. In Alfred J. Menezes and Scott A. Vanstone, editors, *Advances in Cryptology - Crypto '90*, pages 77–93, Berlin, 1990. Springer-Verlag. *Lecture Notes in Computer Science Volume 537*.
12. S. Haber. Long-lived digital integrity using short-lived hash functions. Technical report, HP Laboratories, May 2006.
13. S. Haber and P. Kamat. A content integrity service for long-term digital archives. In *Proceedings of the IS&T Archiving Conference*, 2006. Also published as Technical Memo HPL-2006-54, Trusted Systems Laboratory, HP Laboratories, Princeton.
14. N.A. Lynch and M.R. Tuttle. An introduction to input/output automata. *CWI Quarterly*, 2(3):219–246, September 1989.
15. M. Merritt, F. Modugno, and M.R. Tuttle. Time constrained automata. In *Proceedings of CONCUR 1991*, volume 527 of *LNCS*, pages 408–423, 1991.
16. John Mitchell, Ajith Ramanathan, Andre Scedrov, and Vanessa Teague. A probabilistic polynomial-time process calculus for the analysis of cryptographic protocols. *Theoretical Computer Science*, 353:118–164, 2006.
17. J. Müller-Quade and D. Unruh. Long-term security and universal composability. In *Theory of Cryptography, Proceedings of TCC 2007*, volume 4392 of *LNCS*, pages 41–60. Springer-Verlag, March 2007. Preprint on IACR ePrint 2006/422.
18. Rafail Ostrovsky and Moti Yung. How to withstand mobile virus attacks. In *Proceedings of 10th annual ACM Symposium on Principles of Distributed Computing (PODC-91)*, pages 51–59, August 1991.
19. B. Pfitzmann and M. Waidner. A model for asynchronous reactive systems and its application to secure message transmission. In *IEEE Symposium on Security and Privacy*, pages 184–200, Oakland, CA, May 2001. IEEE Computer Society.
20. R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2(2):250–273, 1995.

Signer( $k : \mathbb{N}, j : SID$ )

### Signature

Input:

tick( $t : \mathbb{T}_k$ )  
signKey( $sk : 2^k$ ) <sub>$j$</sub>   
reqSign( $rid : RID_k,$   
 $m : M_k$ ) <sub>$j$</sub>

Output:

respSign( $rid : RID_k,$   
 $\sigma : \Sigma_k$ ) <sub>$j$</sub>

Internal:

sign( $rid : RID_k, m : M_k$ ) <sub>$j$</sub>

### Transitions

tick( $t$ )

Effect:

if  $j \in \text{alive}(t)$  then  
   $clock := t$   
  if  $awake = \perp$  then  
     $awake := true$   
     $toSign, signed$   
     $:= \text{empty}$   
  else  
     $awake, clock, mySK,$   
     $toSign, signed := \perp$

signKey( $sk$ ) <sub>$j$</sub>

Effect:

if  $awake = true$   
   $\wedge mySK = \perp$   
  then  $mySK := sk$

reqSign( $rid, m$ ) <sub>$j$</sub>

Effect:

if  $awake = true$   
   $\wedge \neg \text{full}(toSign)$   
  then  $toSign :=$   
     $\text{enq}(toSign, \langle rid, m \rangle)$

### Tasks

respSign <sub>$j$</sub>  = {respSign(\*, \*) <sub>$j$</sub> }  
sign <sub>$j$</sub>  = {sign(\*, \*) <sub>$j$</sub> }

### States

$awake : \{true\}_\perp, \text{init } \perp$   
 $clock : (\mathbb{T}_k)_\perp, \text{init } \perp$   
 $mySK : (2^k)_\perp, \text{init } \perp$   
 $toSign : \text{que}_k(RID_k \times M_k)_\perp,$   
 $\text{init } \perp$   
 $signed : \text{que}_k(RID_k \times \Sigma_k)_\perp,$   
 $\text{init } \perp$

sign( $rid, m$ ) <sub>$j$</sub>

local  $\sigma : \Sigma$

Precondition:

$awake = true$   
 $\text{head}(toSign) = \langle rid, m \rangle$   
 $mySK \neq \perp$

Effect:

$toSign := \text{deq}(toSign)$   
 $\sigma \leftarrow \text{Sign}_j(m, mySK)$   
 $signed :=$   
   $\text{enq}(signed, \langle rid, \sigma \rangle)$

respSign( $rid, \sigma$ ) <sub>$j$</sub>

Precondition:

$awake = true$   
 $\text{head}(signed) = \langle rid, \sigma \rangle$

Effect:

$signed := \text{deq}(signed)$

**Fig. 3.** Task-PIOA Code for Signer( $k, j$ )

Verifier( $k : \mathbb{N}, j : SID$ )

### Signature

Input:

tick( $t : \mathbb{T}_k$ )  
 verKey( $vk : 2^k$ ) <sub>$j$</sub>   
 reqVer( $rid : RID_k$ ,  
 $m : M_k, \sigma : \Sigma_k$ ) <sub>$j$</sub>

Output:

respVer( $rid : RID_k$ ,  
 $b : Bool$ ) <sub>$j$</sub>

Internal:

verify( $rid : RID_k$ ,  
 $m : M_k, \sigma : \Sigma_k$ ) <sub>$j$</sub>

### Transitions

tick( $t$ )

Effect:

if  $j \in \text{alive}(t)$  then  
    $clock := t$   
   if  $awake = \perp$  then  
      $awake := true$   
      $toVer, verified$   
        $:= \text{empty}$   
 else  
    $awake, clock, myVK,$   
    $toVer, verified := \perp$

verKey( $vk$ ) <sub>$j$</sub>

Effect:

if  $awake = true$   
 $\wedge myVK = \perp$   
 then  $myVK := vk$

reqVer( $rid, m, \sigma$ ) <sub>$j$</sub>

Effect:

if  $awake = true$   
 $\wedge \neg \text{full}(toVer)$   
 then  $toVer :=$   
    $\text{enq}(toVer, \langle rid, m, \sigma \rangle)$

### Tasks

respVer <sub>$j$</sub>  = {respVer(\*, \*) <sub>$j$</sub> }

verify <sub>$j$</sub>  = {verify(\*, \*, \*) <sub>$j$</sub> }

### States

awake : {true} <sub>$\perp$</sub> , init  $\perp$

clock : ( $\mathbb{T}_k$ ) <sub>$\perp$</sub> , init  $\perp$

myVK : ( $2^k$ ) <sub>$\perp$</sub> , init  $\perp$

toVer : que <sub>$k$</sub> ( $RID_k \times M_k$   
 $\times \Sigma_k$ ) <sub>$\perp$</sub> , init  $\perp$

verified : que <sub>$k$</sub> ( $RID_k \times M_k$   
 $\times \Sigma_k$ ) <sub>$\perp$</sub> , init  $\perp$

verify( $rid, m, \sigma$ ) <sub>$j$</sub>

local  $b : Bool$

Precondition:

$awake = true$   
 $\wedge myVK \neq \perp$   
 $\text{head}(toVer) = \langle rid, m, \sigma \rangle$

Effect:

$toVer := \text{deq}(toVer)$   
 $b := \text{Verify}_j(m, \sigma, myVK)$   
 $verified :=$   
 $\text{enq}(verified, \langle rid, b \rangle)$

respVer( $rid, b$ ) <sub>$j$</sub>

Precondition:

$awake = true$   
 $\text{head}(verified) = \langle rid, b \rangle$

Effect:

$verified := \text{deq}(verified)$

**Fig. 4.** Task-PIOA Code for Verifier( $k, j$ )

SigFunc( $k : \mathbb{N}, j : SID$ )

### Signature

Input:

$I_{\text{Verifier}} \cup I_{\text{Signer}}$

Output:

$O_{\text{Verifier}} \cup O_{\text{Signer}}$

Internal:

$H_{\text{Verifier}} \cup H_{\text{Signer}} \cup \{\text{fail}_j\}$

### Transitions

Same as Signer and Verifier, except the following:

tick( $t$ )

Effect:

if  $j \in \text{alive}(t)$  then

$clock := t$

  if  $awake = \perp$  then

$awake := true$

$toSign, toVer,$

$signed, verified$

$:= \text{empty}$

$history := \emptyset$

$failed := false$

  else

$awake, clock, mySK,$

$myVK, toSign, toVer,$

$signed, history, verified,$

$failed := \perp$

fail <sub>$j$</sub>

Precondition:

$awake = true$

Effect:

$failed := true$

### Tasks

$\mathcal{R}_{\text{Signer}} \cup \mathcal{R}_{\text{Verifier}} \cup \{\{\text{fail}_j\}\}$

### States

All variables of Signer and Verifier

$history : \text{que}_k(M_k)_\perp, \text{init } \perp$   
 $failed : \{\text{true}, \text{false}\}_\perp, \text{init } \perp$

sign( $rid, m$ ) <sub>$j$</sub>

local  $\sigma : \Sigma$

Precondition:

$awake = true$

$\wedge mySK \neq \perp$

$\text{head}(toSign) = \langle rid, m \rangle$

Effect:

$toSign := \text{deq}(toSign)$

$\sigma := \text{Sign}_j(m, mySK)$

$signed :=$

$\text{enq}(signed, \langle rid, \sigma \rangle)$

$history :=$

$\text{enq}(history, m)$

verify( $rid, m, \sigma$ ) <sub>$j$</sub>

Local  $b : Bool$

Precondition:

$awake = true$

$\wedge myVK \neq \perp$

$\text{head}(toVer) = \langle rid, m, \sigma \rangle$

Effect:

$toVer := \text{deq}(toVer)$

$b := (\text{Verify}(m, \sigma, myVK))$

$\wedge (m \in history \vee failed)$

$verified :=$

$\text{enq}(verified, \langle rid, b \rangle)$

**Fig. 5.** Code for SigFunc( $k, j$ )

Dispatcher( $k : \mathbb{N}$ )

### Signature

Input:

$tick(t : \mathbb{T}_k)$   
 $reqCert(rid : RID_k, x : X_k)$   
 $reqCert(rid : RID_k, x : X_k, t : \mathbb{T}_k,$   
 $\sigma_1 : (\Sigma_k)_\perp, \sigma_2 : \Sigma_k, j : SID)$   
 $reqCheck(rid : RID_k, x : X_k, t : \mathbb{T}_k,$   
 $\sigma_1 : (\Sigma_k)_\perp, \sigma_2 : \Sigma_k, j : SID)$   
 $respSign(rid : RID_k, \sigma : \Sigma_k)_j, j \in SID$   
 $respVer(rid : RID_k, b : Bool)_j, j \in SID$

Output:

$reqSign(rid : RID_k, m : M_k)_j, j \in SID$   
 $reqVer(rid : RID_k, m : M_k, \sigma : \Sigma_k)_j, j \in SID$   
 $respCert(rid : RID_k, \sigma : \Sigma_k, j : SID)$   
 $respCert(rid : RID_k, false)$   
 $respCheck(rid : RID_k, b : Bool)$

Internal:

$denyVer(rid : RID_k, op : \{'cert', 'check'\},$   
 $m : M_k, \sigma : \Sigma_k, j : SID)$

### Transitions

$tick(t)$

Effect:

$clock := t$

$reqCert(rid, x)$

Effect:

if  $currCt < b$  then  
 $toSign := enq(toSign, \langle rid, \langle x, clock, \perp \rangle \rangle)$   
 $currCt := currCt + 1$

$reqCert(rid, x, t, \sigma_1, \sigma_2, j)$

Effect:

if  $currCt < b$  then  
 $toVer := enq(toVer,$   
 $\langle rid, 'cert', \langle x, t, \sigma_1 \rangle, \sigma_2, j \rangle)$   
 $currCt := currCt + 1$

$reqCheck(rid, x, t, \sigma_1, \sigma_2, j)$

Effect:

if  $currCt < b$  then  
 $toVer := enq(toVer,$   
 $\langle rid, 'check', \langle x, t, \sigma_1 \rangle, \sigma_2, j \rangle)$   
 $currCt := currCt + 1$

$reqSign(rid, m)_j$

Precondition:

$head(toSign) = \langle rid, m \rangle$   
 $j = pref(clock)$   
 $\neg pendingSign$

Effect:

$pendingSign := true$

### Tasks

$reqSign = \{reqSign(*, *)_*\}$   
 $reqVer = \{reqVer(*, *, *)_*\}$   
 $respCert = \{respCert(*, *, *)\} \cup \{respCert(*, false)\}$   
 $respCheck = \{respCheck(*, *)\}$   
 $denyVer = \{denyVer(*, *, *, *, *)\}$

### States

$clock : \mathbb{T}_k, \text{init } 0$   
 $toSign : que_k(RID_k \times M), \text{init empty}$   
 $toVer : que_k(RID_k \times \{'cert', 'check'\}$   
 $\times M \times \Sigma \times SID), \text{init empty}$   
 $pendingVer, pendingSign : Bool, \text{init false}$   
 $certified : que_k((RID_k \times \Sigma \times SID)$   
 $\cup (RID_k \times \{false\})), \text{init empty}$   
 $checked : que_k(RID_k \times Bool), \text{init empty}$   
 $currCt : \mathbb{N}, \text{init } 0$

$respSign(rid, \sigma_3)_j$

Effect:

if  $pendingSign \wedge (\exists m)(head(toSign) =$   
 $\langle rid, m, j \rangle)$  then  
 choose  $m$  where  $head(toSign) = \langle rid, m, j \rangle$   
 $toSign := deq(toSign)$   
 $pendingSign := false$   
 choose  $x, t$  where  $(\exists \sigma_2)(m = \langle x, t, \sigma_2 \rangle)$   
 $certified := enq(certified, \langle rid, \sigma_3, j \rangle)$

$denyVer(rid, op, m, \sigma_2, j)$

Precondition:

$head(toVer) = \langle rid, op, m, \sigma_2, j \rangle$   
 $j \notin usable(clock)$

Effect:

$toVer := deq(toVer)$   
 if  $op = 'cert'$  then  
 $certified := enq(certified, \langle rid, false \rangle)$   
 else  $checked := enq(checked, \langle rid, false \rangle)$

$reqVer(rid, m, \sigma_2)_j$

Precondition:

$(\exists op)(head(toVer) = \langle rid, op, m, \sigma_2, j \rangle)$   
 $j \in usable(clock)$   
 $\neg pendingVer$

Effect:

$pendingVer := true$

**Fig. 6.** Task-PIOA Code for Dispatcher( $k : \mathbb{N}$ ), Part I

### Transitions

$\text{respVer}(rid, b)_j$

Effect:

```
if  $pendingVer$ 
   $\wedge (\exists op, m, \sigma_2)(head(toVer) =$ 
     $\langle rid, op, m, \sigma_2, j \rangle)$  then
  choose  $op, m, \sigma_2$  where
     $head(toVer) = \langle rid, op, m, \sigma_2, j \rangle$ 
   $toVer := \text{deq}(toVer)$ 
   $pendingVer := false$ 
  if  $op = 'cert' \wedge \neg b$  then
     $certified := \text{enq}(certified, \langle rid, false \rangle)$ 
  if  $op = 'cert' \wedge b$  then
    choose  $x, t$  where  $(\exists \sigma_1)(m = \langle x, t, \sigma_1 \rangle)$ 
     $toSign := \text{enq}(toSign, \langle rid, \langle x, t, \sigma_2 \rangle \rangle)$ 
  if  $op = 'check'$  then
     $checked := \text{enq}(checked, \langle rid, b \rangle)$ 
```

$\text{respCert}(rid, false)$

Precondition:

$head(certified) = \langle rid, false \rangle$

Effect:

```
 $certified := \text{deq}(certified)$ 
 $currCt := currCt - 1$ 
```

$\text{respCert}(rid, \sigma_3, j)$

Precondition:

$head(certified) = \langle rid, \sigma_3, j \rangle$

Effect:

```
 $certified := \text{deq}(certified)$ 
 $currCt := currCt - 1$ 
```

$\text{respCheck}(rid, b)$

Precondition:

$head(checked) = \langle rid, b \rangle$

Effect:

```
 $checked := \text{deq}(checked)$ 
 $currCt := currCt - 1$ 
```

**Fig. 7.** Task-PIOA Code for Dispatcher( $k : \mathbb{N}$ ), Part II