

Task-Structured Probabilistic I/O Automata *

Ran Canetti
IBM T.J. Watson Research Center

Ling Cheung
MIT and Radboud University of Nijmegen

Dilsun Kaynar
Carnegie Mellon University
dilsun@cs.cmu.edu

Moses Liskov
College of William and Mary

Nancy Lynch
MIT

Olivier Pereira
Université catholique de Louvain

Roberto Segala
University of Verona

June 5, 2007

Abstract

Modeling frameworks such as Probabilistic I/O Automata (PIOA) and Markov Decision Processes permit both probabilistic and nondeterministic choices. In order to use these frameworks to express claims about probabilities of events, one needs mechanisms for resolving nondeterministic choices. For PIOAs, nondeterministic choices have traditionally been resolved by schedulers that have perfect information about the past execution. However, these schedulers are too powerful for certain settings, such as cryptographic protocol analysis, where information must sometimes be hidden.

Here, we propose a new, less powerful nondeterminism-resolution mechanism for PIOAs, consisting of *tasks* and *local schedulers*. Tasks are equivalence classes of system actions that are scheduled by oblivious, global task sequences. Local schedulers resolve nondeterminism within system components, based on local information only. The resulting task-PIOA framework yields simple notions of external behavior and implementation, and supports simple compositionality results. We also define a new kind of simulation relation, and show it to be sound for proving implementation. We illustrate the potential of the task-PIOA framework by outlining its use in verifying an Oblivious Transfer protocol.

*This paper presents an extension of the task-PIOA theory first introduced in [CCK⁺05a]. This extension is used in [CCK⁺05b, CCK⁺06c] to carry out a computational analysis of an Oblivious Transfer protocol. An earlier version of the current paper appears as [CCK⁺06a] and an extended abstract appears as [CCK⁺06b].

1 Introduction

The *Probabilistic I/O Automata (PIOA)* modeling framework [Seg95, SL95] is a simple combination of I/O Automata [LT89] and Markov Decision Processes (MDP) [Put94]. As demonstrated in [LSS94, SV99, PSL00], PIOAs are well suited for modeling and analyzing distributed algorithms that use randomness as a computational primitive. In this setting, distributed processes use random choices to break symmetry, in solving problems such as choice coordination [Rab82] and consensus [BO83, AH90]. Each process is modeled as an automaton with randomized transitions, and an entire protocol is modeled as the parallel composition of process automata and automata representing communication channels.

This modeling paradigm combines nondeterministic and probabilistic choices in a natural way. Nondeterminism is used here to model uncertainties in the timing of events in an unpredictable distributed environment. It is also used for modeling distributed algorithms at high levels of abstraction, leaving many details unspecified. This in turn facilitates algorithm verification, because results proven for nondeterministic algorithms apply automatically to an entire family of algorithms, obtained by resolving the nondeterministic choices in particular ways.

In order to formulate and to prove probabilistic properties of distributed algorithms, one needs mechanisms for resolving nondeterministic choices. The most common mechanism is a *perfect-information* event scheduler, which has access to local state and history of all system components and has unlimited computation power. Thus, probabilistic properties of distributed algorithms are typically asserted with respect to worst-case, adversarial schedulers who can choose the next event based on complete knowledge of the past (e.g., [PSL00, Seg95, SL95, BK98]).

One would expect that a similar modeling paradigm, including both probabilistic and nondeterministic choices, would be similarly useful for modeling *cryptographic protocols*, which are special kinds of distributed algorithms that use cryptographic primitives and that guarantee properties such as secrecy and authentication. However, the traditional probabilistic and nondeterministic modeling paradigm does not readily apply to cryptographic protocols. One major reason is that the perfect-information scheduler mechanism used for distributed algorithms is too powerful for this setting. In the presence of nondeterminism (or entropy) the schedulers that are used to resolve nondeterminism may be exploited to create additional channels of informational flow. A scheduler that has perfect information of the history would be able to access sensitive information in the states of the non-corrupted protocol participants such as their random choices, and be able to “divulge” that information to adversarial entities by encoding it in the ordering of events.

Most existing works on cryptographic protocol verification address this issue by requiring that all entities are fully specified up to inputs and coin tosses. This solution does not scale well to large protocols that handle implementation issues explicitly, because a full specification would involve many details that are irrelevant in the security analysis. In this paper, we take a different approach: instead of restricting the presence of nondeterminism, we try to find less powerful mechanisms for resolving nondeterminism. The result is an adaptation of PIOAs called *task-PIOAs*. In this new framework, cryptographic protocols may be specified with nondeterminism, yet without the danger of introducing semantic inconsistencies such as hidden information flow. Using this flexibility, one can simplify the description of practical protocols by leaving inessential choices unspecified. Furthermore, with the help of proof techniques such as probabilistic simulation, these inessential choices may remain unspecified throughout the entire correctness analysis.

Task-PIOAs A task-PIOA is simply a PIOA augmented with a partition of non-input actions into equivalence classes called *tasks*.¹ A task is typically a set of related actions that perform the “same kind of activity” in a protocol. For example, in a protocol with several rounds of message exchange, all the actions that send a round 1 message (with possibly different message contents) would constitute a task. Similarly, if a protocol involves a step in which a random choice is made from a particular domain, we could group all the actions that make a random choice from that domain (yielding possibly different values) into a task.

Tasks are units of scheduling, as for I/O automata; they are scheduled by oblivious, global *task schedule*

¹The terminology of “tasks” and “task partition” traces back to the original I/O Automata framework of Lynch and Tuttle [LT89].

sequences. We think of a task schedule as simply a way of representing the order in which different system activities happen to occur. This order can be determined, for example, by variations in speeds of different system components, or by unpredictable network delays, rather than by a purposeful scheduler entity. This simple, non-adaptive task schedule mechanism eliminates the problem of creating undesirable channels of information flow through schedulers. At first sight, it may seem that our task schedules are insufficient to describe adversarial scheduling patterns of the kind that occur in security protocols. However, we model such patterns in a different way, which we explain in detail later in the paper.

For task-PIOAs, we define notions of *external behavior* and *implementation*, by adapting the trace distribution semantics of Segala [Seg95] to task-based scheduling. We define parallel composition in the usual way and show that our implementation relation is compositional.² We also define a new type of *simulation relation*, which incorporates the notion of tasks, and show that it is sound for proving implementation relations between task-PIOAs. This new definition differs from simulation relations studied earlier [SL95, LSV03], in that it relates probability measures rather than states.

In many cases, including our work on cryptographic protocols (see below), tasks alone suffice for resolving nondeterminism. However, for extra expressive power, we define a second mechanism called *local schedulers*, which uses local information only to resolve nondeterminism within system components. This mechanism is based on earlier work in [CLSV06].

Task-PIOAs are clearly suitable for modeling and analyzing cryptographic protocols; they may also be useful for other kinds of distributed algorithms in which the perfect information assumption is unrealistically strong.

Adversarial Scheduling The standard scheduling mechanism in the cryptographic community is an *adversarial scheduler*, namely, a resource-bounded algorithmic entity that determines the next move adaptively, based on its own view of the computation so far. Clearly, this is weaker than the perfect-information scheduler used for distributed algorithms. It is however stronger than our notion of global task schedule sequences, which are essentially *oblivious schedulers* that fix the entire schedule nondeterministically in advance.

In order to capture the adaptivity of adversarial schedulers within our framework, we separate scheduling concerns into two parts.

1. The adaptive adversarial scheduler is modeled as a system component represented as a task-PIOA, for example, a message delivery service that can eavesdrop on the communications and control the order of message delivery. Such a system component has access to partial information about the execution: it sees information that other components communicate to it during execution, but not “secret information” that is found only in the internal state of these components.
2. Low-level scheduling choices are resolved by a task schedule sequence that is chosen nondeterministically in advance. For example, in a typical protocol, many different parties make independent random choices, and it is inconsequential which of them does so first. Each of these coin tosses would correspond to a single task, which does not contain any information about the actual outcome of the coin toss. A task schedule then fixes a particular order in which the different coin tosses occur.

In short, the high-level adversarial scheduler is responsible for choices that are essential in security analysis, while the low-level schedule of tasks resolves inessential choices. We believe this separation is conceptually meaningful and we illustrate it with a small example in Section 3, where an adaptive adversary uses eavesdropped information to gain advantage. This confirms that adaptive adversarial scheduling can indeed be captured in the task-PIOA framework.

Cryptographic Protocol Analysis In [CCK⁺05b, CCK⁺06c], we apply the task-PIOA framework to analyze the Oblivious Transfer (OT) protocol of Goldreich et al. [GMW87]. This framework can be decomposed into two “layers”: (i) a general foundational layer, not specific to security protocols and (ii) a security

²It is interesting to note that trace distribution semantics is not compositional under perfect-information scheduling. By moving to a less powerful scheduling mechanism, we have also found a solution to the compositionality problem.

layer that follows the general outline of simulation-based security [GMR85, GMW87, GL90, Bea91, MR91, PW94, Can95]. Task-PIOAs serve as the basis of the foundational layer. To express computational limitations, we augment task-PIOAs with additional structures, such as probabilistic Turing machines responsible for computing a next state from any given source state and action. This leads to the notions of *time-bounded task-PIOAs* and *approximate implementation* with respect to time-bounded environments. These are used, for example, to express computational indistinguishability assumptions for cryptographic primitives. Detailed definitions involving time bounds are beyond the scope of this paper. However, for those readers interested in our modeling, we provide a summary of the OT case study in Section 5. An abstract of our general approach to cryptographic protocol modeling can be found in [CLK⁺06].

We observe that, although our correctness proofs for OT are somewhat lengthy, most of the complexity lies in the establishment of simulation relations between specifications at various levels of abstraction. This is consistent with the fact that we use nondeterminism to model implementation freedom (e.g., whether to toss a coin before or after an expected input arrives). Our correctness claims are guaranteed to hold, no matter how an implementer chooses to resolve such choices. Therefore, our modeling is more general than typical models in the literature, where low-level nondeterministic choices are fixed and hard-coded into specifications (e.g. the OT model of [MMS03]).

The correctness theorem that we used in our analysis of OT, stated in terms of our approximate implementation relation, follows the idea of simulation-based security. For that analysis, we use a composition theorem that applies to a constant number of substitutions. In [CCK⁺07] we generalize this composition theorem to any polynomial number (in the security parameter) of substitutions.³ This result complements the fundamental theory of task-PIOAs presented in this paper in that it allows modular (and hence scalable) analysis of cryptographic protocols.

Aside from the OT case study, we have also proposed a novel approach for verifying statistical zero-knowledge (SZK) properties [CMP07]. This approach uses recently developed techniques based on approximate simulation relations [ML06b]. Specifically, statistical indistinguishability is formulated as an implementation relation in the Task-PIOA framework, which can then be proven using approximate simulation relations. This technique separates proof obligations into two categories: those requiring probabilistic reasoning, as well as those that do not. The latter is a good candidate for mechanization. We illustrate the general method by verifying the SZK property of the well-known identification protocol proposed by Girault, Poupard and Stern [GPS06].

1.1 Road Map

Section 2 presents mathematical preliminaries, as well as basic definitions and results for PIOAs. Some detailed constructions appear in Appendix A. Section 3 defines task-PIOAs, task schedules, composition, and implementation, and presents a compositionality result for implementation. Section 4 presents our simulation relation definition and the associated soundness theorem, as well as an example that illustrates the use of simulation relations. Section 5 summarizes our OT protocol case study. Section 6 discusses local schedulers and Section 7 treats related work in greater detail. Concluding remarks follow in Section 8.

2 Preliminaries

We write $\mathbb{R}^{\geq 0}$ and \mathbb{R}^+ for the sets of nonnegative real numbers and positive real numbers, respectively.

Let X be a set. We denote the set of finite sequences and infinite sequences of elements from X by X^* and X^ω , respectively. If ρ is a sequence then we use $|\rho|$ to denote the length of ρ . We use λ to denote the empty sequence (over any set).

If $\rho \in X^*$ and $\rho' \in X^* \cup X^\omega$, then we write $\rho \frown \rho'$ for the concatenation of the sequences ρ and ρ' . Sometimes, when no confusion seems likely, we omit the \frown symbol, writing just $\rho\rho'$.

³The polynomial composition is proven for a slightly stronger variant of our approximate implementation relation $\leq_{neg,pt}$.

2.1 Probability Measures

In this section, we first present basic definitions for probability measures. Then, we define three operations involving probability measures: *flattening*, *lifting*, and *expansion*; we will use these in Section 4 to define our new kind of simulation relation. These three operations have been previously defined in, for example, [LSV03].

Basic Definitions A σ -field over a set X is a set $\mathcal{F} \subseteq 2^X$ that contains the empty set and is closed under complement and countable union. A pair (X, \mathcal{F}) where \mathcal{F} is a σ -field over X , is called a *measurable space*. A measure on a measurable space (X, \mathcal{F}) is a function $\mu : \mathcal{F} \rightarrow [0, \infty]$ that is countably additive: for each countable family $\{X_i\}_i$ of pairwise disjoint elements of \mathcal{F} , $\mu(\cup_i X_i) = \sum_i \mu(X_i)$. A *probability measure* on (X, \mathcal{F}) is a measure on (X, \mathcal{F}) such that $\mu(X) = 1$. A *sub-probability measure* on (X, \mathcal{F}) is a measure on (X, \mathcal{F}) such that $\mu(X) \leq 1$.

A *discrete probability measure* on a set X is a probability measure μ on $(X, 2^X)$, such that, for each $C \subseteq X$, $\mu(C) = \sum_{c \in C} \mu(\{c\})$. A *discrete sub-probability measure* on a set X , is a sub-probability measure μ on $(X, 2^X)$, such that for each $C \subseteq X$, $\mu(C) = \sum_{c \in C} \mu(\{c\})$. We define $\text{Disc}(X)$ and $\text{SubDisc}(X)$ to be, respectively, the set of discrete probability measures and discrete sub-probability measures on X . In the sequel, we often omit the set notation when we refer to the measure of a singleton set.

A *support* of a probability measure μ is a measurable set C such that $\mu(C) = 1$. If μ is a discrete probability measure, then we denote by $\text{supp}(\mu)$ the set of elements that have non-zero measure (thus $\text{supp}(\mu)$ is a support of μ). We let $\delta(x)$ denote the *Dirac measure* for x , the discrete probability measure that assigns probability 1 to $\{x\}$.

Given two discrete measures μ_1, μ_2 on $(X, 2^X)$ and $(Y, 2^Y)$, respectively, we denote by $\mu_1 \times \mu_2$ the *product measure*, that is, the measure on $(X \times Y, 2^{X \times Y})$ such that $\mu_1 \times \mu_2(x, y) = \mu_1(x) \cdot \mu_2(y)$ for each $x \in X, y \in Y$.

If $\{\rho_i\}_{i \in I}$ is a countable family of measures on (X, \mathcal{F}_X) and $\{p_i\}_{i \in I}$ is a family of non-negative values, then the expression $\sum_{i \in I} p_i \rho_i$ denotes a measure ρ on (X, \mathcal{F}_X) such that, for each $C \in \mathcal{F}_X$, $\rho(C) = \sum_{i \in I} p_i \cdot \rho_i(C)$.

A function $f : X \rightarrow Y$ is said to be measurable from $(X, \mathcal{F}_X) \rightarrow (Y, \mathcal{F}_Y)$ if the inverse image of each element of \mathcal{F}_Y is an element of \mathcal{F}_X ; that is, for each $C \in \mathcal{F}_Y$, $f^{-1}(C) \in \mathcal{F}_X$. Note that, if \mathcal{F}_X is 2^X , then any function $f : X \rightarrow Y$ is measurable from $(X, \mathcal{F}_X) \rightarrow (Y, \mathcal{F}_Y)$ for any \mathcal{F}_Y .

Given measurable f from $(X, \mathcal{F}_X) \rightarrow (Y, \mathcal{F}_Y)$ and a measure μ on (X, \mathcal{F}_X) , the function $f(\mu)$ defined on \mathcal{F}_Y by $f(\mu)(C) = \mu(f^{-1}(C))$ for each $C \in \mathcal{F}_Y$ is a measure on (Y, \mathcal{F}_Y) and is called the *image measure* of μ under f . If $\mathcal{F}_X = 2^X$, $\mathcal{F}_Y = 2^Y$, and μ is a sub-probability measure, then the image measure $f(\mu)$ is a sub-probability satisfying $f(\mu)(Y) = \mu(X)$.

Flattening In this and the following two subsections, we define our three operations involving probability measures. The first operation, which we call *flattening*, takes a discrete probability measure over probability measures and “flattens” it into a single probability measure.

Definition 2.1 Let η be a discrete probability measure on $\text{Disc}(X)$. Then the flattening of η , denoted by $\text{flatten}(\eta)$, is the discrete probability measure on X defined by $\text{flatten}(\eta) = \sum_{\mu \in \text{Disc}(X)} \eta(\mu)\mu$.

Lemma 2.2 Let η be a discrete probability measure on $\text{Disc}(X)$ and let f be a function from X to Y . Then $f(\text{flatten}(\eta)) = \text{flatten}(f(\eta))$.

Proof. Recall that $\text{flatten}(\eta)$ is defined to be $\sum_{\mu \in \text{Disc}(X)} \eta(\mu)\mu$. Using the definition of image measures, it is easy to check that f distributes through the summation, so we have

$$f(\text{flatten}(\eta)) = f\left(\sum_{\mu \in \text{Disc}(X)} \eta(\mu)\mu\right) = \sum_{\mu \in \text{Disc}(X)} \eta(\mu)f(\mu) = \sum_{\sigma \in \text{Disc}(Y)} \sum_{\mu \in f^{-1}(\sigma)} \eta(\mu)\sigma.$$

Again by the definition of image measures, we have $f(\eta)(\sigma) = \eta(f^{-1}(\sigma)) = \sum_{\mu \in f^{-1}(\sigma)} \eta(\mu)$. This implies that $f(\text{flatten}(\eta))$ equals $\sum_{\sigma \in \text{Disc}(Y)} f(\eta)(\sigma)\sigma$, which is precisely $\text{flatten}(f(\eta))$. \square

Lemma 2.3 *Let $\{\eta_i\}_{i \in I}$ be a countable family of measures on $\text{Disc}(X)$, and let $\{p_i\}_{i \in I}$ be a family of probabilities such that $\sum_{i \in I} p_i = 1$. Then we have $\text{flatten}(\sum_{i \in I} p_i \eta_i) = \sum_{i \in I} p_i \text{flatten}(\eta_i)$.*

Proof. By the definition of flatten and by rearranging sums. \square

Lifting The second operation, which we call *lifting*, takes a relation R between two domains X and Y and “lifts” it to a relation between discrete measures over X and Y . Informally speaking, a measure μ_1 on X is related to a measure μ_2 on Y if μ_2 can be obtained by “redistributing” the probability masses assigned by μ_1 , in such a way that relation R is respected.

Definition 2.4 *The lifting of R , denoted by $\mathcal{L}(R)$, is the relation from $\text{Disc}(X)$ to $\text{Disc}(Y)$ defined by: $\mu_1 \mathcal{L}(R) \mu_2$ iff there exists a weighting function $w : X \times Y \rightarrow \mathbf{R}^{\geq 0}$ such that the following hold:*

1. For each $x \in X$ and $y \in Y$, $w(x, y) > 0$ implies $x R y$.
2. For each $x \in X$, $\sum_{y \in Y} w(x, y) = \mu_1(x)$.
3. For each $y \in Y$, $\sum_{x \in X} w(x, y) = \mu_2(y)$.

Expansion Finally, we define our third operation, called *expansion*. Expansion is defined in terms of flattening and lifting, and is used directly in our new definition of simulation relations. The *expansion* operation takes a relation between discrete measures on two domains X and Y , and returns a relation of the same kind that relates two measures whenever they can be decomposed into two $\mathcal{L}(R)$ -related measures.

Definition 2.5 *Let R be a relation from $\text{Disc}(X)$ to $\text{Disc}(Y)$. The expansion of R , denoted by $\mathcal{E}(R)$, is a relation from $\text{Disc}(X)$ to $\text{Disc}(Y)$. It is defined by: $\mu_1 \mathcal{E}(R) \mu_2$ iff there exist two discrete measures η_1 and η_2 on $\text{Disc}(X)$ and $\text{Disc}(Y)$, respectively, such that the following hold:*

1. $\mu_1 = \text{flatten}(\eta_1)$.
2. $\mu_2 = \text{flatten}(\eta_2)$.
3. $\eta_1 \mathcal{L}(R) \eta_2$.

Informally speaking, we enlarge R by adding pairs of measures that can be “decomposed” into weighted sums of measures, in such a way that the weights can be “redistributed” in an R -respecting manner. Taking this intuition one step further, the following lemma provides a useful characterization of the expansion relation.

Lemma 2.6 *Let R be a relation on $\text{Disc}(X) \times \text{Disc}(Y)$. Then $\mu_1 \mathcal{E}(R) \mu_2$ iff there exists a countable index set I , a discrete probability measure p on I , and two collections of probability measures, $\{\mu_{1,i}\}_I$ and $\{\mu_{2,i}\}_I$, such that*

1. $\mu_1 = \sum_{i \in I} p(i) \mu_{1,i}$.
2. $\mu_2 = \sum_{i \in I} p(i) \mu_{2,i}$.
3. For each $i \in I$, $\mu_{1,i} R \mu_{2,i}$.

Proof. Suppose that $\mu_1 \mathcal{E}(R) \mu_2$, and let η_1, η_2 and w be the measures and weighting function used in the definition of $\mathcal{E}(R)$. Let $\{(\mu_{1,i}, \mu_{2,i})\}_{i \in I}$ be an enumeration of the pairs for which $w(\mu_{1,i}, \mu_{2,i}) > 0$, and let $p(i)$ be $w(\mu_{1,i}, \mu_{2,i})$. Then $p, \{(\mu_{1,i})\}_{i \in I}$, and $\{(\mu_{2,i})\}_{i \in I}$ satisfy Items 1, 2, and 3.

Conversely, given $p, \{(\mu_{1,i})\}_{i \in I}$, and $\{(\mu_{2,i})\}_{i \in I}$, we define $\eta_1(\mu)$ to be the sum $\sum_{i|\mu=\mu_{1,i}} p(i)$ and $\eta_2(\mu)$ to be $\sum_{i|\mu=\mu_{2,i}} p(i)$. Moreover, define $w(\mu'_1, \mu'_2)$ to be $\sum_{i|\mu'_1=\mu_{1,i}, \mu'_2=\mu_{2,i}} p(i)$. Then, η_1, η_2 and w satisfy the properties required in the definition of $\mathcal{E}(R)$. \square

The next, rather technical lemma gives us a sufficient condition for showing that a pair of functions f and g preserve the relation $\mathcal{E}(R)$; that is, if $\mu_1 \mathcal{E}(R) \mu_2$, then $f(\mu_1) \mathcal{E}(R) g(\mu_2)$. The required condition is that, when μ_1 and μ_2 are decomposed into weighted sums of measures as in the definition of $\mu_1 \mathcal{E}(R) \mu_2$, f and g convert each pair (ρ_1, ρ_2) of R -related probability measures to $\mathcal{E}(R)$ -related probability measures. We will use this lemma in the soundness proof for our new kind of simulation relation (Lemma 4.6), where the two functions f and g apply corresponding sequences of tasks to corresponding measures on executions.

Lemma 2.7 *Let R be a relation from $\text{Disc}(X)$ to $\text{Disc}(Y)$, and let f, g be two endo-functions on $\text{Disc}(X)$ and $\text{Disc}(Y)$, respectively. Suppose that f distributes over convex combinations of measures; that is, for each countable family $\{\rho_i\}_i$ of discrete measures on X and each countable family of probabilities $\{p_i\}_i$ such that $\sum_i p_i = 1$, $f(\sum_i p_i \rho_i) = \sum_i p_i f(\rho_i)$. Similarly for g . Let μ_1 and μ_2 be measures on X and Y , respectively, such that $\mu_1 \mathcal{E}(R) \mu_2$. Let η_1, η_2 , and w be a pair of measures and a weighting function witnessing the fact that $\mu_1 \mathcal{E}(R) \mu_2$. Suppose further that, for any two distributions $\rho_1 \in \text{supp}(\eta_1)$ and $\rho_2 \in \text{supp}(\eta_2)$ with $w(\rho_1, \rho_2) > 0$, we have $f(\rho_1) \mathcal{E}(R) g(\rho_2)$. Then $f(\mu_1) \mathcal{E}(R) g(\mu_2)$.*

Proof. Let W denote the set of pairs (ρ_1, ρ_2) such that $w(\rho_1, \rho_2) > 0$. Note that, by the definition of lifting, $(\rho_1, \rho_2) \in W$ implies $\rho_1 \in \text{supp}(\eta_1)$ and $\rho_2 \in \text{supp}(\eta_2)$. Therefore, by assumption, we have $f(\rho_1) \mathcal{E}(R) g(\rho_2)$ whenever $(\rho_1, \rho_2) \in W$.

Now, for each $(\rho_1, \rho_2) \in W$, choose a pair of measures $(\eta_1)_{\rho_1, \rho_2}, (\eta_2)_{\rho_1, \rho_2}$ and a weighting function w_{ρ_1, ρ_2} as guaranteed by the definition of $f(\rho_1) \mathcal{E}(R) g(\rho_2)$. Let $\eta'_1 = \sum_{(\rho_1, \rho_2) \in W} w(\rho_1, \rho_2) (\eta_1)_{\rho_1, \rho_2}$ and let $\eta'_2 = \sum_{(\rho_1, \rho_2) \in W} w(\rho_1, \rho_2) (\eta_2)_{\rho_1, \rho_2}$. Let $w' = \sum_{(\rho_1, \rho_2) \in W} w(\rho_1, \rho_2) w_{\rho_1, \rho_2}$.

We show that η'_1, η'_2 , and w' satisfy the conditions for $f(\mu_1) \mathcal{E}(R) g(\mu_2)$.

1. $f(\mu_1) = \text{flatten}(\eta'_1)$.

By the definition of η'_1 , $\text{flatten}(\eta'_1) = \text{flatten}(\sum_{(\rho_1, \rho_2) \in W} w(\rho_1, \rho_2) (\eta_1)_{\rho_1, \rho_2})$. By Lemma 2.3, this is in turn equal to $\sum_{(\rho_1, \rho_2) \in W} w(\rho_1, \rho_2) \text{flatten}((\eta_1)_{(\rho_1, \rho_2)})$. By the choice of $(\eta_1)_{(\rho_1, \rho_2)}$, we know that $\text{flatten}((\eta_1)_{(\rho_1, \rho_2)}) = f(\rho_1)$, so we obtain that $\text{flatten}(\eta'_1) = \sum_{(\rho_1, \rho_2) \in W} w(\rho_1, \rho_2) f(\rho_1)$.

We claim that the right side is equal to $f(\mu_1)$: Since $\mu_1 = \text{flatten}(\eta_1)$, by the definition of flattening, $\mu_1 = \sum_{\rho_1 \in \text{Disc}(X)} \eta_1(\rho_1) \rho_1$. Then, by distributivity of f , $f(\mu_1) = \sum_{\rho_1 \in \text{Disc}(X)} \eta_1(\rho_1) f(\rho_1)$. By definition of lifting, $\eta_1(\rho_1) = \sum_{\rho_2 \in \text{Disc}(Y)} w(\rho_1, \rho_2)$.

Therefore, $f(\mu_1) = \sum_{\rho_1 \in \text{Disc}(X)} \sum_{\rho_2 \in \text{Disc}(Y)} w(\rho_1, \rho_2) f(\rho_1)$, and this last expression is equal to $\sum_{(\rho_1, \rho_2) \in W} w(\rho_1, \rho_2) f(\rho_1)$, as needed.

2. $g(\mu_2) = \text{flatten}(\eta'_2)$.

Analogous to the previous case.

3. $\eta'_1 \mathcal{L}(R) \eta'_2$ using w' as a weighting function.

We verify that w' satisfies the three conditions in the definition of a weighting function:

- (a) Let ρ'_1, ρ'_2 be such that $w'(\rho'_1, \rho'_2) > 0$. Then, by definition of w' , there exists at least one pair $(\rho_1, \rho_2) \in R$ such that $w_{\rho_1, \rho_2}(\rho'_1, \rho'_2) > 0$. Since w_{ρ_1, ρ_2} is a weighting function, $\rho'_1 R \rho'_2$ as needed.

(b) By the definition of w' , we have

$$\begin{aligned}
\sum_{\rho'_2 \in \text{Disc}(Y)} w'(\rho'_1, \rho'_2) &= \sum_{\rho'_2 \in \text{Disc}(Y)} \sum_{(\rho_1, \rho_2) \in W} w(\rho_1, \rho_2) w_{\rho_1, \rho_2}(\rho'_1, \rho'_2) \\
&= \sum_{(\rho_1, \rho_2) \in W} \sum_{\rho'_2 \in \text{Disc}(Y)} w(\rho_1, \rho_2) w_{\rho_1, \rho_2}(\rho'_1, \rho'_2) \\
&= \sum_{(\rho_1, \rho_2) \in W} (w(\rho_1, \rho_2) \cdot \sum_{\rho'_2 \in \text{Disc}(Y)} w_{\rho_1, \rho_2}(\rho'_1, \rho'_2)).
\end{aligned}$$

Since w_{ρ_1, ρ_2} is a weighting function, we also have $\sum_{\rho'_2 \in \text{Disc}(Y)} w_{\rho_1, \rho_2}(\rho'_1, \rho'_2) = (\eta_1)_{\rho_1, \rho_2}(\rho'_1)$. This implies $\sum_{\rho'_2 \in \text{Disc}(Y)} w'(\rho'_1, \rho'_2)$ equals $\sum_{(\rho_1, \rho_2)} w(\rho_1, \rho_2) (\eta_1)_{\rho_1, \rho_2}(\rho'_1)$, which is precisely $\eta'_1(\rho'_1)$.

(c) Symmetric to the previous case. □

2.2 Probabilistic I/O Automata

In this subsection, we review basic definitions for Probabilistic I/O Automata.

PIOAs A *probabilistic I/O automaton (PIOA)*, \mathcal{P} , is a tuple (Q, \bar{q}, I, O, H, D) where:

- Q is a countable set of *states*, with *start state* $\bar{q} \in Q$;
- I, O and H are countable and pairwise disjoint sets of actions, referred to as *input, output and internal (hidden) actions*, respectively; and
- $D \subseteq (Q \times (I \cup O \cup H) \times \text{Disc}(Q))$ is a *transition relation*, where $\text{Disc}(Q)$ is the set of discrete probability measures on Q .

An action a is *enabled* in a state q if $(q, a, \mu) \in D$ for some μ . The set $A := I \cup O \cup H$ is called the *action alphabet* of \mathcal{P} . If $I = \emptyset$, then \mathcal{P} is *closed*. The set of *external actions* of \mathcal{P} is $E := I \cup O$, and the set of *locally controlled actions* is $L := O \cup H$.

We assume that \mathcal{P} satisfies the following conditions:

- *Input enabling*: For every state $q \in Q$ and input action $a \in I$, a is enabled in q .
- *Transition determinism*: For every $q \in Q$ and $a \in A$, there is at most one $\mu \in \text{Disc}(Q)$ such that $(q, a, \mu) \in D$. If there is exactly one such μ , it is denoted by $\mu_{q,a}$, and we write $\text{tran}_{q,a}$ for the transition $(q, a, \mu_{q,a})$.

A (non-probabilistic) *execution fragment* of \mathcal{P} is a finite or infinite sequence $\alpha = q_0 a_1 q_1 a_2 \dots$ of alternating states and actions, such that:

- If α is finite, then it ends with a state.
- For every non-final i , there is a transition $(q_i, a_{i+1}, \mu) \in D$ with $q_{i+1} \in \text{supp}(\mu)$.

We write $\text{fstate}(\alpha)$ for q_0 , and, if α is finite, we write $\text{lstate}(\alpha)$ for the last state of α . We use $\text{Frag}(\mathcal{P})$ (resp., $\text{Frag}^*(\mathcal{P})$) to denote the set of all (resp., all finite) execution fragments of \mathcal{P} . An *execution* of \mathcal{P} is an execution fragment beginning from the start state \bar{q} . $\text{Exec}(\mathcal{P})$ (resp., $\text{Exec}^*(\mathcal{P})$) denotes the set of all (resp., finite) executions of \mathcal{P} .

The *trace* of an execution fragment α , written $\text{trace}(\alpha)$, is the restriction of α to the set of external actions of \mathcal{P} . We say that β is a *trace* of \mathcal{P} if there is an execution α of \mathcal{P} with $\text{trace}(\alpha) = \beta$. The symbol \leq denotes the prefix relation on sequences, which applies in particular to execution fragments and traces.

Schedulers and Probabilistic Executions Nondeterministic choices in \mathcal{P} are resolved using a *scheduler*:

Definition 2.8 A scheduler for \mathcal{P} is a function $\sigma : \text{Frag}^*(\mathcal{P}) \rightarrow \text{SubDisc}(D)$ such that $(q, a, \mu) \in \text{supp}(\sigma(\alpha))$ implies $q = \text{lstate}(\alpha)$.

Thus, σ decides (probabilistically) which transition (if any) to take after each finite execution fragment α . Since this decision is a discrete sub-probability measure, it may be the case that σ chooses to *halt* after α with non-zero probability: $1 - \sigma(\alpha)(D) > 0$.

A scheduler σ and a finite execution fragment α generate a measure $\epsilon_{\sigma, \alpha}$ on the σ -field $\mathcal{F}_{\mathcal{P}}$ generated by cones of execution fragments, where the cone $C_{\alpha'}$ of a finite execution fragment α' is the set of execution fragments that have α' as a prefix. The construction of the σ -field is standard and is presented in Appendix A.

Definition 2.9 The measure of a cone, $\epsilon_{\sigma, \alpha}(C_{\alpha'})$, is defined recursively, as:

1. 0, if $\alpha' \not\leq \alpha$ and $\alpha \not\leq \alpha'$;
2. 1, if $\alpha' \leq \alpha$; and
3. $\epsilon_{\sigma, \alpha}(C_{\alpha'}) \mu_{\sigma(\alpha'')}(a, q)$, if α' is of the form $\alpha'' a q$ and $\alpha \leq \alpha''$. Here, $\mu_{\sigma(\alpha'')}(a, q)$ is defined to be $\sigma(\alpha'')(\text{tran}_{\text{lstate}(\alpha''), a}) \mu_{\text{lstate}(\alpha''), a}(q)$, that is, the probability that $\sigma(\alpha'')$ chooses a transition labeled by a and that the new state is q .

Standard measure theoretic arguments ensure that $\epsilon_{\sigma, \alpha}$ is well-defined. Note also that $\epsilon_{\sigma, \alpha}$ is a probability measure. We call the state $\text{fstate}(\alpha)$ the *first state* of $\epsilon_{\sigma, \alpha}$ and denote it by $\text{fstate}(\epsilon_{\sigma, \alpha})$. If α consists of the start state \bar{q} only, we call $\epsilon_{\sigma, \alpha}$ a *probabilistic execution* of \mathcal{P} .

Let μ be a discrete probability measure over $\text{Frag}^*(\mathcal{P})$. We denote by $\epsilon_{\sigma, \mu}$ the measure $\sum_{\alpha} \mu(\alpha) \epsilon_{\sigma, \alpha}$ and we say that $\epsilon_{\sigma, \mu}$ is *generated* by σ and μ . We call the measure $\epsilon_{\sigma, \mu}$ a *generalized probabilistic execution fragment* of \mathcal{P} . If every execution fragment in $\text{supp}(\mu)$ consists of a single state, then we call $\epsilon_{\sigma, \mu}$ a *probabilistic execution fragment* of \mathcal{P} .

We note that the trace function is a measurable function from $\mathcal{F}_{\mathcal{P}}$ to the σ -field generated by cones of traces. Thus, given a probability measure ϵ on $\mathcal{F}_{\mathcal{P}}$, we define the *trace distribution* of ϵ , denoted $\text{tdist}(\epsilon)$, to be the image measure of ϵ under trace. We extend the $\text{tdist}()$ notation to arbitrary measures on execution fragments of \mathcal{P} . We denote by $\text{tdists}(\mathcal{P})$ the set of trace distributions of (probabilistic executions of) \mathcal{P} .

Next we present some basic results about probabilistic executions and trace distributions of PIOAs. In particular, Lemmas 2.10-2.14 give some useful equations involving the probabilities of various sets of execution fragments.

Lemma 2.10 Let σ be a scheduler for PIOA \mathcal{P} , μ be a discrete probability measure on finite execution fragments of \mathcal{P} , and α be a finite execution fragment of \mathcal{P} . Then

$$\epsilon_{\sigma, \mu}(C_{\alpha}) = \mu(C_{\alpha}) + \sum_{\alpha' < \alpha} \mu(\alpha') \epsilon_{\sigma, \alpha'}(C_{\alpha}).$$

Proof. By definition of $\epsilon_{\sigma, \mu}$, $\epsilon_{\sigma, \mu}(C_{\alpha}) = \sum_{\alpha'} \mu(\alpha') \epsilon_{\sigma, \alpha'}(C_{\alpha})$. Since, by definition, $\epsilon_{\sigma, \alpha'}(C_{\alpha}) = 1$ whenever $\alpha \leq \alpha'$, this can be rewritten as

$$\epsilon_{\sigma, \mu}(C_{\alpha}) = \sum_{\alpha': \alpha \leq \alpha'} \mu(\alpha') + \sum_{\alpha' < \alpha} \mu(\alpha') \epsilon_{\sigma, \alpha'}(C_{\alpha}).$$

Observe that $\sum_{\alpha': \alpha \leq \alpha'} \mu(\alpha') = \mu(C_{\alpha})$. Thus, by substitution, we get the statement of the lemma. \square

Lemma 2.11 Let σ be a scheduler for PIOA \mathcal{P} , μ be a discrete probability measure on finite execution fragments of \mathcal{P} , and α be a finite execution fragment of \mathcal{P} . Then

$$\epsilon_{\sigma, \mu}(C_{\alpha}) = \mu(C_{\alpha} - \{\alpha\}) + \sum_{\alpha' \leq \alpha} \mu(\alpha') \epsilon_{\sigma, \alpha'}(C_{\alpha}).$$

Proof. Follows directly from Lemma 2.10 after observing that $\epsilon_{\sigma,\alpha}(C_\alpha) = 1$. \square

Lemma 2.12 *Let σ be a scheduler for PIOA \mathcal{P} , and μ be a discrete measure on finite execution fragments of \mathcal{P} . Let $\alpha = \tilde{\alpha}aq$ be a finite execution fragment of \mathcal{P} . Then*

$$\epsilon_{\sigma,\mu}(C_\alpha) = \mu(C_\alpha) + (\epsilon_{\sigma,\mu}(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})) \sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha},a})\mu_{\tilde{\alpha},a}(q).$$

Proof. By Lemma 2.10 and the definitions of $\epsilon_{\sigma,\alpha'}(C_\alpha)$ and $\mu_{\sigma(\tilde{\alpha})}(a, q)$, we have

$$\begin{aligned} \epsilon_{\sigma,\mu}(C_\alpha) &= \mu(C_\alpha) + \sum_{\alpha' < \alpha} \mu(\alpha') \epsilon_{\sigma,\alpha'}(C_{\tilde{\alpha}}) \sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha},a})\mu_{\tilde{\alpha},a}(q) \\ &= \mu(C_\alpha) + \left(\sum_{\alpha' < \alpha} \mu(\alpha') \epsilon_{\sigma,\alpha'}(C_{\tilde{\alpha}}) \right) (\sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha},a})\mu_{\tilde{\alpha},a}(q)). \end{aligned}$$

Since $\alpha' \leq \tilde{\alpha}$ if and only if $\alpha' < \alpha$, this yields

$$\epsilon_{\sigma,\mu}(C_\alpha) = \mu(C_\alpha) + \left(\sum_{\alpha' \leq \tilde{\alpha}} \mu(\alpha') \epsilon_{\sigma,\alpha'}(C_{\tilde{\alpha}}) \right) (\sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha},a})\mu_{\tilde{\alpha},a}(q)).$$

It suffices to show that $\sum_{\alpha' \leq \tilde{\alpha}} \mu(\alpha') \epsilon_{\sigma,\alpha'}(C_{\tilde{\alpha}}) = \epsilon_{\sigma,\mu}(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})$. But this follows immediately from Lemma 2.11 (with α instantiated as $\tilde{\alpha}$). \square

As a notational convention we introduce a new symbol \perp to denote termination. Given scheduler σ and finite execution fragment α , we write $\sigma(\alpha)(\perp)$ for the probability of terminating after α (namely, $1 - \sigma(\alpha)(D)$).

Lemma 2.13 *Let σ be a scheduler for PIOA \mathcal{P} , μ be a discrete probability measure on finite execution fragments of \mathcal{P} , and α be a finite execution fragment of \mathcal{P} . Then*

$$\epsilon_{\sigma,\mu}(\alpha) = (\epsilon_{\sigma,\mu}(C_\alpha) - \mu(C_\alpha - \{\alpha\})) (\sigma(\alpha)(\perp)).$$

Proof. By definition of $\epsilon_{\sigma,\mu}$, $\epsilon_{\sigma,\mu}(\alpha) = \sum_{\alpha'} \mu(\alpha') \epsilon_{\sigma,\alpha'}(\alpha)$. The sum can be restricted to $\alpha' \leq \alpha$ since for all other α' , $\epsilon_{\sigma,\alpha'}(\alpha) = 0$. Then, since for each $\alpha' \leq \alpha$, $\epsilon_{\sigma,\alpha'}(\alpha) = \epsilon_{\sigma,\alpha'}(C_\alpha) \sigma(\alpha)(\perp)$, we derive $\epsilon_{\sigma,\mu}(\alpha) = \sum_{\alpha' \leq \alpha} \mu(\alpha') \epsilon_{\sigma,\alpha'}(C_\alpha) \sigma(\alpha)(\perp)$. Observe that $\sigma(\alpha)(\perp)$ is a constant with respect to α' , and thus can be moved out of the sum, yielding $\epsilon_{\sigma,\mu}(\alpha) = (\sum_{\alpha' \leq \alpha} \mu(\alpha') \epsilon_{\sigma,\alpha'}(C_\alpha)) (\sigma(\alpha)(\perp))$.

It suffices to show that $\sum_{\alpha' \leq \alpha} \mu(\alpha') \epsilon_{\sigma,\alpha'}(C_\alpha) = \epsilon_{\sigma,\mu}(C_\alpha) - \mu(C_\alpha - \{\alpha\})$. But this follows immediately from Lemma 2.11. \square

Lemma 2.14 *Let σ be a scheduler for PIOA \mathcal{P} , and μ be a discrete probability measure on finite execution fragments of \mathcal{P} . Let α be a finite execution fragment of \mathcal{P} and a be an action of \mathcal{P} that is enabled in $\text{lstate}(\alpha)$. Then*

$$\epsilon_{\sigma,\mu}(C_{\alpha a}) = \mu(C_{\alpha a}) + (\epsilon_{\sigma,\mu}(C_\alpha) - \mu(C_\alpha - \{\alpha\})) \sigma(\alpha)(\text{tran}_{\alpha,a}).$$

Proof. Observe that $C_{\alpha a} = \cup_q C_{\alpha a q}$, where the cones $C - \alpha a q$ are pairwise disjoint. Thus, $\epsilon_{\sigma,\mu}(C_{\alpha a}) = \sum_q \epsilon_{\sigma,\mu}(C_{\alpha a q})$. By Lemma 2.12, the right-hand side is equal to

$$\sum_q (\mu(C_{\alpha a q}) + (\epsilon_{\sigma,\mu}(C_\alpha) - \mu(C_\alpha - \{\alpha\})) \sigma(\alpha)(\text{tran}_{\alpha,a})\mu_{\alpha,a}(q)).$$

Since $\sum_q \mu(C_{\alpha a q}) = \mu(C_{\alpha a})$ and $\sum_q \mu_{\alpha,a}(q) = 1$, this is in turn equal to

$$\mu(C_{\alpha a}) + (\epsilon_{\sigma,\mu}(C_\alpha) - \mu(C_\alpha - \{\alpha\})) \sigma(\alpha)(\text{tran}_{\alpha,a}).$$

Combining the equations yields the result. \square

Finally, we present a result about limits of generalized probabilistic execution fragments.

Definition 2.15 Let ϵ and ϵ' be measures on execution fragments of PIOA \mathcal{P} . Then we say that ϵ is a prefix of ϵ' , denoted by $\epsilon \leq \epsilon'$, if, for each finite execution fragment α of \mathcal{P} , $\epsilon(C_\alpha) \leq \epsilon'(C_\alpha)$.

Proposition 2.16 Let $\epsilon_1 \leq \epsilon_2 \leq \dots$ be a chain of generalized probabilistic execution fragments of a PIOA \mathcal{P} , all generated from the same discrete probability measure μ on finite execution fragments. Then $\lim_{i \rightarrow \infty} \epsilon_i$ is a generalized probabilistic execution fragment of \mathcal{P} generated from μ .

Proof. Let ϵ denote $\lim_{i \rightarrow \infty} \epsilon_i$. For each $i \geq 1$, let σ_i be a scheduler such that $\epsilon_i = \epsilon_{\sigma_i, \mu}$, and for each finite execution fragment α , let $p_\alpha^i = \epsilon_{\sigma_i, \mu}(C_\alpha) - \mu(C_\alpha - \{\alpha\})$. For each finite execution fragment α and each action a , let $p_{\alpha a}^i = \epsilon_{\sigma_i, \mu}(C_{\alpha a}) - \mu(C_{\alpha a})$.

By Lemma 2.14, if a is enabled in $\text{lstate}(\alpha)$ then $p_\alpha^i \sigma_i(\alpha)(\text{tran}_{\alpha, a}) = p_{\alpha a}^i$. Moreover, if $p_{\alpha a}^i \neq 0$, then $\sigma_i(\alpha)(\text{tran}_{\alpha, a}) = p_{\alpha a}^i / p_\alpha^i$.

For each finite execution fragment α , let $p_\alpha = \epsilon(C_\alpha) - \mu(C_\alpha - \{\alpha\})$. For each finite execution fragment α and each action a , let $p_{\alpha a} = \epsilon(C_{\alpha a}) - \mu(C_{\alpha a})$. Define $\sigma(\alpha)(\text{tran}_{\alpha, a})$ to be $p_{\alpha a} / p_\alpha$ if $p_\alpha > 0$; otherwise define $\sigma(\alpha)(\text{tran}_{\alpha, a}) = 0$. By definition of ϵ and simple manipulations, $\lim_{i \rightarrow \infty} p_\alpha^i = p_\alpha$ and $\lim_{i \rightarrow \infty} p_{\alpha a}^i = p_{\alpha a}$. It follows that, if $p_\alpha > 0$, then $\sigma(\alpha)(\text{tran}_{\alpha, a}) = \lim_{i \rightarrow \infty} \sigma_i(\alpha)(\text{tran}_{\alpha, a})$.

It remains to show that σ is a scheduler and that $\epsilon_{\sigma, \mu} = \epsilon$. To show that σ is a scheduler, we must show that, for each finite execution fragment α , $\sigma(\alpha)$ is a sub-probability measure. Observe that, for each $i \geq 1$, $\sum_{\text{tran}} \sigma_i(\alpha)(\text{tran}) = \sum_a \sigma_i(\alpha)(\text{tran}_{\alpha a})$. Similarly, $\sum_{\text{tran}} \sigma(\alpha)(\text{tran}) = \sum_a \sigma(\alpha)(\text{tran}_{\alpha a})$. Since each σ_i is a scheduler, it follows that, for each $i \geq 0$, $\sum_a \sigma_i(\alpha)(\text{tran}_{\alpha a}) \leq 1$. Thus,

$$\lim_{i \rightarrow \infty} \sum_a \sigma_i(\alpha)(\text{tran}_{\alpha a}) \leq \sum_a \lim_{i \rightarrow \infty} \sigma_i(\alpha)(\text{tran}_{\alpha a}) \leq 1.$$

We claim that $\sigma(\alpha)(\text{tran}_{\alpha, a}) \leq \lim_{i \rightarrow \infty} \sigma_i(\alpha)(\text{tran}_{\alpha, a})$, which implies that $\sigma(\alpha)(\text{tran}_{\alpha a}) \leq 1$, as needed. To see this claim, we consider two cases: If $p_\alpha > 0$, then as shown earlier, $\sigma(\alpha)(\text{tran}_{\alpha, a}) = \lim_{i \rightarrow \infty} \sigma_i(\alpha)(\text{tran}_{\alpha, a})$. On the other hand, if $p_\alpha = 0$, then $\sigma(\alpha)(\text{tran}_{\alpha, a})$ is defined to be zero, so that $\sigma(\alpha)(\text{tran}_{\alpha, a}) = 0 \leq \lim_{i \rightarrow \infty} \sigma_i(\alpha)(\text{tran}_{\alpha, a})$.

To show that $\epsilon_{\sigma, \mu} = \epsilon$, we show by induction on the length of a finite execution fragment α that $\epsilon_{\sigma, \mu}(C_\alpha) = \epsilon(C_\alpha)$. For the base case, let α consist of a single state q . By Lemma 2.10, $\epsilon_{\sigma, \mu}(C_q) = \mu(C_q)$, and for each $i \geq 1$, $\epsilon_{\sigma_i, \mu}(C_q) = \mu(C_q)$. Thus, $\epsilon(C_q) = \lim_{i \rightarrow \infty} \epsilon_{\sigma_i, \mu}(C_q) = \mu(C_q)$, as needed.

For the inductive step, let $\alpha = \tilde{\alpha} a q$. By Lemma 2.12,

$$\lim_{i \rightarrow \infty} \epsilon_{\sigma_i, \mu}(C_\alpha) = \lim_{i \rightarrow \infty} (\mu(C_\alpha) + (\epsilon_{\sigma_i, \mu}(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})) \sigma_i(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) \mu_{\tilde{\alpha}, a}(q)).$$

Observe that the left-hand side is $\epsilon(C_\alpha)$. By algebraic manipulation, the right-hand side becomes

$$\mu(C_\alpha) + \left(\left(\lim_{i \rightarrow \infty} \epsilon_{\sigma_i, \mu}(C_{\tilde{\alpha}}) \right) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) \right) \left(\lim_{i \rightarrow \infty} \sigma_i(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) \right) \mu_{\tilde{\alpha}, a}(q).$$

By definition of ϵ , $\lim_{i \rightarrow \infty} \epsilon_{\sigma_i, \mu}(C_{\tilde{\alpha}}) = \epsilon(C_{\tilde{\alpha}})$, and by inductive hypothesis, $\epsilon(C_{\tilde{\alpha}}) = \epsilon_{\sigma, \mu}(C_{\tilde{\alpha}})$. Therefore,

$$\epsilon(C_\alpha) = \mu(C_\alpha) + (\epsilon_{\sigma, \mu}(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})) \left(\lim_{i \rightarrow \infty} \sigma_i(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) \right) \mu_{\tilde{\alpha}, a}(q).$$

Also by Lemma 2.12, we obtain that

$$\epsilon_{\sigma, \mu}(C_\alpha) = \mu(C_\alpha) + (\epsilon_{\sigma, \mu}(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})) \sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) \mu_{\tilde{\alpha}, a}(q).$$

We claim that the right-hand sides of the last two equations are equal. To see this, consider two cases. First, if $p_{\tilde{\alpha}} > 0$, then we have already shown that $\lim_{i \rightarrow \infty} \sigma_i(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) = \sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a})$. Since these two terms are the only difference between the two expressions, the expressions are equal.

On the other hand, if $p_{\tilde{\alpha}} = 0$, then by definition of $p_{\tilde{\alpha}}$, we get that $\epsilon(C_{\tilde{\alpha}}) = \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})$. Then by the induction hypothesis the second terms of the two right-hand sides are both equal to zero, which implies that both expressions are equal to the first term $\mu(C_\alpha)$. Again, the two right-hand sides are equal.

Since the right-hand sides are equal, so are the left-hand sides, that is, $\epsilon_{\sigma, \mu}(C_\alpha) = \epsilon(C_\alpha)$, as needed to complete the inductive step. \square

Composition We define composition of PIOAs as follows.

Definition 2.17 Two PIOAs $\mathcal{P}_i = (Q_i, \bar{q}_i, I_i, O_i, H_i, D_i)$, $i \in \{1, 2\}$, are said to be compatible if $A_i \cap H_j = O_i \cap O_j = \emptyset$ whenever $i \neq j$. In that case, we define their composition $\mathcal{P}_1 \parallel \mathcal{P}_2$ to be the PIOA $(Q_1 \times Q_2, (\bar{q}_1, \bar{q}_2), (I_1 \cup I_2) \setminus (O_1 \cup O_2), O_1 \cup O_2, H_1 \cup H_2, D)$, where D is the set of triples $((q_1, q_2), a, \mu_1 \times \mu_2)$ such that

1. a is enabled in some q_i .
2. For every i , if $a \in A_i$ then $(q_i, a, \mu_i) \in D_i$, otherwise $\mu_i = \delta(q_i)$.

Given a state $q = (q_1, q_2)$ in the composition and $i \in \{1, 2\}$, we use $q \upharpoonright \mathcal{P}_i$ to denote q_i .

The definition of composition can be extended to any finite number of PIOAs rather than just two. Note that if an input of one PIOA is an output of another, then it becomes an output action of the composed automaton.

Hiding We define a hiding operation for PIOAs, which hides output actions.

Definition 2.18 Let $\mathcal{P} = (Q, \bar{q}, I, O, H, D)$ be a PIOA and let $S \subseteq O$. Then $\text{hide}(\mathcal{P}, S)$ is the PIOA \mathcal{P}' that is the same as \mathcal{P} except that $O_{\mathcal{P}'} = O_{\mathcal{P}} - S$ and $H_{\mathcal{P}'} = H_{\mathcal{P}} \cup S$.

3 Task-PIOAs

In this section, we present our definition for task-PIOAs. We introduce task schedules, which are used to generate probabilistic executions. We define composition and hiding operations. We define an implementation relation, which we call \leq_0 . And finally, we state and prove a simple compositionality result. In the next section, Section 4, we define our new simulation relation for task-PIOAs and prove that it is sound for showing implementation relationships.

3.1 Task-PIOA definition

We now augment the PIOA framework with task partitions, our main mechanism for resolving nondeterminism.

Definition 3.1 A task-PIOA is a pair $\mathcal{T} = (\mathcal{P}, R)$ where

- $\mathcal{P} = (Q, \bar{q}, I, O, H, D)$ is a PIOA (satisfying transition determinism).
- R is an equivalence relation on the locally-controlled actions $(O \cup H)$.

For clarity, we sometimes write $R_{\mathcal{T}}$ for R .

The equivalence classes of R are called tasks. A task T is enabled in a state q if some $a \in T$ is enabled in q . It is enabled in a set S of states provided it is enabled in every $q \in S$.

Unless otherwise stated, technical notions for task-PIOAs are inherited from those for PIOAs. Exceptions include the notions of probabilistic executions and trace distributions.

For now, we impose the following action-determinism assumption, which implies that tasks alone are enough to resolve all nondeterministic choices. We will remove this assumption when we introduce local schedulers, in Section 6. To make it easier to remove the action-determinism hypothesis later, we will indicate explicitly, before Section 6, where we are using the action-determinism hypothesis.

- *Action determinism:* For every state $q \in Q$ and task $T \in R$, at most one action $a \in T$ is enabled in q .

Definition 3.2 If $\mathcal{T} = (\mathcal{P}, R)$ is a task-PIOA, then a task schedule for \mathcal{T} is any finite or infinite sequence $\rho = T_1 T_2 \dots$ of tasks in R .

Thus, a task schedule is *static* (or *oblivious*), in the sense that it does not depend on dynamic information generated during execution. Under the action-determinism assumption, a task schedule can be used to generate a unique probabilistic execution, and hence, a unique trace distribution, of the underlying PIOA \mathcal{P} . One can do this by repeatedly scheduling tasks, each of which determines at most one transition of \mathcal{P} .

In general, one could define various classes of task schedules by specifying what dynamic information may be used in choosing the next task. Here, however, we opt for the oblivious version because we intend to model system dynamics separately, via high-level nondeterministic choices (cf. Section 1).

Formally, we define an operation that “applies” a task schedule to a task-PIOA:

Definition 3.3 Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA where $\mathcal{P} = (Q, \bar{q}, I, O, H, D)$. Given $\mu \in \text{Disc}(\text{Frag}^*(\mathcal{P}))$ and a task schedule ρ , $\text{apply}(\mu, \rho)$ is the probability measure on $\text{Frag}(\mathcal{P})$ defined recursively by:

1. $\text{apply}(\mu, \lambda) := \mu$. (λ denotes the empty sequence.)
2. For $T \in R$, $\text{apply}(\mu, T)$ is defined as follows. For every $\alpha \in \text{Frag}^*(\mathcal{P})$, $\text{apply}(\mu, T)(\alpha) := p_1(\alpha) + p_2(\alpha)$, where:
 - $p_1(\alpha) = \mu(\alpha')\eta(q)$ if α is of the form $\alpha' a q$, where $a \in T$ and $(\text{lstate}(\alpha'), a, \eta) \in D$; $p_1(\alpha) = 0$ otherwise.
 - $p_2(\alpha) = \mu(\alpha)$ if T is not enabled in $\text{lstate}(\alpha)$; $p_2(\alpha) = 0$ otherwise.
3. For ρ of the form $\rho' T$, $T \in R$, $\text{apply}(\mu, \rho) := \text{apply}(\text{apply}(\mu, \rho'), T)$.
4. For ρ infinite, $\text{apply}(\mu, \rho) := \lim_{i \rightarrow \infty} (\text{apply}(\mu, \rho_i))$, where ρ_i denotes the length- i prefix of ρ .

In Case (2) above, p_1 represents the probability that α is executed when applying task T at the end of α' . Because of transition-determinism and action-determinism, the transition $(\text{lstate}(\alpha'), a, \eta)$ is unique, and so p_1 is well-defined. The term p_2 represents the original probability $\mu(\alpha)$, which is relevant if T is not enabled after α . It is routine to check that the limit in Case (4) is well-defined. The other two cases are straightforward.

3.2 Example: An Adaptive Adversary

Having defined task-PIOAs in Definition 3.1 and task schedules in Definition 3.2, we illustrate via an example the difference between high- and low-level nondeterministic choices (cf. Section 1) and how they are treated in our formalism.

Consider a toy protocol with one sender and two receivers, who exchange messages via an adversarial entity (Figure 1). In addition to sending messages to Rec_0 and Rec_1 , the party Sender chooses two random bits b and s independently. The first bit b is announced to the adversary Adv and the second bit s is kept secret until Sender receives an acknowledgment from either Rec_0 or Rec_1 . If the acknowledgment from Rec_b arrives *before* the acknowledgment from Rec_{1-b} , Sender reveals s to Adv, otherwise s remains secret. A detailed description of Sender is given in Figure 2.

The adversary Adv acts as a message delivery system between Sender, Rec_0 and Rec_1 . The messages from Sender to Rec_i are delivered whenever they are available, while the acknowledgments from Rec_i to Sender are buffered until Sender announces b . Then Adv delivers ack_b *before* ack_{1-b} . A detailed description of Adv is given in Figure 3.

Finally, a receiver Rec_i simply accepts the message from Sender and responds with an acknowledgment. This is described in Figure 4.

Now, we take a closer look at the output actions of Adv, namely, $\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_0}$, $\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_1}$, $\text{rec}(\text{ack})_{\text{Rec}_0 \rightarrow \text{Sender}}$, and $\text{rec}(\text{ack})_{\text{Rec}_1 \rightarrow \text{Sender}}$. The first two of these actions are used by Adv to just relay

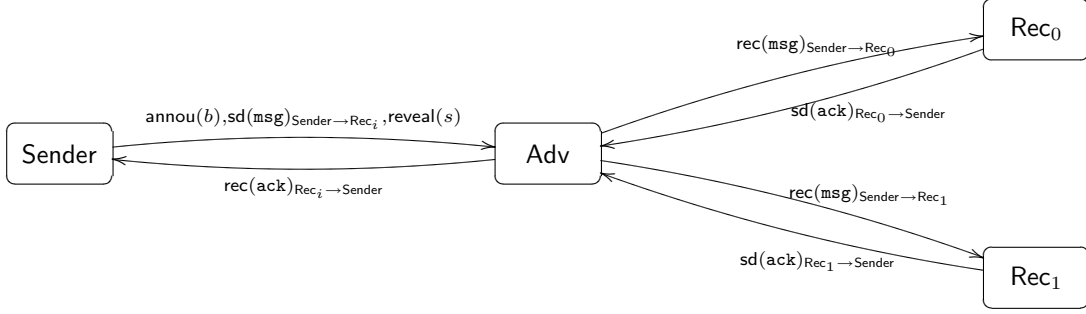


Figure 1: A Toy Protocol

messages received from Sender to Rec_i . The messages can be delivered in any order by Adv once they are received from the Sender. Note, however, that Adv imposes a finer control on the order of occurrence of the the actions $rec(ack)_{Rec_0 \rightarrow Sender}$, and $rec(ack)_{Rec_1 \rightarrow Sender}$. These are the actions that are directly relevant to Adv's ability to learn the secret value s . The conditions that determine whether Adv will send an acknowledgment first to Rec_1 or Rec_2 depend on b , which is generated dynamically by Sender and obtained by Adv as a result of input $annou(x)$. Adv waits until it learns the value of b as a result of $annou(x)$, and then it delivers the acknowledgment from Rec_b . This ensures that Sender will reveal s if the task $reveal(*)$ is scheduled subsequently. In fact, it is easy to check that the following task schedule allows Adv to learn s with probability 1.

choose. $annou(*)$. $sd(msg)_{Sender \rightarrow Rec_0}$. $sd(msg)_{Sender \rightarrow Rec_1}$. $rec(msg)_{Sender \rightarrow Rec_0}$. $rec(msg)_{Sender \rightarrow Rec_1}$.
 $sd(ack)_{Rec_0 \rightarrow Sender}$. $sd(ack)_{Rec_1 \rightarrow Sender}$. $rec(ack)_{Rec_0 \rightarrow Sender}$. $rec(ack)_{Rec_1 \rightarrow Sender}$. $reveal(*)$

The fact that Adv determines the order in which acknowledgments are sent shows that we use an adversarial scheduler Adv to resolve the high-level nondeterminism in this protocol. Also, the fact that Adv uses b in its scheduling decision shows Adv is adaptive. In general, modeling adversarial schedulers as a system component yields a natural way of modeling adaptiveness. (That is, we model the adversary as a task-PIOA that can interact with protocol parties and can compute based on what it sees during execution.)

We now turn to low-level nondeterminism. For instance, the ordering between $sd(msg)_{Sender \rightarrow Rec_0}$ and $sd(msg)_{Sender \rightarrow Rec_1}$ is inessential in the security analysis, provided they are both performed by Sender. Similarly, the ordering between $annou(*)$ and $sd(msg)_{Sender \rightarrow Rec_i}$ is also inessential. All of these are examples of low-level nondeterministic choices and represent implementation freedom in Sender. That is, an actual implementation of Sender may perform $annou(*)$, $sd(msg)_{Sender \rightarrow Rec_0}$ and $sd(msg)_{Sender \rightarrow Rec_1}$ in any order.

Consider the following task schedules.

- (1) $\rho_1 = sd(msg)_{Sender \rightarrow Rec_0}$. choose. $rec(msg)_{Sender \rightarrow Rec_0}$. $sd(msg)_{Sender \rightarrow Rec_1}$. $rec(msg)_{Sender \rightarrow Rec_1}$.
 $sd(ack)_{Rec_1 \rightarrow Sender}$. $annou(*)$. $sd(ack)_{Rec_0 \rightarrow Sender}$. $rec(ack)_{Rec_0 \rightarrow Sender}$. $rec(ack)_{Rec_1 \rightarrow Sender}$. $reveal(*)$
- (2) $\rho_2 = sd(msg)_{Sender \rightarrow Rec_0}$. choose. $rec(msg)_{Sender \rightarrow Rec_0}$. $sd(msg)_{Sender \rightarrow Rec_1}$. $rec(msg)_{Sender \rightarrow Rec_1}$.
 $sd(ack)_{Rec_1 \rightarrow Sender}$. $annou(*)$. $sd(ack)_{Rec_0 \rightarrow Sender}$. $rec(ack)_{Rec_0 \rightarrow Sender}$. $reveal(*)$. $rec(ack)_{Rec_1 \rightarrow Sender}$

It is interesting to note that Adv learns the secret s with probability 1 under ρ_1 , but with probability $\frac{1}{2}$ under ρ_2 . This is because, if $b = 1$, $reveal(*)$ is not enabled in the state resulting from the performance of $rec(ack)_{Rec_0 \rightarrow Sender}$. Nonetheless, Adv is considered to have high advantage in learning s since we define the behavior of automata and the "security" of a protocol that they represent by quantifying over all possible schedulers (see Sections 3.4 and 3.7).

Sender

Signature

Input:
 $\text{rec}(\text{ack})_{\text{Rec}_0 \rightarrow \text{Sender}}, \text{rec}(\text{ack})_{\text{Rec}_1 \rightarrow \text{Sender}}$
Output:
 $\text{annou}(x), x \in \{0, 1\}$
 $\text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_0}, \text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_1}$
 $\text{reveal}(x), x \in \{0, 1\}$
Internal:
choose

Transitions

choose
Precondition:
 $b = \perp \wedge s = \perp$
Effect:
 $b := \text{random}(\text{unif}(\{0, 1\}));$
 $s := \text{random}(\text{unif}(\{0, 1\}))$

$\text{annou}(b)$
Precondition:
 $b \neq \perp$
Effect:
None

$\text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_0}$
Precondition:
True
Effect:
None

Tasks

$\{\text{annou}(x) | x \in \{0, 1\}\}$
 $\{\text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_0}\}, \{\text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_1}\}$
 $\{\text{reveal}(x) | x \in \{0, 1\}\}, \{\text{choose}\}$

States

$b, s \in \{0, 1, \perp\}$, initially \perp
 $c \in \{0, 1, 2\}$, initially 0

$\text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_1}$
Precondition:
True
Effect:
None

$\text{rec}(\text{ack})_{\text{Rec}_0 \rightarrow \text{Sender}}$
Effect:
if $b = 0 \wedge c = 0$ then $c := 1$
else if $b = 1 \wedge c = 0$ then $c := 2$

$\text{rec}(\text{ack})_{\text{Rec}_1 \rightarrow \text{Sender}}$
Effect:
if $b = 1 \wedge c = 0$ then $c := 1$
else if $b = 0 \wedge c = 0$ then $c := 2$

$\text{reveal}(s)$
Precondition:
 $s \neq \perp \wedge c = 1$
Effect:
None

Figure 2: Code for Task-PIOA Sender

3.3 Properties of the apply Function

First, we give some basic properties of the probabilities that arise from the $\text{apply}(\cdot, \cdot)$ function. These turn out to be useful in our results on probabilistic executions and simulation relations.

Lemma 3.4 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. Let μ be a discrete probability measure over finite execution fragments of \mathcal{P} and let T be a task. Let p_1 and p_2 be the functions used in the definition of $\text{apply}(\mu, T)$. Then:*

1. for each state q , $p_1(q) = 0$;
2. for each finite execution fragment α ,

$$\mu(\alpha) = p_2(\alpha) + \sum_{(a,q): \alpha a q \in \text{Frag}^*(\mathcal{P})} p_1(\alpha a q).$$

Proof. Item (1) follows trivially from the definition of $p_1(q)$.
For Item (2), we observe the following facts.

Adv

Signature

Input:

annou(x), $x \in \{0, 1\}$
 $\text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_0}$, $\text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_1}$
 $\text{sd}(\text{ack})_{\text{Rec}_0 \rightarrow \text{Sender}}$, $\text{sd}(\text{ack})_{\text{Rec}_1 \rightarrow \text{Sender}}$
 reveal(x), $x \in \{0, 1\}$

Output:

$\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_0}$, $\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_1}$
 $\text{rec}(\text{ack})_{\text{Rec}_0 \rightarrow \text{Sender}}$, $\text{rec}(\text{ack})_{\text{Rec}_1 \rightarrow \text{Sender}}$

Internal:

None

Transitionsannou(x)

Effect:

 $b := x$ $\text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_0}$

Effect:

if $c_0 = 0$ then $c_0 := 1$ $\text{sd}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_1}$

Effect:

if $c_1 = 0$ then $c_1 := 1$ $\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_0}$

Precondition:

 $c_0 = 1$

Effect:

None

 $\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_1}$

Precondition:

 $c_1 = 1$

Effect:

None

Tasks

$\{\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_0}\}$, $\{\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_1}\}$
 $\{\text{rec}(\text{ack})_{\text{Rec}_0 \rightarrow \text{Sender}}\}$, $\{\text{rec}(\text{ack})_{\text{Rec}_1 \rightarrow \text{Sender}}\}$

States

$b, s \in \{0, 1, \perp\}$, initially \perp
 $c_0, c_1 \in \{0, 1, 2, 3\}$, initially 0

 $\text{sd}(\text{ack})_{\text{Rec}_0 \rightarrow \text{Sender}}$

Effect:

 $c_0 := 2$ $\text{sd}(\text{ack})_{\text{Rec}_1 \rightarrow \text{Sender}}$

Effect:

 $c_1 := 2$ $\text{rec}(\text{ack})_{\text{Rec}_0 \rightarrow \text{Sender}}$

Precondition:

 $(b = 0 \wedge c_0 = 2) \vee (b = 1 \wedge c_1 = 3)$

Effect:

 $c_0 := 3$ $\text{rec}(\text{ack})_{\text{Rec}_1 \rightarrow \text{Sender}}$

Precondition:

 $(b = 1 \wedge c_1 = 2) \vee (b = 0 \wedge c_0 = 3)$

Effect:

 $c_1 := 3$ reveal(x)

Effect:

 $s := x$

Figure 3: Code for Task-PIOA Adv

- If T is not enabled from $\text{lstate}(\alpha)$, then, by definition of p_2 , $\mu(\alpha) = p_2(\alpha)$. Furthermore, for each action a and each state q such that $\alpha a q$ is an execution fragment, we claim that $p_1(\alpha a q) = 0$. Indeed, if $a \notin T$, then the first case of the definition of $p_1(\alpha)$ trivially does not apply; if $a \in T$, then, since T is not enabled from $\text{lstate}(\alpha)$, there is no ρ such that $(\text{lstate}(\alpha), a, \rho) \in D_{\mathcal{P}}$, and thus, again, the first case of the definition of $p_1(\alpha)$ does not apply.
- If T is enabled from $\text{lstate}(\alpha)$, then trivially $p_2(\alpha) = 0$. Furthermore, we claim that $\mu(\alpha) = \sum_{(a,q)} p_1(\alpha a q)$. By action determinism, only one action $b \in T$ is enabled from $\text{lstate}(\alpha)$. By definition of p_1 , $p_1(\alpha a q) = 0$ if $a \neq b$ (either $a \notin T$ or a is not enabled from $\text{lstate}(\alpha)$). Thus,

$$\sum_{(a,q)} p_1(\alpha a q) = \sum_q p_1(\alpha b q) = \sum_q \mu(\alpha) \mu_{\alpha,b}(q).$$

This in turn is equal to $\mu(\alpha)$ since $\sum_q \mu_{\alpha,b}(q) = 1$.

In each case, we get $\mu(\alpha) = p_2(\alpha) + \sum_{(a,q)} p_1(\alpha a q)$, as needed. \square

Rec_i**Signature**

Input:
 $\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_i}$
 Output:
 $\text{sd}(\text{ack})_{\text{Rec}_i \rightarrow \text{Sender}}$
 Internal:
 None

Tasks
 $\{\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_i}\}, \{\text{sd}(\text{ack})_{\text{Rec}_i \rightarrow \text{Sender}}\}$
States
 $c \in \{0, 1\}$, initially 0
Transitions

$\text{rec}(\text{msg})_{\text{Sender} \rightarrow \text{Rec}_i}$
 Effect:
 $c := 1$

$\text{sd}(\text{ack})_{\text{Rec}_i \rightarrow \text{Sender}}$
 Precondition:
 $c \neq 0$
 Effect:
 None

Figure 4: Code for Task-PIOA Rec_i

Lemma 3.5 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. Let μ be a discrete probability measure over finite execution fragments and ρ be a finite sequence of tasks. Then $\text{apply}(\mu, \rho)$ is a discrete probability measure over finite execution fragments.*

Proof. By a simple inductive argument on the length of ρ . The base case is trivial. For the inductive step, it suffices to show that, for each measure ϵ on finite executions fragments and each task T , $\text{apply}(\epsilon, T)$ is a probability measure over finite execution fragments.

Let ϵ' be $\text{apply}(\epsilon, T)$. The fact that ϵ' is a measure on finite execution fragments follows directly by Item (2) of Definition 3.3. To show that ϵ' is in fact a probability measure, we show that $\sum_{\alpha \in \text{Frag}^*(\mathcal{P})} \epsilon'(\alpha) = 1$. By Item (2) of Definition 3.3,

$$\sum_{\alpha \in \text{Frag}^*(\mathcal{P})} \epsilon'(\alpha) = \sum_{\alpha \in \text{Frag}^*(\mathcal{P})} (p_1(\alpha) + p_2(\alpha)).$$

Rearranging terms, we obtain

$$\sum_{\alpha \in \text{Frag}^*(\mathcal{P})} \epsilon'(\alpha) = \sum_q p_1(q) + \sum_{\alpha \in \text{Frag}^*(\mathcal{P})} (p_2(\alpha) + \sum_{(a,q): \alpha a q \in \text{Frag}^*(\mathcal{P})} p_1(\alpha a q)).$$

By Lemma 3.4, the right side becomes $\sum_{\alpha \in \text{Frag}^*(\mathcal{P})} \epsilon(\alpha)$, which equals 1 since ϵ is by assumption a probability measure. Therefore $\sum_{\alpha \in \text{Frag}^*(\mathcal{P})} \epsilon'(\alpha) = 1$, as needed. \square

Lemma 3.6 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA and let T be a task in R . Define $\mu' = \text{apply}(\mu, T)$. Then, for each finite execution fragment α :*

1. *If α consists of a single state q , then $\mu'(C_\alpha) = \mu(C_\alpha)$.*
2. *If $\alpha = \tilde{\alpha} a q$ and $a \notin T$, then $\mu'(C_\alpha) = \mu(C_\alpha)$.*
3. *If $\alpha = \tilde{\alpha} a q$ and $a \in T$, then $\mu'(C_\alpha) = \mu(C_\alpha) + \mu(\tilde{\alpha}) \mu_{\tilde{\alpha}, a}(q)$.*

Proof. Let p_1 and p_2 be the functions used in the definition of $\text{apply}(\mu, T)$, and let α be a finite execution fragment. By definition of a cone and of μ' , $\mu'(C_\alpha) = \sum_{\alpha' | \alpha \leq \alpha'} (p_1(\alpha') + p_2(\alpha'))$. By definition of a cone and Lemma 3.4,

$$\mu(C_\alpha) = \sum_{\alpha' | \alpha \leq \alpha'} (p_2(\alpha') + \sum_{(a,q): \alpha' a q \in \text{Frag}^*(\mathcal{P})} p_1(\alpha' a q)) = \sum_{\alpha' | \alpha \leq \alpha'} (p_1(\alpha') + p_2(\alpha')) - p_1(\alpha).$$

Thus, $\mu'(C_\alpha) = \mu(C_\alpha) + p_1(\alpha)$. We distinguish three cases. If α consists of a single state, then $p_1(\alpha) = 0$ by Lemma 3.4, yielding $\mu'(C_\alpha) = \mu(C_\alpha)$. If $\alpha = \tilde{\alpha}aq$ and $a \notin T$, then $p_1(\alpha) = 0$ by definition, yielding $\mu'(C_\alpha) = \mu(C_\alpha)$. Finally, if $\alpha = \tilde{\alpha}aq$ and $a \in T$, then $p_1(\alpha) = \mu(\tilde{\alpha})\mu_{\tilde{\alpha},a}(q)$ by definition, yielding $\mu'(C_\alpha) = \mu(C_\alpha) + \mu(\tilde{\alpha})\mu_{\tilde{\alpha},a}(q)$. \square

Lemma 3.7 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. Let μ be a discrete probability measure over finite execution fragments, T a task, and $\mu' = \text{apply}(\mu, T)$. Then $\mu \leq \mu'$.*

Proof. Follows directly by Lemma 3.6. \square

Lemma 3.8 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. Let μ be a discrete measure over finite execution fragments and let ρ_1 and ρ_2 be two finite sequences of tasks such that ρ_1 is a prefix of ρ_2 . Then $\text{apply}(\mu, \rho_1) \leq \text{apply}(\mu, \rho_2)$.*

Proof. Simple inductive argument using Lemma 3.7 for the inductive step. \square

The next lemma relates the probability measures on execution fragments that arise as a result when applying a sequence of tasks to a given probability measure μ on execution fragments.

Lemma 3.9 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. Let ρ_1, ρ_2, \dots be a finite or infinite sequence of finite task schedules, and let $\rho = \rho_1\rho_2\dots$ (where juxtaposition denotes concatenation of finite sequences).*

Let μ be a discrete probability measure on finite execution fragments. For each integer i , $1 \leq i \leq |\rho|$, let $\epsilon_i = \text{apply}(\mu, \rho_1\rho_2\dots\rho_i)$, where $\rho_1\dots\rho_i$ denotes the concatenation of the sequences ρ_1 through ρ_i . Let $\epsilon = \text{apply}(\mu, \rho)$. Then the ϵ_i 's form a chain and $\epsilon = \lim_{i \rightarrow \infty} \epsilon_i$.

Proof. The fact that the ϵ_i 's form a chain follows from Lemma 3.7. For the limit property, if the sequence ρ_1, ρ_2, \dots is finite, then the result is immediate. Otherwise, simply observe that the sequence $\epsilon_1, \epsilon_2, \dots$ is a subsequence of the sequence used in the definition of $\text{apply}(\mu, \rho_1\rho_2\dots)$, and therefore, they have the same limit. \square

Lemma 3.10 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. Let μ be a discrete probability measure over finite execution fragments of \mathcal{P} , ρ a task schedule for \mathcal{T} , and q a state of \mathcal{T} . Then $\text{apply}(\mu, \rho)(C_q) = \mu(C_q)$.*

Proof. We prove the result for finite ρ 's by induction on the length of ρ . The infinite case then follows immediately. The base case is trivial since, by definition, $\text{apply}(\mu, \rho) = \mu$. For the inductive step, let $\rho = \rho'T$, and let ϵ be $\text{apply}(\mu, \rho')$. By Definition 3.3, $\text{apply}(\mu, \rho) = \text{apply}(\epsilon, T)$. By induction, $\epsilon(C_q) = \mu(C_q)$. Therefore, it suffices to show $\text{apply}(\epsilon, T)(C_q) = \epsilon(C_q)$.

Let ϵ' be $\text{apply}(\epsilon, T)$. By definition of cone, $\epsilon'(C_q) = \sum_{\alpha: q \leq \alpha} \epsilon'(\alpha)$. By Lemma 3.5, both ϵ and ϵ' are measures over finite execution fragments; therefore we can restrict the sum to finite execution fragments. Let p_1 and p_2 be the two functions used for the computation of $\epsilon'(\alpha)$ according to Item (2) in Definition 3.3. Then $\epsilon'(C_q) = \sum_{\alpha \in \text{Execs}^*(\mathcal{P}): q \leq \alpha} (p_1(\alpha) + p_2(\alpha))$. By rearranging terms, we get $\epsilon'(C_q) = p_1(q) + \sum_{\alpha \in \text{Execs}^*(\mathcal{P}): q \leq \alpha} (p_2(\alpha) + \sum_{(a,s)} p_1(C_{\alpha as}))$. By Lemma 3.4, the right side of the equation above is $\sum_{\alpha: q \leq \alpha} \epsilon(\alpha)$, which is precisely $\epsilon(C_q)$. \square

The next proposition states that $\text{apply}(\cdot, \rho)$ distributes over convex combinations of probability measures. This requires a preliminary lemma.

Lemma 3.11 *Let $\{\mu_i\}_i$ be a countable family of discrete probability measures on finite execution fragments and let $\{p_i\}_i$ be a countable family of probabilities such that $\sum_i p_i = 1$. Let T be a task. Then $\text{apply}(\sum_i p_i \mu_i, T) = \sum_i p_i \text{apply}(\mu_i, T)$.*

Proof. Let p_1 and p_2 be the functions used in the definition of $\text{apply}(\sum_i p_i \mu_i, T)$, and let, for each i , p_1^i and p_2^i be the functions used in the definition of $\text{apply}(\mu_i, T)$. Let α be a finite execution fragment. We show that $p_1(\alpha) = \sum_i p_i p_1^i(\alpha)$ and $p_2(\alpha) = \sum_i p_i p_2^i(\alpha)$. Then

$$\begin{aligned}
\text{apply}\left(\sum_i p_i \mu_i, T\right)(\alpha) &= p_1(\alpha) + p_2(\alpha) && \text{definition of } \text{apply}\left(\sum_i p_i \mu_i, T\right) \\
&= \sum_i p_i p_1^i(\alpha) + \sum_i p_i p_2^i(\alpha) && \text{claims proven below} \\
&= \sum_i p_i (p_1^i(\alpha) + p_2^i(\alpha)) \\
&= \sum_i p_i \text{apply}(\mu_i, T)(\alpha) && \text{definition of } \text{apply}(\mu_i, T)
\end{aligned}$$

To prove our claim about p_1 we distinguish two cases. If α can be written as $\alpha' a q$, where $\alpha' \in \text{supp}(\mu)$, $a \in T$, and $(\text{lstate}(\alpha'), a, \rho) \in D_{\mathcal{P}}$, then, by Definition 3.3, $p_1(\alpha) = (\sum_i p_i \mu_i)(\alpha') \rho(q)$, and, for each i , $p_1^i(\alpha) = \mu_i(\alpha') \rho(q)$. Thus, $p_1(\alpha) = \sum_i p_i p_1^i(\alpha)$ trivially. Otherwise, again by Definition 3.3, $p_1(\alpha) = 0$, and, for each i , $p_1^i(\alpha) = 0$. Thus, $p_1(\alpha) = \sum_i p_i p_1^i(\alpha)$ trivially.

To prove our claim about p_2 we also distinguish two cases. If T is not enabled in $\text{lstate}(\alpha)$, then, by Definition 3.3, $p_2(\alpha) = (\sum_i p_i \mu_i)(\alpha)$, and, for each i , $p_2^i(\alpha) = \mu_i(\alpha)$. Thus, $p_2(\alpha) = \sum_i p_i p_2^i(\alpha)$ trivially. Otherwise, again by Definition 3.3, $p_2(\alpha) = 0$, and, for each i , $p_2^i(\alpha) = 0$. Thus, $p_2(\alpha) = \sum_i p_i p_2^i(\alpha)$ trivially. \square

Proposition 3.12 *Let $\{\mu_i\}_i$ be a countable family of discrete probability measures on finite execution fragments and let $\{p_i\}_i$ be a countable family of probabilities such that $\sum_i p_i = 1$. Let ρ be a finite sequence of tasks. Then, $\text{apply}(\sum_i p_i \mu_i, \rho) = \sum_i p_i \text{apply}(\mu_i, \rho)$.*

Proof. We proceed by induction on the length of ρ . If $\rho = \lambda$, then the result is trivial since $\text{apply}(\cdot, \lambda)$ is defined to be the identity function, which distributes over convex combinations of probability measures. For the inductive step, let ρ be $\rho' T$. By Definition 3.3 and the induction hypothesis,

$$\text{apply}\left(\sum_i p_i \mu_i, \rho' T\right) = \text{apply}\left(\text{apply}\left(\sum_i p_i \mu_i, \rho'\right), T\right) = \text{apply}\left(\sum_i p_i \text{apply}(\mu_i, \rho'), T\right).$$

By Lemma 3.5, each $\text{apply}(\mu_i, \rho')$ is a discrete probability measure over finite execution fragments. By Lemma 3.11, $\text{apply}(\sum_i p_i \text{apply}(\mu_i, \rho'), T) = \sum_i p_i \text{apply}(\text{apply}(\mu_i, \rho'), T)$, and by Definition 3.3, for each i , $\text{apply}(\text{apply}(\mu_i, \rho'), T) = \text{apply}(\mu_i, \rho' T)$. Thus, $\text{apply}(\sum_i p_i \mu_i, \rho' T) = \sum_i p_i \text{apply}(\mu_i, \rho' T)$ as needed. \square

3.4 Probabilistic Executions and Trace Distributions

Here, we show that $\text{apply}(\mu, \rho)$ is a generalized probabilistic execution fragment generated by μ and a scheduler for \mathcal{P} , in the usual sense. Thus, a task schedule for a task-PIOA is a special case of a scheduler for the underlying PIOA.

Theorem 3.13 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. For each probability measure μ on $\text{Frag}_s^*(\mathcal{P})$ and task schedule ρ , there is scheduler σ for \mathcal{P} such that $\text{apply}(\mu, \rho)$ is the generalized probabilistic execution fragment $\epsilon_{\sigma, \mu}$.*

The proof can be found in Appendix B. The idea here is, for any measure μ and task sequence ρ , the probability measure on execution fragments generated by $\text{apply}(\mu, \rho)$ is “standard”, in the sense that it can be obtained from μ and a scheduler as defined in Section 3 for basic PIOAs. Any such $\text{apply}(\mu, \rho)$ is said to be a *generalized probabilistic execution fragment* of the task-PIOA \mathcal{T} . *Probabilistic execution fragments* and *probabilistic executions* are then defined by making the same restrictions as for basic PIOAs.

We write $\text{tdist}(\mu, \rho)$ as shorthand for $\text{tdist}(\text{apply}(\mu, \rho))$, the trace distribution obtained by applying task schedule ρ starting from the measure μ on execution fragments. We write $\text{tdist}(\rho)$ for $\text{tdist}(\text{apply}(\delta(\bar{q}), \rho))$ the trace distribution obtained by applying ρ from the unique start state. (Recall that $\delta(\bar{q})$ denotes the Dirac measure on \bar{q} .) A *trace distribution* of \mathcal{T} is any $\text{tdist}(\rho)$. We use $\text{tdists}(\mathcal{T})$ to denote the set $\{\text{tdist}(\rho) : \rho \text{ is a task schedule for } \mathcal{T}\}$.

3.5 Composition

We define composition of task-PIOAs:

Definition 3.14 *Two task-PIOAs $\mathcal{T}_i = (\mathcal{P}_i, R_i)$, $i \in \{1, 2\}$, are said to be compatible provided the underlying PIOAs are compatible. Then we define their composition $\mathcal{T}_1 \parallel \mathcal{T}_2$ to be the task-PIOA $(\mathcal{P}_1 \parallel \mathcal{P}_2, R_1 \cup R_2)$.*

It is easy to see that $\mathcal{T}_1 \parallel \mathcal{T}_2$ is in fact a task-PIOA. In particular, since compatibility ensures disjoint sets of locally-controlled actions, $R_1 \cup R_2$ is an equivalence relation on the locally-controlled actions of $\mathcal{P}_1 \parallel \mathcal{P}_2$. It is also easy to see that action determinism is preserved under composition. Note that, when two task-PIOAs are composed, no new mechanisms are required to schedule actions of the two components—the tasks alone are enough.

3.6 Hiding

We also define a hiding operator for task-PIOAs. It simply hides output actions:

Definition 3.15 *Let $\mathcal{T} = (\mathcal{P}, R)$ be any task-PIOA, where $\mathcal{P} = (Q, \bar{q}, I, O, H, D)$, and let $S \subseteq O$. Then $\text{hide}(\mathcal{T}, S)$ is the task-PIOA $(\text{hide}(\mathcal{P}, S), R)$, that is, the task-PIOA obtained by hiding S in the underlying PIOA \mathcal{P} , without any change to the task equivalence relation.*

Note that, in the special case where tasks respect the output vs. internal action classification, one can also define a hiding operation that hides all output actions in a set of tasks. We omit the details here.

3.7 Implementation

We now define the notion of external behavior for a task-PIOA and the induced implementation relation between task-PIOAs. Unlike previous definitions of external behavior, the one we use here is not simply a set of trace distributions. Rather, it is a mapping that specifies, for every possible “environment” \mathcal{E} for the given task-PIOA \mathcal{T} , the set of trace distributions that can arise when \mathcal{T} is composed with \mathcal{E} .

Definition 3.16 *Let \mathcal{T} be any task-PIOA and \mathcal{E} be an action-deterministic task-PIOA. We say that \mathcal{E} is an environment for \mathcal{T} if the following hold:*

1. \mathcal{E} is compatible with \mathcal{T} .
2. The composition $\mathcal{T} \parallel \mathcal{E}$ is closed.

Note that \mathcal{E} is allowed to have output actions that are not inputs of \mathcal{T} .

Definition 3.17 *The external behavior of \mathcal{T} , denoted by $\text{extbeh}(\mathcal{T})$, is the total function that maps each environment \mathcal{E} to the set of trace distributions $\text{tdists}(\mathcal{T} \parallel \mathcal{E})$.*

Thus, for each environment, we consider the set of trace distributions that arise from all task schedules. Note that these traces may include new output actions of \mathcal{E} , in addition to the external actions already present in \mathcal{T} .

Our definition of *implementation* says that the lower-level system must “look like” the higher-level system from the perspective of every possible environment. The style of this definition is influenced by common notions in the security protocol literature (e.g., [LMMS98, Can01, PW01]). An advantage of this style of definition is that it yields simple compositionality results (Theorem 3.20). In our case, “looks like” is formalized in terms of inclusion of sets of trace distributions, that is, of external behavior sets.

Definition 3.18 Let $\mathcal{T}_1 = (\mathcal{P}_1, R_1)$ and $\mathcal{T}_2 = (\mathcal{P}_2, R_2)$ be task-PIOAs, and I_i and O_i the input and output actions sets for \mathcal{P}_i , $i \in \{1, 2\}$. Then \mathcal{T}_1 and \mathcal{T}_2 are comparable if $I_1 = I_2$ and $O_1 = O_2$.

Definition 3.19 Let \mathcal{T}_1 and \mathcal{T}_2 be comparable action-deterministic task-PIOAs. Then we say that \mathcal{T}_1 implements \mathcal{T}_2 , written $\mathcal{T}_1 \leq_0 \mathcal{T}_2$, if $\text{extbeh}(\mathcal{T}_1)(\mathcal{E}) \subseteq \text{extbeh}(\mathcal{T}_2)(\mathcal{E})$ for every environment \mathcal{E} for both \mathcal{T}_1 and \mathcal{T}_2 . In other words, we require $\text{tdists}(\mathcal{T}_1 \parallel \mathcal{E}) \subseteq \text{tdists}(\mathcal{T}_2 \parallel \mathcal{E})$ for every \mathcal{E} .

The subscript 0 in the relation symbol \leq_0 refers to the requirement that every trace distribution in $\text{tdists}(\mathcal{T}_1 \parallel \mathcal{E})$ must have an identical match in $\text{tdists}(\mathcal{T}_2 \parallel \mathcal{E})$. For security analysis, we also define another relation $\leq_{\text{neg.pt}}$, which allows “negligible” discrepancies between matching trace distributions [CCK⁺05b, CCK⁺06c].

3.8 Compositionality

Because external behavior and implementation are defined in terms of mappings from environments to sets of trace distributions, a compositionality result for \leq_0 follows easily:

Theorem 3.20 Let $\mathcal{T}_1, \mathcal{T}_2$ be comparable action-deterministic task-PIOAs such that $\mathcal{T}_1 \leq_0 \mathcal{T}_2$, and let \mathcal{T}_3 be an action-deterministic task-PIOA compatible with each of \mathcal{T}_1 and \mathcal{T}_2 . Then $\mathcal{T}_1 \parallel \mathcal{T}_3 \leq_0 \mathcal{T}_2 \parallel \mathcal{T}_3$.

Proof. Let $\mathcal{T}_4 = (\mathcal{P}_4, R_4)$ be any environment (action-deterministic) task-PIOA for both $\mathcal{T}_1 \parallel \mathcal{T}_3$ and $\mathcal{T}_2 \parallel \mathcal{T}_3$. Fix any task schedule ρ_1 for $(\mathcal{T}_1 \parallel \mathcal{T}_3) \parallel \mathcal{T}_4$. Let τ be the trace distribution of $(\mathcal{T}_1 \parallel \mathcal{T}_3) \parallel \mathcal{T}_4$ generated by ρ_1 . It suffices to show that τ is also generated by some task schedule ρ_2 for $(\mathcal{T}_2 \parallel \mathcal{T}_3) \parallel \mathcal{T}_4$.

Note that ρ_1 is also a task schedule for $\mathcal{T}_1 \parallel (\mathcal{T}_3 \parallel \mathcal{T}_4)$, and that ρ_1 generates the same trace distribution τ in the composed task-PIOA $\mathcal{T}_1 \parallel (\mathcal{T}_3 \parallel \mathcal{T}_4)$.

Now, $\mathcal{T}_3 \parallel \mathcal{T}_4$ is an (action-deterministic) environment task-PIOA for each of \mathcal{T}_1 and \mathcal{T}_2 . Since, by assumption, $\mathcal{T}_1 \leq_0 \mathcal{T}_2$, we infer the existence of a task schedule ρ_2 for $\mathcal{T}_2 \parallel (\mathcal{T}_3 \parallel \mathcal{T}_4)$ such that ρ_2 generates trace distribution τ in the task-PIOA $\mathcal{T}_2 \parallel (\mathcal{T}_3 \parallel \mathcal{T}_4)$. Since ρ_2 is also a task schedule for $(\mathcal{T}_2 \parallel \mathcal{T}_3) \parallel \mathcal{T}_4$ and ρ_2 generates τ , this suffices. \square

4 Simulation Relations

Simulation relations provide sufficient conditions for proving that one automaton implements another, according to some precise notion of implementation such as \leq_0 from Definition 3.19. Typically, simulation relations reduce the proof obligations for implementation into conditions relating the start states and conditions relating individual algorithm steps. Checking these individual conditions is generally much easier, and more systematic, than reasoning about entire executions.

In this section, we define a new notion of simulation relation for closed, action-deterministic task-PIOAs, and show that it is sound for proving \leq_0 . Our definition is based on the three operations defined in Section 2.1: flattening, lifting, and expansion.

4.1 Simulation relation definition

We begin with two auxiliary definitions. The first expresses consistency between a probability measure over finite executions and a task schedule. Informally, a measure ϵ over finite executions is said to be consistent with a task schedule ρ if it assigns non-zero probability only to those executions that are possible under the task schedule ρ . We use this condition to avoid extraneous proof obligations in our definition of simulation relation.

Definition 4.1 Let $\mathcal{T} = (\mathcal{P}, R)$ be a closed, action-deterministic task-PIOA and let ϵ be a discrete probability measure over finite executions of \mathcal{P} . Also, let a finite task schedule ρ for \mathcal{T} be given. Then ϵ is consistent with ρ provided that $\text{supp}(\epsilon) \subseteq \text{supp}(\text{apply}(\delta(\bar{q}), \rho))$, where \bar{q} is the start state of \mathcal{P} .

For the second definition, suppose we have two task-PIOAs \mathcal{T}_1 and \mathcal{T}_2 , and a mapping c that takes a finite task schedule ρ and a task T of \mathcal{T}_1 to give a task schedule of \mathcal{T}_2 . The idea is that $c(\rho, T)$ describes how \mathcal{T}_2 matches task T , given that it has already matched the task schedule ρ . Using c , we define a new function $\text{full}(c)$ that, given a task schedule ρ , iterates c on all the elements of ρ , thus producing a “full” task schedule of \mathcal{T}_2 that matches all of ρ .

Definition 4.2 Let $\mathcal{T}_1 = (\mathcal{P}_1, R_1)$ and $\mathcal{T}_2 = (\mathcal{P}_2, R_2)$ be two task-PIOAs, and let $c : (R_1^* \times R_1) \rightarrow R_2^*$ be given. Define $\text{full}(c) : R_1^* \rightarrow R_2^*$ recursively as follows: $\text{full}(c)(\lambda) := \lambda$, and $\text{full}(c)(\rho T) := \text{full}(c)(\rho) \frown c(\rho, T)$ (that is, the concatenation of $\text{full}(c)(\rho)$ and $c(\rho, T)$).

Next, we define our new notion of simulation for task-PIOAs. Note that our simulation relations are relations between probability measures on executions, as opposed to relations between states. Here the use of measures on executions is motivated by certain cases that arise in our OT protocol proof. For example, we wish to match random choices that are made at different points in the low-level and high-level models (see Section 4.3).

Definition 4.3 Let $\mathcal{T}_1 = (\mathcal{P}_1, R_1)$ and $\mathcal{T}_2 = (\mathcal{P}_2, R_2)$ be two comparable task-PIOAs that are closed and action-deterministic. Let R be a relation from $\text{Disc}(\text{Execs}^*(\mathcal{P}_1))$ to $\text{Disc}(\text{Execs}^*(\mathcal{P}_2))$, such that, if $\epsilon_1 R \epsilon_2$, then $\text{tdist}(\epsilon_1) = \text{tdist}(\epsilon_2)$. (That is, the two measures on finite executions yield the same measure on traces.) Then R is a simulation from \mathcal{T}_1 to \mathcal{T}_2 if there exists $c : (R_1^* \times R_1) \rightarrow R_2^*$ such that the following properties hold:

1. Start condition: $\delta(\bar{q}_1) R \delta(\bar{q}_2)$.
2. Step condition: If $\epsilon_1 R \epsilon_2$, $\rho_1 \in R_1^*$, ϵ_1 is consistent with ρ_1 , ϵ_2 is consistent with $\text{full}(c)(\rho_1)$, and $T \in R_1$, then $\epsilon'_1 \mathcal{E}(R) \epsilon'_2$ where $\epsilon'_1 = \text{apply}(\epsilon_1, T)$ and $\epsilon'_2 = \text{apply}(\epsilon_2, c(\rho_1, T))$.

Intuitively, $\epsilon_1 R \epsilon_2$ means that it is possible to simulate from ϵ_2 anything that can happen from ϵ_1 . Furthermore, $\epsilon'_1 \mathcal{E}(R) \epsilon'_2$ means that we can decompose ϵ'_1 and ϵ'_2 into pieces that can simulate each other, and so we can also say that it is possible to simulate from ϵ'_2 anything that can happen from ϵ'_1 . This rough intuition is at the base of the proof of our soundness result, Theorem 4.7.

The lemma below (proved in [CCK⁺05b]) captures a special case of the simulation relation definition we have given above. Any relation that satisfies the hypotheses of this lemma is a simulation relation. We use this special case in Example 4.3, where we demonstrate simulation proofs.

Lemma 4.4 Let $\mathcal{T}_1 = (\mathcal{P}_1, R_1)$ and $\mathcal{T}_2 = (\mathcal{P}_2, R_2)$ be two comparable closed action-deterministic task-PIOAs. Let R be a relation from $\text{Disc}(\text{Execs}^*(\mathcal{P}_1))$ to $\text{Disc}(\text{Execs}^*(\mathcal{P}_2))$, such that, if $\epsilon_1 R \epsilon_2$, then $\text{tdist}(\epsilon_1) = \text{tdist}(\epsilon_2)$. Let $c : (R_1^* \times R_1) \rightarrow R_2^*$. Suppose further that the following conditions hold:

1. Start condition: $\delta(\bar{q}_1) R \delta(\bar{q}_2)$.
2. Step condition: If $\epsilon_1 R \epsilon_2$, $\rho_1 \in R_1^*$, ϵ_1 is consistent with ρ_1 , ϵ_2 is consistent with $\text{full}(c)(\rho_1)$, and $T \in R_1$, then there exist
 - a probability measure p on a countable index set I ,
 - probability measures ϵ'_{1j} , $j \in I$, on finite executions of \mathcal{P}_1 , and
 - probability measures ϵ'_{2j} , $j \in I$, on finite executions of \mathcal{P}_2 ,

such that:

- for each $j \in I$, $\epsilon'_{1j} R \epsilon'_{2j}$,
- $\sum_{j \in I} p(j)(\epsilon'_{1j}) = \text{apply}(\epsilon_1, T)$, and
- $\sum_{j \in I} p(j)(\epsilon'_{2j}) = \text{apply}(\epsilon_2, c(\rho_1, T))$.

Then R is a simulation relation from \mathcal{T}_1 to \mathcal{T}_2 using c .

4.2 Soundness

In this section, we state and prove two soundness results. The first result, Theorem 4.7, says that, for closed task-PIOAs, the existence of a simulation relation implies inclusion of sets of trace distributions.

The proof requires two lemmas. Recall that the definition of simulation relations requires that any two R -related execution distributions must have the same trace distribution. Lemma 4.5 extends this property to the claim that any pair of execution distributions that are related by the expansion of the relation R , $\mathcal{E}(R)$, must also have the same trace distribution. (For the proof, the only property of simulation relations we need is that related execution distributions have the same trace distribution.)

Lemma 4.5 *Let \mathcal{T}_1 and \mathcal{T}_2 be comparable closed action-deterministic task-PIOAs and let R be a simulation from \mathcal{T}_1 to \mathcal{T}_2 . Let ϵ_1 and ϵ_2 be discrete probability measures over finite executions of \mathcal{T}_1 and \mathcal{T}_2 , respectively, such that $\epsilon_1 \mathcal{E}(R) \epsilon_2$. Then $\text{tdist}(\epsilon_1) = \text{tdist}(\epsilon_2)$.*

Proof. Since $\epsilon_1 \mathcal{E}(R) \epsilon_2$, we may choose measures η_1, η_2 and a weighting functions w as in the definition of expansion. Then for all $\rho_1 \in \text{supp}(\eta_1)$, we have $\eta_1(\rho_1) = \sum_{\rho_2 \in \text{supp}(\eta_2)} w(\rho_1, \rho_2)$. Moreover, we have $\epsilon_1 = \text{flatten}(\eta_1)$, therefore

$$\text{tdist}(\epsilon_1) = \sum_{\rho_1 \in \text{supp}(\eta_1)} \eta_1(\rho_1) \text{tdist}(\rho_1) = \sum_{\rho_1 \in \text{supp}(\eta_1)} \sum_{\rho_2 \in \text{supp}(\eta_2)} w(\rho_1, \rho_2) \text{tdist}(\rho_1).$$

Now consider any ρ_1 and ρ_2 with $w(\rho_1, \rho_2) > 0$. By the definition of a weighting function, we may conclude that $\rho_1 R \rho_2$. Since R is a simulation relation, we have $\text{tdist}(\rho_1) = \text{tdist}(\rho_2)$. Thus we may replace $\text{tdist}(\rho_1)$ by $\text{tdist}(\rho_2)$ in the summation above. This yields:

$$\text{tdist}(\epsilon_1) = \sum_{\rho_1 \in \text{supp}(\eta_1)} \sum_{\rho_2 \in \text{supp}(\eta_2)} w(\rho_1, \rho_2) \text{tdist}(\rho_2) = \sum_{\rho_2 \in \text{supp}(\eta_2)} \sum_{\rho_1 \in \text{supp}(\eta_1)} w(\rho_1, \rho_2) \text{tdist}(\rho_2).$$

Using again the fact that w is a weighting function, we can simplify the inner sum above to obtain

$$\text{tdist}(\epsilon_1) = \sum_{\rho_2 \in \text{supp}(\eta_2)} \eta_2(\rho_2) \text{tdist}(\rho_2).$$

This equals $\text{tdist}(\epsilon_2)$ because, by the choice of η_2 , we know that $\epsilon_2 = \text{flatten}(\eta_2)$. \square

The second lemma provides the inductive step needed in the proof of Theorem 4.7.

Lemma 4.6 *Let \mathcal{T}_1 and \mathcal{T}_2 be two comparable closed task-PIOAs and let R be a simulation relation from \mathcal{T}_1 to \mathcal{T}_2 . Furthermore, let c be a mapping witnessing the fact that R is a simulation relation. Let a finite task scheduler ρ_1 of \mathcal{T}_1 be given and set $\rho_2 = \text{full}(c)(\rho_1)$. (Then ρ_2 is a finite task scheduler of \mathcal{T}_2 .) Let ϵ_1 denote $\text{apply}(\delta(\bar{q}_1), \rho_1)$ and let ϵ_2 denote $\text{apply}(\delta(\bar{q}_2), \rho_2)$. Suppose that $\epsilon_1 \mathcal{E}(R) \epsilon_2$.*

Now let T be a task of \mathcal{T}_1 . Let $\epsilon'_1 = \text{apply}(\delta(\bar{q}_1), \rho_1 T)$ and let $\epsilon'_2 = \text{apply}(\delta(\bar{q}_2), \rho_2 c(\rho_1, T))$. Then $\epsilon'_1 \mathcal{E}(R) \epsilon'_2$.

Proof. Let η_1, η_2 and w be the measures and weighting function that witness $\epsilon_1 \mathcal{E}(R) \epsilon_2$. Observe that $\epsilon'_1 = \text{apply}(\epsilon_1, T)$ and $\epsilon'_2 = \text{apply}(\epsilon_2, c(\rho_1, T))$.

We apply Lemma 2.7: define the function f on discrete distributions on finite executions of \mathcal{T}_1 by $f(\epsilon) = \text{apply}(\epsilon, T)$, and the function g on discrete distributions on finite executions of \mathcal{T}_2 by $g(\epsilon) = \text{apply}(\epsilon, c(\rho_1, T))$. We show that the hypothesis of Lemma 2.7 is satisfied, so we can invoke Lemma 2.7 to conclude that $\epsilon'_1 \mathcal{E}(R) \epsilon'_2$.

Distributivity of f and g follows directly by Proposition 3.12. Let μ_1, μ_2 be two measures such that $w(\mu_1, \mu_2) > 0$. We must show that $f(\mu_1) \mathcal{E}(R) g(\mu_2)$. Since w is a weighting function for $\epsilon_1 \mathcal{E}(R) \epsilon_2$, $\mu_1 R \mu_2$. Observe that $\text{supp}(\mu_1) \subseteq \text{supp}(\epsilon_1)$ and $\text{supp}(\mu_2) \subseteq \text{supp}(\epsilon_2)$; thus, μ_1 is consistent with ρ_1 and μ_2 is consistent with ρ_2 . By the step condition for R , $\text{apply}(\mu_1, T) \mathcal{E}(R) \text{apply}(\mu_2, c(\rho_1, T))$. Observe that $\text{apply}(\mu_1, T) = f(\mu_1)$ and that $\text{apply}(\mu_2, c(\rho_1, T)) = g(\mu_2)$. Thus, $f(\mu_1) \mathcal{E}(R) g(\mu_2)$, as needed. \square

The following theorem, Theorem 4.7, is the main soundness result. The proof simply puts the pieces together, using Lemma 3.9 (which says that the probabilistic execution generated by an infinite task scheduler can be seen as the limit of the probabilistic executions generated by some of the finite prefixes of the task scheduler), Lemma 4.6 (the step condition), Lemma 4.5 (related probabilistic executions have the same trace distribution), and Lemma A.8 (limit commutes with $tdist$).

Theorem 4.7 *Let \mathcal{T}_1 and \mathcal{T}_2 be comparable task-PIOAs that are closed and action-deterministic. If there exists a simulation relation from \mathcal{T}_1 to \mathcal{T}_2 , then $tdists(\mathcal{T}_1) \subseteq tdists(\mathcal{T}_2)$.*

Proof. Let R be the assumed simulation relation from \mathcal{T}_1 to \mathcal{T}_2 . Let ϵ_1 be the probabilistic execution of \mathcal{T}_1 generated by \bar{q}_1 and a (finite or infinite) task schedule, $T_1 T_2 \dots$. For each $i > 0$, define ρ_i to be $c(T_1 \dots T_{i-1}, T_i)$. Let ϵ_2 be the probabilistic execution generated by \bar{q}_2 and the concatenation $\rho_1 \rho_2 \dots$. It is sufficient to prove $tdist(\epsilon_1) = tdist(\epsilon_2)$.

For each $j \geq 0$, let $\epsilon_{1,j} = \text{apply}(\bar{q}_1, T_1 \dots T_j)$, and $\epsilon_{2,j} = \text{apply}(\bar{q}_2, \rho_1 \dots \rho_j)$. Then by Lemma 3.9, for each $j \geq 0$, $\epsilon_{1,j} \leq \epsilon_{1,j+1}$ and $\epsilon_{2,j} \leq \epsilon_{2,j+1}$; moreover, $\lim_{j \rightarrow \infty} \epsilon_{1,j} = \epsilon_1$ and $\lim_{j \rightarrow \infty} \epsilon_{2,j} = \epsilon_2$. Also, for every $j \geq 0$, $\text{apply}(\epsilon_{1,j}, T_{j+1}) = \epsilon_{1,j+1}$ and $\text{apply}(\epsilon_{2,j}, \rho_{j+1}) = \epsilon_{2,j+1}$.

Observe that $\epsilon_{1,0} = \delta(\bar{q}_1)$ and $\epsilon_{2,0} = \delta(\bar{q}_2)$. The start condition for a simulation relation and a trivial expansion imply that $\epsilon_{1,0} \mathcal{E}(R) \epsilon_{2,0}$. Then by induction, using Lemma 4.6 for the definition of a simulation relation in proving the inductive step, for each $j \geq 0$, $\epsilon_{1,j} \mathcal{E}(R) \epsilon_{2,j}$. Then, by Lemma 4.5, for each $j \geq 0$, $tdist(\epsilon_{1,j}) = tdist(\epsilon_{2,j})$.

By Lemma A.8, $tdist(\epsilon_1) = \lim_{j \rightarrow \infty} tdist(\epsilon_{1,j})$, and $tdist(\epsilon_2) = \lim_{j \rightarrow \infty} tdist(\epsilon_{2,j})$. Since for each $j \geq 0$, $tdist(\epsilon_{1,j}) = tdist(\epsilon_{2,j})$, we conclude that $tdist(\epsilon_1) = tdist(\epsilon_2)$, as needed. \square

The second soundness result, Corollary 4.8, asserts soundness for (not necessarily closed) task-PIOAs, with respect to the \leq_0 relation.

Corollary 4.8 *Let \mathcal{T}_1 and \mathcal{T}_2 be two comparable action-deterministic task-PIOAs. Suppose that, for every environment \mathcal{E} for both \mathcal{T}_1 and \mathcal{T}_2 , there exists a simulation relation R from $\mathcal{T}_1 \parallel \mathcal{E}$ to $\mathcal{T}_2 \parallel \mathcal{E}$. Then $\mathcal{T}_1 \leq_0 \mathcal{T}_2$.*

Proof. Immediate by Theorem 4.7 and the definition of \leq_0 . \square

4.3 Example: Trapdoor vs. Rand

The following example, taken from our Oblivious Transfer case study, is a key motivation for generalizing prior notions of simulation relations. We consider two closed task-PIOAs, Trapdoor and Rand. Rand simply chooses a number in $\{1, \dots, n\}$ randomly, from the uniform distribution (using a choose internal action), and then outputs the chosen value k (using a report(k) output action). Trapdoor, on the other hand, first chooses a random number, then applies a known permutation f to the chosen number, and then outputs the result. (The name Trapdoor refers to the type of permutation f that is used in the OT protocol.)

More precisely, neither Rand nor Trapdoor has any input actions. Rand has output actions report(k), $k \in [n] = \{1, \dots, n\}$ and internal action choose. It has tasks Report = {report(k) : $k \in [n]$ }, and Choose = {choose}. Its state contains one variable $zval$, which assumes values in $[n] \cup \{\perp\}$, initially \perp . The choose action is enabled when $zval = \perp$, and has the effect of setting $zval$ to a number in $[n]$, chosen uniformly at random. The report(k) action is enabled when $zval = k$, and has no effect on the state (so it may happen repeatedly). Precondition/effect code for Rand appears in Figure 5, and a diagram appears in Figure 6.

Trapdoor has the same actions as Rand, plus internal action compute. It has the same tasks as Rand, plus the task Compute = {compute}. Trapdoor's state contains two variables, y and z , each of which takes on values in $[n] \cup \{\perp\}$, initially \perp . The choose action is enabled when $y = \perp$, and sets y to a number in $[n]$, chosen uniformly at random. The compute action is enabled when $y \neq \perp$ and $z = \perp$, and sets $z := f(y)$. The report(k) action behaves exactly as in Rand. Precondition/effect code for Trapdoor appears in Figure 7, and a diagram appears in Figure 8.

Rand

Signature

Input:
None
Output:
 $\text{report}(k), k \in \{1, \dots, n\}$
Internal:
choose

Tasks

Report = $\{\text{report}(k) : k \in \{1, \dots, n\}\}$
Choose = $\{\text{choose}\}$

States

$zval \in \{1, \dots, n\} \cup \{\perp\}$, initially \perp

Transitions

choose
Precondition:
 $zval = \perp$
Effect:
 $zval := \text{random}(\text{unif}(\{1, \dots, n\}))$

$\text{report}(k)$
Precondition:
 $zval = k$
Effect:
None

Figure 5: Code for Task-PIOA Rand




Figure 6: Task-PIOA Rand

We want to use a simulation relation to prove that $\text{tdists}(\text{Trapdoor}) \subseteq \text{tdists}(\text{Rand})$. To do so, it is natural to allow the steps that define z to correspond in the two automata, which means the choose steps of Trapdoor (which define y) do not have corresponding steps in Rand. Note that, between the choose and compute in Trapdoor, a randomly-chosen value appears in the y component of the state of Trapdoor, but no such value appears in the corresponding state of Rand. Thus, the desired simulation relation should allow the correspondence between a probability measure on states of Trapdoor and a single state of Rand.

Simulation relation: We are able to express this correspondence using a simulation relation in the sense of Definition 4.3. Let $\epsilon_1 \in \text{Disc}(\text{Execs}^*(\text{Trapdoor}))$ and $\epsilon_2 \in \text{Disc}(\text{Execs}^*(\text{Rand}))$ such that $\text{tdist}(\epsilon_1) = \text{tdist}(\epsilon_2)$.⁴ Then we say that ϵ_1 and ϵ_2 are related by R whenever the following conditions hold:

1. For every $s \in \text{supp}(\text{lstate}(\epsilon_1))$ and $u \in \text{supp}(\text{lstate}(\epsilon_2))$, $s.zval = u.zval$.
2. For every $u \in \text{supp}(\text{lstate}(\epsilon_2))$, if $u.zval = \perp$ then either
 - (a) $\text{lstate}(\epsilon_1).yval$ is the Dirac distribution on \perp , or else
 - (b) $\text{lstate}(\epsilon_1).yval$ is the uniform distribution on $[n]$.

We define the task correspondence mapping c as follows.⁵

- $c(\rho, \text{Choose}) = \lambda$.

⁴In the extended abstract [CCK⁺06b], this condition is missing.

⁵In the extended abstract [CCK⁺06b], the definition of c contains a small error. Namely, in the second clause, $c(\rho, \text{Compute})$ is set to Choose even if ρ does not contain Choose.

Trapdoor

Signature

Input:
None
Output:
 $\text{report}(k), k \in \{1, \dots, n\}$
Internal:
choose, compute

Tasks

Report = $\{\text{report}(k) : k \in \{1, \dots, n\}\}$
Choose = {choose}, Compute = {compute}

States

$yval \in \{1, \dots, n\} \cup \{\perp\}$, initially \perp
 $zval \in \{1, \dots, n\} \cup \{\perp\}$, initially \perp

Transitions

choose
Precondition:
 $yval = \perp$
Effect:
 $yval := \text{random}(\text{unif}(\{1, \dots, n\}))$

report(k)
Precondition:
 $zval = k$
Effect:
None

compute
Precondition:
 $yval \neq \perp; zval = \perp$
Effect:
 $zval := f(yval)$

Figure 7: Code for Task-PIOA Trapdoor




Figure 8: Task-PIOA Trapdoor

- If ρ contains Choose, then $c(\rho, \text{Compute}) = \text{Choose}$; otherwise, $c(\rho, \text{Compute}) = \lambda$.
- $c(\rho, \text{Report}) = \text{Report}$.

Proving the simulation relation: We now prove that the relation R defined above is a simulation relation in the sense of Definition 4.3. We do this by showing that R satisfies the two conditions in Lemma 4.4.

The following two invariant properties are used in the simulation proof. They both follow from easy inductive arguments.

Lemma 4.9 *In every reachable state s of Trapdoor, if $s.zval \neq \perp$ then $s.yval \neq \perp$.*

Lemma 4.10 *Suppose ρ is a finite task schedule for Trapdoor, and ϵ is a discrete distribution on finite executions of Trapdoor that is consistent with ρ . Let $s \in \text{supp}(\text{lstate}(\epsilon))$ be given. Suppose further that $\text{Choose} \in \rho$, that is, the task Choose occurs in the schedule ρ . Then $s.yval \neq \perp$. If $\text{Choose} \notin \rho$, then $s.yval = \perp$.*

Lemma 4.11 *The relation R is a simulation relation from Trapdoor to Rand using the mapping*

$$c : R_{\text{Trapdoor}}^* \times R_{\text{Trapdoor}} \rightarrow R_{\text{Rand}}^*$$

defined above.

Proof. We show that R satisfies the two conditions in Lemma 4.4.

Start condition: The Dirac measures on executions consisting of the unique start states s and u of, respectively, Trapdoor and Rand are R -related. Properties 1 and 2 of R hold since for every $s \in \text{supp}(\text{lstate}(\epsilon_1))$ and $u \in \text{supp}(\text{lstate}(\epsilon_2))$, $s.yval = s.zval = u.zval = \perp$.

Step condition: Suppose $\epsilon_1 R \epsilon_2$, $\rho_1 \in R_1^*$, ϵ_1 is consistent with ρ_1 , ϵ_2 is consistent with $\text{full}(c)(\rho_1)$. We now successively consider each task T in R_1 .

1. $T = \text{Choose}$.

Fix any pair of states $s \in \text{supp}(\text{lstate}(\epsilon_1))$ and $u \in \text{supp}(\text{lstate}(\epsilon_2))$. We consider two cases, according to the value of $s.zval$.

- (a) $s.zval \neq \perp$. By Property 1 of R , for every $v \in \text{supp}(\text{lstate}(\epsilon_1))$, we have that $v.zval \neq \perp$. By the invariant expressed in Lemma 4.9, for every $v \in \text{supp}(\text{lstate}(\epsilon_1))$, we have that $v.yval \neq \perp$. As a result, T is disabled in every state in $\text{supp}(\text{lstate}(\epsilon_1))$.

Let I be the singleton index $\{1\}$, let p be the Dirac measure on 1 and let $\epsilon'_{11} = \epsilon'_1$ and $\epsilon'_{21} = \epsilon'_2$. By Definition 3.3 we have $\epsilon'_1 = \epsilon_1$, $\epsilon'_2 = \epsilon_2$. Since $\epsilon_1 R \epsilon_2$, we have $\epsilon'_{11} R \epsilon'_{21}$, as needed. The trace distribution equivalence condition $\text{tdist}(\epsilon'_1) = \text{tdist}(\epsilon'_2)$ also holds since $\text{tdist}(\epsilon_1) = \text{tdist}(\epsilon_2)$. The summations in the step condition of Lemma 4.4 clearly hold.

- (b) $s.zval = \perp$. By Property 1 of R , for every $v \in \text{supp}(\text{lstate}(\epsilon_1)) \cup \text{supp}(\text{lstate}(\epsilon_2))$, we have that $v.zval = \perp$. Following Property 2 of R , we distinguish two cases.

- i. $\text{lstate}(\epsilon_1).yval$ is the Dirac distribution on \perp . In this case, T is enabled in every state in $\text{supp}(\text{lstate}(\epsilon_1))$.

Let I be the singleton index $\{1\}$, let p be the Dirac measure on 1 and let $\epsilon'_{11} = \epsilon'_1$ and $\epsilon'_{21} = \epsilon'_2 = \epsilon_2$. We now show that $(\epsilon'_1, \epsilon_2) \in R$. By Definition 3.3 and by the effect of the unique action in T , we observe that $\text{supp}(\text{lstate}(\epsilon'_1))$ is equal to $\text{supp}(\text{lstate}(\epsilon_1))$, except that $\text{supp}(\text{lstate}(\epsilon'_1))$ projected on $yval$ is now the uniform distribution on $[n]$. As a result, Property 1 of R holds since the $zval$ component is not modified by the application of T to ϵ_1 and λ to ϵ_2 , and Property 2b of R holds since the effect of the *choose* action in Trapdoor is to select $yval$ according to a uniform distribution. The summations in the step condition clearly hold.

- ii. $\text{lstate}(\epsilon_1).yval$ is uniformly distributed on $[n]$. In this case, T is disabled in every state in $\text{supp}(\text{lstate}(\epsilon_1))$, and the treatment is similar to that of Case 1a.

In all cases, the only action that may take place is the choose action, and it is an internal action. So, the trace condition of R is trivially satisfied.

2. $T = \text{Compute and Choose} \in \rho_1$.

Fix any pair of states $s \in \text{supp}(\text{lstate}(\epsilon_1))$ and $u \in \text{supp}(\text{lstate}(\epsilon_2))$. Lemma 4.10 guarantees that $s.yval \neq \perp$. We consider two cases, according to the value of $s.zval$.

- (a) $s.zval \neq \perp$. Property 1 of R guarantees that $v.zval \neq \perp$ for every v in $\text{supp}(\text{lstate}(\epsilon_1)) \cup \text{supp}(\text{lstate}(\epsilon_2))$. As a result T is disabled in every state in $\text{supp}(\text{lstate}(\epsilon_1))$, and $c(\rho_1, T)$ is disabled in every state $\text{supp}(\text{lstate}(\epsilon_2))$. The treatment of this case is then the same as the one of Case 1a of the Choose task.

- (b) $s.zval = \perp$. By the same observation as in the previous paragraph, we obtain that T is enabled in every state in $\text{supp}(\text{lstate}(\epsilon_1))$, and $c(\rho, T)$ is enabled every state in $\text{supp}(\text{lstate}(\epsilon_2))$. Also, using Property 2 of R , we obtain that the projection of $\text{lstate}(\epsilon_1)$ on $yval$ is the uniform distribution on $[n]$.

We define the probability measures needed to show the step correspondence. Let p be the uniform probability measure on the index set $I = \{1 \cdots n\}$. That is, $p(j) = \frac{1}{n}$ for each $j \in I$. For each j , we define the probability measure ϵ'_{1j} as follows. The support $\text{supp}(\epsilon'_{1j})$ is the set of executions

$\alpha \in \text{supp}(\epsilon'_1)$ such that $\text{lstate}(\alpha).zval = j$. For each $\alpha \in \text{supp}(\epsilon'_{1j})$ of the form α' compute q , let $\epsilon'_{1j}(\alpha) = \epsilon_1(\alpha')$. We define ϵ'_{2j} analogously from ϵ_2 , except that we now consider $\alpha \in \text{supp}(\epsilon'_{2j})$ of the form α' choose q .

Now fix $j \in I$; we show that $(\epsilon'_{1j}, \epsilon'_{2j}) \in R$. To do this, we establish the properties of R for ϵ'_{1j} and ϵ'_{2j} , and show trace distribution equivalence for ϵ'_{1j} and ϵ'_{2j} .

Consider any pair of states $s' \in \text{supp}(\text{lstate}(\epsilon'_{1j}))$ and $u' \in \text{supp}(\text{lstate}(\epsilon'_{2j}))$. By definition of ϵ'_{1j} and ϵ'_{2j} , we know that Property 1 of R holds. Also, since $u'.zval \neq \perp$, Property 2 of R holds too. Since choose and compute are internal actions, trace distribution equivalence holds.

Eventually, since the projection of $\text{lstate}(\epsilon_1)$ on $yval$ is the uniform distribution on $[n]$, since f is a permutation, and since the effect of the choose action of Rand is to select $zval$ according to the uniform distribution on $[n]$, the summations of the step condition of Lemma 4.4 hold.

3. $T = \text{Compute and Choose} \notin \rho_1$.

Fix any pair of states $s \in \text{supp}(\text{lstate}(\epsilon_1))$ and $u \in \text{supp}(\text{lstate}(\epsilon_2))$. Lemma 4.10 guarantees that $s.yval = \perp$. As a result, T is disabled in all states in $\text{supp}(\text{lstate}(\epsilon_1))$. Since $c(\rho_1, T) = \lambda$, the treatment of this case is then the same as the one of Case 1a of the Choose task.

4. $T = \text{Report}$. Fix any pair of states $s \in \text{supp}(\text{lstate}(\epsilon_1))$ and $u \in \text{supp}(\text{lstate}(\epsilon_2))$. We consider two cases, according to the value of $s.zval$.

- (a) $s.zval \neq \perp$. In this case, Property 1 guarantees that the projection of $\text{supp}(\text{lstate}(\epsilon_1)) \cup \text{supp}(\text{lstate}(\epsilon_2))$ on $zval$ is a unique value j . As a result, there is a unique action $a = \text{report}(j) \in T \cup c(\rho_1, T)$ that is enabled in all states in $\text{supp}(\text{lstate}(\epsilon_1)) \cup \text{supp}(\text{lstate}(\epsilon_2))$.

Let I be the singleton index $\{1\}$, let p be the Dirac measure on 1 and let $\epsilon'_{11} = \epsilon'_1$ and $\epsilon'_{21} = \epsilon'_2$. We show that $(\epsilon'_{11}, \epsilon'_{21}) \in R$. First, since the action a has no effect, Properties 1 and 2 of R trivially hold for $(\epsilon'_{11}, \epsilon'_{21})$. Then, since the same action a is executed in both systems, trace distribution equivalence for ϵ'_1 and ϵ'_2 holds.

- (b) $s.zval = \perp$. In this case, Property 1 guarantees that the projection of $\text{supp}(\text{lstate}(\epsilon_1)) \cup \text{supp}(\text{lstate}(\epsilon_2))$ on $zval$ is equal to \perp . Then, T and $c(\rho_1, T)$ are disabled in every state in $\text{supp}(\text{lstate}(\epsilon_1)) \cup \text{supp}(\text{lstate}(\epsilon_2))$, and the treatment of this case is the same as that of Case 1a of the Choose task.

□

5 Application to Security Protocols

In [CCK⁺05b, CCK⁺06c], we use the task-PIOAs to model and analyze the Oblivious Transfer (OT) protocol of Goldreich et al. [GMW87]. Our analysis takes into account computational issues such as resource-bounded adversaries and computational indistinguishability assumptions. Therefore, it requires extensions to the task-PIOA theory presented in this paper. These extensions, and the full details of our analysis, are beyond the scope of this paper. Our purpose in this section is to discuss how task-PIOAs can be used in modeling and analysis of cryptographic protocols based on a concrete, nontrivial example.

Oblivious Transfer In the OT problem, two input bits (x_0, x_1) are submitted to a Transmitter Trans and a single input bit i to a Receiver Rec. After engaging in an OT protocol, Rec should output only the single bit x_i . Rec should not learn the other bit x_{1-i} , and Trans should not learn i ; moreover, an eavesdropping adversary should not, by observing the protocol messages, be able to learn anything about the inputs or the progress of the protocol. OT has been shown to be “complete” for multi-party secure computation, in the sense that, using OT as the only cryptographic primitive, one can construct protocols that securely realize any functionality.

The protocol of [GMW87] uses trap-door permutations (and hard-core predicates) as an underlying cryptographic primitive. It uses three rounds of communication: First, Trans chooses a random trap-door permutation f and sends it to Rec. Second, Rec chooses two random numbers (y_0, y_1) and sends (z_0, z_1) to Trans, where $z_i = f(y_i)$ and $z_{1-i} = y_{1-i}$ (here i is the input of Rec.) Third, Trans applies the same transformation to each of z_0 and z_1 and sends the results back as (b_0, b_1) . Finally, Rec decodes and outputs the correct bit. The protocol uses cryptographic primitives and computational hardness in an essential way, therefore its security is inherently only computational and its analysis requires the modeling of computational assumptions.

Our Analysis Our analysis follows the *trusted party* paradigm of [GMW87], with a formalization of secure emulation that is close in spirit to [PW00, Can01]. Our modeling involves two systems of automata:

- The *real system (RS)* consists of two automata representing protocol parties (Trans and Rec) and one automaton representing an adversarial communication service (Adv), which has access to all messages sent during the execution of the protocol. In *RS*, typical tasks include “choose random (y_0, y_1) ”, “send round 1 message”, and “deliver round 1 message”, as well as arbitrary tasks of environment and adversary automata. (The environment and adversary automata are purposely under-specified, so that our results are as general as possible.) Note that these tasks do not specify exactly what transition occurs. For example, the “choose” task does not specify the chosen values of (y_0, y_1) . And the “send” task does not specify the message contents—these are computed by Trans, based on its own internal state.
- The *ideal system (IS)* consists of an ideal *Oblivious Transfer functionality* automaton, and a *simulator* automaton, which interacts with the functionality and tries to mimic the behavior of the real system such as the messages between the protocol parties.

Then we prove that *RS* implements *IS*. The proof consists of four cases, depending on which parties are corrupted.⁶ In the two cases where Trans is corrupted, we can show that *RS* implements *IS* perfectly, using \leq_0 . In the cases where Trans is not corrupted, we can show implementation only in a “computational” sense, namely, (i) for resource-bounded adversaries, (ii) up to negligible difference in probabilities, and (iii) under computational indistinguishability assumptions. Modeling these aspects requires additions to the task-PIOA framework of this paper, namely, defining a *time-bounded* version of task-PIOAs, and defining a variation, $\leq_{neg,pt}$, of the \leq_0 relation, which describes approximate implementation with respect to polynomial-time environments. Similar relations were defined in [LMMS98, PW01]. Our simulation relations (extended with time bounds on schedule lengths) are also shown to be sound for proving $\leq_{neg,pt}$.

We also provide models for computational objects, namely, trap-door functions and hard-core predicates. Part of the specification for such objects is that their behavior should look “approximately random” to outside observers; we formalize this in terms of $\leq_{neg,pt}$.

The correctness proofs proceed by levels of abstraction, relating each pair of models at successive levels using $\leq_{neg,pt}$. In the case where only Rec is corrupted, all but one of the relationships between levels are proved using simulation relations as defined in this paper (and so, they guarantee \leq_0). The only exception is between a level in which a hard-core bit is used, and a higher level in which the hard-core bit is replaced by a random bit. Showing this correspondence relies on our $\leq_{neg,pt}$ -based definition of the cryptographic primitive, and on composition results for time-bounded task-PIOAs. Since this type of reasoning is isolated to one correspondence, the methods of this paper in fact suffice to accomplish most of the work of verifying OT.

Each of our system models, at each level, includes an explicit adversary component automaton, which acts as a message delivery service that can eavesdrop on communications and control the order of message delivery. The behavior of this adversary is arbitrary, subject to polynomial time constraints on its computational capabilities. In our models, the adversary is the same at all levels, so our simulation relations relate

⁶In [CCK⁺05b], only one case is treated in full detail—when only Rec is corrupted. We prove all four cases in [CCK⁺05a], but using a less general definition of task-PIOAs than the one used here and in [CCK⁺05b], and with non-branching adversaries.

the adversary states at consecutive levels directly, using the identity function. This treatment allows us to consider arbitrary adversaries without examining their structure in detail (they can do anything, but must do the same thing at all levels).

The following patterns arise in our simulation relation proofs. They are the motivations for our new definition of simulation relations, which has the expansion capability and relates measures to measures.

1. We often correspond random choices at two levels of abstraction—for instance, when the adversary makes a random choice, from the same state, at both levels. We would like our simulation relation to relate the individual outcomes of the choices at the two levels, matching up the states in which the same result is obtained. Modeling this correspondence uses the expansion feature.
2. The Trapdoor vs. Rand example described in Section 4 occurs in our OT proof. Here, the low-level system chooses a random y and then computes $z = f(y)$ using a trap-door permutation f . The higher level system simply chooses the value of z randomly, without using value y or permutation f . In order to correspond the steps that define z in both automata, we need to correspond the point between random choice of y and the computation of z in Trapdoor to the initial state of Rand. This correspondence relates measures to measures and uses expansion.
3. In another case, a lower-level system chooses a random value y and then computes a new value by applying XOR to y and an input value. The higher level system just chooses a random value. We establish a correspondence between the two levels using the fact that XOR preserves the uniform distribution. This correspondence again relates measures to measures and uses expansion.

6 Local Schedulers

With the action-determinism assumption, our task mechanism is enough to resolve all nondeterminism. However, action determinism limits expressive power. Now we remove this assumption and add a second mechanism for resolving the resulting additional nondeterminism, namely, a *local scheduler* for each component task-PIOA. A local scheduler for a given component can be used to resolve nondeterministic choices among actions in the same task, using only information about that component. Here, we define one type of local scheduler, which uses only the current state, and indicate how our results for the action-deterministic case carry over to this setting.

Our notion of local scheduler is simply a “sub-automaton”: We could add more expressive power by allowing the local scheduler to depend on the past execution. This could be formalized in terms of an explicit function of the past execution, or perhaps in terms of a refinement mapping or other kind of simulation relation.

Definition 6.1 We say that task-PIOA $\mathcal{T}' = (\mathcal{P}', R')$ is a sub-task-PIOA of task-PIOA $\mathcal{T} = (\mathcal{P}, R)$ provided that all components are identical except that $D' \subseteq D$, where D and D' are the sets of discrete transitions of \mathcal{P} and \mathcal{P}' , respectively. Thus, the only difference is that \mathcal{T}' may have a smaller set of transitions.

Definition 6.2 A local scheduler for a task-PIOA \mathcal{T} is any action-deterministic sub-task-PIOA of \mathcal{T} . A probabilistic system is a pair $\mathcal{M} = (\mathcal{T}, \mathcal{S})$, where \mathcal{T} is a task-PIOA and \mathcal{S} is a set of local schedulers for \mathcal{T} .

Definition 6.3 A probabilistic execution of a probabilistic system $\mathcal{M} = (\mathcal{T}, \mathcal{S})$ is defined to be any probabilistic execution of any task-PIOA $S \in \mathcal{S}$.

We next define composition for probabilistic systems.

Definition 6.4 If $\mathcal{M}_1 = (\mathcal{T}_1, \mathcal{S}_1)$ and $\mathcal{M}_2 = (\mathcal{T}_2, \mathcal{S}_2)$ are two probabilistic systems, and \mathcal{T}_1 and \mathcal{T}_2 are compatible, then their composition $\mathcal{M}_1 \parallel \mathcal{M}_2$ is the probabilistic system $(\mathcal{T}_1 \parallel \mathcal{T}_2, \mathcal{S})$, where \mathcal{S} is the set of local schedulers for $\mathcal{T}_1 \parallel \mathcal{T}_2$ of the form $S_1 \parallel S_2$, for some $S_1 \in \mathcal{S}_1$ and $S_2 \in \mathcal{S}_2$.

Definition 6.5 If $\mathcal{M} = (\mathcal{T}, \mathcal{S})$ is a probabilistic system, then an environment for \mathcal{M} is any environment (action-deterministic task-PIOA) for \mathcal{T} . If $\mathcal{M} = (\mathcal{T}, \mathcal{S})$ is a probabilistic system, then the external behavior of \mathcal{M} , $\text{extbeh}(\mathcal{M})$, is the total function that maps each environment task-PIOA \mathcal{E} for \mathcal{M} to the set $\bigcup_{S \in \mathcal{S}} \text{tdists}(S \parallel \mathcal{E})$.

Thus, for each environment, we consider the set of trace distributions that arise from two choices: of a local scheduler of \mathcal{M} and of a global task schedule ρ .

Definition 6.6 Two probabilistic systems $(\mathcal{T}_1, \mathcal{S}_1)$ and $(\mathcal{T}_2, \mathcal{S}_2)$ are comparable if \mathcal{T}_1 and \mathcal{T}_2 are comparable task-PIOAs.

We define an implementation relation for comparable probabilistic systems in terms of inclusion of sets of trace distributions for each probabilistic system based on an environment task-PIOA:

Definition 6.7 If $\mathcal{M}_1 = (\mathcal{T}_1, \mathcal{S}_1)$ and $\mathcal{M}_2 = (\mathcal{T}_2, \mathcal{S}_2)$ are comparable probabilistic systems (i.e., \mathcal{T}_1 and \mathcal{T}_2 are comparable), then \mathcal{M}_1 implements \mathcal{M}_2 , written $\mathcal{M}_1 \leq_0 \mathcal{M}_2$, provided that $\text{extbeh}(\mathcal{M}_1)(\mathcal{E}) \subseteq \text{extbeh}(\mathcal{M}_2)(\mathcal{E})$ for every environment (action-deterministic) task-PIOA \mathcal{E} for both \mathcal{M}_1 and \mathcal{M}_2 .

We obtain a sufficient condition for implementation of probabilistic systems, in which each local scheduler for the low-level system always corresponds to the same local scheduler of the high-level system.

Theorem 6.8 Let $\mathcal{M}_1 = (\mathcal{T}_1, \mathcal{S}_1)$ and $\mathcal{M}_2 = (\mathcal{T}_2, \mathcal{S}_2)$ be two comparable probabilistic systems. Suppose there is a total function f from \mathcal{S}_1 to \mathcal{S}_2 such that, for every $S_1 \in \mathcal{S}_1$, $S_1 \leq_0 f(S_1)$. Then $\mathcal{M}_1 \leq_0 \mathcal{M}_2$.

We also obtain a compositionality result for probabilistic systems. The proof is similar to that of Theorem 3.20, for the action-deterministic case.

Theorem 6.9 Let $\mathcal{M}_1, \mathcal{M}_2$ be comparable probabilistic systems such that $\mathcal{M}_1 \leq_0 \mathcal{M}_2$, and let \mathcal{M}_3 be a probabilistic system compatible with each of \mathcal{M}_1 and \mathcal{M}_2 . Then $\mathcal{M}_1 \parallel \mathcal{M}_3 \leq_0 \mathcal{M}_2 \parallel \mathcal{M}_3$.

Proof. Let $\mathcal{T}_4 = (\mathcal{P}_4, R_4)$ be any environment (action-deterministic) task-PIOA for both $\mathcal{M}_1 \parallel \mathcal{M}_3$ and $\mathcal{M}_2 \parallel \mathcal{M}_3$. Let \mathcal{M}_4 be the trivial probabilistic system $(\mathcal{T}_4, \{\mathcal{T}_4\})$. Fix any task schedule ρ_1 for $(\mathcal{T}_1 \parallel \mathcal{T}_3) \parallel \mathcal{T}_4$ and local scheduler \mathcal{P}'_{13} of $\mathcal{M}_1 \parallel \mathcal{M}_3$. Let τ be the trace distribution of $(\mathcal{T}_1 \parallel \mathcal{T}_3) \parallel \mathcal{T}_4$ generated by ρ_1 and \mathcal{P}'_{13} . It suffices to show that τ is also generated by some task schedule ρ_2 for $(\mathcal{T}_2 \parallel \mathcal{T}_3) \parallel \mathcal{T}_4$, local scheduler \mathcal{P}'_{23} of $\mathcal{M}_2 \parallel \mathcal{M}_3$, and \mathcal{P}_4 .

Note that ρ_1 is also a task schedule for $\mathcal{T}_1 \parallel (\mathcal{T}_3 \parallel \mathcal{T}_4)$. Since \mathcal{P}'_{13} is a local scheduler of $\mathcal{M}_1 \parallel \mathcal{M}_3$, it is (by definition) of the form $\mathcal{P}'_1 \parallel \mathcal{P}'_3$, where $\mathcal{P}'_1 \in \mathcal{S}_1$ and $\mathcal{P}'_3 \in \mathcal{S}_3$. Let $\mathcal{P}'_{34} = \mathcal{P}'_3 \parallel \mathcal{P}_4$. Then \mathcal{P}'_{34} is a local scheduler of $\mathcal{M}_3 \parallel \mathcal{M}_4$. Then, ρ_1, \mathcal{P}'_1 , and \mathcal{P}'_{34} generate the same trace distribution τ in the composed task-PIOA $\mathcal{T}_1 \parallel (\mathcal{T}_3 \parallel \mathcal{T}_4)$.

Define \mathcal{T}_5 to be the task-PIOA $\mathcal{T}_3 \parallel \mathcal{T}_4$. Note that \mathcal{T}_5 is an environment task-PIOA for each of \mathcal{T}_1 and \mathcal{T}_2 . Define the probabilistic system \mathcal{M}_5 to be $(\mathcal{T}_5, \{\mathcal{P}'_{34}\})$, that is, we consider just a singleton set of local schedulers, containing the one scheduler we are actually interested in.

Now, by assumption, $\mathcal{M}_1 \leq_0 \mathcal{M}_2$. Therefore, there exists a task schedule ρ_2 for $\mathcal{T}_2 \parallel \mathcal{T}_5$ and a local scheduler \mathcal{P}'_2 for \mathcal{P}_2 such that ρ_2, \mathcal{P}'_2 , and \mathcal{P}'_{34} generate the same trace distribution τ in the task-PIOA $\mathcal{T}_2 \parallel \mathcal{T}_5$. Note that ρ_2 is also a task schedule for $(\mathcal{T}_2 \parallel \mathcal{T}_3) \parallel \mathcal{T}_4$. Let $\mathcal{P}'_{23} = \mathcal{P}'_2 \parallel \mathcal{P}'_3$. Then \mathcal{P}'_{23} is a local scheduler of $\mathcal{M}_2 \parallel \mathcal{M}_3$. Also, \mathcal{P}'_4 is a local scheduler of \mathcal{M}_4 . Then $\rho_2, \mathcal{P}'_{23}$ and \mathcal{P}'_4 also generate τ , which suffices to show the required implementation relationship. \square

7 Related Work

There are numerous models in the literature that combine nondeterministic and probabilistic choices (see [SdV04] for a survey). However, very few tackle the issue of partial-information scheduling, as we do.

Exceptions include [dA99], which uses partitions on the state space to obtain partial-information schedules, and [CH05], which models local-oblivious scheduling. The former is essentially within the framework of *partially observable MDPs (POMDPs)*, originally studied in the context of reinforcement learning [KLC98]. In these frameworks, basic units of scheduling are individual actions, as opposed to equivalent classes of actions. That allows the adversary some access to random choices generated during execution, therefore these frameworks are not suitable for cryptographic modeling.

Our general approach to cryptographic protocol verification is inspired by the Interactive Turing Machine (ITM) framework used in [Can01]. There, protocol participants are modeled as ITMs and messages as bit strings written on input and output tapes. ITMs are purely probabilistic and are activated sequentially via message deliveries. This framework is well suited for computational analysis: a typical argument reduces the correctness of a protocol to the computational hardness (or indistinguishability) assumptions associated with the underlying cryptographic primitives. However, a complete and rigorous analysis is often impractical, because it involves too many low-level machine details. Indeed, in the community of computational cryptography, protocols are typically described using an informal high-level language, and proof sketches are given in terms of these informal descriptions.

Thus, our motivation is to provide a framework in which protocols can be specified clearly and concisely, and computational proofs can be carried out formally and at a high level of abstraction. This motivation is certainly not new: several other research groups have taken conventional concurrency modeling frameworks and added features intended for computational cryptographic analysis. This includes the *Probabilistic Polynomial-time process Calculus (PPC)* framework of Mitchell et al. [LMMS98, MMS03, MRST06] and the *Reactive Simulatability (RSIM)* framework of Pfitzmann et al. [PW00, PW01, BPW04]. However, the semantic foundations of concurrent computation used in these frameworks differ from task-PIOAs in some fundamental ways. These differences arise in part because we aim to exploit the benefits of nondeterminism to a greater extent. Below we give a brief summary of these two frameworks and compare them against task-PIOAs.

PPC On a conceptual level, the PPC and task-PIOA frameworks handle concurrency and nondeterminism in a similar way. Processes with nondeterministic choices are definable in both frameworks, and a global scheduler is used to resolve nondeterminism.⁷ In [MRST06], the scheduler is a probabilistic polynomial-time function that selects an action label from the set of enabled actions. Typically, action labels in PPC correspond to the types of protocol messages, as opposed to the messages themselves. This is similar to our distinction between tasks and actions.

Despite these similarities, a major difference exists between the operational semantics of PPC and our task-scheduling mechanism. Our task schedules are oblivious sequences of tasks. Since tasks may represent either internal computations or external communications, a task schedule resolves choices involving both types of activities with no priority restrictions. In PPC, however, internal computations are prioritized over secure communications, which are in turn prioritized over insecure communications. The probabilistic polynomial-time scheduler is invoked when no more internal computations are possible.

Like our task schedules, the scheduler function of PPC is an entity distinct from the adversary. To faithfully model adversarial scheduling, one must construct the PPC adversary contexts in such a way that message timing is controlled by the adversary and not by the scheduler function. This point is observed in [MMS03] and it corresponds to our proposal that high-level nondeterminism is resolved by an adversary automaton. Thus, the real difference between PPC’s scheduling mechanism and our own lies in the handling of low-level nondeterminism. In our view, task-based scheduling is a simple and sufficient method for that purpose.

The PPC framework differs from our task-PIOA framework in another respect, namely, the use of observational equivalence and probabilistic bisimulation as the main semantic relations. Both of these are symmetric relations, whereas our implementation and simulation relations are asymmetric, expressing the idea that a system P can emulate another system Q but the converse is not necessarily true. The asymmetry of our definitions arises from our quantification over schedules: we assert that “for every schedule of P ,

⁷The authors have also developed a sequential version of PPC [DKMR05], with a semantics akin to the ITM and RSIM frameworks.

there is a schedule of Q that yields equivalent behavior.” This is analogous to the traditional formulation for non-probabilistic systems, where implementation means “every behavior of P is a behavior of Q ,” with no requirement imposed in the other direction. In our experience with distributed algorithms, this asymmetry can be used to make specifications more simple, by keeping irrelevant details unspecified in the abstract specification. At the same time, it produces guarantees that are more general, because correctness statements will hold no matter how an implementer chooses to fill in the unspecified details.

In PPC, symmetric semantic relations are sufficient because process expressions typically do not contain low-level nondeterminism. For example, under the grammar of PPC, there is no natural way to express a process P that performs actions a and b in either order and then proceeds as process P' . This is because PPC allows nondeterministic choices only via parallel composition, but one cannot prepend $a||b$ to a separate process P' . In contrast, this type of choices is very common in our models (cf. our OT model in [CCK⁺05b]). Using a simulation relation, we are able to handle the two cases (i.e., whether a precedes b or vice versa) separately but without duplicating any arguments involving P' .

To sum up, the PPC and task-PIOAs frameworks share similar goals but have different features and strengths. Having a process algebraic foundation, PPC proofs are based on formal deduction, hence very amenable to mechanization. However, process expressions constructed from a formal syntax are less flexible compared to our automata-based specifications. These differences are inherent and, because of them, one framework may be more appropriate than the other for a given application.

RSIM Next we consider the RSIM framework of Pfitzmann et al. There a protocol is modeled by a network of interrupt-driven probabilistic state machines, with special “buffer” machines to capture message delays, and special “clock ports” to control the scheduling of message delivery. Each individual machine is purely probabilistic; that is, it is fully-specified up to inputs and/or random choices generated during execution. In particular, if a system of machines is closed (i.e., every input of every machine is provided by some machine in the system), it is fully specified up to coin tosses.

For a closed RSIM system, a sequential activation scheme is used to define a unique probabilistic run for each possible initial state of the system. Under this scheme, machines are switched one at a time, and the current active machine M_1 selects the next machine M_2 by scheduling a message intended for M_2 . (This is possible only when M_1 “owns” the clock port to an input buffer of M_2 .) If no machines are active, a special “master scheduler” machine is triggered by default.

Thus, in order to capture nondeterministic choices within the RSIM framework, one must associate explicit inputs to each schedulable event and then quantify over different machines that provide these scheduling inputs. That is, each scheduler machine corresponds to one particular schedule. In contrast, nondeterministic choices can be expressed in a task-PIOA without the use of explicit inputs, and we quantify over task schedules to capture the possible ways of resolving nondeterminism.

These technical differences are again the result of our divergent views on nondeterminism. The RSIM framework is designed with high-level nondeterminism in mind. Therefore, message buffering and delivery are made explicit, while internal computations are entirely encapsulated within single-step transitions. As in PPC, there is no natural way (aside from adding inputs and scheduler machines) to express nondeterministic choices that correspond to implementation freedom.

Aside from practical consequences, the different treatments of nondeterminism also affect the meaning of security definitions. For example, in the definition of reactive simulatability, the user and adversary are fixed *after* all other machines are determined. Essentially, this allows the worst possible adversary for every schedule of the system. In our security definitions [CCK⁺06c, CCK⁺05b], the environment⁸ and adversary are fixed *before* the task schedules. Therefore, we consider instead the worst possible schedule for each given adversary. In light of the separation result in [CCLP07], we believe RSIM and task-PIOA definitions are *not* equivalent, because RSIM is based on sequential scheduling.

⁸The role of “user” in RSIM is comparable to the role of “environment” in our definitions.

8 Conclusions

We have extended the traditional PIOA model with a task mechanism, which provides a systematic way of resolving nondeterministic scheduling choices without using dynamically generated information. In the resulting Task-PIOA framework, we develop basic machinery for verification, including a compositional trace-based semantics and a new kind of simulation relation that is sound for proving implementation. The utility of these tools are demonstrated in the Oblivious Transfer case study, which is outlined in this paper and presented in full in [CCK⁺05b]. We have also proposed a further extension, in which local nondeterminism are resolved by schedulers that use local information only.

Although our development is motivated by concerns in cryptographic protocol analysis, the notion of partial-information scheduling is interesting in its own right. For example, some distributed algorithms are designed using partial-information scheduling assumptions, because the problems they address are provably unsolvable under perfect-information scheduling [Cha96, Asp03]. Also, one may argue that partial-information scheduling is more realistic in models of large distributed systems, in which basic scheduling decisions are made locally, and not by any centralized mechanism.

There are many more interesting problems in our general project of cryptographic protocol verification. Below we discuss some of them.

Computational Assumptions We would like to express standard computational assumptions in terms of implementation relations between task-PIOAs. Judging from our results on hard-core predicates [CCK⁺05b], decisional assumptions can be formulated quite easily by comparing the accept probabilities of a distinguisher environment. To handle hardness assumptions, we need to formalize the usual “conversion” argument: any successful adversary can be converted into an efficient algorithm for solving a difficult problem.

Statistical Indistinguishability Another problem is to formalize the notion of statistical indistinguishability. This yields a third type of implementation relation, in addition to perfect implementation \leq_0 and the computational approximate implementation $\leq_{neg,pt}$. In [ML06a], metric distances between trace distributions are used to define approximate implementations and approximate simulations for individual task-PIOAs. These definitions are extended to families of task-PIOAs, thereby obtaining a notion of statistical approximate implementation relation [CMP07].

Simulation-Based Security Definitions Our security definition falls under the category of simulation-based security. In [DKMR05], various simulation-based security notions, including universally composable security [Can01] and reactive simulatability [PW00, PW01], are translated into the sequential version of PPC. A semantic hierarchy of these notions is given, based on the placement of a master process and the ability to forward communication between processes. Interestingly, there is no obvious way to place our notion of security in this hierarchy. This is because there is no inherent notion of master process in our definitions. Moreover, our task schedules are determined *after* the adversary, simulator, and environment automata have been fixed. This additional quantification is not found in other security definitions. We investigate the role of scheduling in simulation-based security and show that different choices of scheduling mechanisms give rise to differences in the meanings of security definitions in [CCLP07].

Composition Theorem In [CCK⁺07], we prove that our notion of secure emulation is preserved under a polynomial number of protocol substitutions. While the result is standard, the actual proof in our framework requires a nontrivial adaptation of the usual hybrid argument. This reveals some interesting implications of our definitions. In particular, we allow a time-bounded task-PIOA to continue its computation indefinitely, even though each individual transition is of bounded complexity. A task schedule thus takes on two roles: resolving nondeterminism and limiting the number of individual transitions. This separation between single step complexity and schedule length complexity is a design decision, because eventually we are interested in proving long-term correctness properties. As a result, our definition of secure emulation involves an

additional quantification over schedule length bounds, which makes the hybrid argument more delicate than usual.

Mechanized Proofs The simulation proofs in our OT case study are quite similar in character to typical analysis of distributed algorithms. Namely, we establish correspondence between two specifications by examining their respective behaviors at comparable stages in the protocol. Invariant-style properties are often invoked to prove these correspondences. This type of reasoning is a good candidate for mechanization. In fact, we are quite confident that the (non-probabilistic) invariant properties in our proofs can be proven in a straightforward fashion using a theorem prover such as PVS. Since our language is based on I/O automata, translation into PVS (without probabilities) can be done systematically using the TEMPO Toolset [Too] and the TAME interface [Arc]. It will be very interesting to explore the possibility of mechanizing a greater portion of our simulation proofs. This may involve the development of new reasoning techniques to handle sets of traces (as opposed to individual traces or trace distributions).

Acknowledgments

We thank Frits Vaandrager for collaboration in early stages of this project, and Michael Backes, Anupam Datta, Joshua Guttman, Jon Herzog, Susan Hohenberger, Ralf Kuesters, John Mitchell, Birgit Pfitzmann and Andre Scedrov for technical discussions that helped us in clarifying our ideas and their connections to other work in analysis of cryptographic protocols. We thank Silvio Micali for impressing upon us the importance of adaptive adversarial schedulers in the cryptographic setting. We thank Sayan Mitra both for technical discussions and for help in typesetting the paper.

Canetti's work on this project was supported by NSF CyberTrust Grant #0430450. Cheung was supported by DFG/NWO bilateral cooperation project 600.050.011.01 Validation of Stochastic Systems (VOSS) and by NSF Award #CCR-0326277. Part of this work was done when she was at Radboud University Nijmegen. Kaynar and Lynch were supported by DARPA/AFOSR MURI Award #F49620-02-1-0325, MURI AFOSR Award #SA2796PO 1-0000243658, NSF Awards #CCR-0326277 and #CCR-0121277, and USAF, AFRL Award #FA9550-04-1-0121, and Kaynar was supported by US Army Research Office grant #DAAD19-01-1-0485. Pereira was supported by the Belgian National Fund for Scientific Research (F.R.S.-FNRS).

References

- [AH90] J. Aspnes and M. Herlihy. Fast randomized consensus using shared memory. *Journal of Algorithms*, 11(3):441–461, 1990.
- [Arc] M. Archer. TAME support for refinement proofs. Available at <http://chacs.nrl.navy.mil/personnel/archer.html>.
- [Asp03] J. Aspnes. Randomized protocols for asynchronous consensus. *Distributed Computing*, 16(2-3):165–175, 2003.
- [Bea91] D. Beaver. Secure multiparty protocols and zero-knowledge proof systems tolerating a faulty minority. *Journal of Cryptology*, 4(2):75–122, 1991.
- [BK98] C. Baier and M. Kwiatkowska. Model checking for a probabilistic branching time logic with fairness. *Distributed Computing*, 11(3):125–155, 1998.
- [BO83] M. Ben-Or. Another advantage of free choice: completely asynchronous agreement protocols. In *Proc. 2nd ACM Symposium on Principles of Distributed Computing*, pages 27–30, Montreal, Quebec, Canada, August 1983.
- [BPW04] M. Backes, B. Pfitzmann, and M. Waidner. Secure asynchronous reactive systems. *Cryptology ePrint Archive Report 2004/082*, 2004.

- [Can95] R. Canetti. *Studies in Secure Multi-Party Computation and Applications*. PhD thesis, Weizmann Institute, Israel, 1995.
- [Can01] R. Canetti. Universally composable security: a new paradigm for cryptographic protocols. In *Proc. 42nd IEEE Symposium on Foundations of Computing*, pages 136–145, 2001.
- [CCK⁺05a] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala. Using probabilistic I/O automata to analyze an oblivious transfer protocol. Technical Report MIT-LCS-TR-1001a or MIT-CSAIL-TR-2005-055, MIT CSAIL, 2005.
- [CCK⁺05b] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala. Using task-structured probabilistic I/O automata to analyze an oblivious transfer protocol. Cryptology ePrint Archive, Report 2005/452, 2005. Available at <http://eprint.iacr.org/>.
- [CCK⁺06a] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala. Task-structured probabilistic I/O automata. Technical Report MIT-CSAIL-TR-2006-023, MIT CSAIL, 2006.
- [CCK⁺06b] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala. Task-structured probabilistic I/O automata. In *Proceedings of the 8th International Workshop on Discrete Event Systems (WODES'06)*, 2006. Ann Arbor, Michigan, July 2006.
- [CCK⁺06c] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala. Time-bounded task-PIOAs: a framework for analyzing security protocols. In *Proceedings of the 20th International Symposium on Distributed Computing (DISC '06)*, 2006. Stockholm, Sweden, September 2006.
- [CCK⁺07] R. Canetti, L. Cheung, D. Kaynar, N. Lynch, and Olivier Pereira. Compositional security for Task-PIOAs. In *Proceedings of the 20th IEEE Computer Security Foundations Symposium (CSF-20)*, 2007. To appear.
- [CCLP07] R. Canetti, L. Cheung, N. Lynch, and O. Pereira. On the role of scheduling in simulation-based security. In *Proceedings of the 7th International Workshop on Issues in the Theory of Security (WITS '07)*, 2007. Available at <http://eprint.iacr.org/2007/102>.
- [CH05] L. Cheung and M. Hendriks. Causal dependencies in parallel composition of stochastic processes. Technical Report ICIS-R05020, Institute for Computing and Information Sciences, University of Nijmegen, 2005.
- [Cha96] T.D. Chandra. Polylog randomized wait-free consensus. In *Proc. 15th ACM Symposium on Principles of Distributed Computing*, pages 166–175, 1996.
- [CLK⁺06] R. Canetti, L. Cheung, D. Kaynar, N. Lynch, M. Liskov, O. Pereira, and R. Segala. Using task-structured probabilistic I/O automata to analyze cryptographic protocol. In V. Cortier and S. Kremer, editors, *Proceedings of Workshop on Formal and Computational Cryptography (FCC '06)*, pages 34–39, 2006.
- [CLSV06] L. Cheung, N. Lynch, R. Segala, and F. Vaandrager. Switched PIOA: Parallel composition via distributed scheduling. *Theoretical Computer Science*, 365(1–2):83–108, 2006.
- [CMP07] Ling Cheung, Sayan Mitra, and Olivier Pereira. Verifying statistical zero knowledge with approximate implementations. Cryptology ePrint Archive, Report 2007/195, 2007. Available at <http://eprint.iacr.org/2007/195>.
- [dA99] L. de Alfaro. The verification of probabilistic systems under memoryless partial-information policies is hard. In *Proc. PROBMIV 99*, pages 19–32, 1999.

- [DKMR05] A. Datta, R. Küsters, J.C. Mitchell, and A. Ramanathan. On the relationships between notions of simulation-based security. In *Proceedings TCC 2005*, pages 476–494, 2005.
- [GL90] S. Goldwasser and L. Levin. Fair computation of general functions in presence of immoral majority. In Alfred J. Menezes and Scott A. Vanstone, editors, *Advances in Cryptology - Crypto '90*, pages 77–93, Berlin, 1990. Springer-Verlag. Lecture Notes in Computer Science Volume 537.
- [GMR85] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing (STOC'85)*, pages 291–304, 1985.
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *Proc. 19th ACM Symposium on Theory of Computing (STOC)*, pages 218–229, 1987.
- [GPS06] Marc Girault, Guillaume Poupard, and Jacques Stern. On the fly authentication and signature schemes based on groups of unknown order. *Journal of Cryptology*, 19(4):463–487, 2006.
- [KLC98] L.P. Kaelbling, M.L. Littman, and A.R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.
- [LMMS98] P. Lincoln, J.C. Mitchell, M. Mitchell, and A. Scedrov. A probabilistic poly-time framework for protocol analysis. In *ACM Conference on Computer and Communications Security*, pages 112–121, 1998.
- [LSS94] N.A. Lynch, I. Saias, and R. Segala. Proving time bounds for randomized distributed algorithms. In *Proc. 13th ACM Symposium on the Principles of Distributed Computing*, pages 314–323, 1994.
- [LSV03] N.A. Lynch, R. Segala, and F.W. Vaandrager. Compositionality for probabilistic automata. In *Proc. 14th International Conference on Concurrency Theory (CONCUR 2003)*, volume 2761 of *LNCS*, pages 208–221. Springer-Verlag, 2003.
- [LT89] N.A. Lynch and M.R. Tuttle. An introduction to input/output automata. *CWI Quarterly*, 2(3):219–246, September 1989.
- [ML06a] S. Mitra and N.A. Lynch. Probabilistic timed I/O automata with continuous state spaces. Preliminary version available at http://theory.lcs.mit.edu/~mitras/research/csptioa_preprint.pdf, May 2006.
- [ML06b] Sayan Mitra and Nancy Lynch. Approximate implementation relations for probabilistic I/O automata, October 2006. To appear in ENTCS.
- [MMS03] P. Mateus, J.C. Mitchell, and A. Scedrov. Composition of cryptographic protocols in a probabilistic polynomial-time process calculus. In *Proc. 14th International Conference on Concurrency Theory (CONCUR 2003)*, pages 323–345, 2003.
- [MR91] S. Micali and P. Rogaway. Secure computation. In Joan Feigenbaum, editor, *Advances in Cryptology - Crypto '91*, pages 392–404, Berlin, 1991. Springer-Verlag. Lecture Notes in Computer Science Volume 576.
- [MRST06] J. Mitchell, A. Ramanathan, A. Scedrov, and V. Teague. A probabilistic polynomial-time process calculus for the analysis of cryptographic protocols. *Theoretical Computer Science*, 353:118–164, 2006.
- [PSL00] A. Pogosyants, R. Segala, and N.A. Lynch. Verification of the randomized consensus algorithm of Aspnes and Herlihy: a case study. *Distributed Computing*, 13(3):155–186, 2000.

- [Put94] M.L. Puterman. *Markov Decision Process – Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, 1994.
- [PW94] Birgit Pfitzmann and Michael Waidner. A general framework for formal notions of “secure” system. Technical report, Hildesheimer Informatik-Berichte 11/94, Institut für Informatik, Universität Hildesheim., April 1994.
- [PW00] B. Pfitzmann and M. Waidner. Composition and integrity preservation of secure reactive systems. In *Proc. 7th ACM Conference on Computer and Communications Security (CCS 2000)*, pages 245–254, 2000.
- [PW01] B. Pfitzmann and M. Waidner. A model for asynchronous reactive systems and its application to secure message transmission. In *Proc. IEEE Symposium on Research in Security and Privacy*, pages 184–200, 2001.
- [Rab82] M. Rabin. The choice coordination problem. *Acta Informatica*, 17:121–134, 1982.
- [SdV04] A. Sokolova and E.P. de Vink. Probabilistic automata: system types, parallel composition and comparison. In *Validation of Stochastic Systems*, volume 2925 of *LNCS*, pages 1–43. Springer-Verlag, 2004.
- [Seg95] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, June 1995. Available as Technical Report MIT/LCS/TR-676.
- [SL95] R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2(2):250–273, 1995.
- [SV99] M.I.A. Stoelinga and F.W. Vaandrager. Root contention in IEEE 1394. In *Proc. 5th International AMAST Workshop on Formal Methods for Real-Time and Probabilistic Systems*, volume 1601 of *LNCS*, pages 53–74. Springer-Verlag, 1999.
- [Too] TAME Tempo Toolset. Available at <http://www.veromodo.com/Veromodo/Tempo-Toolset.html>.

A σ -Fields of Execution Fragments and Traces

In order to define probability measures on executions and traces, we need appropriate σ -fields. We begin with a σ -field over the set of execution fragments of a PIOA \mathcal{P} :

Definition A.1 *The cone of a finite execution fragment α , denoted by C_α , is the set $\{\alpha' \in \text{Frag}(\mathcal{P}) \mid \alpha \leq \alpha'\}$. Then $\mathcal{F}_\mathcal{P}$ is the σ -field generated by the set of cones of finite execution fragments of \mathcal{P} .*

A probability measure on execution fragments of \mathcal{P} is then simply a probability measure on the σ -field $\mathcal{F}_\mathcal{P}$.

Since Q , I , O , and H are countable, $\text{Frag}^*(\mathcal{P})$ is countable, and hence the set of cones of finite execution fragments of \mathcal{P} is countable. Therefore, any union of cones is measurable. Moreover, for each finite execution fragment α , the set $\{\alpha\}$ is measurable since it can be expressed as the intersection of C_α with the complement of $\bigcup_{\alpha': \alpha < \alpha'} C_{\alpha'}$. Thus, any set of finite execution fragments is measurable; in other words, the discrete σ -field of finite executions is included in $\mathcal{F}_\mathcal{P}$.

We often restrict our attention to probability measures on finite execution fragments, rather than those on arbitrary execution fragments. Thus, we define:

Definition A.2 *Let ϵ be a probability measure on execution fragments of \mathcal{P} . We say that ϵ is finite if $\text{Frag}^*(\mathcal{P})$ is a support for ϵ .*

Since any set of finite execution fragments is measurable, any finite probability measure on execution fragments of \mathcal{P} can also be viewed as a discrete probability measure on $\text{Frag}^*(\mathcal{P})$. Formally, given any finite probability measure ϵ on execution fragments of \mathcal{P} , we obtain a discrete probability measure $\text{finite}(\epsilon)$ on $\text{Frag}^*(\mathcal{P})$ by simply defining $\text{finite}(\epsilon)(\alpha) = \epsilon(\{\alpha\})$ for every finite execution fragment α of \mathcal{P} . The difference between $\text{finite}(\epsilon)$ and ϵ is simply that the domain of ϵ is $\mathcal{F}_{\mathcal{P}}$, whereas the domain of $\text{finite}(\epsilon)$ is $\text{Execs}^*(\mathcal{P})$. Henceforth, we will ignore the distinction between $\text{finite}(\epsilon)$ and ϵ .

Definition A.3 A chain of probability measures on execution fragments of PIOA \mathcal{P} is an infinite sequence, $\epsilon_1, \epsilon_2, \dots$ of probability measures on execution fragments of \mathcal{P} such that, for each $i \geq 0$, $\epsilon_i \leq \epsilon_{i+1}$ (cf. Definition 2.15). Given a chain $\epsilon_1, \epsilon_2, \dots$ of probability measures on execution fragments of \mathcal{P} , we define a new function ϵ on the σ -field generated by cones of execution fragments of \mathcal{P} as follows: for each finite execution fragment α ,

$$\epsilon(C_\alpha) = \lim_{i \rightarrow \infty} \epsilon_i(C_\alpha).$$

Standard measure theoretic arguments ensure that ϵ can be extended uniquely to a probability measure on the σ -field generated by the cones of finite execution fragments. Furthermore, for each $i \geq 0$, $\epsilon_i \leq \epsilon$. We call ϵ the limit of the chain, and we denote it by $\lim_{i \rightarrow \infty} \epsilon_i$.

If α is a finite execution fragment of a PIOA \mathcal{P} and a is an action of \mathcal{P} , then $C_{\alpha a}$ denotes the set of execution fragments of \mathcal{P} that start with αa . The cone construction can also be used to define a σ -field of traces:

Definition A.4 The cone of a finite trace β , denoted by C_β , is the set $\{\beta' \in E^* \cup E^\omega \mid \beta \leq \beta'\}$, where \leq denotes the prefix ordering on sequences. The σ -field of traces of \mathcal{P} is simply the σ -field generated by the set of cones of finite traces of \mathcal{P} .

Again, the set of cones is countable and the discrete σ -field on finite traces is included in the σ -field generated by cones of traces. We often refer to a probability measure on the σ -field generated by cones of traces of a PIOA \mathcal{P} as simply a probability measure on traces of \mathcal{P} .

Definition A.5 Let τ be a probability measure on traces of \mathcal{P} . We say that τ is finite if the set of finite traces is a support for τ . Any finite probability measure on traces of \mathcal{P} can also be viewed as a discrete probability measure on the set of finite traces.

Definition A.6 Let τ and τ' be probability measures on traces of PIOA \mathcal{P} . Then we say that τ is a prefix of τ' , denoted by $\tau \leq \tau'$, if, for each finite trace β of \mathcal{P} , $\tau(C_\beta) \leq \tau'(C_\beta)$.

Definition A.7 A chain of probability measures on traces of PIOA \mathcal{P} is an infinite sequence, τ_1, τ_2, \dots of probability measures on traces of \mathcal{P} such that, for each $i \geq 0$, $\tau_i \leq \tau_{i+1}$. Given a chain τ_1, τ_2, \dots of probability measures on traces of \mathcal{P} , we define a new function τ on the σ -field generated by cones of traces of \mathcal{P} as follows: for each finite trace β ,

$$\tau(C_\beta) = \lim_{i \rightarrow \infty} \tau_i(C_\beta).$$

Then τ can be extended uniquely to a probability measure on the σ -field of cones of finite traces. Furthermore, for each $i \geq 0$, $\tau_i \leq \tau$. We call τ the limit of the chain, and we denote it by $\lim_{i \rightarrow \infty} \tau_i$.

Recall from Section 2.2 the definition of the trace distribution $\text{tdist}(\epsilon)$ of a probability measure ϵ on execution fragments. Namely, $\text{tdist}(\epsilon)$ is the image measure of ϵ under the measurable function trace .

Lemma A.8 Let $\epsilon_1, \epsilon_2, \dots$ be a chain of measures on execution fragments, and let ϵ be $\lim_{i \rightarrow \infty} \epsilon_i$. Then $\lim_{i \rightarrow \infty} \text{tdist}(\epsilon_i) = \text{tdist}(\epsilon)$.

Proof. It suffices to show that, for any finite trace β , $\lim_{i \rightarrow \infty} \text{tdist}(\epsilon_i)(C_\beta) = \text{tdist}(\epsilon)(C_\beta)$. Fix a finite trace β .

Let Θ be the set of minimal execution fragments whose trace is in C_β . Then $\text{trace}^{-1}(C_\beta) = \cup_{\alpha \in \Theta} C_\alpha$, where all the cones are pairwise disjoint. Therefore, for $i \geq 0$, $\text{tdist}(\epsilon_i)(C_\beta) = \sum_{\alpha \in \Theta} \epsilon_i(C_\alpha)$, and $\text{tdist}(\epsilon)(C_\beta) = \sum_{\alpha \in \Theta} \epsilon(C_\alpha)$.

Since we have monotone limits here (that is, our limits are also suprema), limits commute with sums and our goal can be restated as showing:

$$\sum_{\alpha \in \Theta} \lim_{i \rightarrow \infty} \epsilon_i(C_\alpha) = \sum_{\alpha \in \Theta} \epsilon(C_\alpha).$$

Since $\lim_{i \rightarrow \infty} \epsilon_i = \epsilon$, we have $\lim_{i \rightarrow \infty} \epsilon_i(C_\alpha) = \epsilon(C_\alpha)$ for each finite execution fragment α . Therefore, the two sums above are in fact equal. \square

The lstate function is a measurable function from the discrete σ -field of finite execution fragments of \mathcal{P} to the discrete σ -field of states of \mathcal{P} . If ϵ is a probability measure on execution fragments of \mathcal{P} , then we define the lstate distribution of ϵ , $\text{lstate}(\epsilon)$, to be the image measure of ϵ under the function lstate .

B Task Schedules

The following theorem states that, for any measure μ and task sequence ρ , the probability measure on execution fragments generated by $\text{apply}(\mu, \rho)$ is “standard”, in the sense that it can be obtained from μ and a scheduler as defined in Section 3 for basic PIOAs.

Theorem B.1 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. For each probability measure μ on $\text{Frag}_s^*(\mathcal{P})$ and task schedule ρ , there is scheduler σ for \mathcal{P} such that $\text{apply}(\mu, \rho)$ is the generalized probabilistic execution fragment $\epsilon_{\sigma, \mu}$.*

The proof of Theorem 3.13 uses several auxiliary lemmas. We give the full proof in Appendix B. The first lemma is about applying λ , the empty sequence of tasks. It is used in the base case of the inductive proof for Lemma B.4, which involves applying any finite sequence of tasks.

Lemma B.2 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. Let μ be a discrete probability measure over finite execution fragments. Then $\text{apply}(\mu, \lambda)$ is a generalized probabilistic execution fragment generated by μ .*

Proof. Follows directly from the definitions, by defining a scheduler σ such that $\sigma(\alpha)(\text{tran}) = 0$ for each finite execution fragment α and each transition tran . \square

The next lemma provides the inductive step needed for Lemma B.4.

Lemma B.3 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. If ϵ is a generalized probabilistic execution fragment generated by a measure μ , then, for each task T , $\text{apply}(\epsilon, T)$ is a generalized probabilistic execution fragment generated by μ .*

Proof. Suppose ϵ is generated by μ together with a scheduler σ (that is, $\epsilon_{\sigma, \mu} = \epsilon$). Let ϵ' be $\text{apply}(\epsilon, T)$. For each finite execution fragment α , let $D(\text{lstate}(\alpha))$ denote the set of transitions of D with source state $\text{lstate}(\alpha)$. For each $\text{tran} \in D$, let $\text{act}(\text{tran})$ denote the action that occurs in tran . Now we define a new scheduler σ' as follows: given finite execution fragment α and $\text{tran} \in D$,

- if $\epsilon'(C_\alpha) - \mu(C_\alpha - \{\alpha\}) = 0$, then $\sigma'(\alpha)(\text{tran}) = 0$;

- otherwise, if $\text{tran} \in D(\text{lstate}(\alpha))$ and $\text{act}(\text{tran}) \in T$, then

$$\sigma'(\alpha)(\text{tran}) = \frac{\epsilon(C_\alpha) - \mu(C_\alpha - \{\alpha\})}{\epsilon'(C_\alpha) - \mu(C_\alpha - \{\alpha\})} (\sigma(\alpha)(\text{tran}) + \sigma(\alpha)(\perp));$$

- otherwise,

$$\sigma'(\alpha)(\text{tran}) = \frac{\epsilon(C_\alpha) - \mu(C_\alpha - \{\alpha\})}{\epsilon'(C_\alpha) - \mu(C_\alpha - \{\alpha\})} \sigma(\alpha)(\text{tran}).$$

We first argue that σ' , thus defined, is a scheduler. Let a finite execution fragment α be given. If the first clause applies, then $\sigma'(\alpha)$ is 0 everywhere, hence is a sub-probability measure. Assume otherwise. By action- and transition-determinism, there is at most one tran with $\text{tran} \in D(\text{lstate}(\alpha))$ and $\text{act}(\text{tran}) \in T$. Let Y denote $\{\text{tran}\}$ if such tran exists and \emptyset otherwise. Then we have the following.

$$\begin{aligned} & \sum_{\text{tran} \notin Y} \sigma(\alpha)(\text{tran}) + \sum_{\text{tran} \in Y} (\sigma(\alpha)(\text{tran}) + \sigma(\alpha)(\perp)) \\ &= \left(\sum_{\text{tran} \in D} \sigma(\alpha)(\text{tran}) \right) + \sigma(\alpha)(\perp) && Y \text{ is either empty or a singleton} \\ &= 1 && \sigma \text{ is a scheduler} \end{aligned}$$

Furthermore, by Lemma 3.7, we know that $\epsilon(C_\alpha) \leq \epsilon'(C_\alpha)$, thus the fraction $\frac{\epsilon(C_\alpha) - \mu(C_\alpha - \{\alpha\})}{\epsilon'(C_\alpha) - \mu(C_\alpha - \{\alpha\})}$ is at most 1. Putting the pieces together, we have

$$\sum_{\text{tran} \in D} \sigma'(\alpha)(\text{tran}) = \frac{\epsilon(C_\alpha) - \mu(C_\alpha - \{\alpha\})}{\epsilon'(C_\alpha) - \mu(C_\alpha - \{\alpha\})} \cdot \left(\sum_{\text{tran} \notin Y} \sigma(\alpha)(\text{tran}) + \sum_{\text{tran} \in Y} (\sigma(\alpha)(\text{tran}) + \sigma(\alpha)(\perp)) \right) \leq 1.$$

Next, we prove by induction on the length of a finite execution fragment α that $\epsilon_{\sigma', \mu}(C_\alpha) = \epsilon'(C_\alpha)$.

For the base case, let $\alpha = q$. By Lemma 2.10,

$$\epsilon_{\sigma', \mu}(C_q) = \mu(C_q) = \epsilon_{\sigma, \mu}(C_q).$$

By the choice of σ , the last expression equals $\epsilon(C_q)$, which in turn is equal to $\epsilon'(C_q)$ by virtue of Lemma 3.10. Thus, $\epsilon_{\sigma', \mu}(C_q) = \epsilon'(C_q)$, as needed.

For the inductive step, let $\alpha = \tilde{\alpha}aq$. By Lemma 2.10 and the definition of the measure of a cone, we get

$$\epsilon_{\sigma', \mu}(C_\alpha) = \mu(C_\alpha) + \sum_{\alpha' < \alpha} \mu(\alpha') \epsilon_{\sigma', \alpha'}(C_\alpha) = \mu(C_\alpha) + \sum_{\alpha' \leq \tilde{\alpha}} \mu(\alpha') \epsilon_{\sigma', \alpha'}(C_{\tilde{\alpha}}) \mu_{\sigma'(\tilde{\alpha})}(a, q).$$

We know that a is enabled from $\text{lstate}(\tilde{\alpha})$, because α is an execution fragment of \mathcal{P} . Thus, $\text{tran}_{\tilde{\alpha}, a}$ and $\mu_{\tilde{\alpha}, a}$ are defined. By expanding $\mu_{\sigma'(\tilde{\alpha})}(a, q)$ in the equation above, we get

$$\epsilon_{\sigma', \mu}(C_\alpha) = \mu(C_\alpha) + \sum_{\alpha' \leq \tilde{\alpha}} \mu(\alpha') \epsilon_{\sigma', \alpha'}(C_{\tilde{\alpha}}) \sigma'(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) \mu_{\tilde{\alpha}, a}(q). \quad (1)$$

We distinguish three cases.

1. $\epsilon'(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) = 0$.

By inductive hypothesis, $\epsilon_{\sigma', \mu}(C_{\tilde{\alpha}}) = \epsilon'(C_{\tilde{\alpha}})$. Then by Lemma 2.12, $\epsilon_{\sigma', \mu}(C_\alpha) = \mu(C_\alpha)$. It is therefore sufficient to show that $\epsilon'(C_\alpha) = \mu(C_\alpha)$.

By Lemma 3.7, $\epsilon(C_{\tilde{\alpha}}) \leq \epsilon'(C_{\tilde{\alpha}})$. Thus, using $\epsilon'(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) = 0$, we get $\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) \leq 0$. On the other hand, from Lemma 2.11 and the fact that $\epsilon = \epsilon_{\sigma, \mu}$, we have $\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) \geq 0$.

$\{\tilde{\alpha}\} \geq 0$. Thus, $\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) = 0$. Now, using Lemma 2.12 and the fact that $\epsilon_{\sigma, \mu} = \epsilon$ and $\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) = 0$, we get $\epsilon(C_{\alpha}) = \mu(C_{\alpha})$.

Since $C_{\tilde{\alpha}} - \{\tilde{\alpha}\}$ is a union of cones, we may use Lemma 3.7 to obtain $\mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) \leq \epsilon(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})$. Adding $\epsilon(\{\tilde{\alpha}\})$ on both sides, we get $\mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) + \epsilon(\{\tilde{\alpha}\}) \leq \epsilon(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) + \epsilon(\{\tilde{\alpha}\}) = \epsilon(C_{\tilde{\alpha}})$. Since $\epsilon(C_{\tilde{\alpha}}) = \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})$, the previous inequalities imply $\epsilon(C_{\tilde{\alpha}}) + \epsilon(\{\tilde{\alpha}\}) \leq \epsilon(C_{\tilde{\alpha}})$, therefore $\epsilon(\{\tilde{\alpha}\}) = 0$. By Lemma 3.6 (Items (2) and (3)), we have $\epsilon'(C_{\alpha}) = \epsilon(C_{\alpha}) = \mu(C_{\alpha})$, as needed.

2. $\epsilon'(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) > 0$ and $a \notin T$.

By Equation (1) and the definition of σ' , we know that $\epsilon_{\sigma', \mu}(C_{\alpha})$ equals

$$\mu(C_{\alpha}) + \sum_{\alpha' \leq \tilde{\alpha}} \mu(\alpha') \epsilon_{\sigma', \alpha'}(C_{\tilde{\alpha}}) \frac{\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})}{\epsilon'(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})} \sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) \mu_{\tilde{\alpha}, a}(q).$$

Observe that in the sum above only the factors $\mu(\alpha') \epsilon_{\sigma', \alpha'}(C_{\tilde{\alpha}})$ are not constant with respect to the choice of α' . By Lemma 2.11, $\sum_{\alpha' \leq \tilde{\alpha}} \mu(\alpha') \epsilon_{\sigma', \alpha'}(C_{\tilde{\alpha}}) = \epsilon_{\sigma', \mu}(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})$. By the inductive hypothesis, $\epsilon_{\sigma', \mu}(C_{\tilde{\alpha}}) = \epsilon'(C_{\tilde{\alpha}})$. Thus, replacing $\sum_{\alpha' \leq \tilde{\alpha}} \mu(\alpha') \epsilon_{\sigma', \alpha'}(C_{\tilde{\alpha}})$ with $\epsilon'(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})$ and simplifying the resulting expression, we obtain

$$\epsilon_{\sigma', \mu}(C_{\alpha}) = \mu(C_{\alpha}) + (\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})) \sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) \mu_{\tilde{\alpha}, a}(q).$$

By definition, $\epsilon = \epsilon_{\sigma, \mu}$. Therefore, by Lemma 2.12, the right side of the equation above is $\epsilon(C_{\alpha})$. Moreover, $\epsilon(C_{\alpha}) = \epsilon'(C_{\alpha})$ by Lemma 3.6, Item (2). Thus, $\epsilon_{\sigma', \mu}(C_{\alpha}) = \epsilon'(C_{\alpha})$, as needed.

3. $\epsilon'(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\}) > 0$ and $a \in T$.

As in the previous case, $\epsilon_{\sigma', \mu}(C_{\alpha})$ equals

$$\mu(C_{\alpha}) + (\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})) (\sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) + \sigma(\tilde{\alpha})(\perp)) \mu_{\tilde{\alpha}, a}(q).$$

Also shown in the previous case, we have

$$\epsilon(C_{\alpha}) = \mu(C_{\alpha}) + (\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})) \sigma(\tilde{\alpha})(\text{tran}_{\tilde{\alpha}, a}) \mu_{\tilde{\alpha}, a}(q).$$

Therefore,

$$\epsilon_{\sigma', \mu}(C_{\alpha}) = \epsilon(C_{\alpha}) + (\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})) \sigma(\tilde{\alpha})(\perp) \mu_{\tilde{\alpha}, a}(q).$$

By definition, $\epsilon = \epsilon_{\sigma, \mu}$. Applying Lemma 2.13, we substitute $\epsilon(\tilde{\alpha})$ for $(\epsilon(C_{\tilde{\alpha}}) - \mu(C_{\tilde{\alpha}} - \{\tilde{\alpha}\})) \sigma(\tilde{\alpha})(\perp)$. Now we have

$$\epsilon_{\sigma', \mu}(C_{\alpha}) = \epsilon(C_{\alpha}) + \epsilon(\tilde{\alpha}) \mu_{\tilde{\alpha}, a}(q).$$

The desired result now follows from Lemma 3.6, Item (3). □

Now we can show that applying any finite sequences of tasks to a probability measure on finite execution fragments leads to a generalized probabilistic execution fragment.

Lemma B.4 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. For each probability measure μ on finite execution fragments and each finite sequence of tasks ρ , $\text{apply}(\mu, \rho)$ is a generalized probabilistic execution fragment generated by μ .*

Proof. Simple inductive argument using Lemma B.2 for the base case and Lemma B.3 for the inductive step. □

And now we consider infinite sequences of tasks.

Lemma B.5 *Let $\mathcal{T} = (\mathcal{P}, R)$ be an action-deterministic task-PIOA. For each measure μ on finite execution fragments and each infinite sequence of tasks ρ , $\text{apply}(\mu, \rho)$ is a generalized probabilistic execution fragment generated by μ .*

Proof. For each $i \geq 0$, let ρ_i denote the length- i prefix of ρ and let ϵ_i be $\text{apply}(\mu, \rho_i)$. By Lemmas B.4 and 3.8, the sequence $\epsilon_0, \epsilon_1, \dots$ is a chain of generalized probabilistic execution fragments generated by μ . By Proposition 2.16, $\lim_{i \rightarrow \infty} \epsilon_i$ is a generalized probabilistic execution fragment generated by μ . This suffices, since $\text{apply}(\mu, \rho)$ is $\lim_{i \rightarrow \infty} \epsilon_i$ by definition. \square

This completes the proof of Theorem 3.13.

Proof (Theorem 3.13). Follows directly from Lemmas B.4 and B.5. \square