Computer Science and Artificial Intelligence Laboratory

Technical Report

MIT-CSAIL-TR-2008-068 November 22, 2008

Modeling Computational Security in Long-Lived Systems, Version 2

Ran Canetti, Ling Cheung, Dilsun Kaynar, Nancy Lynch, and Olivier Pereira

massachusetts institute of technology, cambridge, ma 02139 usa — www.csail.mit.edu

Modeling Computational Security in Long-Lived Systems, Version 2

Ran Canetti[1,2], Ling Cheung[2], Dilsun Kaynar[3], Nancy Lynch[2], and Olivier Pereira[4]

[1] IBM T. J. Watson Research Center
[2] Massachusetts Institute of Technology
[3] Carnegie Mellon University
[4] 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 [13, 7, 21, 9, 12, 5, 4, 10]).

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* [20]. 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* [3, 8], 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* [3, 11] is defined analogously. Yet another task of the same flavor is timestamping [6, 14, 15]. 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 work: A first suggestion for an approach might be to use existing models, such as the PPT calculus [18], the Reactive Simulatability [5], or the Universally Composable security frameworks [9], 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 [19]. 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 [19] 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 probabilitic automata model, the Task-PIOA model [10], 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*, $\leq_{\mathsf{neg,pt}}$, 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 no failure steps occur starting from some time t, then the ideal system starts following the specified ideal behavior.

Our long-term implementation relation $\leq_{\mathsf{neg,pt}}$ 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 failures 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 $\leq_{\mathsf{neg,pt}}$ 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 $\leq_{\mathsf{neg,pt}}$, in verifying protocols using several levels of abstraction.

We show that $\leq_{\mathsf{neg,pt}}$ 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. [6, 14, 15], 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 amout 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.

Note: This paper is a revision of a previous paper that appeared in Concur 2008 [1], and its full technical report version [2]. The current paper uses a new, stronger definition of the long-lived implementation relation, based on conditional probability.

2 Task-PIOAs

We build our new framework using task-PIOAs [10], which are a version of Probabilistic Automata [22], 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.

Notation: Given a set S, let $\mathsf{Disc}(S)$ 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* $\mathsf{type}(v)$, which is the set of all possible values of v. We assume that $\mathsf{type}(v)$ 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 $\mathsf{type}(v)$. The set of all valuations for V is denoted $val(V)$. 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 $S_\perp := S \cup \{\perp\}$.

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

(i) V is a set of *state variables* and $S \subseteq val(V)$ is a set of *states*;
(ii) $s^{\mathsf{init}} \in S$ 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) $\Delta \subseteq S \times (I \cup O \cup H) \times \mathsf{Disc}(S)$ is a *transition relation*.

The set $Act := I \cup O \cup H$ 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 $\alpha = q_0 a_1 q_1 a_2 \ldots$ of alternating states and actions where $q_0 = s^{\mathsf{init}}$ 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 $q_{i+1} \in \mathsf{Support}(\mu)$. 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 μ. We require that $\mathcal{A}$ satisfy 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 $\mu \in \mathsf{Disc}(S)$ with $\langle s, a, \mu \rangle \in \Delta$. We write $\Delta(s, a)$ for such μ, 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 $V_i \cap V_j = Act_i \cap H_j = O_i \cap O_j = \emptyset$ whenever $i \neq j$. In that case, we define their *composition* $\mathcal{A}_1 \| \mathcal{A}_2$ to be $\langle V_1 \cup V_2, S_1 \times S_2, \langle s_1^{\mathsf{init}}, s_2^{\mathsf{init}} \rangle, (I_1 \cup I_2) \setminus (O_1 \cup O_2), O_1 \cup O_2, H_1 \cup H_2, \Delta \rangle$, where Δ 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 [16, 10]. 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 $\mathsf{Clock}(\mathbb{T})$, which has a $\mathsf{tick}(t)$ output action for every t in some discrete time domain $\mathbb{T}$. For concreteness, we assume that $\mathbb{T} = \mathbb{N}$, and write simply Clock. Clock has a single task tick, consisting of all $\mathsf{tick}(t)$ actions. These clock ticks are produced in order, for $t = 1, 2, \ldots$. In Section 3, we will define a mechanism that will ensure that each $\mathsf{tick}(t)$ occurs exactly at real time t.

$\mathsf{Clock}(\mathbb{T})$

Signature

Input:
none
Output:
$\mathsf{tick}(t : \mathbb{T}), t > 0$

Transitions

$\mathsf{tick}(t)$
Precondition:
$count = t - 1$
Effect:
$count := t$

Tasks

$\mathsf{tick} = \{\mathsf{tick}(*)\}$

States

$count \in \mathbb{T}$, initially 0

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

Operations: Given compatible task-PIOAs $\mathcal{A}_1$ and $\mathcal{A}_2$, we define their *composition* to be $\langle \mathcal{A}_1 \| \mathcal{A}_2, \mathcal{R}_1 \cup \mathcal{R}_2 \rangle$. 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 $\mathcal{A} = \langle V, S, s^{\mathsf{init}}, I, O, H, \Delta \rangle$ and $S \subseteq O$, $\mathsf{hide}(\mathcal{A}, S)$ is the task-PIOA given by $\langle V, S, s^{\mathsf{init}}, I, O', H', \Delta \rangle$, 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 $\rho = T_1, T_2, \ldots$ of tasks in $\mathcal{R}$. This induces a well-defined run of $\mathcal{A}$ as follows.

(i) From the start state s^{init}, we *apply* the first task T_1: due to action- and transition-determinism, T_1 specifies at most one transition from s^{init}; if such a transition exists, it is taken, otherwise nothing happens.

(ii) Repeat with remaining T_i's.

Such a run gives rise to a unique *probabilistic execution*, which is a probability distribution over executions in $\mathcal{A}$. We denote this probabilistic execution by $\mathsf{Execs}(\mathcal{A}, \rho)$.

For finite ρ, let $\mathsf{lstate}(\mathcal{A}, \rho)$ denote the state distribution of $\mathcal{A}$ after executing according to ρ. A state s is said to be *reachable* under ρ if $\mathsf{lstate}(\mathcal{A}, \rho)(s) > 0$. Moreover, the probabilistic execution induces a unique *trace distribution* $\mathsf{tdist}(\mathcal{A}, \rho)$, which is a probability distribution over the set of traces of $\mathcal{A}$. We refer to [10] 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 τ for a closed task-PIOA $\langle \mathcal{A}, \mathcal{R} \rangle$ to be a finite or infinite sequence $\langle T_1, t_1 \rangle, \langle T_2, t_2 \rangle, \ldots$ such that: $T_i \in \mathcal{R}$ and $t_i \in \mathbb{R}_{\geq 0}$ for every i, and $t_1, t_2, \ldots$ is non-decreasing. Given a timed task schedule $\tau = \langle T_1, t_1 \rangle, \langle T_2, t_2 \rangle, \ldots$ and $t \in \mathbb{R}_{\geq 0}$, let $\mathsf{trunc}_{\geq t}(\tau)$ denote the result of removing all pairs $\langle T_i, t_i \rangle$ with $t_i \geq t$.

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

- If τ is empty, then $\mathsf{ltime}(\tau) := 0$.
- If $t_1, t_2, \ldots$ is bounded, then $\mathsf{ltime}(\tau) := \lim_{i \to \infty} t_i$, otherwise $\mathsf{ltime}(\tau) := \infty$.

Following [17], 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 $\langle \mathsf{rate}, \mathsf{burst}, \mathsf{lb}, \mathsf{ub} \rangle$ such that: (i) $\mathsf{rate}, \mathsf{burst}, \mathsf{lb} : \mathcal{R} \to \mathbb{R}_{\geq 0}$, (ii) $\mathsf{ub} : \mathcal{R} \to \mathbb{R}^{\infty}_{>0}$, and (iii) for all $T \in \mathcal{R}$, $\mathsf{lb}(T) \leq \mathsf{ub}(T)$. To ensure that rate and ub can be satisfied simultaneously, we require $\mathsf{rate}(T) \geq 1/\mathsf{ub}(T)$ whenever $\mathsf{rate}(T) \neq 0$ and $\mathsf{ub}(T) \neq \infty$. From this point on, we assume that every task-PIOA is associated with a particular bound map.

Given a timed schedule τ and a task T, let $\mathsf{proj}_T(\tau)$ 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 τ is *valid* for the interval I (under a bound map $\langle \mathsf{rate}, \mathsf{burst}, \mathsf{lb}, \mathsf{ub} \rangle$) if the following hold for every task T.

(i) If the pair $\langle T, t \rangle$ appears in τ, then $t \in I$.

(ii) If $\mathsf{lb}(T) > 0$, then: (a) if $\langle T, t \rangle$ is the first element of $\mathsf{proj}_T(\tau)$, then $t \geq \mathsf{lb}(T)$; (b) for every interval I' of a non-negative real length less than $\mathsf{lb}(T)$, $\mathsf{proj}_T(\tau)$ contains at most one element $\langle T, t \rangle$ with $t \in I'$.

(iii) If $\mathsf{ub}(T) \neq \infty$, then, for every interval $I' \subseteq I$ of a non-negative real length greater than $\mathsf{ub}(T)$, $\mathsf{proj}_T(\tau)$ 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, $\mathsf{proj}_T(\tau)$ contains at most $\mathsf{rate}(T) \cdot d + \mathsf{burst}(T)$ elements $\langle T, t \rangle$ with $t \in I'$.

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

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

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

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

This is well defined since the task partition of $\mathcal{A}_1 \| \mathcal{A}_2$ 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 $\mathsf{lb}(\mathsf{tick}) = \mathsf{ub}(\mathsf{tick}) = 1$. 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 $\langle \mathsf{tick}, 1\rangle, \langle \mathsf{tick}, 2\rangle, \ldots$ 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 $v \in V$, $\mathsf{type}(v) \subseteq \{0,1\}^p$.
(ii) The name of every action, task, and variable of $\mathcal{A}$ has length at most p.
(iii) The initial state s^{init} is p-bounded.
(iv) There exists a deterministic Turing machine M_{enable} satisfying: for every p-bounded state s, M_{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_{enable} is at most p bits long, and M_{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 τ and every state s reachable under τ, there are at most p variables v such that $s.v \neq \perp$.*

Proof. By the definition of step bounds, we have s^{init} is p-bounded. For a state s' reachable under schedule τ', let s be a state immediately preceding s' in the probabilistic execution induced by τ'. Thus s is reachable under some prefix of

τ. 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. □

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 [10, Lemma 4.2]. We also note that the hiding operator introduced in Section 2 preserves step bounds.

Lemma 2. *Suppose $\{\mathcal{A}_i | 1 \leq i \leq b\}$ is a compatible set of task-PIOAs, where each A_i has step bound $p_i \in \mathbb{N}$. The composition $\|_{i=1}^{b} \mathcal{A}_i$ has step bound $c_{\mathsf{comp}} \cdot \sum_{i=1}^{b} p_i$, where c_{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 $\mathcal{A}$, $\mathsf{rate}(T)$ and $\mathsf{burst}(T)$ are both at most p.
(iii) For every $t \in \mathbb{N}$, let S_t denote the set of states s of $\mathcal{A} \| \mathsf{Clock}$ such that s is reachable under some valid schedule τ and $s.count = t$. There are at most p tasks T such that T is enabled in some $s \in S_t$. (Here, $s.count$ is the value of variable $count$ of Clock in state s).

We say that $\mathcal{A}$ is *quasi-p-bounded* if $\mathcal{A}$ is of the form $\mathcal{A}' \| \mathsf{Clock}$ where $\mathcal{A}'$ is p-bounded.

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 $p : \mathbb{N} \to \mathbb{N}$, 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 σ 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 σ is said to be *valid* for m and vk if $\mathsf{Verify}(m, \sigma, vk) = 1$.

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, $\mathsf{KeyGen}(k, j)$, $\mathsf{Signer}(k, j)$, and $\mathsf{Verifier}(k, j)$ for every pair $\langle k, j \rangle$, representing the key generator, signer, and verifier, respectively. We assume a function $\mathsf{alive} : \mathbb{T} \to 2^{SID}$ such that, for every t, $\mathsf{alive}(t)$ is the set of services alive at discrete time t. The lifetime of each service j is then given by $\mathsf{aliveTimes}(j) := \{t \in \mathbb{T} | j \in \mathsf{alive}(t)\}$; we assume this to be a finite set of consecutive numbers.

For every value k of the security parameter, we assume the following finite domains: RID_k (request identifiers), M_k (messages to be signed) and Σ_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 $\mathsf{KeyGen}(k, j)$, $\mathsf{Signer}(k, j)$, and $\mathsf{Verifier}(k, j)$ has a set of input actions $\mathsf{tick}(t)$, $t \in \mathbb{T}_k$, 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 $clock$. Since $clock$ can produce $\mathsf{tick}(t)$ outputs for arbitrary $t \in \mathbb{T}$, this means that these new components do not receive all of $clock$'s inputs, but only those with $t \in \mathbb{T}_k$.

$\mathsf{KeyGen}(k, j)$ 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 $\mathsf{signKey}(sk)_j$ and $\mathsf{verKey}(vk)_j$. The signing key goes to $\mathsf{Signer}(k, j)$, while the verification key goes to $\mathsf{Verifier}(k, j)$.

The code for $\mathsf{KeyGen}(k, j)$ is given in Figure 2. As we mentioned before, the $\mathsf{tick}(t)$ 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 $\mathsf{chooseKeys}$ action uses KeyGen_j to choose the key pair, and is enabled only when j is awake and the keys are currently $\perp$. Note that the KeyGen algorithm is indexed by j, because different services may use different algorithms. The same applies to Sign_j in $\mathsf{Signer}(k, j)$ and Verify_j in $\mathsf{Verifier}(k, j)$. The $\mathsf{signKey}$ and verKey actions output the keys, and they are enabled only when j is awake and the keys have been chosen.

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

Signature

Input:
$\mathsf{tick}(t : \mathbb{T}_k)$
Output:
$\mathsf{signKey}(sk : 2^k)_j$
$\mathsf{verKey}(vk : 2^k)_j$
Internal:
$\mathsf{chooseKeys}_j$

Tasks

$\mathsf{verKey}_j = \{\mathsf{verKey}(*)_j\}$
$\mathsf{signKey}_j = \{\mathsf{signKey}(*)_j\}$
$\mathsf{chooseKeys}_j = \{\mathsf{chooseKeys}_j\}$

States

$awake : \{\mathsf{true}\}_\perp$, init $\perp$
$clock : (\mathbb{T}_k)_\perp$, init $\perp$
$mySK : (2^k)_\perp$, init $\perp$
$myVK : (2^k)_\perp$, init $\perp$

Transitions

$\mathsf{tick}(t)$
Effect:
  if $j \in \mathsf{alive}(t)$ then
    $clock := t$
    if $awake = \perp$ then
      $awake := true$
  else
    $awake, clock, mySK, myVK := \perp$

$\mathsf{chooseKeys}_j$
Precondition:
  $awake = true$
  $mySK = myVK = \perp$
Effect:
  $\langle mySK, myVK \rangle \leftarrow \mathsf{KeyGen}_j(1^k)$

$\mathsf{signKey}(sk)_j$
Precondition:
  $awake = true$
  $sk = mySK \neq \perp$
Effect:
  none

$\mathsf{verKey}(vk)_j$
Precondition:
  $awake = true$
  $vk = myVK \neq \perp$
Effect:
  none

Fig. 2. Task-PIOA Code for $\mathsf{KeyGen}(k, j)$

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

The data type 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 $\mathsf{Signer}(k, j)$ has a bounded amount of memory. For concreteness, we assume here that p is the constant function $\underline{1}$ for the queues $toSign$ and $signed$.

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

Assuming the algorithms KeyGen_j, Sign_j and Verify_j are polynomial time, it not hard to check that the composite $\mathsf{KeyGen}(k, j) \| \mathsf{Signer}(k, j) \| \mathsf{Verifier}(k, j)$ has step bound $p(k)$ for some polynomial p. If $\mathsf{rate}(T)$ and $\mathsf{burst}(T)$ are at

most $p(k)$ for every T, then the composite is $p(k)$-bounded. The family $\{\mathsf{KeyGen}(k,j)\|\mathsf{Signer}(k,j)\|\mathsf{Verifier}(k,j)\}_{k\in\mathbb{N}}$ 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 acc and task $\mathsf{acc} = \{\mathsf{acc}\}$. Let τ be a timed task schedule for $\mathcal{A}$. The *acceptance probability* of $\mathcal{A}$ under τ is: $\mathbf{P}_{\mathsf{acc}}(\mathcal{A},\tau) := \Pr[\beta \text{ contains } \mathsf{acc} : \beta \leftarrow_{\mathsf{R}} \mathsf{tdist}(\mathcal{A},\tau)]$; that is, the probability that a trace drawn from the distribution $\mathsf{tdist}(\mathcal{A},\tau)$ contains the action acc. If $\mathcal{A}$ is not necessarily closed, we include a closing environment. A task-PIOA Env is an *environment* for $\mathcal{A}$ if it is compatible with $\mathcal{A}$ and $\mathcal{A}\|\mathsf{Env}$ is closed. From here on, we assume that every environment has output action 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 [14, 15] 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 A^2. We assume that each task in F^1 and F^2 has upper bound ∞.

Given $t \in \mathbb{R}_{\geq 0}$ and an environment Env for both $\mathcal{A}^1$ and $\mathcal{A}^2$, we consider two experiments. In the first experiment, Env interacts with $\mathcal{A}^1$ according to some valid task schedule τ_1 of $\mathcal{A}^1\|\mathsf{Env}$, where τ_1 does not contain any tasks from F^1 from time t onwards. In the second experiment, Env interacts with $\mathcal{A}^2$ according to some valid task schedule τ_2 of $\mathcal{A}^2\|\mathsf{Env}$, where τ_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) from time t onwards, then it simulates $\mathcal{A}^2$, also acting ideally from time t onwards.

More specifically, we require that, for any valid τ_1, there exists a valid τ_2 as above such that the two executions are identical before time t from the point of view of the environment. That is, the probabilistic execution is the same before time t. Moreover, for any particular execution α before time t, the continuations of α in the two experiments are *computationally indistinguishable*, namely, the difference in acceptance probabilities is negligible provided Env is computationally bounded.

If τ is a schedule of $\mathcal{A}\|\mathcal{B}$, then we define $\mathsf{proj}_{\mathcal{B}}(\tau)$ to be the result of removing all $\langle T_i, t_i\rangle$ where T_i is *not* a task of $\mathcal{B}$.

Let $\mathsf{Execs}_{\mathcal{B}}(\mathcal{A}\|\mathcal{B},\tau)$ denote the distribution of executions of $\mathcal{B}$ when $\mathcal{A}\|\mathcal{B}$ executes under schedule τ. Define $\mathbf{P}_{\mathsf{acc}}(\mathcal{A}\|\mathcal{B},\tau,t,\alpha)$, where α is in the support of $\mathsf{Execs}_{\mathcal{B}}(\mathcal{A}\|\mathcal{B},\mathsf{trunc}_{\geq t}(\tau))$, to be the probability of acceptance for an execution chosen randomly from those in the support of $\mathsf{Execs}(\mathcal{A}\|\mathcal{B},\tau)$ whose portion before time t projects on $\mathcal{B}$ to yield α. More precisely, this is the conditional probability that an execution α' contains acc, where α' is chosen from the distribution $\mathsf{Execs}(\mathcal{A}\|\mathcal{B},\tau)$ conditioned on the event $\alpha'' \upharpoonright \mathcal{B} = \alpha$, where α'' is the prefix of α' derived from $\mathsf{trunc}_{\geq t}(\tau)$.

Definition 1. *Let $\mathcal{A}^1$ and $\mathcal{A}^2$ be comparable task-PIOAs that are both compatible with* Clock*. Let F^1 and F^2 be sets of tasks of, respectively, $\mathcal{A}^1$ and $\mathcal{A}^2$, such that for any $T \in (F^1 \cup F^2)$, $\mathsf{ub}(T) = \infty$. 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 quasi-$p$-bounded environment Env*, and every valid timed schedule τ_1 for $\mathcal{A}^1 \| \mathsf{Env}$ 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 τ_2 for $\mathcal{A}^2 \| \mathsf{Env}$ for the interval $[0, t+q]$ such that:*

(i) $\mathsf{proj}_{\mathsf{Env}}(\tau_1) = \mathsf{proj}_{\mathsf{Env}}(\tau_2)$*;*
(ii) *τ_2 does not contain any pairs of the form $\langle T_i, t_i \rangle$ where $T_i \in F^2$ and $t_i \geq t$;*
(iii) $\mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_1)) = \mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^2 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_2))$*; and*
(iv) *For every α in the support of $\mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_1))$,*

$$|\, \mathbf{P}_{\mathsf{acc}}(\mathcal{A}^1 \| \mathsf{Env}, \tau_1, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\mathcal{A}^2 \| \mathsf{Env}, \tau_2, t, \alpha)| \leq \epsilon.$$

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, F^2, and F^3 be sets of tasks of $\mathcal{A}^1$, $\mathcal{A}^2$, and $\mathcal{A}^3$, respectively, such that for any $T \in F^1 \cup F^2 \cup F^3$, $\mathsf{ub}(T) = \infty$. 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 $t \in \mathbb{R}_{\geq 0}$, Env a quasi-$p$-bounded environment, and a valid timed schedule τ_1 for $\mathcal{A}^1 \| \mathsf{Env}$ for the interval $[0, t+q]$ be given, where τ_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 τ_2 for $\mathcal{A}^2 \| \mathsf{Env}$ according to the assumption $(\mathcal{A}^1, F^1) \leq_{p,q,\epsilon_1} (\mathcal{A}^2, F^2)$. Using τ_2, choose τ_3 for $\mathcal{A}^3 \| \mathsf{Env}$ according to the assumption $(\mathcal{A}^2, F^2) \leq_{p,q,\epsilon_2} (\mathcal{A}^3, F^3)$.

Clearly, we have

(i) $\mathsf{proj}_{\mathsf{Env}}(\tau_1) = \mathsf{proj}_{\mathsf{Env}}(\tau_2) = \mathsf{proj}_{\mathsf{Env}}(\tau_3)$;
(ii) τ_3 does not contain any pairs of the form $\langle T_i, t_i \rangle$ where $T_i \in F^3$ and $t_i \geq t$; and
(iii) $\mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_1)) = \mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^2 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_2)) = \mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^3 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_3))$.

For Condition (iv), let α be any execution in the support of $\mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}\, \tau_1)$. Then

$$\begin{aligned}
&|\, \mathbf{P}_{\mathsf{acc}}(\mathcal{A}^1 \| \mathsf{Env}, \tau_1, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\mathcal{A}^3 \| \mathsf{Env}, \tau_3, t, \alpha)| \\
&\leq |\, \mathbf{P}_{\mathsf{acc}}(\mathcal{A}^1 \| \mathsf{Env}, \tau_1, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\mathcal{A}^2 \| \mathsf{Env}, \tau_2, t, \alpha)| \\
&\quad + |\, \mathbf{P}_{\mathsf{acc}}(\mathcal{A}^2 \| \mathsf{Env}, \tau_2, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\mathcal{A}^3 \| \mathsf{Env}, \tau_3, t, \alpha)| \\
&\leq \epsilon_1 + \epsilon_2.
\end{aligned}$$

□

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 $\bar{F}^1$ and $\bar{F}^2$ be families of sets of tasks such that each $(\bar{F}^1)_k$ is a set of tasks of $(\bar{\mathcal{A}}^1)_k$ and each $(\bar{F}^2)_k$ is a set of tasks of $(\bar{\mathcal{A}}^2)_k$, such that each task in each of those sets has upper bound ∞. Let $\epsilon : \mathbb{N} \to \mathbb{R}_{\geq 0}$ and $p, q : \mathbb{N} \to \mathbb{N}$ be given. Then 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 $\leq_{\mathsf{neg,pt}}$. Formally, a function $\epsilon : \mathbb{N} \to \mathbb{R}_{\geq 0}$ 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, ϵ 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 $\bar{F}^2$ of $\bar{\mathcal{A}}^1$ and $\bar{\mathcal{A}}^2$, respectively, we say that $(\bar{\mathcal{A}}^1, \bar{F}^1) \leq_{\mathsf{neg,pt}} (\bar{\mathcal{A}}^2, \bar{F}^2)$ 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 ϵ is a negligible function.

Lemma 4 (Transitivity of $\leq_{\mathsf{neg,pt}}$). *Let $\bar{\mathcal{A}}^1$, $\bar{\mathcal{A}}^2$, and $\bar{\mathcal{A}}^3$ be comparable task-PIOA families. Let $\bar{F}^1$, $\bar{F}^2$, and $\bar{F}^3$ be task set families of $\bar{\mathcal{A}}^1$, $\bar{\mathcal{A}}^2$, and $\bar{\mathcal{A}}^3$, respectively (satisfying the upper bound condition). Suppose $(\bar{\mathcal{A}}^1, \bar{F}^1) \leq_{\mathsf{neg,pt}} (\bar{\mathcal{A}}^2, \bar{F}^2)$ and $(\bar{\mathcal{A}}^2, \bar{F}^2) \leq_{\mathsf{neg,pt}} (\bar{\mathcal{A}}^3, \bar{F}^3)$. Then $(\bar{\mathcal{A}}^1, \bar{F}^1) \leq_{\mathsf{neg,pt}} (\bar{\mathcal{A}}^3, \bar{F}^3)$.*

Proof. Given polynomials p and q, choose negligible functions ϵ_1 and ϵ_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* SigFunc, and show that it is implemented, in the sense of our $\leq_{\mathsf{neg,pt}}$ definition, by the real signature service of Section 4.

As with KeyGen, Signer, and Verifier, each instance of SigFunc is parameterized with a security parameter k and an identifier j. The code for $\mathsf{SigFunc}(k, j)$ appears in Figure 5. It is very similar to the composition of $\mathsf{Signer}(k, j)$ and $\mathsf{Verifier}(k, j)$. The important difference is that $\mathsf{SigFunc}(k, j)$ maintains an additional variable $history$, which records the set of signed messages. In addition, $\mathsf{SigFunc}(k, j)$ has an internal action fail_j, which sets a boolean flag $failed$. If $failed = \mathsf{false}$, then $\mathsf{SigFunc}(k, j)$ uses $history$ to answer verification requests: a signature is rejected if the submitted message is not in $history$, even if Verify_j returns 1. If $failed = \mathsf{true}$, then $\mathsf{SigFunc}(k, j)$ 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, $\mathsf{rate}(T)$ and $\mathsf{burst}(T)$ are bounded by $p(k)$ for some polynomial p. We assume that the same bound applies to $\mathsf{SigFunc}(k, j)$. Since $\mathsf{aliveTimes}(j)$ 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 $\mathsf{SigFunc}(k, j)$, 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* $\mathsf{Sig} = \langle \mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify} \rangle$ *is* complete *if* $\mathsf{Verify}(m, \sigma, vk) = 1$ *whenever* $\langle sk, vk \rangle \leftarrow \mathsf{KeyGen}(1^k)$ *and* $\sigma \leftarrow \mathsf{Sign}(sk, m)$. *We say that* 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* KeyGen *to obtain* $\langle sk, vk \rangle$ *and gives the forger* vk.

Query *The forger submits message* m. *The challenger responds with signature* $\sigma \leftarrow \mathsf{Sign}(m, sk)$. *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* $\mathsf{Verify}(m^*, \sigma^*, vk) = 1$.

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

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

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

To prove Theorem 1, we show that, for every time point t, the environment cannot distinguish $\mathsf{RealSig}(j)_k$ from $\mathsf{IdealSig}(j)_k$ with high probability between time t and $t + q(k)$, where q is a polynomial. This holds even when the task $\{fail_j\}$ 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 $\mathsf{SigFunc}(k, j)$ 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 quasi-$p(k)$-bounded environments.) Applying the definition of EUF-CMA security, we obtain the desired negligible bound on distinguishing probability.

Proof. Unwinding the definition of $\leq_{\mathsf{neg,pt}}$ using the given failure sets, we need to show the following: For every pair of polynomials p and q, there is a negligible function ϵ such that, for every $k \in \mathbb{N}$, $t \in \mathbb{R}_{\geq 0}$, quasi-$p(k)$-bounded environment Env for $\mathsf{RealSig}(j)_k$, and valid schedule τ_1 for $\mathsf{RealSig}(j)_k \| \mathsf{Env}$ for the interval $[0, t + q(k)]$, there is a valid schedule τ_2 for $\mathsf{IdealSig}(j)_k \| \mathsf{Env}$ such that:

(i) $\mathsf{proj}_{\mathsf{Env}}(\tau_1) = \mathsf{proj}_{\mathsf{Env}}(\tau_2)$;

(ii) τ_2 does not contain any pairs of the form $\langle \mathsf{fail}_j, t_i \rangle$ where $t_i \geq t$;

(iii) $\mathsf{Execs}_{\mathsf{Env}}(\mathsf{RealSig}(j)_k \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_1)) = \mathsf{Execs}_{\mathsf{Env}}(\mathsf{IdealSig}(j)_k \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_2))$; and
(iv) For every α in the support of $\mathsf{Execs}_{\mathsf{Env}}(\mathsf{RealSig}(j)_k \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_1))$,
$|\,\mathbf{P}_{\mathsf{acc}}(\mathsf{RealSig}(j)_k \| \mathsf{Env}, \tau_1, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\mathsf{IdealSig}(j)_k \| \mathsf{Env}, \tau_2, t, \alpha)| \leq \epsilon(k)$.

Since Sig is complete, we observe that the difference between the acceptance probabilities of the two automata compared in Condition (iv) can only be non-zero if Env succeeds in producing a forged signature (that is, a valid signature for a message that was not previously signed by the Sign or SigFunc automata) and in having this signature rejected when the verify and respVer actions of SigFunc execute.

Fix polynomials p and q. We must obtain a negligible ϵ bound that satisfies the four conditions above for every k, t, Env, and valid τ_1, for some corresponding τ_2. To define ϵ, we rely on the EUF-CMA security of Sig. However, here we must bound, not the success probability of one specific forger, as in the EUF-CMA definition, but the success probability of all forgers that satisfy the fixed polynomial p and q bounds, for every time t and every schedule τ_1.

Define t_l to be the time point marking the beginning of j's lifetime: $t_l = min(\mathsf{aliveTimes}(j))$. We know that both $\mathsf{RealSig}(j)_k$ and $\mathsf{IdealSig}(j)_k$ are dormant before time t_l.

For every $k \in \mathbb{N}$, we define a quasi-$p(k)$-bounded environment $(\mathsf{Env}_{max})_k$ for $\mathsf{RealSig}_k$, a time $(t_{max})_k \leq t_l$, a schedule $(\tau_{1max})_k$ for $\mathsf{RealSig}(j)_k \| (\mathsf{Env}_{max})_k$ that is valid for interval $[0, (t_{max})_k + q(k)]$, and an execution $(\alpha_{max})_k$ in the support of $\mathsf{Execs}_{(\mathsf{Env}_{max})_k}(\mathsf{RealSig}(j)_k \| (\mathsf{Env}_{max})_k, \mathsf{trunc}_{\geq (t_{max})_k}((\tau_{1max})_k))$, satisfying the following property: For every quasi-$p(k)$-bounded environment Env for $\mathsf{RealSig}_k$, every time $t \leq t_l$, every valid schedule τ_1 for $\mathsf{RealSig}_k \| \mathsf{Env}$ for the interval $[0, t+q(k)]$, and every execution α in the support of $\mathsf{Execs}_{\mathsf{Env}}(\mathsf{RealSig}(j)_k \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_1)$:

$$
\begin{aligned}
&|\,\mathbf{P}_{\mathsf{acc}}(\mathsf{RealSig}(j)_k \| \mathsf{Env}, \tau_1, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\mathsf{IdealSig}(j)_k \| \mathsf{Env}, \tau_1, t, \alpha)| \\
&\leq |\,\mathbf{P}_{\mathsf{acc}}(\mathsf{RealSig}(j)_k \| (\mathsf{Env}_{max})_k, (\tau_{1max})_k, (t_{max})_k, (\alpha_{max})_k) \\
&\qquad - \mathbf{P}_{\mathsf{acc}}(\mathsf{IdealSig}(j)_k \| (\mathsf{Env}_{max})_k, (\tau_{1max})_k, (t_{max})_k, (\alpha_{max})_k)|.
\end{aligned}
$$

To see that such $(\mathsf{Env}_{max})_k$, $(t_{max})_k$, $(\tau_{1max})_k$, and $(\alpha_{max})_k$ exist, it is enough to observe:

- The set of quasi-$p(k)$-bounded environments is finite (up to isomorphism).
- The set of *untimed* versions of the candidate τ_1 schedules is finite; this follows from the rate restriction and the task enabling properties (properties (ii) and (iii) in the definition of p-boundedness).
- The set of candidate α executions is finite.
- The probability of acceptance depends only on the untimed version of the τ_1 schedule and on the (untimed) execution α.

Based on these observations, we first define $(\mathsf{Env}_{max})_k$, an untimed schedule ρ, and $(\alpha_{max})_k$ that yield the maximum difference between the acceptance probabilities. Then we fix $(\tau_{1max})_k$ and $(t_{max})_k$ to be any valid timed schedule and real time that yield ρ and α.

We use $(\mathsf{Env}_{max})_k$, $(\tau_{1max})_k$, $(t_{max})_k$, and $(\alpha_{max})_k$, for all k, to define the needed negligible function ϵ. We do this by defining a probabilistic polynomial-time (non-uniform) forger $G = \{G_k\}_{k \in \mathbb{N}}$ for Sig, in such a way that each G_k essentially emulates an execution of the automaton $\mathsf{IdealSig}(j)_k \| (\mathsf{Env}_{max})_k$ with schedule $(\tau_{1max})_k$, beginning with $(\alpha_{max})_k$.

First, we require some observations about $(\alpha_{max})_k$. Since $(\alpha_{max})_k$ is in the support of $\mathsf{Execs}_{(\mathsf{Env}_{max})_k}(\mathsf{RealSig}(j)_k \| (\mathsf{Env}_{max})_k, \mathsf{trunc}_{\geq (t_{max})_k}((\tau_{1max})_k))$, $(\alpha_{max})_k$ is equal to $\alpha' \restriction (\mathsf{Env}_{max})_k$, where α' is some execution in the support of $\mathsf{Execs}(\mathsf{RealSig}(j)_k \| (\mathsf{Env}_{max})_k, \mathsf{trunc}_{\geq (t_{max})_k}((\tau_{1max})_k))$. Since $(t_{max})_k \leq t_l$, α' contains only trivial steps of $\mathsf{RealSig}(j)_k$—no locally-controlled steps, and only self-loops for inputs. That is, the only interesting activity in α' occurs in $(\mathsf{Env}_{max})_k$, and in particular, probabilistic branching involves only the state of $(\mathsf{Env}_{max})_k$. So $(\alpha_{max})_k$ completely determines execution α'. Furthermore, since only trivial steps of $\mathsf{RealSig}(j)_k$ occur in such an α', $\mathsf{RealSig}(j)_k$ could be replaced by $\mathsf{IdealSig}(j)_k$: $(\alpha_{max})_k$ completely determines an execution in the support of $\mathsf{Execs}(\mathsf{IdealSig}(j)_k \| (\mathsf{Env}_{max})_k, \mathsf{trunc}_{\geq (t_{max})_k}((\tau_{1max})_k))$.

Now we define G_k. G_k successively reads all the tasks in the schedule $(\tau_{1max})_k$, and uses them to internally emulate an execution of $\mathsf{IdealSig}(j)_k \| (\mathsf{Env}_{max})_k$, with the following exceptions:

1. When G_k emulates the $\{\mathsf{verKey}(*)\}$ task, it replaces the verification algorithm obtained when emulating the $\{\mathsf{chooseKeys}\}$ task with the one provided by Sig in the EUF-CMA game;
2. When G_k emulates the $\{\mathsf{sign}(*,*)\}$ task, it obtains signatures by using the signing oracle available in the EUF-CMA game;
3. When G_k emulates any task with an associated time strictly less than t_{max}, it determines the next state of $(\mathsf{Env}_{max})_k$ from α rather than according to the probabilistic branching allowed by the transition.

Furthermore, G_k stores a list of all messages that the emulated $(\mathsf{Env}_{max})_k$ asked to sign, and checks whether $(\mathsf{Env}_{max})_k$ 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. (Although t_l may be very large, it does not depend on k, and so does not cause a violation of the polynomial-time requirement.)

We also observe that the first two proposed exceptions in the emulation of the execution of $\mathsf{IdealSig}(j)_k \| (\mathsf{Env}_{max})_k$ do not change the distribution of the messages that $(\mathsf{Env}_{max})_k$ 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 (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 the existence of a negligible function ϵ bounding the success probability of G. It follows that:

$$\begin{aligned}&|\,\mathbf{P}_{\mathsf{acc}}(\mathsf{RealSig}(j)_k \| (\mathsf{Env}_{max})_k, (\tau_{1max})_k, (t_{max})_k, (\alpha_{max})_k)\\ &\quad - \mathbf{P}_{\mathsf{acc}}(\mathsf{IdealSig}(j)_k \| (\mathsf{Env}_{max})_k, (\tau_{1max})_k, (t_{max})_k, (\alpha_{max})_k)| \le \epsilon(k).\end{aligned}$$

We fix this function ϵ for the rest of our proof.

It remains to show that, for every $k \in \mathbb{N}$, $t \in \mathbb{R}_{\ge 0}$, quasi-$p(k)$-bounded environment Env for $\mathsf{RealSig}(j)_k$, and valid schedule τ_1 for $\mathsf{RealSig}(j)_k \| \mathsf{Env}$ for the interval $[0, t + q(k)]$, there is a valid schedule τ_2 for $\mathsf{IdealSig}(j)_k, \| \mathsf{Env}$ satisfying the four required conditions. Fix k, Env, t, and τ_1. We consider two cases.

First, suppose that $t_l < t$. We obtain τ_2 by inserting $\langle \{\mathsf{fail}_j\}, t_l \rangle$ immediately after $\langle \mathsf{tick}, t_l \rangle$. This sets the *failed* flag in $\mathsf{SigFunc}(k, j)$ to true immediately after *awake* becomes true. Notice that, if *failed* = true, the verify transition bypasses the check $m \in history$ (Figure 5). In other words, $\mathsf{SigFunc}(k, j)$ answers verify requests in exactly the same way as $\mathsf{Verifier}(k, j)$, using the Verify algorithm only. Furthermore, it is easy to check that *failed* remains true as long as $\mathsf{SigFunc}(k, j)$ is alive. Therefore, $\mathsf{IdealSig}(j)_k$ has exactly the same visible behavior as $\mathsf{RealSig}(j)_k$ and Conditions (i) through (iv) are satisfied. (For Condition (iv), we obtain a bound of 0, which implies the needed bound of $\epsilon(k)$.)

Second, suppose that $t \le t_l$. Define $\tau_2 := \tau_1$. Since both $\mathsf{RealSig}(j)_k$ and $\mathsf{IdealSig}(j)_k$ are dormant during $[0, t]$, Condition (i) is immediate and Condition (ii) holds because fail_j is not a task of $\mathsf{RealSig}(j)_k$. Condition (iii) also must hold. For Condition (iv), observe that, for every every execution α in the support of $\mathsf{Execs}_{\mathsf{Env}}(\mathsf{RealSig}(j)_k \| \mathsf{Env}, \mathsf{trunc}_{\ge t}(\tau_1))$,

$$\begin{aligned}&|\,\mathbf{P}_{\mathsf{acc}}(\mathsf{RealSig}(j)_k \| \mathsf{Env}, \tau_1, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\mathsf{IdealSig}(j)_k \| \mathsf{Env}, \tau_1, t, \alpha)|\\ &\le |\,\mathbf{P}_{\mathsf{acc}}(\mathsf{RealSig}(j)_k \| (\mathsf{Env}_{max})_k, (\tau_{1max})_k, (t_{max})_k, (\alpha_{max})_k)\\ &\qquad - \mathbf{P}_{\mathsf{acc}}(\mathsf{IdealSig}(j)_k \| (\mathsf{Env}_{max})_k, (\tau_{1max})_k, (t_{max})_k, (\alpha_{max})_k)|\\ &\le \epsilon(k),\end{aligned}$$

as needed. □

7 Composition Theorems

In practice, cryptographic services are seldom used in isolation. Usually, 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 begins. An example of such a collection is the signature services in the timestamping protocol of [15, 14], 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.[5] This highlights a unique aspect of our implementation relation: essentially, from any point t on the real time line, 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 (Parallel Composition Theorem). *Let $\mathcal{A}_1^1, \mathcal{A}_2^1, \ldots$ and $\mathcal{A}_1^2, \mathcal{A}_2^2, \ldots$ be two infinite sequences of task-PIOAs, with $\mathcal{A}_i^1$ comparable to $\mathcal{A}_i^2$ for every i. Suppose that $\mathcal{A}_1^{\alpha_1}, \mathcal{A}_2^{\alpha_2}, \ldots$ are pairwise compatible for any combination of $\alpha_i \in \{1, 2\}$. Let $b \in \mathbb{N}$, and let $\hat{\mathcal{A}}^1$ and $\hat{\mathcal{A}}^2$ denote $\|_{i=1}^b \mathcal{A}_i^1$ and $\|_{i=1}^b \mathcal{A}_i^2$, respectively. Let r be a nondecreasing function, $r : \mathbb{N} \to \mathbb{N}$ such that, for every i, both $\mathcal{A}_i^1$ and $\mathcal{A}_i^2$ are $r(i)$-bounded.*

For each i, let F_i^1 and F_i^2 be sets of tasks of $\mathcal{A}_i^1$ and $\mathcal{A}_i^2$, respectively, all with infinite upper bounds. Let $\hat{F}^1$ and $\hat{F}^2$ denote $\bigcup_{i=1}^b F_i^1$ and $\bigcup_{i=1}^b F_i^2$, respectively.

Let $p, q \in \mathbb{N}$ and $\epsilon \in \mathbb{R}_{\geq 0}$. Suppose that $(\mathcal{A}_i^1, F_i^1) \leq_{p,q,\epsilon} (\mathcal{A}_i^2, F_i^2)$ for every i.

Let $p' \in \mathbb{N}$ and $\epsilon' \in \mathbb{R}_{\geq 0}$, with $p = c_{\mathsf{comp}} \cdot (b \cdot r(b) + p')$ (where c_{comp} is the constant factor for parallel composition of task-PIOAs), and $\epsilon' = b \cdot \epsilon$. Then $(\hat{\mathcal{A}}^1, \hat{F}^1) \leq_{p',q,\epsilon'} (\hat{\mathcal{A}}^2, \hat{F}^2)$.

Proof. Let $t \in \mathbb{R}_{\geq 0}$ be given. Let Env be a quasi-p'-bounded environment and let τ_0 be a valid timed task schedule for $\hat{\mathcal{A}}^1 \| \mathsf{Env}$ for the interval $[0, t + q]$ where τ_0 contains no tasks from $\hat{F}_1$ occuring at time t or later. We must find τ_b for $\hat{\mathcal{A}}^2 \| \mathsf{Env}$ such that:

(i) $\mathsf{proj}_{\mathsf{Env}}(\tau_0) = \mathsf{proj}_{\mathsf{Env}}(\tau_b)$;
(ii) τ_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 $t_i \geq t$;
(iii) $\mathsf{Execs}_{\mathsf{Env}}(\hat{\mathcal{A}}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_0)) = \mathsf{Execs}_{\mathsf{Env}}(\hat{\mathcal{A}}^2 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_b))$; and
(iv) For every α in the support of $\mathsf{Execs}_{\mathsf{Env}}(\hat{\mathcal{A}}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}\, \tau_0)$, $|\, \mathbf{P}_{\mathsf{acc}}(\hat{\mathcal{A}}^1 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\hat{\mathcal{A}}^2 \| \mathsf{Env}, \tau_b, t, \alpha)| \leq \epsilon$.

For each i, $0 \leq i \leq b$, let H_i denote $\mathcal{A}_1^2 \| \ldots \| \mathcal{A}_i^2 \| \mathcal{A}_{i+1}^1 \| \ldots \| \mathcal{A}_b^1$. In particular, $H_0 = \|_{i=1}^b \mathcal{A}_i^1$ and $H_b = \|_{i=1}^b \mathcal{A}_i^2$. Also, for each i, $0 \leq i \leq b$, let Env_i denote $\mathcal{A}_1^2 \| \ldots \| \mathcal{A}_{i-1}^2 \| \mathcal{A}_{i+1}^1 \| \ldots \| \mathcal{A}_b^1 \| \mathsf{Env}$. Note that every Env_i is quasi-p-bounded and is an environment for $\mathcal{A}_i^1$ and $\mathcal{A}_i^2$. In fact, we have $H_{i-1} \| \mathsf{Env} = \mathcal{A}_i^1 \| \mathsf{Env}_i$ and $H_i \| \mathsf{Env} = \mathcal{A}_i^2 \| \mathsf{Env}_i$.

Since τ_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 $(\mathcal{A}_1^1, F_1^1) \leq_{p,q,\epsilon} (\mathcal{A}_1^2, F_2^2)$ and τ_0 is a valid schedule for $\mathcal{A}_1^1 \| \mathsf{Env}_1$ in which no tasks from F_1^1 occur from time t onwards, we may choose a valid schedule τ_1 for $\mathcal{A}_1^2 \| \mathsf{Env}_1$ for the interval $[0, t + q]$ such that

(i) $\mathsf{proj}_{\mathsf{Env}_1}(\tau_0) = \mathsf{proj}_{\mathsf{Env}_1}(\tau_1)$;
(ii) τ_1 does not contain any pairs of the form $\langle T_i, t_i \rangle$ where $T_i \in F_1^2$ and $t_i \geq t$;
(iii) $\mathsf{Execs}_{\mathsf{Env}_1}(\mathcal{A}_1^1 \| \mathsf{Env}_1, \mathsf{trunc}_{\geq t}(\tau_0)) = \mathsf{Execs}_{\mathsf{Env}_1}(\mathcal{A}_1^2 \| \mathsf{Env}_1, \mathsf{trunc}_{\geq t}(\tau_1))$; and
(iv) For every α' in the support of $\mathsf{Execs}_{\mathsf{Env}_1}(\mathcal{A}_1^1 \| \mathsf{Env}_1, \mathsf{trunc}_{\geq t}(\tau_0))$,
$|\, \mathbf{P}_{\mathsf{acc}}(\mathcal{A}_1^1 \| \mathsf{Env}_1, \tau_0, t, \alpha') - \mathbf{P}_{\mathsf{acc}}(\mathcal{A}_1^2 \| \mathsf{Env}_1, \tau_1, t, \alpha')| \leq \epsilon$.

Repeating this argument, we choose valid schedules $\tau_2, \ldots, \tau_b$ for $H_2 \| \mathsf{Env}, \ldots, H_b \| \mathsf{Env}$, respectively, all satisfying the appropriate four conditions.

Since Env is part of every Env_i, Condition (i) guarantees that $\mathsf{proj}_{\mathsf{Env}}(\tau_0) = \mathsf{proj}_{\mathsf{Env}}(\tau_b)$. Using both Conditions (i) and (ii), we can infer that τ_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 $t_i \geq t$. Since Env is part of every Env_i, Condition (iii) guarantees that

$$\mathsf{Execs}_{\mathsf{Env}}(H_0 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_0)) = \mathsf{Execs}_{\mathsf{Env}}(H_b \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_b)).$$

It remains to show Condition (iv). For this, fix α in the support of $\mathsf{Execs}_{\mathsf{Env}}(\hat{\mathcal{A}}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_0))$; we must show that $|\, \mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_b \| \mathsf{Env}, \tau_b, t, \alpha)| \leq \epsilon$.

We first show that $|\, \mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_1 \| \mathsf{Env}, \tau_1, t, \alpha)| \leq \epsilon$. To show this, we first decompose the probability $\mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha)$ using the total probability theorem, as the sum of terms of the form $p_{\alpha'} \cdot \mathbf{P}_{\mathsf{acc}}(A_1^1 \| \mathsf{Env}_1, \tau_0, t, \alpha')$, where α' is in the support of $\mathsf{Execs}_{\mathsf{Env}_1}(\mathcal{A}_1^1 \| \mathsf{Env}_1, \mathsf{trunc}_{\geq t}(\tau_0))$. Here, $p_{\alpha'}$ is the conditional

[5] 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.

probability that an execution's portion before time t projects on Env_1 to yield α', given that it projects on Env to yield α. More precisely $p_{\alpha'}$ is the conditional probability that an execution α'' in the support of $\mathsf{Execs}(H_0 \| \mathsf{Env}, \tau_0)$ satisfies $\alpha''' \upharpoonright \mathsf{Env}_1 = \alpha'$, where α''' is the prefix of α'' derived from $\mathsf{trunc}_{\geq t}(\tau)$, conditioned on the event that $\alpha''' \upharpoonright \mathsf{Env} = \alpha$.

Similarly, $\mathbf{P}_{\mathsf{acc}}(H_1 \| \mathsf{Env}, \tau_1, t, \alpha) = \Sigma_{\alpha'} p_{\alpha'} \cdot \mathbf{P}_{\mathsf{acc}}(A_1^2 \| \mathsf{Env}_1, \tau_1, t, \alpha')$. Then $|\mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_1 \| \mathsf{Env}, \tau_1, t, \alpha)|$ is equal to $|\Sigma_{\alpha'}(p_{\alpha'} \cdot \mathbf{P}_{\mathsf{acc}}(A_1^1 \| \mathsf{Env}_1, \tau_0, t, \alpha')) - \Sigma_{\alpha'}(p_{\alpha'} \cdot \mathbf{P}_{\mathsf{acc}}(A_1^2 \| \mathsf{Env}_1, \tau_1, t, \alpha'))|$, $\leq \Sigma_{\alpha'} p_{\alpha'} |\mathbf{P}_{\mathsf{acc}}(A_1^1 \| \mathsf{Env}_1, \tau_0, t, \alpha') - \mathbf{P}_{\mathsf{acc}}(A_1^2 \| \mathsf{Env}_1, \tau_1, t, \alpha')|$, $\leq \Sigma_{\alpha'} p_{\alpha'} \cdot \epsilon$ by the assumed Condition iv, $= \epsilon$. Therefore, $|\mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_1 \| \mathsf{Env}, \tau_1, t, \alpha)| \leq \epsilon$, as claimed.

By analogous arguments, $|\mathbf{P}_{\mathsf{acc}}(H_i \| \mathsf{Env}, \tau_i, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_{i+1} \| \mathsf{Env}, \tau_{i+1}, t, \alpha)| \leq \epsilon$ for every i, $0 \leq i \leq b-1$. Therefore,

$$
\begin{aligned}
&|\mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_b \| \mathsf{Env}, \tau_b, t, \alpha)| \\
&\leq |\mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_1 \| \mathsf{Env}, \tau_1, t, \alpha)| + \ldots \\
&\quad + |\mathbf{P}_{\mathsf{acc}}(H_i \| \mathsf{Env}, \tau_i, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_{i+1} \| \mathsf{Env}, \tau_{i+1}, t, \alpha)| + \ldots \\
&\quad + |\mathbf{P}_{\mathsf{acc}}(H_{b-1} \| \mathsf{Env}, \tau_{b-1}, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_b \| \mathsf{Env}, \tau_b, t, \alpha)| \\
&\leq b \cdot \epsilon = \epsilon'.
\end{aligned}
$$

Thus, $|\mathbf{P}_{\mathsf{acc}}(\hat{\mathcal{A}}^1 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\hat{\mathcal{A}}^2 \| \mathsf{Env}, \tau_b, t, \alpha)| \leq b \cdot \epsilon = \epsilon'$, as needed for Condition iv. □

Using Theorem 2, it is not hard to prove a polynomial composition theorem for $\leq_{\mathsf{neg,pt}}$. The theorem contains a technicality: instead of simply assuming $\leq_{\mathsf{neg,pt}}$ relationships for all the components, we assume a slightly stronger property, in which the same negligible function ϵ is assumed for all of the components; that is, ϵ is not allowed to depend on the component index i.

Theorem 3 (Parallel Composition Theorem for $\leq_{\mathsf{neg,pt}}$). *Let $\bar{\mathcal{A}}_1^1, \bar{\mathcal{A}}_2^1, \ldots$ and $\bar{\mathcal{A}}_1^2, \bar{\mathcal{A}}_2^2, \ldots$ be two infinite sequences of task-PIOA families, with $\bar{\mathcal{A}}_i^1$ comparable to $\bar{\mathcal{A}}_i^2$ for every i. Suppose that $\bar{\mathcal{A}}_1^{\alpha_1}, \bar{\mathcal{A}}_2^{\alpha_2}, \ldots$ are pairwise compatible for any combination of $\alpha_i \in \{1, 2\}$. Let b be any polynomial, and for each k, let $(\widehat{\mathcal{A}}^1)_k$ and $(\widehat{\mathcal{A}}^2)_k$ denote $\|_{i=1}^{b(k)} (\bar{\mathcal{A}}_i^1)_k$ and $\|_{i=1}^{b(k)} (\bar{\mathcal{A}}_i^2)_k$, respectively. Let r and s be polynomials, $r, s : \mathbb{N} \to \mathbb{N}$, such that r is nondecreasing, and for every i, k, both $(\bar{\mathcal{A}}_i^1)_k$ and $(\bar{\mathcal{A}}_i^2)_k$ are bounded by $s(k) \cdot r(i)$.*

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, where all these tasks have infinite upper bounds. Let $(\hat{F}^1)_k$ and $(\hat{F}^2)_k$ denote $\bigcup_{i=1}^{b(k)} (\bar{F}_i^1)_k$ and $\bigcup_{i=1}^{b(k)} (\bar{F}_i^2)_k$, respectively.

Assume:

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

where p, q are polynomials and ϵ is a negligible function. Then $(\widehat{\mathcal{A}}^1, \hat{F}^1) \leq_{\mathsf{neg,pt}} (\widehat{\mathcal{A}}^2, \hat{F}^2)$.

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

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

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

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 ϵ, 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 $((\widehat{\mathcal{A}}^1)_k, (\hat{F}^1)_k) \leq_{p'(k),q(k),\epsilon'(k)} ((\widehat{\mathcal{A}}^2)_k, (\hat{F}^2)_k)$, as needed. □

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.[6] In Definition 3, we formalize the idea of an entity $\mathcal{A}$ being active during a particular 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 every $t \notin [t_1, t_2]$, environment Env for $\mathcal{A}$ of the form $\mathsf{Env}' \| \mathsf{Clock}$, valid schedule τ for $\mathcal{A} \| \mathsf{Env}$ for $[0, t]$, and state s reachable under τ, no locally controlled actions of $\mathcal{A}$ are enabled in s, and $s.v = \perp$ for every variable v of $\mathcal{A}$.*

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 Env be an environment for both $\mathcal{A}^1$ and $\mathcal{A}^2$, of the form $\mathsf{Env}' \| \mathsf{Clock}$. Let $t \in \mathbb{R}_{\geq 0}$ and $q \in \mathbb{N}$. Suppose τ_1 is a valid schedule for $\mathcal{A}^1 \| \mathsf{Env}$ for the interval $[0, t+q]$, and τ_2 is a valid schedule for $\mathcal{A}^2 \| \mathsf{Env}$ for $[0, t+q]$, satisfying:*

- $\mathsf{proj}_{\mathsf{Env}}(\tau_1) = \mathsf{proj}_{\mathsf{Env}}(\tau_2)$;
- $\mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_1)) = \mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^2 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_2))$.

Assume further that either $t_2 < t$ or $t_1 > t + q$. Then for every α in the support of $\mathsf{Execs}_{\mathsf{Env}}(\mathcal{A}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}\, \tau_1)$, $\mathbf{P}_{\mathsf{acc}}(\mathcal{A}^1 \| \mathsf{Env}, \tau_1, t, \alpha) = \mathbf{P}_{\mathsf{acc}}(\mathcal{A}^2 \| \mathsf{Env}, \tau_2, t, \alpha))$.

Proof. 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 $[t, t+q]$. Since τ_1 and τ_2 agree on the tasks of Env, and the execution α of Env just before time t is identical in the two experiments, the probability that Env outputs acc during $[t, t+q]$ must be identical. Also, since the executions before t are the same in the two experiments, the probability that Env outputs acc during $[0, t)$ is the same. Therefore, the acceptance probabilities are the same for the entire interval $[0, t+q]$, as needed.

□

Definition 4 (Sequentiality). *Let $\mathcal{A}_1, \mathcal{A}_2, \ldots$ be pairwise compatible task-PIOAs. We say that $\mathcal{A}_1, \mathcal{A}_2, \ldots$ are* sequential with respect to *the the nondecreasing sequence $t_1, t_2, \ldots$ of nonnegative reals provided that for every i, $\mathcal{A}_i$ is restricted to $[t_i, t_{i+1}]$.*

Note the slight technicality that each $\mathcal{A}_i$ may overlap with $\mathcal{A}_{i+1}$ at the boundary time t_{i+1}.

Theorem 4 (Sequential Composition Theorem). *Let $\mathcal{A}_1^1, \mathcal{A}_2^1, \ldots$ and $\mathcal{A}_1^2, \mathcal{A}_2^2, \ldots$ be two infinite sequences of task-PIOAs, with $\mathcal{A}_i^1$ comparable to $\mathcal{A}_i^2$ for every i. Suppose that $\mathcal{A}_1^{\alpha_1}, \mathcal{A}_2^{\alpha_2}, \ldots$ are pairwise compatible for any combination of $\alpha_i \in \{1, 2\}$. Let $L \in \mathbb{N}$, and let $\hat{\mathcal{A}}^1$ and $\hat{\mathcal{A}}^2$ denote $\|_{i=1}^{L} \mathcal{A}_i^1$ and $\|_{i=1}^{L} \mathcal{A}_i^2$, respectively. Let $\hat{p} \in \mathbb{N}$, and assume that both $\hat{\mathcal{A}}^1$ and $\hat{\mathcal{A}}^2$ are $\hat{p}$-bounded.*

Assume that both $\mathcal{A}_1^1, \ldots, \mathcal{A}_L^1$ and $\mathcal{A}_1^2, \ldots, \mathcal{A}_L^2$ are sequential with respect to the same nondecreasing sequence of reals $t_1, t_2, \ldots$. Assume that $b \in \mathbb{N}$ is an upper bound on the number of t_i's that fall into a single closed interval of length q.

For each i, let F_i^1 and F_i^2 be sets of tasks of $\mathcal{A}_i^1$ and $\mathcal{A}_i^2$, respectively, all with infinite upper bounds. Let $\hat{F}^1$ and $\hat{F}^2$ denote $\bigcup_{i=1}^{L} F_i^1$ and $\bigcup_{i=1}^{L} F_i^2$, respectively.

Let $p, q \in \mathbb{N}$ and $\epsilon \in \mathbb{R}_{\geq 0}$. Suppose that $(\mathcal{A}_i^1, F_i^1) \leq_{p,q,\epsilon} (\mathcal{A}_i^2, F_i^2)$ for every i.

Let $p' \in \mathbb{N}$ and $\epsilon' \in \mathbb{R}_{\geq 0}$, with $p \geq c_{\mathsf{comp}} \cdot (\hat{p} + p')$ (where c_{comp} is the constant factor for parallel composition), and $\epsilon' \geq (b+2) \cdot \epsilon$. Then $(\hat{\mathcal{A}}^1, \hat{F}^1) \leq_{p',q,\epsilon'} (\hat{\mathcal{A}}^2, \hat{F}^2)$.

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 lifetime of each component is $\frac{q}{3}$, then b is 5.[7] 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.

[6] 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.

[7] Recall that two components may be active simultaneously at the boundary time.

Proof. The proof is similar to that of Theorem 2. Let $t \in \mathbb{R}_{\geq 0}$ be given. Let Env be a quasi-p'-bounded environment and let τ_0 be a valid timed task schedule for $\hat{\mathcal{A}}1 \| \mathsf{Env}$ for the interval $[0, t+q]$ where τ_0 has no tasks from $\hat{F}^1$ occuring at time t or later. We must find τ_L for $\hat{\mathcal{A}}^2 \| \mathsf{Env}$ such that:

(i) $\mathsf{proj}_{\mathsf{Env}}(\tau_0) = \mathsf{proj}_{\mathsf{Env}}(\tau_L)$;
(ii) τ_L does not contain any pairs of the form $\langle T_i, t_i \rangle$ where $T_i \in \hat{F}^2$ and $t_i \geq t$;
(iii) $\mathsf{Execs}_{\mathsf{Env}}(\hat{\mathcal{A}}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_0)) = \mathsf{Execs}_{\mathsf{Env}}(\hat{\mathcal{A}}^2 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_L))$; and
(iv) For every α in the support of $\mathsf{Execs}_{\mathsf{Env}}(\hat{\mathcal{A}}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_0))$, $|\, \mathbf{P}_{\mathsf{acc}}(\hat{\mathcal{A}}^1 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(\hat{\mathcal{A}}^2 \| \mathsf{Env}, \tau_L, t, \alpha)| \leq \epsilon'$.

For each i, $0 \leq i \leq L$, let H_i denote $\mathcal{A}_1^2 \| \ldots \| \mathcal{A}_i^2 \| \mathcal{A}_{i+1}^1 \| \ldots \| \mathcal{A}_L^1$, and let $\mathsf{Env}_i := \mathcal{A}_1^2 \| \ldots \| \mathcal{A}_{i-1}^2 \| \mathcal{A}_{i+1}^1 \| \ldots \| \mathcal{A}_L^1 \| \mathsf{Env}$. Note that Env_i is quasi-p-bounded; therefore we may choose τ_{i+1} using τ_i and the assumption that $(\mathcal{A}_i^1, F_i^1) \leq_{p,q,\epsilon} (\mathcal{A}_i^2, F_i^2)$.

Since Env is part of Env_i for every i, Conditions (i) and (iii) are clearly satisfied at every replacement step. Condition (ii) is satisfied because the following hold at every step i.

– The new task schedule τ_{i+1} does not contain tasks from F_{i+1}^2.
– Condition (i) guarantees that τ_{i+1} does not contain tasks from $\bigcup_{j=1}^{i} F_j^2$.

It remains to show Condition (iv). Thus, fix α in the support of $\mathsf{Execs}_{\mathsf{Env}}(\hat{\mathcal{A}}^1 \| \mathsf{Env}, \mathsf{trunc}_{\geq t}(\tau_0))$; we must show that $|\, \mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_L \| \mathsf{Env}, \tau_L, t, \alpha)| \leq \epsilon'$. Arguing as before, we obtain that

$$
\begin{aligned}
&|\, \mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_L \| \mathsf{Env}, \tau_L), t, \alpha| \\
&\leq |\, \mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_1 \| \mathsf{Env}, \tau_1, t, \alpha)| + \ldots \\
&\quad + |\, \mathbf{P}_{\mathsf{acc}}(H_i \| \mathsf{Env}, \tau_i, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_{i+1} \| \mathsf{Env}, \tau_{i+1}, t, \alpha)| + \ldots \\
&\quad + |\, \mathbf{P}_{\mathsf{acc}}(H_{L-1} \| \mathsf{Env}, \tau_{L-1}, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_L \| \mathsf{Env}, \tau_L, t, \alpha)|.
\end{aligned}
$$

As before, each of the terms in the summation is at most ϵ. However, this is not sufficient here to prove the needed bound, because we have exponentially many terms. To prove the bound, we argue that most of the terms in the summation are in fact 0.

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 $t_{l+b} > t+q$.

If $i < l-1$ or $i \geq l+b$, then we can apply Lemma 5 to conclude that $\mathbf{P}_{\mathsf{acc}}(\mathcal{A}_i^1 \| \mathsf{Env}_i, \tau_i, t, \alpha)$ is in fact equal to $\mathbf{P}_{\mathsf{acc}}(\mathcal{A}_i^2 \| \mathsf{Env}_i, \tau_{i+1}, t, \alpha)$, that is, that $\mathbf{P}_{\mathsf{acc}}(H_i \| \mathsf{Env}, \tau_i, t, \alpha) = \mathbf{P}_{\mathsf{acc}}(H_{i+1} \| \mathsf{Env}, \tau_{i+1}, t, \alpha)$. Summing over all indices i, we have $|\, \mathbf{P}_{\mathsf{acc}}(H_0 \| \mathsf{Env}, \tau_0, t, \alpha) - \mathbf{P}_{\mathsf{acc}}(H_L \| \mathsf{Env}, \tau_L, t, \alpha)| \leq (b+2) \cdot \epsilon = \epsilon'$, as needed.

□

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

Theorem 5 (Sequential Composition Theorem for $\leq_{\mathsf{neg,pt}}$). *Let $\bar{\mathcal{A}}_1^1, \bar{\mathcal{A}}_2^1, \ldots$ and $\bar{\mathcal{A}}_1^2, \bar{\mathcal{A}}_2^2, \ldots$ be two infinite sequences of task-PIOA families, with $\bar{\mathcal{A}}_i^1$ comparable to $\bar{\mathcal{A}}_i^2$ for every i. Suppose that $\bar{\mathcal{A}}_1^{\alpha_1}, \bar{\mathcal{A}}_2^{\alpha_2}, \ldots$ are pairwise compatible for any combination of $\alpha_i \in \{1, 2\}$. Let $L : \mathbb{N} \to \mathbb{N}$ be an exponential function and, for each k, let $(\widehat{\mathcal{A}}^1)_k$ and $(\widehat{\mathcal{A}}^2)_k$ denote $\|_{i=1}^{L(k)} (\bar{\mathcal{A}}_i^1)_k$ and $\|_{i=1}^{L(k)} (\bar{\mathcal{A}}_i^2)_k$, respectively. Let $\hat{p}$ be a polynomial such that both $\widehat{\mathcal{A}}^1$ and $\widehat{\mathcal{A}}^2$ are $\hat{p}$-bounded.*

Suppose there exists an increasing sequence of nonnegative reals $t_1, t_2, \ldots$ such that, for each k, both $(\bar{\mathcal{A}}_1^1)_k, \ldots, (\bar{\mathcal{A}}_{L(k)}^1)_k$ and $(\bar{\mathcal{A}}_1^2)_k, \ldots, (\bar{\mathcal{A}}_{L(k)}^2)_k$ are sequential for $t_1, t_2, \ldots$ Assume there is a constant real number c such that consecutive t_i's are at least c apart.

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, where all these tasks have infinite upper bounds. Let $(\hat{F}^1)_k$ and $(\hat{F}^2)_k$ denote $\bigcup_{i=1}^{L(k)} (\bar{F}_i^1)_k$ and $\bigcup_{i=1}^{L(k)} (\bar{F}_i^2)_k$, respectively.

Assume:

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

where p, q are polynomials and ϵ is a negligible function. Then $(\widehat{\mathcal{A}}^1, \hat{F}^1) \leq_{\mathsf{neg,pt}} (\widehat{\mathcal{A}}^2, \hat{F}^2)$.

Proof. Let polynomials p', q be given and define $p := c_{\mathsf{comp}} \cdot (\hat{p} + p')$, where c_{comp} is the constant factor for composing task-PIOAs in parallel. Choose ϵ 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 $((\widehat{\mathcal{A}}^1)_k, (\hat{F}^1)_k) \leq_{p'(k),q(k),\epsilon'(k)} ((\widehat{\mathcal{A}}^2)_k, (\hat{F}^2)_k$, as needed. □

Next, we present a corollary to Theorem 5, which provides a composition result for d-bounded concurrent systems, for d any positive integer.

Definition 5 (d-Bounded Concurrency). *Let $\mathcal{A}_1, \mathcal{A}_2, \ldots$ be pairwise compatible task-PIOAs, d a positive integer. We say that $\mathcal{A}_1, \mathcal{A}_2, \ldots$ are* d-bounded-concurrent with respect to *sequences $l_1, l_2, \ldots$ and $r_1, r_2, \ldots$ of nonnegative reals provided that:*

(i) $0 \leq l_1 \leq l_2 \leq \ldots$, and for every i, $l_i \leq r_i$.
(ii) For every positive real t, t is in the interior of at most d of the intervals $[l_i, r_i]$, that is, $|\{i : l_i < t < r_i\}| \leq d$.
(iii) For every i, $\mathcal{A}_i$ is restricted to $[l_i, r_i]$,

Corollary 1 (d-Bounded Composition Theorem for $\leq_{\mathsf{neg,pt}}$). *Let $\bar{\mathcal{A}}_1^1, \bar{\mathcal{A}}_2^1, \ldots$ and $\bar{\mathcal{A}}_1^2, \bar{\mathcal{A}}_2^2, \ldots$ be two infinite sequences of task-PIOA families, with $\bar{\mathcal{A}}_i^1$ comparable to $\bar{\mathcal{A}}_i^2$ for every i. Suppose that $\bar{\mathcal{A}}_1^{\alpha_1}, \bar{\mathcal{A}}_2^{\alpha_2}, \ldots$ are pairwise compatible for any combination of $\alpha_i \in \{1, 2\}$. Let $L : \mathbb{N} \to \mathbb{N}$ be an exponential function and, for each k, let $(\widehat{\mathcal{A}}^1)_k$ and $(\widehat{\mathcal{A}}^2)_k$ denote $\|_{i=1}^{L(k)} (\bar{\mathcal{A}}_i^1)_k$ and $\|_{i=1}^{L(k)} (\bar{\mathcal{A}}_i^2)_k$, respectively. Let $\hat{p}$ be a polynomial such that both $\widehat{\mathcal{A}}^1$ and $\widehat{\mathcal{A}}^2$ are $\hat{p}$-bounded.*

Let d be a positive integer. Suppose that $l_1, l_2, \ldots$ and $r_1, r_2, \ldots$ are two sequences of nonnegative reals, and for each k, both $(\bar{\mathcal{A}}_1^1)_k, (\bar{\mathcal{A}}_2^1)_k, \ldots$ and $(\bar{\mathcal{A}}_1^2)_k, (\bar{\mathcal{A}}_2^2)_k, \ldots$ are d-bounded concurrent with respect to $l_1, l_2, \ldots$ and $r_1, r_2, \ldots$. Let c be a positive real, and suppose that $l_i + c \leq r_i$ for every i.

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, where all these tasks have infinite upper bounds. Let $(\hat{F}^1)_k$ and $(\hat{F}^2)_k$ denote $\bigcup_{i=1}^{L(k)} (\bar{F}_i^1)_k$ and $\bigcup_{i=1}^{L(k)} (\bar{F}_i^2)_k$, respectively.

Assume:

$$\forall p, q\ \exists \epsilon\ \forall i\ (\bar{\mathcal{A}}_i^1, \bar{F}_i^1) \leq_{p,q,\epsilon} (\bar{\mathcal{A}}_i^2, \bar{F}_i^2), \tag{3}$$

where p, q are polynomials and ϵ is a negligible function. Then $(\widehat{\mathcal{A}}^1, \hat{F}^1) \leq_{\mathsf{neg,pt}} (\widehat{\mathcal{A}}^2, \hat{F}^2)$.

Proof. By induction on d. The base case, $d = 1$, follows easily from Theorem 5, where the increasing sequence $t_1, t_2, \ldots$ is simply the sequence of left interval endpoints $l_1, l_2, \ldots$.

For the inductive step, we assume that the result holds for all values up to $d - 1$, and show it for d. We extract a pair of sequences of task-PIOA families to which we can apply Theorem 5, in such a way that the remaining pair of sequences of task-PIOA families satisfy the inductive hypothesis. To extract these sequences, we select a subset $I = \{i_1, i_2, \ldots\}$ of the indices, with $i_1 < i_2 < \ldots$, and consider the task-PIOA families associated with the indices in I.

We construct the subset I as follows: Let $i_1 = 1$. Then for each $j > 1$ in turn, define m_j and i_j as follows: Let $m_j = min\{l_i : l_i \geq r_{i_{j-1}}\}$, that is, the smallest left endpoint of any interval that is greater than or equal to the right endpoint of the previously-chosen interval, and let i_j be the smallest index with $l_{i_j} = m_j$.

Now we consider the two sequences of task-PIOA families associated with the indices in I, $\bar{\mathcal{A}}_{i_1}^1, \bar{\mathcal{A}}_{i_2}^1, \ldots$ and $\bar{\mathcal{A}}_{i_1}^2, \bar{\mathcal{A}}_{i_2}^2, \ldots$. We apply Theorem 5 to these two sequences, and conclude that the compositions of these families are related by $\leq_{\mathsf{neg,pt}}$. More precisely, for every k, define $I(k) = I \cap \{i : i \leq L(k)\}$. Define task-PIOA families $\hat{\mathcal{B}}^1$ and $\hat{\mathcal{B}}^2$, where for every k, $(\hat{\mathcal{B}}^1)_k = \|_{i \in I(k)} (\bar{\mathcal{A}}_i^1)_k$ and $(\hat{\mathcal{B}}^2)_k = \|_{i \in I(k)} (\bar{\mathcal{A}}_i^2)_k$. Also define failure-task-set families $\hat{G}^1$ and $\hat{G}^2$, where $(\hat{G}^1)_k = \bigcup_{i \in I(k)} (\bar{F}_i^1)_k$ and $(\hat{G}^2)_k = \bigcup_{i \in I(k)} (\bar{F}_i^2)_k$. Observe that, for every k, the sequences $(\bar{\mathcal{A}}_{i_1}^1)_k, (\bar{\mathcal{A}}_{i_2}^1)_k, \ldots$ and $(\bar{\mathcal{A}}_{i_1}^2)_k, (\bar{\mathcal{A}}_{i_2}^2)_k, \ldots$ are both sequential for $l_{i_1}, l_{i_2}, \ldots$. Then Theorem 5 implies that $(\hat{\mathcal{B}}^1, \hat{G}^1) \leq_{\mathsf{neg,pt}} (\hat{\mathcal{B}}^2, \hat{G}^2)$.

Let $J = \mathbb{N} - I$ be the set of non-selected indices. We claim that J satisfies $d - 1$-bounded concurrency; namely, for every positive real t, t is in the interior of at most $d - 1$ intervals $[l_i, r_i]$ for $i \in J$.

To see this, we argue by contradiction: Consider any time t that falls into the interior of d of the intervals for indices in J. Then t cannot also be in the interior of an interval for an index in I, since that would mean that t is in the interior of at least $d + 1$ intervals overall, which violates the d-bounded-concurrency assumption. Similarly, t cannot be either a left or right endpoint of any interval for an index in I, since in either case, a slight perturbation of t would be in the

interior of $d+1$ intervals overall. It follows that t must lie in the "gap" between intervals for indices i_{j-1} and i_j, for some j. But then we claim that at least one of the d intervals for indices in J containing t in its interior must have its left endpoint $\geq r_{i_{j-1}}$: if not, then all of these intervals would overlap the interval for i_{j-1} by more than just one point, again violating d-bounded concurrency. But this claim violates the choice of l_{i_j} as the smallest left endpoint $\geq r_{i_{j-1}}$.

It follows that the pair of subsequences of task-PIOA families associated with the indices in J satisfy the assumptions for the inductive hypothesis. So by the conclusion of the inductive hypothesis, the two compositions of these families of task-PIOAs are related by $\leq_{\mathsf{neg,pt}}$. More precisely, for every k, define $J(k) = J \cap \{i : i \leq L(k)\}$. Define task-PIOA families $\hat{\mathcal{C}}^1$ and $\hat{\mathcal{C}}^2$, where for every k, $(\hat{\mathcal{C}}^1)_k = \|_{i \in J(k)} (\bar{\mathcal{A}}_i^1)_k$ and $(\hat{\mathcal{C}}^2)_k = \|_{i \in J(k)} (\bar{\mathcal{A}}_i^2)_k$. Also define failure-task-set families $\hat{H}^1$ and $\hat{H}^2$, where $(\hat{H}^1)_k = \bigcup_{i \in J(k)} (\bar{F}_i^1)_k$ and $(\hat{H}^2)_k = \bigcup_{i \in J(k)} (\bar{F}_i^2)_k$. Then the inductive hypothesis implies that $(\hat{\mathcal{C}}^1, \hat{H}^1) \leq_{\mathsf{neg,pt}} (\hat{\mathcal{C}}^2, \hat{H}^2)$.

Finally, we combine the claims $(\hat{\mathcal{B}}^1, \hat{G}^1) \leq_{\mathsf{neg,pt}} (\hat{\mathcal{B}}^2, \hat{G}^2)$ and $(\hat{\mathcal{C}}^1, \hat{H}^1) \leq_{\mathsf{neg,pt}} (\hat{\mathcal{C}}^2, \hat{H}^2)$ using Theorem 3, to conclude the final result, $(\widehat{\mathcal{A}}^1, \hat{F}^1) \leq_{\mathsf{neg,pt}} (\widehat{\mathcal{A}}^2, \hat{F}^2)$. Note that, in applying Theorem 3, we need that the two negligible functions implicit in the claims $(\hat{\mathcal{B}}^1, \hat{G}^1) \leq_{\mathsf{neg,pt}} (\hat{\mathcal{B}}^2, \hat{G}^2)$ and $(\hat{\mathcal{C}}^1, \hat{H}^1) \leq_{\mathsf{neg,pt}} (\hat{\mathcal{C}}^2, \hat{H}^2)$ are the same. However, since we are composing only two task-PIOA families, we can simply use the maximum of the two negligible functions. The r and s bounds follow from the fact that we are composing only two families and each of these is polynomially bounded.

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 bounded concurrent composition corollary (Corollary 1), we prove that the real protocol implements the ideal protocol with respect to the long-term implementation relation $\leq_{\mathsf{neg,pt}}$. 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, Haber's 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 alive and $\mathsf{aliveTimes}$ (cf. Section 4), we assume the following.

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

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

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

The task-PIOA code for the component $\mathsf{Dispatcher}$ appears in Figure 7. As a convention, we use σ_1, σ_2 and σ_3 to denote previous, current, and new signatures, respectively.

Concrete Time Scheme: Let d be a positive natural number. Each service j is in $\mathsf{alive}(t)$ for $t = (j-1)d, \ldots, (j+2)d-1$, so j is alive in the real time interval $[(j-1)d, (j+2)d]$. Thus, at any real time t, at most three services are concurrently alive; more precisely, t lies in the interior of the intervals for at most three services. Moreover, service j is preferred for signing for discrete times $(j-1)d, \ldots, jd-1$, that is, for real times in the interval $[(j-1)d, jd-1]$, and is usable for discrete times $(j-1)d, \ldots, (j+1)d-1$, that is, for real times in the interval $[(j-1)d, (j+1)d-1]$. Between real times $(j+1)d$ and $(j+2)d$, service 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 $\mathsf{RealSig}(j)_k = \mathsf{hide}(\mathsf{KeyGen}(k,j) \| \mathsf{Signer}(k,j) \| \mathsf{Verifier}(k,j), \mathsf{signKey}_j)$ and $\mathsf{IdealSig}(j)_k = \mathsf{hide}(\mathsf{KeyGen}(k,j) \| \mathsf{SigFunc}(k,j), \mathsf{signKey}_j)$. Here we define

$$\mathsf{Real}_k = \|_{j \in SID_k} \mathsf{RealSig}(j)_k, \mathsf{Ideal}_k = \|_{j \in SID_k} \mathsf{IdealSig}(j)_k, \text{ and}$$

$$\mathsf{RealSigSys}_k := \mathsf{Dispatcher}_k \| \mathsf{Real}_k, \mathsf{IdealSigSys}_k := \mathsf{Dispatcher}_k \| \mathsf{Ideal}_k.$$

Also, define $\overline{\mathsf{Real}} := \{\mathsf{Real}_k\}_{k \in \mathbb{N}}$, $\overline{\mathsf{Ideal}} := \{\mathsf{Ideal}_k\}_{k \in \mathbb{N}}$, $\overline{\mathsf{RealSigSys}} := \{\mathsf{RealSigSys}_k\}_{k \in \mathbb{N}}$ and $\overline{\mathsf{IdealSigSys}} := \{\mathsf{IdealSigSys}_k\}_{k \in \mathbb{N}}$.

Our goal is to show that

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

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

First, we observe that certain components of the real and ideal systems are restricted to certain time intervals, in the sense of Definition 3.

Lemma 6. *Suppose $k \in \mathbb{N}$, $j \in SID_k$. Then $\mathsf{RealSig}(j)_k$ and $\mathsf{IdealSig}(j)_k$ are restricted to $[(j-1)d, (j+2)d]$.*

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

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

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

The proof for $\mathsf{IdealSig}(j)_k$ is analogous. □

Lemma 7. *For every k, both $\mathsf{RealSig}(1)_k, \mathsf{RealSig}(2)_k, \ldots$ and $\mathsf{IdealSig}(1)_k, \mathsf{IdealSig}(2)_k, \ldots$ are 3-bounded-concurrent.*

Proof. Follows from Lemma 6.

Lemma 8. *The task-PIOA families $\overline{\mathsf{Real}}$ and $\overline{\mathsf{Ideal}}$ are polynomially bounded.*

Theorem 6. *Assume the concrete time scheme described above and assume that every signature scheme used in the timestamping protocol is complete and existentially unforgeable. Then $(\overline{\mathsf{RealSigSys}}, \emptyset) \leq_{\mathsf{neg,pt}} (\overline{\mathsf{IdealSigSys}}, \bar{F})$, where $\bar{F}_k := \bigcup_{j \in SID_k} \{\{\mathsf{fail}_j\}\}$ for every k.*

Proof. We apply Corollary 1 to the two sequences $\overline{\mathsf{RealSig}}(1), \overline{\mathsf{RealSig}}(2), \ldots$ and $\overline{\mathsf{IdealSig}}(1), \overline{\mathsf{IdealSig}}(2), \ldots$. It is easy to see that for each $j \in SID$, $\overline{\mathsf{RealSig}}(j)$ is comparable to $\overline{\mathsf{IdealSig}}(j)$, and that the needed compatibility conditions are satisfied. The number of components in Real_k is bounded by the cardinality of the set SID_k. Since SID_k is the set of $p(k)$-bit numbers for some polynomial p, the size of SID_k is bounded by some exponential in k. We use this exponential for the L bound in Corollary 1. Lemma 8 implies that conditions on the complexity bounds are met. Lemma 7 yields the needed sequences of positive reals for 3-bounded concurrency.

Theorem 1 implies that $(\overline{\mathsf{RealSig}}(j), \emptyset) \leq_{\mathsf{neg,pt}} (\overline{\mathsf{IdealSig}}(j), \{\mathsf{fail}_j\})$ for every $j \in SID$. We need a stronger statement here: that, for every pair of polynomials p and q, there exists a *single* negligible function ϵ such that $(\overline{\mathsf{RealSig}}(j), \emptyset) \leq_{p,q,\epsilon} (\overline{\mathsf{IdealSig}}(j), \{\mathsf{fail}_j\})$ for every $j \in SID$. That is, we require that the negligible function be independent of j. In our particular example, this independence follows because all of the $\overline{\mathsf{RealSig}}(j)$ are identical

except for the parameter j, and likewise for all of the $\overline{\mathsf{IdealSig}}(j)$.[8] Thus, we can apply Theorem 5, which shows that $(\overline{\mathsf{Real}}, \emptyset) \leq_{\mathsf{neg,pt}} (\overline{\mathsf{Ideal}}, \bar{F})$.

Then, we apply Theorem 3 to $\overline{\mathsf{Dispatcher}} \| \overline{\mathsf{Real}}$ and $\overline{\mathsf{Dispatcher}} \| \overline{\mathsf{Ideal}}$. In order to apply this theorem we first observe that $\overline{\mathsf{Dispatcher}}$ is comparable to $\overline{\mathsf{Dispatcher}}$, and for each $j \in SID$, $\overline{\mathsf{RealSig}}(j) \in \overline{\mathsf{Real}}$ is comparable to $\overline{\mathsf{IdealSig}}(j) \in \overline{\mathsf{Ideal}}$. Observe also that compatibility conditions are satisfied.

It is also obvious that for every pair of polynomials p and q, $(\overline{\mathsf{Dispatcher}}, \emptyset) \leq_{p,q,0} (\overline{\mathsf{Dispatcher}}, \emptyset)$, and we just showed that there is a negligible function ϵ such that $(\overline{\mathsf{Real}}, \emptyset) \leq_{p,q,\epsilon} (\overline{\mathsf{Ideal}}, \bar{F})$. The fact that each of the composed families is polynomially bounded, and that we are only composing of two of them, provides the r and s bounds and guarantees the uniformity condition (1) required for Theorem 3 (we can simply select the larger of the bounds of the individual families). Those observations are sufficient to apply Theorem 3, which yields the result.

Abstract long-lived timestamp service: It is possible to define a somewhat more abstract specification for a long-lived timestamp service—one that does not include explicit representations of individual short-lived services—and to show that our ideal level system model implements this specification, in the sense of $\leq_{\mathsf{neg,pt}}$. The abstract specification would, for example, include global sets of signing and verification keys instead of individual $mySK$ and $myVK$ variables, and a global table of issued certificates instead of individual $history$ queues. Old entries in the table that are not recertified quickly enough would be garbage-collected, in order to keep the model polynomial-bounded. Otherwise, the specification would be essentially the same as our ideal system model.

Given the close correspondence between our ideal system model and the new abstract specification, it should be straightforward to show that the two models are related by $\leq_{\mathsf{neg,pt}}$. Then transitivity of $\leq_{\mathsf{neg,pt}}$ (Lemma 4) can be used to show that our real system model also implements the new abstract specification, in the sense of $\leq_{\mathsf{neg,pt}}$.

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, $\leq_{\mathsf{neg,pt}}$, which expresses security in any polynomial-length interval of time, despite prior security violations. We have proved polynomial parallel composition and exponential sequential composition theorems for $\leq_{\mathsf{neg,pt}}$. 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 remains to carry out the details of defining a higher-level abstract functionality specification for a long-lived timestamp service, and to use $\leq_{\mathsf{neg,pt}}$ 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).

[8] In other examples, this independence might not follow, e.g., because not all of the services are identical. In such cases, we would have to add an independence assumption.

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.
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):111–159, 2008.
9. Ran Canetti, Ling Cheung, Dilsun Kaynar, Nancy Lynch, and Olivier Pereira. Modeling computational security in long-lived systems. In *Proceedings of the 19th International Conference on Concurrency Theory*, number 5201 in LNCS, pages 114–130. Springer Berlin/Heidelberg, 2008.
10. Ran Canetti, Ling Cheung, Dilsun Kaynar, Nancy Lynch, and Olivier Pereira. Modeling computational security in long-lived systems. Cryptology ePrint Archive, Report 2007/406, 2008. Available at `http://eprint.iacr.org/2007/406/`.
11. 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 LNCS, pages 255–271. Springer, 2003.
12. O. Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, 2001 (reprint of 2003).
13. S. Goldwasser and L. Levin. Fair computation of general functions in presence of immoral majority. In A. J. Menezes and S. A. Vanstone, editors, *Advances in Cryptology - Crypto '90*, number 536 in LNCS, pages 77–93. Springer-Verlag, 1990.
14. S. Haber. Long-lived digital integrity using short-lived hash functions. Technical report, HP Laboratories, May 2006.
15. 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.
16. N.A. Lynch and M.R. Tuttle. An introduction to input/output automata. *CWI Quarterly*, 2(3):219–246, September 1989.
17. M. Merritt, F. Modugno, and M.R. Tuttle. Time constrained automata. In *Proceedings of CONCUR 1991*, volume 527 of *LNCS*, pages 408–423, 1991.
18. 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.
19. 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.
20. 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.
21. 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.
22. R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2(2):250–273, 1995.

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

Signature

Input:
- $\mathsf{tick}(t : \mathbb{T}_k)$
- $\mathsf{signKey}(sk : 2^k)_j$
- $\mathsf{reqSign}(rid : RID_k, m : M_k)_j$

Output:
- $\mathsf{respSign}(rid : RID_k, \sigma : \Sigma_k)_j$

Internal:
- $\mathsf{sign}(rid : RID_k, m : M_k)_j$

Transitions

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

$\mathsf{signKey}(sk)_j$
Effect:
- if $awake = true \wedge mySK = \bot$ then $mySK := sk$

$\mathsf{reqSign}(rid, m)_j$
Effect:
- if $awake = true \wedge \neg\, \mathsf{full}(toSign)$ then $toSign := \mathsf{enq}(toSign, \langle rid, m \rangle)$

Tasks

$\mathsf{respSign}_j = \{\mathsf{respSign}(*, *)_j\}$
$\mathsf{sign}_j = \{\mathsf{sign}(*, *)_j$

States

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

$\mathsf{sign}(rid, m)_j$
local $\sigma : \Sigma$
Precondition:
- $awake = true$
- $\mathsf{head}(toSign) = \langle rid, m \rangle$
- $mySK \neq \bot$

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

$\mathsf{respSign}(rid, \sigma)_j$
Precondition:
- $awake = true$
- $head(signed) = \langle rid, \sigma \rangle$

Effect:
- $signed := \mathsf{deq}(signed)$

Fig. 3. Task-PIOA Code for $\mathsf{Signer}(k, j)$

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

Signature

Input:
- $\mathsf{tick}(t : \mathbb{T}_k)$
- $\mathsf{verKey}(vk : 2^k)_j$
- $\mathsf{reqVer}(rid : RID_k, m : M_k, \sigma : \Sigma_k)_j$

Output:
- $\mathsf{respVer}(rid : RID_k, b : Bool)_j$

Internal:
- $\mathsf{verify}(rid : RID_k, m : M_k, \sigma : \Sigma_k)_j$

Tasks

$\mathsf{respVer}_j = \{\mathsf{respVer}(*, *)_j\}$
$\mathsf{verify}_j = \{\mathsf{verify}(*, *, *)_j\}$

States

$awake : \{\mathsf{true}\}_\perp$, init $\perp$
$clock : (\mathbb{T}_k)_\perp$, init $\perp$
$myVK : (2^k)_\perp$, init $\perp$
$toVer : \mathsf{que}_k(RID_k \times M_k \times \Sigma_k)_\perp$, init $\perp$
$verified : \mathsf{que}_k(RID_k \times M_k \times \Sigma_k)_\perp$, init $\perp$

Transitions

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

$\mathsf{verKey}(vk)_j$
Effect:
- if $awake = true \wedge myVK = \perp$ then $myVK := vk$

$\mathsf{reqVer}(rid, m, \sigma)_j$
Effect:
- if $awake = true \wedge \neg\, \mathsf{full}(toVer)$ then $toVer := \mathsf{enq}(toVer, \langle rid, m, \sigma \rangle)$

$\mathsf{verify}(rid, m, \sigma)_j$
local $b : Bool$
Precondition:
- $awake = true \wedge myVK \neq \perp$
- $\mathsf{head}(toVer) = \langle rid, m, \sigma \rangle$

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

$\mathsf{respVer}(rid, b)_j$
Precondition:
- $awake = true$
- $\mathsf{head}(verified) = \langle rid, b \rangle$

Effect:
- $verified := \mathsf{deq}(verified)$

Fig. 4. Task-PIOA Code for $\mathsf{Verifier}(k, j)$

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

Signature

Input:
$I_{\mathsf{Verifier}} \cup I_{\mathsf{Signer}}$
Output:
$O_{\mathsf{Verifier}} \cup O_{\mathsf{Signer}}$
Internal:
$H_{\mathsf{Verifier}} \cup H_{\mathsf{Signer}} \cup \{\mathsf{fail}_j\}$

Transitions

Same as Signer and Verifier, except the following:

$\mathsf{tick}(t)$
Effect:
  if $j \in \mathsf{alive}(t)$ then
    $clock := t$
    if $awake = \bot$ then
      $awake := true$
      $toSign, toVer,$
      $signed, verified$
        $:=$ empty
      $history := \emptyset$
      $failed :=$ false
  else
    $awake, clock, mySK,$
    $myVK, toSign, toVer,$
    $signed, history, verified,$
    $failed := \bot$

fail_j
Precondition:
  $awake =$ true
Effect:
  $failed :=$ true

Tasks

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

States

All variables of Signer
and Verifier
$history : \mathsf{que}_k(M_k)_\bot$, init $\bot$
$failed : \{\mathsf{true}, \mathsf{false}\}_\bot$, init $\bot$

$\mathsf{sign}(rid, m)_j$
local $\sigma : \Sigma$
Precondition:
  $awake = true$
    $\wedge mySK \neq \bot$
  $\mathsf{head}(toSign) = \langle rid, m \rangle$
Effect:
  $toSign := \mathsf{deq}(toSign)$
  $\sigma := \mathsf{Sign}_j(m, mySK)$
  $signed :=$
    $\mathsf{enq}(signed, \langle rid, \sigma \rangle)$
  $history :=$
    $\mathsf{enq}(history, m)$

$\mathsf{verify}(rid, m, \sigma)_j$
Local $b : Bool$
Precondition:
  $awake = true$
    $\wedge myVK \neq \bot$
  $\mathsf{head}(toVer) = \langle rid, m, \sigma \rangle$
Effect:
  $toVer := \mathsf{deq}(toVer)$
  $b := (\mathsf{Verify}(m, \sigma, myVK)$
    $\wedge (m \in history \vee failed))$
  $verified :=$
    $\mathsf{enq}(verified, \langle rid, b \rangle)$

Fig. 5. Code for $\mathsf{SigFunc}(k, j)$

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

Signature

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

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

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

Tasks

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

States

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

Transitions

$\mathsf{tick}(t)$
Effect:
- $clock := t$

$\mathsf{reqCert}(rid, x)$
Effect:
- if $currCt < b$ then
  - $toSign := \mathsf{enq}(toSign, \langle rid, \langle x, clock, \perp\rangle\rangle)$
  - $currCt := currCt + 1$

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

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

$\mathsf{reqSign}(rid, m)_j$
Precondition:
- $\mathsf{head}(toSign) = \langle rid, m\rangle$
- $j = \mathsf{pref}(clock)$
- $\neg pendingSign$

Effect:
- $pendingSign := true$

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

$\mathsf{denyVer}(rid, op, m, \sigma_2, j)$
Precondition:
- $head(toVer) = \langle rid, op, m, \sigma_2, j\rangle$
- $j \notin \mathsf{usable}(clock)$

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

$\mathsf{reqVer}(rid, m, \sigma_2)_j$
Precondition:
- $(\exists op)(\mathsf{head}(toVer) = \langle rid, op, m, \sigma_2, j\rangle$
- $j \in \mathsf{usable}(clock)$
- $\neg pendingVer$

Effect:
- $pendingVer := true$

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

Transitions

$\mathsf{respVer}(rid, b)_j$
Effect:
  if $pendingVer$
    $\wedge(\exists op, m, \sigma_2)(head(toVer) =$
      $\langle rid, op, m, \sigma_2, j\rangle)$ then
    choose op, m, σ_2 where
      $\mathsf{head}(toVer) = \langle rid, op, m, \sigma_2, j\rangle$
    $toVer := \mathsf{deq}(toVer)$
    $pendingVer := false$
    if $op =' cert' \wedge \neg b$ then
      $certified := \mathsf{enq}(certified, \langle rid, false\rangle)$
    if $op =' cert' \wedge b$ then
      choose x, t where $(\exists \sigma_1)(m = \langle x, t, \sigma_1\rangle)$
      $toSign := \mathsf{enq}(toSign, \langle rid, \langle x, t, \sigma_2\rangle\rangle)$
    if $op =' check'$ then
      $checked := \mathsf{enq}(checked, \langle rid, b\rangle)$

$\mathsf{respCert}(rid, false)$
Precondition:
  $\mathsf{head}(certified) = \langle rid, false\rangle$
Effect:
  $certified := \mathsf{deq}(certified)$
  $currCt := currCt - 1$

$\mathsf{respCert}(rid, \sigma_3, j)$
Precondition:
  $head(certified) = \langle rid, \sigma_3, j\rangle$
Effect:
  $certified := \mathsf{deq}(certified)$
  $currCt := currCt - 1$

$\mathsf{respCheck}(rid, b)$
Precondition:
  $\mathsf{head}(checked) = \langle rid, b\rangle$
Effect:
  $checked := \mathsf{deq}(checked)$
  $currCt := currCt - 1$

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