

# Switched PIOA: Parallel Composition via Distributed Scheduling

Ling Cheung<sup>a,\*</sup>,<sup>1</sup>

<sup>a</sup>*Department of Computer Science, University of Nijmegen,  
P.O. Box 9010, 6500 GL Nijmegen, The Netherlands*

Nancy Lynch<sup>b</sup>,<sup>2</sup>

<sup>b</sup>*MIT Computer Science and Artificial Intelligence Laboratory,  
Cambridge, MA 02139, USA*

Roberto Segala<sup>c</sup>,<sup>3</sup>

<sup>c</sup>*Dipartimento di Informatica, Università di Verona,  
Strada Le Grazie 15, 37134 Verona, Italy*

Frits Vaandrager<sup>a</sup>,<sup>1</sup>

---

## Abstract

This paper presents the framework of switched probabilistic input/output automata (or switched PIOA), augmenting the original PIOA framework with an explicit control exchange mechanism. Using this mechanism, we model a network of processes passing a single token among them, so that the location of this token determines which process is scheduled to make the next move. This token structure therefore implements a distributed scheduling scheme: scheduling decisions are always made by the (unique) active component.

Distributed scheduling allows us to draw a clear line between local and global non-deterministic choices. We then require that local non-deterministic choices are resolved using strictly local information. This eliminates unrealistic schedules that arise under the more common centralized scheduling scheme. As a result, we are able to prove that our trace-style semantics is compositional.

*Key words:* Switched Probabilistic Input/Output Automata, Trace-style Semantics, Parallel Composition, Distributed Scheduling, Compositionality

---

## 1 Introduction

Over the past few decades, a large number of modeling frameworks have been adopted for the purpose of verifying and analyzing stochastic systems. Some of these frameworks, for example, *continuous-time Markov chains* [Ste94] and *labeled Markov processes* [BDEP02], are designed to handle continuous probability distributions, thus finding many applications in the area of performance and reliability analysis [Hav98]. Others, such as *discrete-time Markov chains* [KS76] and *probabilistic automata* [Seg95], deal with discrete probability distributions and are popular in the verification of distributed algorithms and communication protocols [Agg94,LSS94,PSL00,SV99].

Designers of such frameworks are often presented with two challenges:

- (i) defining a sensible notion of *parallel composition*;
- (ii) defining a sensible notion of *semantic equivalence (or preorder)* that is compositional with respect to the proposed parallel operator.

Both notions are important tools for verification. Parallel composition underlies the so-called modular approaches to system development and analysis, where large and complex systems are decomposed into smaller and more tangible subsystems. Semantic equivalence, on the other hand, allows us move across different level of abstraction, from high-level abstract specifications to low-level detailed implementations. A successful combination of parallel composition and process semantics provides great flexibility in model construction and correctness analysis, thus increasing the appeal of the particular modeling framework.

In this paper, we focus on systems that exhibit both non-deterministic behavior and stochastic behavior, while the latter is restricted to discrete probability distributions. Our goal is to develop a compositional framework for modeling these systems. In particular, we are interested in compositionality of trace-style semantics. This has proven to be a surprisingly difficult problem and, resorting to distributed scheduling among parallel components, we offer a solution quite unlike most existing proposals.

---

\* Preliminary versions of this paper appeared as [CLSV04a] and [CLSV04b].

\* Corresponding author.

*Email address:* lcheung@cs.ru.nl (Ling Cheung).

<sup>1</sup> Supported by DFG/NWO bilateral cooperation project 600.050.011.01 Validation of Stochastic Systems.

<sup>2</sup> Supported by DARPA/AFOSR MURI Award #F49620-02-1-0325, NSF Award #CCR-0326277, and USAF, AFRL Award #FA9550-04-1-0121.

<sup>3</sup> Supported by MURST project Constraint-based Verification of Reactive Systems.

A good part of this introduction is devoted to discussions of related literature (Sections 1.1 and 1.2). A disinterested reader may wish to begin with Section 1.3, where we illustrate via a simple example the difficulty with trace-style semantics.

### 1.1 Process Semantics

We focus on systems with both non-deterministic and probabilistic choices, as opposed to purely probabilistic systems. This is because non-determinism is essential in modeling lack of information in either the object system or the external environment. Moreover, as we shall discuss shortly, the interleaving interpretation of parallel composition relies on the presence of non-deterministic choices. Finally, non-deterministic choices can be used to model implementation freedom, making our framework more widely applicable.

In the literature, one can find a great variety of probabilistic process semantics, most of which are extensions of familiar semantic notions for labeled transition systems. Earlier proposals include probabilistic bisimulation [LS91] and testing preorder [YL92], followed by probabilistic simulation [SL95,LSV03], observational testing preorder [SV03] and many others.

Overall, semantics of a more branching character, such as bisimulation and simulation, have been more common than their linear counterparts, such as trace distribution preorder [Seg95]. This is perhaps due to the presence of both non-deterministic and probabilistic choices. In order to define a linear semantics in this setting, one often resorts to the so-called *adversaries*<sup>1</sup> to resolve all non-deterministic choices in a system. Once an adversary is specified, the system becomes purely probabilistic and can be analyzed as a discrete-time Markov chain. Process behavior is then defined by quantifying over all possible adversaries.

In comparison, branching-style semantics are easier to define and more pleasant to work with. For instance, in order to establish bisimilarity between two processes, one simply defines a binary relation on states (or probability distributions on states) and proves that the proposed relation satisfies certain transfer properties. Most importantly, these transfer properties are typically *local*, concerning only the states in relation and their near successors.

Despite the apparent advantages of branching-style semantics, we remain interested in linear, trace-style semantics, because they better capture the idea of externally visible behavior. As shown in [SAGG<sup>+</sup>93,DGRV00], one is often

---

<sup>1</sup> These are called *policies* in the setting of *Markov decision processes*, as employed in planning and optimization.

willing to say that a low-level automaton *implements* a high-level one, even if there exists no bisimulation relation between them. In other words, trace-style equivalence is useful when bisimilarity is considered too fine. Moreover, trace-style semantics are central to *black-box testing*, where we have no convenient access to actual architectures of the system in question.

## 1.2 Parallel Composition

A fundamental idea in concurrency theory is the *interleaving* interpretation of parallel composition:

- (i) every atomic step of a composite system is an atomic step of one of its components (more in case of synchronization);
- (ii) the scheduling among components is arbitrary, up to some appropriate fairness constraints.

In particular, the parallel composition of two independent actions is interpreted as a non-deterministic choice between the two possible interleavings of these actions. This interpretation is generally regarded as a simplifying assumption, reducing the complexity of single-step evolution.

Most existing proposals of parallel composition for stochastic processes adopt the interleaving assumption, thereby necessitating the use of (some form of) adversaries to resolve non-deterministic choices among parallel components. Below we attempt to summarize a few prominent approaches.

- *Parameterized composition* [JLY01,DHK98]. Each (binary) composition operator  $\parallel_p$  is parameterized with a real number  $p \in [0, 1]$ , indicating the bias towards the left process. Sometimes a family of such operators are considered, with  $p$  ranging over some subset of  $[0, 1]$ . Each  $\parallel_p$  is essentially a *static* adversary, resolving the choice between two processes in the same manner at every step.
- *Real-time delay* [WSS94]. Each state  $s$  of a process is associated with a delay parameter  $\delta_s$ . Upon entering a state, every process draws a real-time delay from an exponential distribution with parameter  $\delta_s$ . Among a group of parallel processes, the process with the shortest delay performs the next move. Given delay parameters of all components, one can use specific properties of exponential distributions to calculate the bias towards each component. Therefore, this approach essentially uses *state-dependent* adversaries to resolve non-deterministic choices arising from parallel composition.
- *Compose-and-schedule* [DHK98,Seg95]. Nondeterministic choices remain unresolved in the composition of parallel processes. Eventually, a possible behavior of the composite is obtained by specifying a *history-dependent*

adversary, which has access to internal history of every component and is responsible for resolving *local* non-deterministic choices (i.e., those within each component) as well as *global* ones (i.e., those between parallel components).

Clearly, the last approach is the most robust, in that scheduling decisions may depend on dynamic behaviors of the entire system. Here we pay a hefty price for such expressivity: trace-style semantics is not compositional [Seg95]. Put simply, trace-style semantics abstracts away from internal branching structures of processes. Yet a powerful adversary can observe differences in internal branching and is therefore capable of exposing these differences when equivalent processes are composed in parallel with the same probabilistic context. We shall return to this point in Section 1.3 and give a concrete example (Figures 1 and 2).

Moving to the less robust approach of real-time delay, one can in fact achieve compositionality for trace-style semantics [WSS94]. However, this approach relies on the assumption that delay patterns of processes can be universally characterized by exponential distributions. In the end, it is unclear whether the theory is applicable outside specific areas such as hybrid systems and queuing networks.

Finally, we mention an approach that takes us away from the realm of interleaving semantics. In the models of [dAHJ01,vGSS95], components may make simultaneous moves, even if they are not involved in action synchronization. Assuming independence of coin tosses, the probability of a composite move can be calculated by simply multiplying the probabilities of all atomic moves involved. In this setting, it is also possible to obtain a compositional trace-style semantics [dAHJ01]. Nonetheless, synchronous models are not suitable for a large class of problems in which simultaneous executions are not possible. Obvious examples include mutual exclusion and distributed consensus algorithms. This takes us back to the challenge of scheduling via adversaries.

### 1.3 Observational Powers of Adversaries

As promised, we give a simple example in which a trace-style semantics fails to be compositional. In particular, we take the trace distribution semantics of [Seg95], where each possible behavior is a probability space on the set of traces and is induced by a history-dependent adversary.

As their names suggest, automaton **Early** forces the adversary to choose between  $b$  and  $c$  as it chooses one of the two available  $a$ -transitions, whereas in automaton **Late** the adversary may postpone this decision until after the  $a$ -transition. Clearly, these two automata have the same set of trace distri-




Fig. 1. Probabilistic automata Early, Late and Coin

butions, but they can be distinguished by composing with the context **Coin**. This context has a probabilistic  $a$ -transition leading to a uniform distribution on two states, one of which enables a  $d$ -transition while the other enables an  $e$ -transition.

The composed system **Late**  $\parallel$  **Coin** has a trace distribution that assigns probability  $\frac{1}{2}$  to each of these traces:  $adb$  and  $aec$  (Figure 2). This is induced by an adversary that chooses the  $b$ -transition in **Late** if and only if the random choice in **Coin** results in the left state. Such total correlations between actions  $d$  and  $b$ , and between actions  $e$  and  $c$ , cannot be achieved by the composite **Early**  $\parallel$  **Coin**, therefore the two composites are not trace distribution equivalent.




Fig. 2. Non-substitutivity of trace distribution equivalence

Inspired by this example, we showed that the coarsest pre-congruence refining trace distribution preorder coincides with the probabilistic simulation preorder [LSV03]. In other words, the observational power given to adversaries forces us into the realm of branching-style semantics, where internal branching structures can be used to distinguish processes.

In the present paper, we follow a different direction: rather than taking the largest pre-congruence induced by trace distribution preorder, we attempt to weaken the observational power of adversaries. Notice, in the composition mechanism of probabilistic automata, nondeterministic choices are resolved after the two automata are composed, allowing the adversary to make decisions in one component using state information of the other. This sort of “information leakage” is precisely the source of difficulty in compositionality. We therefore aim at a framework in which global non-determinism is clearly separated from local non-determinism. (Recall that the former arises from uncertainty in a distributed environment, while the latter from uncertainty within components.) The challenge is then to achieve this separation without sacrificing the flexibility to treat a composite of multiple components as yet a single component.

In fact, adversary models of various strengths have been studied in the setting of randomized distributed algorithms [Asp03,AB04], where correctness and complexity of algorithms depend crucially upon the particular choice of adversary model. In the formal methods community, however, this issue has not received much attention. Some initial steps along these lines can be found in [CH05].

#### 1.4 Distributed Scheduling

We propose a composition mechanism where local scheduling decisions are based on strictly local information, while global scheduling conflicts are eliminated using a control-passage mechanism. Note that the term *control* is used here in the spirit of “control flow” in sequential programming: a component is said to possess the control of a system if it is scheduled to actively perform the next action. This should not be confused with the notion of controllers for plants, as in control theory.

Intuitively, we model a network of processes passing a single token among them, with the property that a process enables a locally controlled transition (i.e., non-input) only if it possesses the token. Thus, the location of this unique token determines which process is scheduled to make the next move. We call this model *switched probabilistic input/output automata* (or *switched PIOA* for short). It augments the *probabilistic input/output automata (PIOA)* model [WSS94,PSL00,BPW04] with additional structures and axioms for control exchange.

In particular, we add a predicate **active** on the set of states, indicating whether an automaton is active or inactive. We require that locally controlled actions are enabled only if the automaton is active. In other words, an inactive automaton must be quiescent and can only accept inputs from the environment.

This activity status can be changed only by performing special control input and control output actions. Control inputs correspond to an incoming token, thus switching the automaton from inactive mode to active mode. And vice versa for control outputs. We make sure that all such control synchronizations are “handshakes”: at most two components may participate in a transition labeled by a control action. Together with an appropriate initialization condition, this ensures that at most one component is active at any point of an execution.

In this framework, scheduling decisions are always made locally: each process is equipped with a local scheduler, which has access to local history and is responsible for resolving local non-deterministic choices. Among other things, the local scheduler chooses when to give up the activity token and to whom

the token is sent. This is precisely the sense in which our scheduling scheme is *distributed*: global scheduling is performed collectively by all local schedulers. This scheme eliminates the need for adversaries such as the one in Figure 2 and allows us to give a compositional trace-style semantics (Definition 10 and Theorem 29).

Distributed scheduling (as opposed to centralized scheduling) has been a mainstream approach in the area of security analysis [BPW04,Can01], where information flow is a sensitive issue. Compared with the *interactive Turing machines* of [Can01] and *asynchronous reactive systems* of [BPW04], our framework provides much better modeling flexibility, as we allow local non-deterministic choices to accommodate for lack of information and implementation freedom. However, we must admit this is an unfair comparison, because the two frameworks mentioned above are highly specialized for delicate reasonings in computational cryptography.

For those who may be skeptical of distributed scheduling, we argue that centralized scheduling can be implemented in our framework by modeling adversaries explicitly via an *arbiter* automaton. In other words, processes do not exchange control among each others directly, but they do so via the arbiter. This arbiter observes the whole system by way of action synchronization and it makes scheduling decisions accordingly. Since the input signature of such an arbiter is completely flexible, we have a convenient means to specify what information is available for inter-component scheduling. This will be further discussed in Section 8.

### 1.5 Overview

This introduction is followed by Section 2, which contains basic mathematical preliminaries. Section 3 presents a new formulation of the PIOA framework, combining *reactive* and *generative* system types [DHK98,vGSS95,SdV04] in the presence of input/output (i/o) distinction. Then, in Section 4, we define *i/o schedulers* for PIOAs and derive an external behavior semantics based on *execution trees* and *likelihood assignments*.

Starting from Section 5, we focus on switched PIOAs and distributed scheduling. First we introduce a set of switch axioms, formalizing the notion of control exchange. Section 6 then defines parallel composition for switched PIOAs. The main technical contribution of this paper is presented in Section 7: the external behavior semantics for switched PIOAs is compositional (Theorem 29).

Section 8 describes *controllable PIOAs* and arbiters, which can be used to implement various centralized scheduling schemes. Concluding discussions follow in Section 9.

## 2 Preliminaries

Let two sets  $X$  and  $Y$  be given. A function  $\mu : X \rightarrow [0, 1]$  is called a *discrete (probability) distribution* on  $X$  if  $\sum_{x \in X} \mu(x) = 1$ . The *support* of  $\mu$ , denoted  $\text{Supp}(\mu)$ , is the set  $\{x \in X \mid \mu(x) > 0\}$ . We write  $\text{Disc}(X)$  for the set of all discrete distributions on  $X$ . Given  $x \in X$ , the *Dirac distribution* on  $x$ , denoted  $\text{Dirac}(x)$ , assigns probability 1 to  $x$ .

If  $\mu$  is a discrete distribution on  $X$  and  $Y$  is a superset of  $X$ , we shall freely regard  $\mu$  as a discrete distribution on  $Y$ , where  $\mu(y) := 0$  for all  $y \in Y \setminus X$ .

We write  $X \times Y$  for the *Cartesian product* of  $X$  and  $Y$ , where the projection maps are denoted  $\pi_L$  and  $\pi_R$ , respectively. The product of an indexed family  $\{X_i \mid i \in \mathcal{I}\}$  of sets is denoted  $\prod_{i \in \mathcal{I}} X_i$ , with projection maps  $\pi_i$ . Throughout this paper, we assume that the index set  $\mathcal{I}$  is non-empty and finite. Also, when  $\mathcal{I}$  is clear from context, we write  $\vec{x}$  for a typical member of  $\prod_{i \in \mathcal{I}} X_i$ .

Given a family  $\{\mu_i \mid i \in \mathcal{I}\}$  where each  $\mu_i$  is a discrete distribution on  $X_i$ , we form the *product distribution*, denoted  $\prod_{i \in \mathcal{I}} \mu_i$ , as follows:

$$\left(\prod_{i \in \mathcal{I}} \mu_i\right)(\vec{s}) := \prod_{i \in \mathcal{I}} \mu_i(s_i).$$

This is easily shown to be a discrete distribution on  $\prod_{i \in \mathcal{I}} X_i$ . Conversely, given any distribution  $\mu$  on  $\prod_{i \in \mathcal{I}} X_i$  and  $i \in \mathcal{I}$ , one can form the  *$i$ th-projection* of  $\mu$ , denoted  $\pi_i(\mu)$ , by:

$$\pi_i(\mu)(s) := \sum_{\vec{t}: t_i = s} \mu(\vec{t}).$$

We have the obvious identity:  $\pi_i(\prod_{j \in \mathcal{I}} \mu_j) = \mu_i$ .

Finally, the set of all *partial functions* from  $X$  to  $Y$  is denoted  $X \rightarrow Y$ . For each  $f \in X \rightarrow Y$ , we write  $\text{dom}(f)$  for the *domain* of  $f$  and  $f(x) = \perp$  whenever  $x \notin \text{dom}(f)$ . The symbol  $\emptyset$  denotes the empty function (as well as the empty set).

## 3 Probabilistic Input/Out Automata

In this section, we define the basic framework of probabilistic input/output automata, following the tradition of Input/Output Automata (IOA) of Lynch and Tuttle [LT89]. Variations of this framework have appeared in many places (e.g. [BPW04, PSL00, WSS94]), yet the actual definitions diverge significantly. Among other goals, this paper aims to provide a concise and unifying formulation.

We assume a fixed, countably infinite alphabet  $Act$  of action symbols. Inspired by [vGSS95], we define reactive and generative transition structures as follows.

**Definition 1** Let  $S$  be a set of states and let  $X \subseteq Act$  be given.

- (i) A reactive transition structure on  $\langle S, X \rangle$  is a function  $\mathbf{R} : S \times X \rightarrow \mathcal{P}(\text{Disc}(S))$ .
- (ii) A generative transition structure on  $\langle S, X \rangle$  is a function  $\mathbf{G} : S \rightarrow \mathcal{P}(\text{Disc}(X \times S))$ .

A state  $s \in S$  blocks action  $a \in X$  if  $\mathbf{R}(\langle s, a \rangle) = \emptyset$ . It is said to be quiescent if  $\mathbf{G}(s) = \emptyset$ .

A reactive transition structure  $\mathbf{R}$  describes a system that reacts to input signals. Given a state  $s$  and an action  $a$ ,  $\mathbf{R}(\langle s, a \rangle)$  yields a set of discrete distributions on  $S$ . Thus we allow non-deterministic choices over possible distributions on end states, while each such distribution specifies an effect of randomization on system evolution. We use variables  $\mu, \nu$ , etc., for these state distributions. Figure 3 below illustrates two such reactive systems.




Fig. 3. Examples of Reactive Transition Systems

On the other hand, a generative transition structure  $\mathbf{G}$  describes a system that evolves in an active fashion. That is, every state  $s$  enables a (possibly empty) set of *transition bundles*, where each bundle is a discrete distribution on  $Act \times S$ . Again, we have non-deterministic choices over bundles, while each bundle specifies a random choice over next transitions. We use variables  $f, g$ , etc., for these transition bundles. Figure 4 below illustrates two such generative systems.




Fig. 4. Examples of Generative Transition Systems

Now we introduce the notion of probabilistic i/o automata as a combination of reactive and generative system types, in the presence of i/o distinction. Notice that, we impose i/o distinction not only on the action signature, but also on the transition structure.

**Definition 2** A probabilistic i/o automaton (PIOA)  $A$  is a tuple

$$\langle S_A, s_A^0, I_A, O_A, H_A, \mathbf{R}_A, \mathbf{G}_A \rangle$$

where:

- (1)  $S_A$  is a set of states with initial state  $s_A^0 \in S_A$ ;
- (2)  $\{I_A, O_A, H_A\}$  are pairwise disjoint subsets of  $Act$ , referred to as: input, output and hidden actions, respectively;
- (3)  $\mathbf{R}_A$  is a reactive transition structure on  $\langle S_A, I_A \rangle$  and  $\mathbf{G}_A$  is a generative transition structure on  $\langle S_A, O_A \cup H_A \rangle$ .

The automaton  $A$  is said to be *closed* if  $I_A$  is empty and *open* otherwise. As usual, input and output actions are *visible*, while output and hidden actions are *locally controlled*. The union  $I_A \cup O_A \cup H_A$  is often denoted by  $Act_A$ . Notice that we omit the input enabling axiom of IOA (i.e., all inputs are accepted at every state). This flexibility facilitates our introduction of switched PIOAs in Section 5.

In a typical automata-theoretic setting, an execution (or a path) is a sequence of states and actions in alternating fashion satisfying the obvious reachability conditions. Our version, called an execution branch, is enriched with additional information from the reactive and generative transition structures.

**Definition 3** Let  $A$  be a PIOA and let  $s \in S_A$  be given. We use joint recursion to define the set of execution branches from  $s$ , denoted  $\mathbf{Bran}(s)$ , together with the function  $\mathbf{last} : \mathbf{Bran}(s) \rightarrow S_A$ .

- The length-one sequence containing  $s$  (written  $\underline{s}$ ) is in  $\mathbf{Bran}(s)$  and is called the empty branch, where  $\mathbf{last}(\underline{s}) := s$ .
- For all  $r \in \mathbf{Bran}(s)$ ,  $a \in I_A$ ,  $\mu \in \mathbf{R}_A(\langle \mathbf{last}(r), a \rangle)$  and  $s' \in \text{Supp}(\mu)$ , we have  $r.a.\mu.s' \in \mathbf{Bran}(s)$ . Moreover,  $\mathbf{last}(r.a.\mu.s') := s'$ .
- For all  $r \in \mathbf{Bran}(s)$ ,  $f \in \mathbf{G}_A(\mathbf{last}(r))$  and  $\langle a, s' \rangle \in \text{Supp}(f)$ , we have  $r.f.a.s' \in \mathbf{Bran}(s)$ . Moreover,  $\mathbf{last}(r.f.a.s') := s'$ .

The *trace* of a branch  $r$  is defined in the usual way:

- $\text{tr}(\underline{s}) := \epsilon$ ,
- $\text{tr}(r.a.\mu.s') := \text{tr}(r).a$  (in this case  $a \in I_A$ ), and
- $\text{tr}(r.f.a.s')$  is  $\text{tr}(r).a$  if  $a \in O_A$  and  $\text{tr}(r)$  if  $a \in H_A$ .

We write  $\mathbf{Bran}(A)$  for  $\mathbf{Bran}(s^0)$ .

Notice that execution branches are always finite, because  $\mathbf{Bran}(s)$  is given by a recursive definition. An infinite branch from  $s$  is simply an infinite subset of  $\mathbf{Bran}(s)$  that is linearly ordered by the prefix ordering on sequences, which is

denoted  $\sqsubseteq$ . We write  $\text{Bran}^{\leq\omega}(s)$  for the set of finite and infinite branches from  $s$ . Similarly,  $\text{Bran}^{\leq\omega}(A) := \text{Bran}^{\leq\omega}(s^0)$ .

It is often convenient to speak of reachability with non-zero probability, abstracting away from the actual probability distributions. Given  $s, s' \in S_A$  and  $a \in \text{Act}_A$ , we say that  $s'$  is *reachable* (in one step) from  $s$  via action  $a$ , denoted  $s \xrightarrow{a} s'$ , just in case:

- $s.a.\mu.s' \in \text{Bran}(s)$  for some  $\mu$ , or
- $s.f.a.s' \in \text{Bran}(s)$  for some  $f$ .

Similarly, a state  $s$  is *reachable* if there exists  $r \in \text{Bran}(A)$  such that  $\text{last}(r) = s$ .

Given two PIOAs  $A$  and  $B$  with the same action signature, one can speak of  $A$  being a *sub-automaton* of  $B$ . Intuitively, it means  $A$  can be obtained from  $B$  by removing certain states and/or transitions. This is made precise in Definition 4 below.

**Definition 4** *Suppose  $A$  and  $B$  are PIOAs with the same action signature  $\{I, O, H\}$ . We say that  $A$  is a sub-automaton of  $B$ , denoted  $A \subseteq B$ , if*

- $S_A \subseteq S_B$  and  $s_A^0 = s_B^0$ ;
- for all  $s \in S_A$  and  $a \in I$ ,  $\mathbf{R}_A(\langle s, a \rangle) \subseteq \mathbf{R}_B(\langle s, a \rangle)$  and  $\mathbf{G}_A(s) \subseteq \mathbf{G}_B(s)$ .

## 4 Probabilistic Systems and Their External Behaviors

As we argued in Section 1.3, the full observational power of history-dependent adversaries leads to unrealistic schedules in a parallel composition (Figure 2). To exclude these schedules, we pair a PIOA with a set of acceptable schedules, forming a *probabilistic system* (Definition 6 below).

First, we make explicit the notion of schedules in an i/o setting.

**Definition 5** *Let  $A$  be a PIOA. An input scheduler  $\sigma$  for  $A$  is a partial function*

$$\sigma : \text{Bran}(A) \times I \rightarrow \text{Disc}(S_A)$$

*such that: for all  $\langle r, a \rangle \in \text{Bran}(A) \times I$ , if  $\mathbf{R}_A(\langle \text{last}(r), a \rangle)$  is non-empty, then  $\sigma(\langle r, a \rangle)$  is defined and is in  $\mathbf{R}_A(\langle \text{last}(r), a \rangle)$ . An output scheduler  $\rho$  for  $A$  is a partial function*

$$\rho : \text{Bran}(A) \rightarrow \text{Disc}((O_A \cup H_A) \times S_A)$$

*such that: for all  $r \in \text{Bran}(A)$ , if  $\rho(r)$  is defined, then  $\rho(r) \in \mathbf{G}_A(\text{last}(r))$ . An i/o scheduler for  $A$  is then a pair  $\langle \sigma, \rho \rangle$  where  $\sigma$  is an input scheduler for  $A$*

and  $\rho$  is an output scheduler for  $A$ .

I/O schedulers remove non-deterministic choices in  $A$ . The input scheduler  $\sigma$  specifies the reactive schedule: given a finite history  $r$  and an input signal  $a$  that is not blocked by  $\text{last}(r)$ ,  $\sigma$  selects a distribution from  $\mathbf{R}_A(\langle \text{last}(r), a \rangle)$ . Similarly, the output scheduler  $\rho$  specifies the generative schedule: given a finite history  $r$ ,  $\rho$  selects a bundle from  $\mathbf{G}_A(\text{last}(r))$  if  $\text{last}(r)$  is not quiescent. Notice that the output scheduler has slightly more freedom compared to its input counterpart: it may *halt* the execution by setting  $\rho(r)$  to  $\perp$ , even if  $\mathbf{G}_A(\text{last}(r))$  is non-empty.

From the notion of i/o schedulers, it is now straightforward to define probabilistic systems.

**Definition 6** A probabilistic system  $\mathcal{A}$  is a pair  $\langle A, \mathcal{S} \rangle$ , where  $A$  is a PIOA and  $\mathcal{S}$  is a set of i/o schedulers for  $A$ . Such a system is full if  $\mathcal{S}$  is the set of all i/o schedulers for  $A$ . Moreover,  $\mathcal{A}$  is called a switched probabilistic system if  $A$  is a switched PIOA.

In the rest of this section, we define a trace-style notion of external behavior for probabilistic systems. In particular, we derive a *likelihood assignment* from each triple  $\langle A, \sigma, \rho \rangle$ , where  $A$  is a PIOA and  $\langle \sigma, \rho \rangle$  is an i/o scheduler for  $A$ . This is analogous to the notion of *trace distributions* in [Seg95], where a trace distribution is obtained from a probabilistic automaton (without i/o distinction) together with a randomized, history-dependent scheduler.

#### 4.1 Execution Trees

First we define the *execution tree* induced by such a triple  $\langle A, \sigma, \rho \rangle$ .

**Definition 7** Let  $A$  be a PIOA and let  $\langle \sigma, \rho \rangle$  be an i/o scheduler for  $A$ . The execution tree generated by  $\langle A, \sigma, \rho \rangle$  is the function  $Q_{\sigma, \rho} : \text{Bran}(A) \rightarrow [0, 1]$  defined recursively by:

- $Q_{\sigma, \rho}(s_A^0) = 1$ ;
- given  $r'$  of the form  $r.a.\mu.s'$ ,
  - $Q_{\sigma, \rho}(r') := Q_{\sigma, \rho}(r) \cdot \mu(s')$ , if  $\mu = \sigma(\langle r, a \rangle)$ ;
  - $Q_{\sigma, \rho}(r') := 0$ , otherwise;
- given  $r'$  of the form  $r.f.a.s'$ ,
  - $Q_{\sigma, \rho}(r') := Q_{\sigma, \rho}(r) \cdot f(\langle a, s' \rangle)$ , if  $f = \rho(r)$ ;
  - $Q_{\sigma, \rho}(r') := 0$ , otherwise.

If  $A$  is closed, then the input scheduler  $\sigma$  must be the empty function. In that case, we write  $Q_\rho$  for  $Q_{\sigma, \rho}$ . We claim that  $Q_\rho$  induces a probability space over

the sample space  $\Omega_A := \mathbf{Bran}^{<\omega}(A)$ . The construction is completely standard, so we provide an outline below and refer the reader to e.g. [Seg95] for details.

- (i) Each  $r \in \mathbf{Bran}(A)$  generates a *cone* of executions as follows:  $\mathbf{C}_r := \{r' \in \mathbf{Bran}^{<\omega}(A) \mid r \sqsubseteq r'\}$ .
- (ii) Let  $\mathcal{F}_A$  denote the smallest  $\sigma$ -field on  $\Omega_A$  generated by the collection  $\{\mathbf{C}_r \mid r \in \mathbf{Bran}(A)\}$ .
- (iii) Construct a (unique) probability measure  $\mathbf{m}_\rho$  on  $\mathcal{F}_A$  such that  $\mathbf{m}_\rho[\mathbf{C}_r] = Q_\rho(r)$  for all  $r$  in  $\mathbf{Bran}(A)$ .

In this way,  $Q_\rho$  gives rise to the probability space  $(\Omega_A, \mathcal{F}_A, \mathbf{m}_\rho)$ .

For open PIOAs, however, an execution tree does not always induce a probability measure, because it does not take into account the probabilities with which various inputs are provided by the environment. For example, if  $r'$  is of the form  $r.a.\mu.s'$ , the value  $Q_{\sigma,\rho}(r')$  is computed from  $Q_{\sigma,\rho}(r)$  and  $\mu(s')$ , neither of which contains information about the probability of  $a$  being provided as an input.

Nonetheless, the notion of execution trees is an important technical tool in our development. It gives great flexibility in manipulating open components, which are typically part of a parallel composition forming a closed PIOA. In the end, we are assured that any probabilistic statement about the final, closed composite is meaningful (i.e., based on a well-defined probability measure).

## 4.2 Likelihood Assignments

Likelihood assignments are behavioral abstractions of execution trees. Roughly speaking, the probability of observing a certain trace  $\alpha \in Act^{<\omega}$  is the probability of the automaton executing *any* branch with trace  $\alpha$ . This can be computed by summing the probabilities of all such branches in the execution tree. As we mentioned at the end of Section 4.1, execution trees of open PIOAs need not always induce probability measures. That is the reason we opt for the term “likelihood”, rather than “probability”. Nonetheless, the method of abstraction is completely analogous.

To begin, we need the notion of minimal branches: a branch  $r \in \mathbf{Bran}(s)$  is said to be *minimal* if every proper prefix of  $r$  in  $\mathbf{Bran}(s)$  has a strictly shorter trace. Notice, the empty branch is minimal. For non-empty  $r$ , it is minimal if and only if its last action label is visible. We write  $\mathbf{Bran}_{\min}(s)$  for the set of minimal branches in  $\mathbf{Bran}(s)$ . For each  $\alpha \in Act_A^{<\omega}$ , let  $\mathbf{tr}_{\min}^{-1}(\alpha)$  denote the set of minimal branches of  $A$  with trace  $\alpha$ .

Minimality is important because distinct minimal branches with the same

trace always represent mutually exclusive events, whereas distinct branches (not necessarily minimal) with the same trace may be prefix-related, i.e., one event is strictly included in the other. With this notion at hand, we define a lifting of the trace operator  $\text{tr} : \text{Bran}(A) \rightarrow \text{Act}_A^{<\omega}$ . Given a function  $Q : \text{Bran}(A) \rightarrow [0, 1]$ , we define  $\text{tr}(Q) : \text{Act}_A^{<\omega} \rightarrow [0, 1]$  by

$$\text{tr}(Q)(\alpha) := \sum_{r \in \text{tr}_{\min}^{-1}(\alpha)} Q(r).$$

Using this lifted trace operator, it is straightforward to derive likelihood assignments.

**Definition 8** *Let  $A$  be a PIOA and let  $\langle \sigma, \rho \rangle$  be an i/o scheduler for  $A$ . The likelihood assignment induced by  $\langle A, \sigma, \rho \rangle$ , denoted  $\mathbb{L}_{\sigma, \rho}$ , is the function  $\text{tr}(Q_{\sigma, \rho}) : \text{Act}_A^{<\omega} \rightarrow [0, 1]$ .*

As with execution trees, we omit the input scheduler  $\sigma$  whenever  $A$  is closed. In that case, each  $\mathbb{L}_\rho$  induces a probability measure on the sample space  $\Omega := \text{Act}_A^{<\omega}$ . The  $\sigma$ -field  $\mathcal{F}$  on  $\Omega$  is generated by the collection  $\{\mathbf{C}_\alpha \mid \alpha \in \text{Act}_A^{<\omega}\}$ , where  $\mathbf{C}_\alpha := \{\alpha' \in \Omega \mid \alpha \sqsubseteq \alpha'\}$ . The measure  $\mathbf{m}^\rho$  on  $\mathcal{F}$  is uniquely determined by the equations  $\mathbf{m}^\rho[\mathbf{C}_\alpha] = \mathbb{L}_\rho(\alpha)$  for all  $\alpha \in \text{Act}_A^{<\omega}$ .

In the literature, some authors define probabilistic executions (resp. trace distributions) to be the probability spaces  $\langle \Omega_A, \mathcal{F}_A, \mathbf{m}_\rho \rangle$  (resp.  $\langle \Omega, \mathcal{F}, \mathbf{m}^\rho \rangle$ ). These definitions no longer apply when we move to a setting with open inputs. Therefore, we propose the notions of execution trees and likelihood assignments, generalizing probabilistic executions and trace distributions, respectively. Again we refer to [Seg95] for these alternative definitions and measure theoretic proofs.

We are now ready to define external behavior for probabilistic systems. The implementation relation is simply behavioral inclusion.

**Definition 9** *Let  $\mathcal{A} = \langle A, \mathcal{S} \rangle$  be a probabilistic system. An external behavior of  $\mathcal{A}$  is a likelihood assignment  $\mathbb{L}_{\sigma, \rho}$  induced by some  $\langle \sigma, \rho \rangle \in \mathcal{S}$ . We write  $\text{ExtBeh}(\mathcal{A})$  for the set of all external behaviors of  $\mathcal{A}$ .*

**Definition 10** *Probabilistic systems  $\mathcal{A} = \langle A, \mathcal{S} \rangle$  and  $\mathcal{B} = \langle B, \mathcal{T} \rangle$  are said to be comparable if:*

- $\text{active}_A(s_A^0) = \text{active}_B(s_B^0)$  and
- $I_A = I_B$ ,  $O_A = O_B$ , and  $\text{Sync}_A = \text{Sync}_B$ .

*Given such comparable  $\mathcal{A}$  and  $\mathcal{B}$ , we say that  $\mathcal{A}$  implements  $\mathcal{B}$  if  $\text{ExtBeh}(\mathcal{A}) \subseteq \text{ExtBeh}(\mathcal{B})$ .*

This concludes our treatment of external behavior for PIOAs in general. Start-

ing in the next section, we focus on switched PIOAs and parallel composition via distributed scheduling.

## 5 Switched PIOAs

We now augment the PIOA model of Section 3 with additional structures and axioms, yielding the notion of switched PIOAs. These changes are prompted by our proposal of distributed scheduling (cf. Section 1.4). Namely, we use a token structure to eliminate global scheduling conflicts, ensuring that

- (i) at any point of an execution, at most one component is active;
- (ii) the currently active component always selects the next active component.

In order to implement this token structure, we must distinguish between *active* and *inactive* states of an automaton. Moreover, we designate special *control actions* and impose five *switch axioms*, formalizing our intuitions about control passage among components. This leads to Definition 11 below.

For technical simplicity, we assume that  $Act$  is partitioned into two sets:  $BAct$  (*basic actions*) and  $CAct$  (*control actions*). Both sets are assumed to be countably infinite.

**Definition 11** *A switched PIOA is given by a PIOA  $A$ , together with a function  $\text{active}_A : S_A \rightarrow \{0, 1\}$  and a set  $\text{Sync}_A \subseteq O_A \cap CAct$  of synchronized control actions such that the following (universally quantified) axioms are satisfied.*

- (S1)  $\text{active}_A(s) = 0 \Rightarrow . \mathbf{G}_A(s) = \emptyset \wedge \forall a \in I_A. \mathbf{R}_A(\langle a, s \rangle) \neq \emptyset$
- (S2)  $\text{active}_A(s) = 1 \Rightarrow . \forall a \in I_A. \mathbf{R}_A(\langle a, s \rangle) = \emptyset$
- (S3)  $(s \xrightarrow{a} s' \wedge a \in I_A \cap CAct) \Rightarrow \text{active}_A(s') = 1$
- (S4)  $(s \xrightarrow{a} s' \wedge a \in (O_A \cap CAct) \setminus \text{Sync}_A) \Rightarrow \text{active}_A(s') = 0$
- (S5)  $(s \xrightarrow{a} s' \wedge a \in BAct \cup H_A \cup \text{Sync}_A) \Rightarrow \text{active}_A(s) = \text{active}_A(s')$

To increase readability, we classify the action symbols of  $A$  as follows:

- $BI_A := I_A \cap BAct$  (*basic inputs*);
- $BO_A := O_A \cap BAct$  (*basic outputs*);
- $CI_A := I_A \cap CAct$  (*control inputs*);
- $CO_A := (O_A \cap CAct) \setminus \text{Sync}_A$  (*control outputs*).

Essentially, we have a partition  $\{BI_A, BO_A, H_A, CI_A, CO_A, \text{Sync}_A\}$  of  $Act_A$ . We say that  $A$  is *initially active* if  $\text{active}_A(s^0) = 1$ . Otherwise, it is *initially inactive*.

The first two axioms constrain the behavior of  $A$  based on its activity status. Essentially, Axiom (S1) says that an inactive automaton is a reactive machine, therefore all inactive states of  $A$  must be quiescent and satisfy the usual input enabling assumption. On the other hand, an active automaton is a generative machine, therefore Axiom (S2) requires all active states of  $A$  to be input blocking.

The last three axioms specify how the various types of actions change the activity status of an automaton. Axioms (S3) and (S4) say that control inputs lead to active states and control outputs to inactive states. Axiom (S5) says that no other actions may change the activity status.

Together, these five axioms describe an “activity cycle” for the automaton  $A$ :

- (i) while in inactive mode,  $A$  does not enable locally controlled transitions, although it may still receive inputs from its environment;
- (ii) when  $A$  receives a control input it moves into active mode, where it may perform hidden or output transitions, possibly followed by a control output;
- (iii) via this control output  $A$  returns to inactive mode.

This is captured in Lemma 12 below.

**Lemma 12** *Let  $A$  be a switched PIOA and let  $s, s'$  in  $S_A$  and  $a \in Act_A$  be given. Suppose that  $s \xrightarrow{a} s'$ .*

- (1) *If  $a \in BI_A$ , then  $\text{active}_A(s) = \text{active}_A(s') = 0$ .*
- (2) *If  $a \in CI_A$ , then  $\text{active}_A(s) = 0$  and  $\text{active}_A(s') = 1$ .*
- (3) *If  $a \in BO_A \cup H_A \cup Sync_A$ , then  $\text{active}_A(s) = \text{active}_A(s') = 1$ .*
- (4) *If  $a \in CO_A$ , then  $\text{active}_A(s) = 1$  and  $\text{active}_A(s') = 0$ .*

**PROOF.** For Item (1), note that  $a \in I_A$ . By the definition of  $s \xrightarrow{a} s'$ , we may choose distribution  $\mu \in \mathbf{R}_A(\langle s, a \rangle)$  such that  $s' \in \text{Supp}(\mu)$ . Therefore, by Axiom (S2), we know that  $\text{active}_A(s) = 0$ . Applying Axiom (S5) we have  $\text{active}_A(s) = \text{active}_A(s') = 0$ . Item (3) follows similarly from Axioms (S1) and (S5).

For Item (2), we first use Axiom (S2) to argue that  $\text{active}_A(s) = 0$ . Moreover, Axiom (S3) implies  $\text{active}_A(s') = 1$ . Item (4) follows similarly from Axioms (S1), (S4).  $\square$

To give some concrete examples of switched PIOAs, we return to automata Early, Late and Coin of Figure 1. Their adaptations to the switched PIOA framework are illustrated in Figure 5 below. We have chosen to assign actions

$b$  and  $c$  to the basic output signature of **Early'** and **Late'**, whereas  $a$ ,  $d$  and  $e$  are basic outputs of **Coin'**. Following conventions in process algebra, we use  $?$  to indicate input actions and  $!$  to indicate output actions.




Fig. 5. Adaptations of **Early**, **Late** and **Coin**

Due to the additional predicate **active**, the state spaces have been doubled. Active states are drawn in the foreground and inactive ones in the background. Thus, **Early'** and **Late'** are initially inactive and **Coin'** is initially active.

Each two-headed arrow indicates a control output from active to inactive and a control input from inactive to active. We assume that **Early'** and **Late'** have a sole control input **go** and a sole control output **done**; and vice versa for **Coin'**. For a clearer picture, we have omitted the names of control actions, as well as non-essential input loops.

## 6 Parallel Composition

In this section, we define parallel composition in an incremental fashion. First we do so for PIOAs (Section 6.1), specifying the composite transition structures. As usual, this definition is based on action synchronization and does not attempt to resolve non-deterministic choices among parallel components. In Section 6.2, we extend this composition operator to switched PIOAs, taking care that the composite still satisfies all switch axioms.

Then, departing from the “compose-and-schedule” approach (cf. Section 1.2), we describe how to compose i/o schedulers for compatible switched PIOAs to form a single i/o scheduler for their composite (Section 6.3). This extends easily to parallel composition for probabilistic systems. Thus, our approach can be described as “schedule-and-compose”, where parallel composition is imposed *after* local schedules have been completely specified.

Unlike most composition operators in the literature, our definitions do *not* involve normalization mechanisms, which collect and redistribute deadlock probabilities. Instead, we take advantage of i/o distinction and use probabilistic input enabling to make sure that deadlocks never occurs.

### 6.1 Composing PIOAs

Let us begin with an example to illustrate how we intend to compose reactive and generative transition structures. Consider automata  $A$ ,  $B$  and  $C$  in Figure 6 and assume that action  $a$  is in the signatures of all three automata, while  $b$  is in the signatures of  $B$  and  $C$  only.




Fig. 6. Automata  $A$ ,  $B$  and  $C$

First we consider  $A\|B$ . Here both  $A$  and  $B$  are reactive, therefore their composite is constructed in a straightforward manner via synchronization of shared actions. In particular, if input  $a$  is provided, then both  $A$  and  $B$  react and move to corresponding new states. If, on the other hand,  $b$  is provided, then only  $B$  reacts and  $A$  simply stutters (i.e., no transition takes place). This is illustrated in Figure 7 on the left.




Fig. 7. Parallel Composites  $A\|B$  and  $A\|B\|C$

Next we add  $C$  to the parallel composition. Now the composite exhibits generative behavior, because both actions  $a$  and  $b$  are locally controlled by  $C$ . In  $A\|B\|C$ , these action each take place with probability  $\frac{1}{2}$ , just as in  $C$ . If  $a$  is chosen, then all three components participate in the transition. Otherwise,  $b$  is chosen and only  $B$  and  $C$  participate. This is illustrated in Figure 7 on the right.

Despite its simplicity, Figure 7 demonstrates our basic idea of parallel composition: in each step of the composite, at most one component behaves actively, while all others react to the action performed by the active component. In the rest of this section, we try to formalize this simple idea in the general setting of PIOAs, where components may exhibit non-deterministic behavior.

We start with the notion of compatibility: two PIOAs  $A$  and  $B$  are said to be *compatible* if  $O_A \cap O_B = Act_A \cap H_B = Act_B \cap H_A = \emptyset$ .

Let  $\{A_i \mid 1 \leq i \leq n\}$  denote a set of pairwise compatible PIOAs and, for readability, we replace all subscripts  $A_i$  with  $i$ . (The same convention will be adopted throughout this paper.) The *parallel composite*, denoted  $\uparrow\uparrow_{i=1}^n A_i$ , is the PIOA  $B$  with the following state space and action signature:

- (1)  $S_B := \prod_{i=1}^n S_i$  with  $s_B^0 := \langle s_1^0, \dots, s_n^0 \rangle$ ;
- (2)  $I_B := \bigcup_{i=1}^n I_i \setminus \bigcup_{i=1}^n O_i$ ,  $O_B := \bigcup_{i=1}^n O_i$ , and  $H_B := \bigcup_{i=1}^n H_i$ ;

The reactive transition structure  $\mathbf{R}_B$  and the generative transition structure  $\mathbf{G}_B$  are given in Definition 13 and Definition 14, respectively.

**Definition 13** *Let  $\vec{s} \in S_B$  and  $a \in I_B$  be given. We define  $\mathbf{R}_B(\langle \vec{s}, a \rangle) \subseteq \text{Disc}(S_B)$  to be the set of all discrete distributions of the form  $\prod_{i=1}^n \mu_i$  for some family  $\vec{\mu} \in \prod_{i=1}^n \text{Disc}(S_i)$  satisfying:*

- if  $a \notin I_i$ , then  $\mu_i = \text{Dirac}(s_i)$ ;
- otherwise,  $\mu_i \in \mathbf{R}_i(\langle s_i, a \rangle)$ .

In other words, each process  $A_i$  stutters if the given input  $a$  is not in the signature of  $A_i$ . Otherwise,  $A_i$  reacts to this input by

- (i) first choosing *non-deterministically* a distribution  $\mu_i$  from  $\mathbf{R}_i(\langle s_i, a \rangle)$ ;
- (ii) then choosing *randomly* a state  $t_i$  according to  $\mu_i$ .

We assume that processes evolve independently, therefore a product construction on state distributions  $\mu_i$  yields a typical member of  $\mathbf{R}_B(\langle \vec{s}, a \rangle)$ .

The definition of  $\mathbf{G}_B$  for  $B = \uparrow\uparrow_{i=1}^n A_i$  is slightly more complicated, where exactly one component  $B_j$  is generative and all others are reactive.

**Definition 14** *Let  $\vec{s} \in S_B$  and  $1 \leq j \leq n$  be given. Let  $N_j$  denote the index set  $(O_B \cup H_B) \times \{i \mid 1 \leq i \leq n, i \neq j\}$ . Suppose we have a transition bundle  $g_j \in \mathbf{G}_j(s_j)$  and a family  $\vec{\mu} \in \prod_{\langle a, i \rangle \in N_j} \text{Disc}(S_i)$  of state distributions so that: for all  $\langle a, i \rangle \in N_j$ ,*

- if  $a \notin I_i$ , then  $\mu_{a,i} = \text{Dirac}(s_i)$ ;
- otherwise,  $\mu_{a,i} \in \mathbf{R}_i(\langle s_i, a \rangle)$ .

Then  $g_j$  and  $\vec{\mu}$  are said to generate the following distribution  $f$  on  $(O_B \cup H_B) \times S_B$ : for all  $\langle a, \vec{t} \rangle$ ,

$$f(\langle a, \vec{t} \rangle) := g_j(\langle a, t_j \rangle) \cdot \prod_{i \neq j} \mu_{a,i}(t_i).$$

With slight abuse of notation, we write  $f = g_j \times \prod_{\langle a, i \rangle \in N_j} \mu_{a,i}$ .

We define  $\mathbf{G}_B^j(\vec{s}) \subseteq \text{Disc}((O_B \cup H_B) \times S_B)$  to be the set of all bundles  $f$  so that

$f$  is generated by some  $g_j \in \mathbf{G}_j(s_j)$  and some  $\vec{\mu} \in \prod_{\langle a, i \rangle \in N_j} \text{Disc}(S_i)$  satisfying the conditions above. Then  $\mathbf{G}_B(\vec{s}) := \bigcup_{1 \leq j \leq n} \mathbf{G}_B^j(\vec{s})$ .

Here the unique active component  $A_j$  chooses *non-deterministically* a transition bundle  $g_j$  enabled from  $s_j$ . Once  $g_j$  is specified, a pair  $\langle a, t_j \rangle$  is chosen *randomly* according to  $g_j$ . The other processes  $A_i$  either stutter or react to the action performed by  $A_j$ , whichever dictated by their action signatures. Note that the choice of the family  $\vec{\mu}$  is *non-deterministic* and is independent from the particular pair  $\langle a, t_j \rangle$  drawn from  $g_j$ .

Lemma 15 below shows that the new bundles  $f$  constructed in Definition 14 are in fact well-defined discrete distributions.

**Lemma 15** *The bundle  $f$  in Definition 14 is well-defined.*

**PROOF.** We need to verify that  $f$  is a discrete distribution on  $(O_B \cup H_B) \times S_B$ . First consider fixed  $a \in O_j \cup H_j$ . By the definition of  $f$ , we have

$$\sum_{\vec{t} \in S_B} f(\langle a, \vec{t} \rangle) = \sum_{t_1 \in S_1} \dots \sum_{t_n \in S_n} g_j(\langle a, t_j \rangle) \cdot \prod_{i \neq j} \mu_{a,i}(t_i).$$

We can rearrange the sums and factor out  $g_j(\langle a, t_j \rangle)$  to obtain:

$$\sum_{t_j \in S_j} g_j(\langle a, t_j \rangle) \cdot \left( \sum_{t_1 \in S_1} \dots \sum_{t_{j-1} \in S_{j-1}} \sum_{t_{j+1} \in S_{j+1}} \dots \sum_{t_n \in S_n} \prod_{i \neq j} \mu_{a,i}(t_i) \right).$$

Since each  $\mu_{a,i}$  is a discrete distribution on  $S_i$ , an easy inductive argument shows that

$$\sum_{t_1 \in S_1} \dots \sum_{t_{j-1} \in S_{j-1}} \sum_{t_{j+1} \in S_{j+1}} \dots \sum_{t_n \in S_n} \prod_{i \neq j} \mu_{a,i}(t_i) = 1.$$

Then we have  $\sum_{\vec{t} \in S_B} f(\langle a, \vec{t} \rangle) = \sum_{t_j \in S_j} g_j(\langle a, t_j \rangle)$ .

Now notice that  $f(\langle a, \vec{t} \rangle) = 0$  whenever  $a \notin O_j \cup H_j$ . Therefore,

$$\begin{aligned} \sum_{\langle a, \vec{t} \rangle \in (O_B \cup H_B) \times S_B} f(\langle a, \vec{t} \rangle) &= \sum_{a \in O_j \cup H_j} \sum_{\vec{t} \in S_B} f(\langle a, \vec{t} \rangle) \\ &= \sum_{a \in O_j \cup H_j} \sum_{t_j \in S_j} g_j(\langle a, t_j \rangle) && \text{calculation above} \\ &= 1 && g_j \text{ discrete distribution} \end{aligned}$$

Therefore  $f$  is a discrete distribution on  $(O_B \cup H_B) \times S_B$ .  $\square$

This completes the definition of parallel composition for PIOAs. We write  $\uparrow\uparrow^n$  for the  $n$ -ary composition operator and, when  $n = 2$ , we omit the superscript

and use infix notation. Due to symmetries in our definitions, it is easy to see that  $\uparrow\uparrow$  is commutative. We claim that  $\uparrow\uparrow$  is also associative, because both  $(A \uparrow\uparrow B) \uparrow\uparrow C$  and  $A \uparrow\uparrow (B \uparrow\uparrow C)$  are isomorphic to  $\uparrow\uparrow^3 \{A, B, C\}$ . We omit the details.

## 6.2 Composing Switched PIOAs

As usual, we need an appropriate notion of compatibility: switched PIOAs  $A$  and  $B$  are said to be *compatible* if

- they are compatible as PIOAs;
- $Act_A \cap Sync_B = Act_B \cap Sync_A = CI_A \cap CI_B = \emptyset$ ;
- at most one of them is initially active.

Since switched PIOAs are special cases of PIOAs, one may apply the operator  $\uparrow\uparrow$  of Section 6.1 to compatible switched PIOAs. Unfortunately, the result does not always satisfy all switch axioms. We give a simple example.

Consider automata  $D$  and  $E$  in Figure 8 below and assume that all actions shown are control actions.



Fig. 8. Automata  $D$  and  $E$

If from the initial state the composite  $D \uparrow\uparrow E$  receives an input signal  $a$ , then  $D$  moves into an active state,  $s_1$ , and  $E$  remains at its initial state. This is shown in Figure 9. In state  $\langle s_1, s_E^0 \rangle$ , the composite is considered active, because  $D$  is. However, an input transition with label  $b$  is still enabled, violating Axiom (S2). Moreover, suppose in fact an input signal  $b$  is received from state  $\langle s_1, s_E^0 \rangle$ . Then in the resulting state  $\langle s_1, s_2 \rangle$  both  $D$  and  $E$  are active. This state violates Axiom (S4), because a single control output (say  $c$ ) is not sufficient to deactivate both components (Figure 9).

$$\langle s_D^0, s_E^0 \rangle \xrightarrow{a?} \langle s_1, s_E^0 \rangle \xrightarrow{b?} \langle s_1, s_2 \rangle \xrightarrow{c!} \langle s_D^0, s_2 \rangle$$

Fig. 9. A Potential Execution of  $D \uparrow\uparrow E$

This is a counterintuitive scenario: if the environment of  $D \uparrow\uparrow E$  is itself a switched PIOA, then it should have become inactive after providing the first control input  $a$ , thus unable to provide the second control input  $b$ . In fact, it is

shown in [CLSV04b] that any state with more than one active components is unreachable, provided the closing environment is also a switched PIOA. (The proof involves lengthy inductive arguments and is omitted here.)

This example suggests that, when switched PIOAs are composed using the PIOA parallel operator  $\uparrow\uparrow$ , the resulting state space and reactive transition structure both contain too many elements. Therefore, we are prompted to consider an appropriate sub-automaton with fewer states and fewer input transitions.

**Definition 16** *Let  $\{A_i \mid i \in I\}$  be a set of pairwise compatible switched PIOAs. The parallel composite,  $\parallel_{i=1}^n A_i$ , is based on the sub-automaton  $B$  of  $\uparrow\uparrow_{i=1}^n A_i$  obtained by*

- (i) *removing all states in which more than one  $A_i$ 's are active;*
- (ii) *removing all input transitions from states in which at least one  $S_i$  is active.*

Moreover,  $Sync_B := \bigcup_{1 \leq i \leq n} Sync_i \cup \bigcup_{1 \leq i, j \leq n} (CI_i \cap CO_j)$ , and  $active_B(\vec{s}) := 0$  if and only if  $active_i(s_i) = 0$  for all  $i$ .

Although the signature of  $B = \parallel_{i=1}^n A_i$  is completely specified in Definition 16, it is instructive to provide a list of explicit identities.

**Lemma 17** *The following equalities hold:*

- $BI_B = \bigcup_{1 \leq i \leq n} BI_i \setminus \bigcup_{1 \leq i \leq n} BO_i$ ;
- $CI_B = \bigcup_{1 \leq i \leq n} CI_i \setminus \bigcup_{1 \leq i \leq n} CO_i$ ;
- $BO_B = \bigcup_{1 \leq i \leq n} BO_i$ ;
- $CO_B = \bigcup_{1 \leq i \leq n} CO_i \setminus \bigcup_{1 \leq i \leq n} CI_i$ .

**PROOF.** By definition,  $I_B = \bigcup_{1 \leq i \leq n} I_i \setminus \bigcup_{1 \leq i \leq n} O_i$ . Since  $BAct$  and  $CAct$  are disjoint, we have the desired properties about  $BI_B$  and  $CI_B$ .

Similarly,  $O_B = \bigcup_{1 \leq i \leq n} O_i$ , therefore  $BO_B = \bigcup_{1 \leq i \leq n} BO_i$  and  $O_B \cap CAct = \bigcup_{1 \leq i \leq n} CO_i$ . Applying the definitions of  $CO_B$  and  $Sync_B$ , we have  $CO_B = \bigcup_{1 \leq i \leq n} CO_i \setminus \bigcup_{1 \leq i \leq n} CI_i$ .  $\square$

To show that such  $B$  is a well-defined PIOA, we need to verify (i)  $s_B^0 \in S_B$  and (ii)  $S_B$  is closed under the reduced transition structures. Clearly, the first claim holds by the definition of compatibility. The second is confirmed by Lemmas 18 and 19 below.

For convenience, we partition  $S_B$  into two sets:

- $S_{B,0}$  is the set of all  $\vec{s}$  such that  $\text{active}_i(s_i) = 0$  for all  $i$ ;
- $S_{B,1}$  is the set of all  $\vec{s}$  such that  $\text{active}_i(s_i) = 1$  for exactly one  $i$ .

**Lemma 18** *Let  $\vec{s} \in S_B$  and  $a \in I_B$  be given. For all  $\mu \in \mathbf{R}_B(\langle \vec{s}, a \rangle)$ :*

- $a \in BI_B$  implies  $\text{Supp}(\mu) \subseteq S_{B,0}$ ;
- $a \in CI_B$  implies  $\text{Supp}(\mu) \subseteq S_{B,1}$ .

**PROOF.** By definition,  $\mathbf{R}_B(\langle \vec{s}, a \rangle)$  is empty whenever  $\vec{s} \in S_{B,1}$ . Therefore we may assume that  $\vec{s} \in S_{B,0}$ . Let  $\mu \in \mathbf{R}_B(\langle \vec{s}, a \rangle)$  and  $\vec{s}' \in \text{Supp}(\mu)$  be given.

First assume  $a \in BI_B$ . For every  $i$ , if  $a \notin \text{Act}_i$ , it must be the case that  $s_i = s'_i$  and hence  $\text{active}_i(s_i) = \text{active}_i(s'_i) = 0$ . Otherwise, we have  $a \in BI_i$  and we may apply Lemma 12 to conclude that  $\text{active}_i(s'_i) = 0$ . Therefore  $\vec{s}' \in S_{B,0}$ .

Now assume  $a \in CI_B$ . By compatibility,  $a \in \text{Act}_j$  for exactly one  $j$ . Choose such  $j$ . By Lemma 12, we know  $\text{active}_j(s'_j) = 1$ . For all other  $i$ ,  $a \notin \text{Act}_i$  and hence  $\text{active}_i(s_i) = \text{active}_i(s'_i) = 0$ . This proves  $\vec{s}' \in S_{B,1}$ .  $\square$

**Lemma 19** *Let  $\vec{s} \in S_B$  and  $f \in \mathbf{G}_B(\vec{s})$  be given. For every  $\langle a, \vec{s}' \rangle \in \text{Supp}(f)$ :*

- If  $a \in BO_B \cup \text{Sync}_B \cup H_B$ , then  $\vec{s}' \in S_{B,1}$ ;
- If  $a \in CO_B$ , then  $\vec{s}' \in S_{B,0}$ .

**PROOF.** By Axiom (1), we know that  $\mathbf{G}_i(s_i)$  is empty for every  $i$  with  $\text{active}_i(s_i) = 0$ . This implies  $s \in S_{B,1}$ , because otherwise  $\mathbf{G}_B(\vec{s})$  would be empty. Let  $j$  be the unique index with  $\text{active}_j(s_j) = 1$  and choose  $g_j \in \mathbf{G}_j(s_j)$  such that  $f$  is generated by  $g_j$ . By Definition 14,  $a$  must be in  $O_j \cup H_j$ . We have the following cases.

- (1)  $a \in H_j \cup \text{Sync}_j$ . Compatibility of switched PIOAs requires that  $a \notin \text{Act}_i$  for all  $i \neq j$ . This implies, for all  $i \neq j$ ,  $s_i = s'_i$  and hence  $\text{active}_i(s'_i) = \text{active}_i(s_i) = 0$ . On the other hand, we may apply Lemma 12 to  $A_j$  and conclude that  $\text{active}_i(s'_j) = \text{active}_i(s_j) = 1$ . Therefore,  $\vec{s}' \in S_{B,1}$ .
- (2)  $a \in BO_j$ . For every  $i$  such that  $a \notin \text{Act}_i$ , we know that  $s_i = s'_i$  and hence  $\text{active}_i(s_i) = \text{active}_i(s'_i) = 0$ . For every  $i$  such that  $i \neq j$  and  $a \in \text{Act}_i$ , it must be the case that  $a \in BI_i$ , so we apply Lemma 12 to conclude that  $\text{active}_i(s'_i) = 0$ . As in the previous case, we know  $\text{active}_i(s'_j) = 1$ . Therefore,  $\vec{s}' \in S_{B,1}$ .
- (3)  $a \in CO_j \cap CI_k$  for some  $k \neq j$ . By Lemma 12, we have  $\text{active}_j(s'_j) = 0$  and  $\text{active}_k(s'_k) = 1$ . By the compatibility of switched PIOAs, there is at most one such  $k$ . For all other indices  $i$ ,  $\text{active}_i(s_i) = \text{active}_i(s'_i) = 0$ . Again we conclude  $\vec{s}' \in S_{B,1}$ .

- (4)  $a \in CO_B$ . By the definition of  $CO_B$ , we know that  $a \notin Act_i$  for all  $i \neq j$ . Hence  $\mathbf{active}_i(s_i) = \mathbf{active}_i(s'_i) = 0$  for all  $i \neq j$ . By Lemma 12, we have  $\mathbf{active}_j(s'_j) = 0$ . Thus,  $\vec{s}' \in S_{B,0}$ .  $\square$

It remains to show that  $B$  satisfies all switch axioms.

**Lemma 20** *The PIOA  $B$ , together with  $\mathbf{active}_B$  and  $\mathit{Sync}_B$ , satisfies Axioms (S1) through (S5) in Definition 11.*

**PROOF.** Note that  $\mathbf{active}_B(\vec{s}) = 0$  if and only if  $\vec{s} = S_{B,0}$ . For Axiom (S1), let  $\vec{s} \in S_{B,0}$  and  $a \in I_B$  be given. Applying Axiom (S1) on each component, we know that  $\mathbf{G}_i(s_i)$  is empty for every  $i$  with  $\mathbf{active}_i(s_i) = 0$  and hence  $\mathbf{G}_B(\vec{s}) = \emptyset$ . On the other hand, for all  $i$  with  $a \in I_i$ , Axiom (S2) requires  $\mathbf{R}_i(\langle s_i, a \rangle)$  is non-empty. Hence  $\mathbf{R}_B(\langle \vec{s}, a \rangle)$  is non-empty. This proves that  $B$  satisfies Axiom (S1).

Axiom (S2) follows from the definition of  $\mathbf{R}_B$ . Axioms (S3) through (S5) follow from Lemmas 18 and 19.  $\square$

We adopt the same notational conventions as with  $\uparrow\uparrow$ . Namely,  $\|$  <sup>$n$</sup>  denotes the  $n$ -ary operator and  $\|$  denotes the (infix) binary operator. Again commutativity is trivial. associativity. For associativity, it is easy to see that  $(A \| B) \| C$  has the same state space as  $\|$ <sup>3</sup>  $\{A, B, C\}$ . Similarly for  $A \| (B \| C)$ . The transition structures are isomorphic because they are based on parallel composition of PIOAs, which is associative.

### 6.3 Composing I/O Schedulers

The goal of this section is to extend the parallel operator  $\|$  to probabilistic systems, therefore we consider composition of i/o schedulers. For that end, we need some basic notions of projection.

In Section 2, we described projection operators for discrete distributions on a product space. Extending the same idea, we define projection on composite transition bundles.

**Definition 21** *Let  $\{A_i \mid 1 \leq i \leq n\}$  be a set of pairwise compatible PIOAs and let  $B$  denote  $\uparrow\uparrow_{i=1}^n A_i$ . Let  $\vec{s} \in S_B$  and  $f \in \mathbf{G}_B(\vec{s})$  be given. Let  $j$  be the unique index such that  $\pi_L(\mathbf{Supp}(f)) \subseteq O_j \cup H_j$  (equivalently,  $f \in \mathbf{G}_B^j(\vec{s})$ ). The  $j$ th-projection of  $f$ , denoted  $\pi_j(f)$ , is the discrete distribution on  $(O_j \cup H_j) \times S_j$*

given by:

$$\pi_j(f)(\langle a, t \rangle) := \sum_{\vec{i} \in S_B : t_j = t} f(\langle a, \vec{t} \rangle).$$

For every  $a \in \pi_L(\text{Supp}(f))$  and  $i \neq j$ , the  $\langle a, i \rangle$ th-projection of  $f$ , denoted  $\pi_{a,i}(f)$ , is the discrete distribution on  $S_i$  given by:

$$\pi_{a,i}(f)(t) := \frac{\sum_{\vec{i} \in S_B : t_i = t, t_j = u} f(\langle a, \vec{t} \rangle)}{\pi_j(f)(\langle a, u \rangle)},$$

where  $u$  is any state in  $S_j$  such that  $\pi_j(f)(\langle a, u \rangle) \neq 0$ .

Lemmas 22 and 23 below show that these projection operators are in fact well-defined.

**Lemma 22** *The distribution  $\pi_j(f)$  in Definition 21 is well-defined and is in  $\mathbf{G}_j(s_j)$ .*

**PROOF.** By the definition of  $\mathbf{G}_B(\vec{s})$ , we may choose  $g_j \in \mathbf{G}_j(s_j)$  such that  $f$  is generated by  $g_j$ . It suffices to show  $\pi_j(f) = g_j$ . Let  $\langle a, t \rangle \in (O_j \cup H_j) \times S_j$  be given. By definition,

$$\pi_j(f)(\langle a, t \rangle) = \sum_{\vec{i} \in S_B : t_j = t} f(\langle a, \vec{t} \rangle) = \sum_{\vec{i} \in S_B : t_j = t} g_j(\langle a, t_j \rangle) \cdot \prod_{i \neq j} \mu_{a,i}(t_i).$$

We can rearrange the sums and factor out  $g_j(\langle a, t_j \rangle)$  to obtain:

$$\pi_j(f)(\langle a, t \rangle) = g_j(\langle a, t \rangle) \cdot \left( \sum_{t_1 \in S_1} \dots \sum_{t_{j-1} \in S_{j-1}} \sum_{t_{j+1} \in S_{j+1}} \dots \sum_{t_n \in S_n} \prod_{i \neq j} \mu_{a,i}(t_i) \right).$$

Since every  $\mu_{a,i}$  is a discrete distribution on  $S_i$ , the second factor equals 1. Hence  $\pi_j(f)(\langle a, t \rangle) = g_j(\langle a, t \rangle)$ .  $\square$

**Lemma 23** *The distribution  $\pi_{a,i}(f)$  in Definition 21 is well-defined. Moreover, if  $a \in I_i$ , then  $\pi_{a,i}(f) \in \mathbf{R}_i(\langle s_i, a \rangle)$ ; otherwise,  $\pi_{a,i}(f) = \text{Dirac}(s_i)$ .*

**PROOF.** By the definition of  $\mathbf{G}_B(\vec{s})$ , we may choose  $\mu_{a,i} \in \text{Disc}(S_i)$  and  $g_j \in \mathbf{G}_j(s_j)$  such that  $f$  is generated by  $\mu_{a,i}$  and  $g_j$ . It suffices to show  $\pi_{a,i}(f) = \mu_{a,i}$ . Let  $t \in S_i$  be given. By definition,  $\pi_{a,i}(f)(t)$  equals

$$\frac{\sum_{\vec{i} \in S_B : t_i = t, t_j = u} f(\langle a, \vec{t} \rangle)}{\pi_j(f)(\langle a, u \rangle)} = \frac{\sum_{\vec{i} \in S_B : t_i = t, t_j = u} (g_j(\langle a, t_j \rangle) \cdot \prod_{k \neq j} \mu_{a,k}(t_k))}{\pi_j(f)(\langle a, u \rangle)}.$$

Factoring out  $g_j(\langle a, u \rangle)$  and  $\mu_{a,i}(t)$ , the numerator becomes

$$g_j(\langle a, u \rangle) \cdot \mu_{a,i}(t) \cdot \sum_{\vec{i} \in S_B : t_i = t, t_j = u} \prod_{k \neq i, j} \mu_{a,k}(t_k).$$

Again the third factor is easily seen to be 1 and hence the numerator equals  $g_j(\langle a, u \rangle) \cdot \mu_{a,i}(t)$ . Moreover, we saw in the proof of Lemma 22 that  $\pi_j(f) = g_j$ , therefore the denominator equals  $g_j(\langle a, u \rangle)$ . Now we have

$$\pi_{a,i}(f)(t) = \frac{g_j(\langle a, u \rangle) \cdot \mu_{a,i}(t)}{g_j(\langle a, u \rangle)} = \mu_{a,i}(t).$$

Notice we haven't used any additional assumption on  $u$ , therefore the equality holds regardless of the choice of  $u$ .  $\square$

Given these projection operators on transition bundles, it is straightforward to define projection on execution branches.

**Definition 24** *Let  $\{A_i \mid 1 \leq i \leq n\}$  be a set of pairwise compatible PIOAs and let  $B$  denote  $\uparrow\uparrow_{i=1}^n A_i$ . Let  $\vec{s} \in S_B$  and  $1 \leq i \leq n$  be given. We define, recursively, the  $i$ -th projection operator on  $\text{Bran}(\vec{s})$  as follows:*

- $\pi_i(\langle s_1^0, \dots, s_n^0 \rangle) := s_i^0$ ;
- $\pi_i(r.a.\mu.\vec{t})$  equals
  - $\pi_i(r).a.\pi_i(\mu).t_i$ , if  $a \in I_i$ ;
  - $\pi_i(r)$ , otherwise;
- $\pi_i(r.f.a.\vec{t})$  equals
  - $\pi_i(r).\pi_i(f).a.t_i$ , if  $i$  is the unique index with  $\pi_L(\text{Supp}(f)) \subseteq O_i \cup H_i$ ;
  - $\pi_i(r).a.\pi_{a,i}(f).t_i$ , if  $a \in I_i$ ;
  - $\pi_i(r)$ , otherwise.

These projected branches are well-defined by virtue of Lemma 25 below.

**Lemma 25** *Let  $1 \leq i \leq n$  be given. For all  $q \in \text{Bran}(\vec{s})$ , we have*

- (1)  $\pi_i(\text{last}(q)) = \text{last}(\pi_i(q))$ ;
- (2) if  $q$  is of the form  $r.a.\mu.\vec{t}$  and  $a \in I_i$ , then  $\pi_i(\mu) \in \mathbf{R}_i(\langle \text{last}(\pi_i(r)), a \rangle)$  and  $t_i \in \text{Supp}(\pi_i(\mu))$ ;
- (3) if  $q$  is of the form  $r.f.a.\vec{t}$  and  $a \in O_i \cup H_i$ , then  $\pi_i(f) \in \mathbf{G}_i(\text{last}(\pi_i(r)))$  and  $\langle a, t_i \rangle \in \text{Supp}(\pi_i(f))$ ;
- (4) if  $q$  is of the form  $r.f.a.\vec{t}$  and  $a \in I_i$ , then  $\pi_{a,i}(f) \in \mathbf{R}_i(\langle \text{last}(\pi_i(r)), a \rangle)$  and  $t_i \in \text{Supp}(\pi_{a,i}(f))$ ;

**PROOF.** We proceed by induction on the length of  $r$ . The base case is trivial.

Consider a branch of the form  $r.a.\mu.\vec{t}$  and let  $\vec{u}$  denote  $\text{last}(r)$ . By the induction hypothesis, we have  $\pi_i(\text{last}(r)) = u_i = \text{last}(\pi_i(r))$ . Recall that  $\mu = \prod_{i=1}^n \pi_i(\mu_i)$ . We have two cases.

- $a \in I_i$ . Then by Definition 13 we have  $\pi_i(\mu) \in \mathbf{R}_i(\langle u_i, a \rangle)$ . Since  $\vec{t} \in \text{Supp}(\mu)$ , it must be that  $t_i \in \text{Supp}(\pi_i(\mu))$ . Moreover,  $\pi_i(\text{last}(q)) = t_i = \text{last}(\pi_i(q))$ .
- $a \notin I_i$ . Then by Definition 13 we have  $\pi_i(\mu) = \text{Dirac}(u_i)$ . Since  $\vec{t} \in \text{Supp}(\mu)$ , it must be that  $t_i = u_i$ . Therefore  $\pi_i(\text{last}(q)) = t_i = u_i = \text{last}(\pi_i(r)) = \text{last}(\pi_i(q))$ .

Now we consider a branch of the form  $r.f.a.\vec{t}$ . Again, let  $\vec{u}$  denote  $\text{last}(r)$  and we have  $\pi_i(\text{last}(r)) = u_i = \text{last}(\pi_i(r))$  by the induction hypothesis. By Definition 14 we may choose unique  $j$  such that  $f = g_j \times \prod_{\langle a, i \rangle \in N_j} \mu_{a,i}$  for some  $g_j \in \mathbf{G}_j(u_j)$  and family  $\{\mu_{a,i}\}_{\langle a, i \rangle \in N_j} \in \prod_{\langle a, i \rangle \in N_j} \text{Disc}(S_i)$ . We have three cases.

- $i = j$ . Then we have  $\pi_i(f) = g_i \in \mathbf{G}_i(u_i)$ . Since  $\langle a, \vec{t} \rangle \in \text{Supp}(f)$ , it must be that  $\langle a, t_i \rangle \in \text{Supp}(g_i) = \text{Supp}(\pi_i(f))$ . Moreover,  $\pi_i(\text{last}(q)) = t_i = \text{last}(\pi_i(q))$ .
- $i \neq j$  and  $a \in I_i$ . Then by Definition 14 we have  $\pi_{a,i}(f) \in \mathbf{R}_i(\langle u_i, a \rangle)$ . Since  $\langle a, \vec{t} \rangle \in \text{Supp}(f)$ , it must be that  $t_i \in \text{Supp}(\pi_{a,i}(f))$ . Moreover,  $\pi_i(\text{last}(q)) = t_i = \text{last}(\pi_i(q))$ .
- $i \neq j$  and  $a \notin I_i$ . Then by Definition 14 we have  $\pi_{a,i}(f) = \text{Dirac}(u_i)$ . Since  $\langle a, \vec{t} \rangle \in \text{Supp}(f)$ , it must be that  $t_i = u_i$ . Then  $\pi_i(\text{last}(q)) = t_i = u_i = \text{last}(\pi_i(r)) = \text{last}(\pi_i(q))$ .  $\square$

We are now ready to consider composition of i/o schedulers for switched PIOAs.

**Definition 26** *Let  $\{A_i \mid 1 \leq i \leq n\}$  be a set of pairwise compatible switched PIOAs and let  $B$  denote  $\prod_{i=1}^n A_i$ . Suppose we have, for each  $i$ , an i/o scheduler  $\langle \sigma_i, \rho_i \rangle$  for  $A_i$ . These i/o schedulers are said to generate the following i/o scheduler  $\langle \sigma, \rho \rangle$  for  $B$ . Let  $r \in \text{Bran}(B)$  be given and let  $\vec{s}$  denote  $\text{last}(r)$ .*

- If  $\text{active}_B(\vec{s}) = 1$ , then  $\sigma(\langle r, a \rangle) := \perp$  for all  $a \in I_B$ .
- If  $\text{active}_B(\vec{s}) = 0$ , then for all  $a \in I_B$ ,  $\sigma(\langle r, a \rangle) := \prod_{i=1}^n \mu_i$ , where  $\mu_i$  equals  $\text{Dirac}(s_i)$  whenever  $a \notin I_i$  and  $\sigma_i(\langle \pi_i(r), a \rangle)$  otherwise.
- If  $\text{active}_B(\vec{s}) = 0$ , then  $\rho(r) := \perp$ .
- If  $\text{active}_B(\vec{s}) = 1$ , then  $\rho(r) \neq \perp$  if and only if  $\rho_j(\pi_j(r)) \neq \perp$ , where  $j$  is the unique index with  $\text{active}_j(s_j) = 1$ . In that case,  $\rho(r)$  is the bundle  $f = \rho_j(\pi_j(r)) \times \prod_{\langle a, i \rangle \in N_j} \mu_{a,i}$ , where  $\mu_{a,i}$  equals  $\text{Dirac}(s_i)$  whenever  $a \notin I_i$  and  $\sigma_i(\langle \pi_i(r), a \rangle)$  otherwise.

**Lemma 27** *The i/o scheduler  $\langle \sigma, \rho \rangle$  in Definition 26 is well-defined.*

**PROOF.** Let  $r \in \text{Bran}(B)$  and  $a \in I_B$  be given. Let  $\vec{s}$  denote  $\text{last}(r)$ . First we consider the case where  $\mathbf{R}_B(\langle \text{last}(r), a \rangle)$  is non-empty. Since  $B$  satisfies

Axiom (S2), it must be the case that  $\mathbf{active}_B(\vec{s}) = 0$  and hence  $\mathbf{active}_i(s_i) = 0$  for all  $i$ .

By Axiom (S1),  $\mathbf{R}_i(\langle s_i, a \rangle)$  is non-empty for all  $a \in I_i$ . By the definition of input schedulers, this implies  $\sigma_i(\langle \pi_i(r), a \rangle)$  is defined and is in  $\mathbf{R}_i(\langle s_i, a \rangle)$ . By the definition of  $\mathbf{R}_B$ , we have that  $\prod_{i=1}^n \mu_i$  is in  $\mathbf{R}_B(\langle \vec{s}, a \rangle)$ . This proves that  $\sigma(\langle r, a \rangle)$  is in  $\mathbf{R}_B(\langle \mathbf{last}(r), a \rangle)$  whenever  $\mathbf{R}_B(\langle \mathbf{last}(r), a \rangle)$  is non-empty.

Now assume that  $\mathbf{R}_B(\langle \mathbf{last}(r), a \rangle)$  is empty. By Axiom (S1), we may conclude that  $\mathbf{active}_B(\vec{s}) = 1$ , in which case  $\sigma(\langle r, a \rangle)$  is by definition undefined for all  $a \in I_B$ . This completes the proof that  $\sigma$  is a well-defined input scheduler for  $B$ .

For the output scheduler  $\rho$ , we need to show that  $\rho(r) \in \mathbf{G}_B(\mathbf{last}(r))$  whenever  $\rho(r)$  is defined. Therefore, we may focus on the case in which  $\mathbf{active}_B(\vec{s}) = 1$ . By the definition of  $S_B$ , there is in fact unique  $j$  with  $\mathbf{active}_j(s_j) = 1$ . Assume without loss that  $\rho_j(\pi_j(r))$  is defined. By the definition of output schedulers,  $\rho_j(\pi_j(r)) \in \mathbf{G}_j(s_j)$ .

Moreover, we know that  $\mathbf{active}_i(s_i) = 0$  for all  $i \neq j$ . Fix  $a \in O_B \cup H_B$  and  $i \neq j$ . By Axiom (S1),  $\mathbf{R}_i(\langle s_i, a \rangle)$  is non-empty whenever  $a \in I_i$ . This implies that  $\sigma_i(\langle \pi_i(r), a \rangle)$  is defined and is in  $\mathbf{R}_i(\langle s_i, a \rangle)$ . Therefore, the family  $\{\mu_{a,i}\}_{\langle a, i \rangle \in N_j}$  satisfies the conditions in Definition 14 and thus the bundle  $f$  generated by  $\rho_j(\pi_j(r))$  and  $\{\mu_{a,i}\}_{\langle a, i \rangle \in N_j}$  is in  $\mathbf{G}_B(\mathbf{last}(r))$ . This completes the proof that  $\rho$  is a well-defined input scheduler for  $B$ .  $\square$

Notice that Definition 26 and the proof of Lemma 27 rely on the definition of  $\parallel$  and switch axioms, therefore they do not apply to PIOAs in general. Roughly speaking, the parallel composition mechanism for PIOAs does not attempt to resolve global non-determinism, therefore it is not possible to combine two local schedules to form a single global schedule. The token structure of switched PIOAs serves precisely the purpose of eliminating such global non-determinism.

Definition 26 induces to a very natural notion of composition for switched probabilistic systems.

**Definition 28** *Let  $\{\mathcal{A}_i \mid 1 \leq i \leq n\}$  be a set of probabilistic systems where  $\mathcal{A}_i = \langle A_i, \mathcal{S}_i \rangle$  and  $\{\mathcal{A}_i \mid 1 \leq i \leq n\}$  are pairwise compatible switched PIOAs. The parallel composite, denoted  $\parallel_{i=1}^n \mathcal{A}_i$ , is the probabilistic system  $\mathcal{B} = \langle B, \mathcal{T} \rangle$  defined as follows:*

- *the underlying switched PIOA is  $B = \parallel_{i=1}^n A_i$ ;*
- *the set  $\mathcal{T}$  of i/o schedulers contains precisely those  $\langle \sigma, \rho \rangle$  generated by some family  $\{\langle \sigma_i, \rho_i \rangle\}_{1 \leq i \leq n} \in \prod_{i=1}^n \mathcal{S}_i$ .*

Again, we adopt notational conventions as in the case of  $\parallel$  for switched PIOAs. Commutativity and associativity follow similarly.

Before ending this section, let us briefly revisit automata  $\text{Early}'$ ,  $\text{Late}'$  and  $\text{Coin}'$  of Figure 5. Consider the full probabilistic systems induced by these automata (i.e., each automaton is paired with all possible local i/o schedulers). We claim that, when  $\text{Late}'$  and  $\text{Coin}'$  are composed using Definition 28, it is no longer possible to obtain the schedule depicted in Figure 2. This is because the local i/o scheduler of  $\text{Late}'$  must choose between  $b$  and  $c$  without “knowing” the random outcome in  $\text{Coin}'$ . Extending this intuition, it is not hard to show that  $\text{Early}' \parallel \text{Coin}'$  and  $\text{Late}' \parallel \text{Coin}'$  are equivalent in our external behavior semantics.

## 7 Compositionality

We proceed to state and prove our main theorem: the external behavior semantics for switched probabilistic systems (Definition 9) is compositional with respect to the composition operator introduced in Definition 28.

**Theorem 29** *Let  $\mathcal{A} = \langle A, \mathcal{S} \rangle$ ,  $\mathcal{C} = \langle C, \mathcal{U} \rangle$  and  $\mathcal{D} = \langle D, \mathcal{V} \rangle$  be switched probabilistic systems. Assume that  $\mathcal{A}$  and  $\mathcal{D}$  are comparable and  $\text{ExtBeh}(\mathcal{A}) \subseteq \text{ExtBeh}(\mathcal{D})$ . Moreover, assume that  $\mathcal{C}$  is compatible with both  $\mathcal{A}$  and  $\mathcal{D}$ . Then  $\text{ExtBeh}(\mathcal{A} \parallel \mathcal{C}) \subseteq \text{ExtBeh}(\mathcal{D} \parallel \mathcal{C})$ .*

To prove this theorem, we need quite a few auxiliary results. Recall from Definition 8 that likelihood assignments are defined in terms of minimal execution branches. We will start with a pasting result on minimal branches in a parallel composite (Section 7.1, Lemma 33). Then, in Section 7.2, we consider pasting results for execution trees and likelihood assignments. That lays sufficient ground for the proof of Theorem 29 in Section 7.3.

Throughout the rest of this section, let  $A_1$  and  $A_2$  be compatible switched PIOAs and define  $B := A_1 \parallel A_2$ . Moreover, let  $\langle \sigma_1, \rho_1 \rangle$  and  $\langle \sigma_2, \rho_2 \rangle$  be i/o schedulers for  $A_1$  and  $A_2$ , respectively, and let  $\langle \sigma, \rho \rangle$  denote the i/o scheduler for  $B$  generated by  $\langle \sigma_1, \rho_1 \rangle$  and  $\langle \sigma_2, \rho_2 \rangle$  (cf. Definition 26).

### 7.1 Minimal Execution Branches

First we make an observation about the activity status of end states of non-minimal branches. This is essentially a corollary of Lemma 12.

**Lemma 30** *Let  $A$  be any switched PIOA and let  $s$  be a state in  $A$ . For every non-minimal branch  $r$  in  $\text{Bran}(s)$ ,  $\text{active}_A(\text{last}(r)) = 1$ .*

**PROOF.** Since  $r$  is non-minimal, it must be non-empty and of the form  $q.f.a.t$  where  $a \in H_A$ . Then we have  $\text{last}(q) \xrightarrow{a} t$ . By Lemma 12, we know that  $\text{active}_A(\text{last}(r)) = \text{active}_A(t) = 1$ .  $\square$

Lemma 31 below says, when we project a minimal branch in  $B$  onto one of its components, the result is always minimal. Lemma 32 states that, given  $r_1 \in \text{Bran}_{\min}(A_1)$  and  $r_2 \in \text{Bran}_{\min}(A_2)$  with matching traces, we can “zip” them together in a unique way to form a minimal branch in  $B$ . Finally, Lemma 33 states that, given a fixed trace  $\alpha$ , there is a bijective correspondence between  $\text{tr}_{\min}^{-1}(\alpha)$  in  $B$  and the Cartesian product of  $\text{tr}_{\min}^{-1}(\pi_1(\alpha))$  in  $A_1$  and  $\text{tr}_{\min}^{-1}(\pi_2(\alpha))$  in  $A_2$ .

**Lemma 31** *For every minimal branch  $r$  in  $\text{Bran}(B)$ , both  $\pi_1(r)$  and  $\pi_2(r)$  are minimal.*

**PROOF.** Without loss of generality, we consider only  $\pi_1(r)$ . Recall that empty branches are always minimal, so we may focus on non-empty branches.

Consider a minimal branch of the form  $r.a.\mu.\vec{t}$  and let  $\vec{s}$  denote  $\text{last}(r)$ . Notice that,  $a$  must be in  $I_B$ , hence in  $I_1 \cup I_2$ . There are two cases:

- $a \in I_1$ . Then  $\pi_1(r.a.\mu.\vec{t}) = \pi_1(r).a.\pi_1(\mu).t_1$ , which is minimal because  $a$  is visible.
- $a \notin I_1$ . Then  $\pi_1(r.a.\mu.\vec{t}) = \pi_1(r)$ . Moreover, note that  $\mu \in \mathbf{R}_B(\langle \vec{s}, a \rangle)$ . Therefore, by Axiom (2), we know that  $\text{active}_B(\vec{s}) = 0$ . This implies  $\text{active}_1(\text{last}(\pi_1(r))) = \text{active}_1(s_1) = 0$ . Therefore by Lemma 30 we know  $\pi_1(r)$  is minimal.

Now we consider a minimal branch of the form  $r.f.a.\vec{t}$  and again let  $\vec{s}$  denote  $\text{last}(r)$ . In this case,  $a$  must be in  $O_B$ , hence in  $O_1 \cup O_2$ . Here we have three cases.

- $a \in O_1$ . Then  $\pi_1(r.f.a.\vec{t}) = \pi_1(r).\pi_1(f).a.t_1$ , which is minimal because  $a$  is visible.
- $a \in I_1$ . Then  $\pi_1(r.f.a.\vec{t}) = \pi_1(r).a.\pi_{a,1}(f).t_1$ , which is minimal because  $a$  is visible.
- $a \notin I_1$ . Then  $\pi_1(r.f.a.\vec{t}) = \pi_1(r)$ . Moreover, note that  $f$  must be generated by some  $g_2 \in \mathbf{G}_2(s_2)$ . Therefore, by Axiom (2), we know that  $\text{active}_2(s_2) = 1$ . By the definition of  $S_B$ , we have  $\text{active}_1(\text{last}(\pi_1(r))) = \text{active}_1(s_1) = 0$ . Again, by Lemma 30, we know  $\pi_1(r)$  is minimal.  $\square$

**Lemma 32** *Let  $\alpha \in (I_B \cup O_B)^{<\omega}$  be given. Let  $p$  be a minimal branch of  $A_1$  such that  $\text{tr}(p) = \pi_1(\alpha)$ . Similarly for  $q$  in  $A_2$ . There is a unique minimal branch  $r$  of  $B$  such that  $\pi_1(r) = p$ ,  $\pi_2(r) = q$ , and  $\text{tr}(r) = \alpha$ .*

**PROOF.** We proceed by induction on the length of  $\alpha$ . If  $\alpha$  is empty, then, by minimality,  $p$  and  $q$  are both empty. Take  $r$  to be the empty branch in  $B$ .

Consider  $\alpha a$ . Let  $p'$  be a minimal branch of  $A_1$  with trace  $\pi_1(\alpha a)$  and let  $p$  denote the unique minimal prefix of  $p'$  with trace  $\pi_1(\alpha)$ . Similarly for  $q \sqsubseteq q'$  in  $A_2$ . By induction hypothesis, choose a unique minimal branch  $r$  such that  $\pi_1(r) = p$ ,  $\pi_2(r) = q$ , and  $\text{tr}(r) = \alpha$ .

First assume that  $a$  is in  $O_1 \cup H_1$ . We have two cases.

- $a \notin I_2$ . Then  $\pi_2(\alpha) = \pi_2(\alpha a)$ . Therefore  $q = q'$  and we take  $r'$  to be the unique extension of  $r$  in which  $A$  follows  $p'$  and  $B$  idles after  $q$ .
- $a \in I_2$ . Then  $q'$  ends with an  $a$ -transition. Let  $q_0$  be the one-step prefix of  $q'$ . By Lemma 12, we know that  $\text{active}_2(\text{last}(q_0)) = 0$ . By Lemma 30,  $q_0$  is minimal and hence coincides with  $q$ . Take  $r'$  to be the unique extension of  $r$ , in which  $A_1$  follows  $p'$  and  $A_2$  idles after  $r$  until the last step (i.e., the  $a$ -step).

The case in which  $a$  is locally controlled by  $A_2$  is symmetric. It remains to consider the case where  $a$  is an input of  $B$ . Again, if  $a$  is not in the signature of  $A_1$ , then  $p = p'$ ; otherwise,  $a \in I_1$  and we apply Lemma 12 and Lemma 30 to conclude that  $p$  is the one-step prefix of  $p'$ . Similarly for  $q$  and  $q'$ . Take  $r'$  to be the unique (one-step) extension of  $r$  in which (1)  $A_i$  takes an  $a$ -step after  $r$ , if  $a \in I_i$ ; (2)  $A_i$  idles after  $r$  otherwise;  $\square$

**Lemma 33** *Let  $X$  denote  $\text{tr}_{\min}^{-1}(\alpha)$  in  $B$ . Let  $Y$  and  $Z$  denote  $\text{tr}_{\min}^{-1}(\pi_1(\alpha))$  in  $A_1$  and  $\text{tr}_{\min}^{-1}(\pi_2(\alpha))$  in  $A_2$ , respectively. There exists an isomorphism  $\text{zip} : Y \times Z \rightarrow X$  whose inverse is  $\langle \pi_1, \pi_2 \rangle$ .*

**PROOF.** By Lemma 31 and Lemma 32.  $\square$

## 7.2 Execution Trees and Likelihood Assignments

For the rest of this section, let  $Q, Q_1$  and  $Q_2$  be abbreviations for  $Q_{\sigma, \rho}, Q_{\sigma_1, \rho_1}$  and  $Q_{\sigma_2, \rho_2}$ , respectively. Similarly for  $\mathbb{L}, \mathbb{L}_1$  and  $\mathbb{L}_2$ . Lemma 34 below says an execution tree of the parallel composite can be obtained as a pointwise product of the execution trees of the components. Lemma 35 then combines Lemma 33 and Lemma 34 to show the analogous result for likelihood assignments.

**Lemma 34** For every  $r$  in  $\text{Bran}(B)$ , we have  $Q(r) = Q_1(\pi_1(p)) \cdot Q_2(\pi_2(p))$ .

**PROOF.** If  $r$  is empty,  $Q(r) = 1 = Q_1(\pi_1(r)) \cdot Q_2(\pi_2(r))$ .

Consider  $r' = r.a.\mu.\vec{t}$  and let  $\vec{s}$  denote  $\text{last}(r)$ . By Definition 13,  $\mu$  is of the form  $\mu_1 \times \mu_2$ , where  $\mu_i = \text{Dirac}(s_i)$  whenever  $a \notin I_i$ . Define  $c_i$  to be 0 if  $a \in I_i$  but  $\mu_i \neq \sigma_i(\langle \pi_i(r), a \rangle)$ . Otherwise,  $c_i$  is 1. Then we have

$$\begin{aligned} Q(r') &= Q(r) \cdot \mu(\vec{t}) \cdot c_1 \cdot c_2 && \text{definitions } \sigma, Q \\ &= Q_1(\pi_1(r)) \cdot Q_2(\pi_2(r)) \cdot c_1 \cdot \mu_1(t_1) \cdot c_2 \cdot \mu_2(t_2) && \text{I.H.} \\ &= Q_1(\pi_1(r')) \cdot Q_2(\pi_2(r')) && \text{definitions } Q_1, Q_2 \end{aligned}$$

Next we consider  $r' = r.f.a.\vec{t}$  and also let  $\vec{s}$  denote  $\text{last}(r)$ . Without loss of generality, assume that  $f$  is generated by some  $g_1$  and  $\{\mu_{b,2}\}_{(b,2) \in N_1}$ . Notice that, if  $b \notin I_2$ , then  $\mu_{b,2}$  must be  $\text{Dirac}(s_2)$ .

Now define  $c_1$  to be 0 if  $g_1 \neq \rho_1(\pi_1(r))$  and 1 otherwise. Similarly, define  $c_2$  to be 0 if  $a \in I_2$  but  $\mu_{a,2} \neq \sigma_2(\langle \pi_2(r), a \rangle)$ . Otherwise,  $c_2$  is 1. Similar to the previous case, we have

$$\begin{aligned} Q(r') &= Q(r) \cdot f(\langle a, \vec{t} \rangle) \cdot c_1 \cdot c_2 && \text{definitions } \rho, Q \\ &= Q_1(\pi_1(r)) \cdot Q_2(\pi_2(r)) \cdot c_1 \cdot g_1(\langle a, t_1 \rangle) \cdot c_2 \cdot \mu_{a,2}(t_2) && \text{definition } f \text{ and I.H.} \\ &= Q_1(\pi_1(r')) \cdot Q_2(\pi_2(r')) && \text{definitions } Q_1, Q_2 \end{aligned}$$

□

**Lemma 35** Let  $\alpha \in (I_B \cup O_B)^{<\omega}$  be given. We have  $\mathbb{L}(\alpha) = \mathbb{L}_1(\pi_1(\alpha)) \cdot \mathbb{L}_2(\pi_2(\alpha))$ .

**PROOF.** Let  $X$  denote  $\text{tr}_{\min}^{-1}(\alpha)$  in  $B$ . Let  $Y$  and  $Z$  denote  $\text{tr}_{\min}^{-1}(\pi_1(\alpha))$  in  $A_1$  and  $\text{tr}_{\min}^{-1}(\pi_2(\alpha))$  in  $A_2$ , respectively. We have

$$\begin{aligned} \mathbb{L}(\alpha) &= \sum_{r \in X} Q(r) && \text{definition } \mathbb{L} \\ &= \sum_{r \in X} Q_1(\pi_1(r)) \cdot Q_2(\pi_2(r)) && \text{Lemma 34} \\ &= \sum_{p \in Y, q \in Z} Q_1(p) \cdot Q_2(q) && \text{Lemma 33} \\ &= \left( \sum_{p \in Y} Q_1(p) \right) \cdot \left( \sum_{q \in Z} Q_2(q) \right) && \text{factorization} \\ &= \mathbb{L}_1(\pi_1(\alpha)) \cdot \mathbb{L}_2(\pi_2(\alpha)). && \text{definition } \mathbb{L}_1 \text{ and } \mathbb{L}_2 \end{aligned}$$

□

### 7.3 Main Proof

**PROOF.** [Proof of Theorem 29] First note that, if  $\mathcal{A}$  and  $\mathcal{D}$  are comparable and  $\mathcal{C}$  is compatible with both  $\mathcal{A}$  and  $\mathcal{D}$ , then  $\mathcal{A} \parallel \mathcal{C}$  is comparable to  $\mathcal{D} \parallel \mathcal{C}$ .

Let  $\mathbb{L} \in \text{ExtBeh}(\mathcal{A} \parallel \mathcal{C})$  be given. We need to show that  $\mathbb{L}$  is also in  $\text{ExtBeh}(\mathcal{D} \parallel \mathcal{C})$ . Let  $\langle \sigma, \rho \rangle$  be an i/o scheduler for  $\mathcal{A} \parallel \mathcal{C}$  such that  $\mathbb{L} = \text{tr}(Q_{\sigma, \rho})$ . By the definition of  $\parallel$  for probabilistic systems, we may choose  $\langle \sigma_A, \rho_A \rangle \in \mathcal{S}$  and  $\langle \sigma_C, \rho_C \rangle \in \mathcal{U}$  so that they generate  $\langle \sigma, \rho \rangle$ . Let  $\mathbb{L}_A$  and  $\mathbb{L}_C$  denote  $\text{tr}(Q_{\sigma_A, \rho_A})$  and  $\text{tr}(Q_{\sigma_C, \rho_C})$ , respectively.

On the other hand, we know that  $\mathbb{L}_A \in \text{ExtBeh}(\mathcal{A}) \subseteq \text{ExtBeh}(\mathcal{D})$ . Therefore, we may choose  $\langle \sigma_D, \rho_D \rangle \in \mathcal{V}$  such that  $\mathbb{L}_D := \text{tr}(Q_{\sigma_D, \rho_D}) = \mathbb{L}_A$ . Let  $\langle \sigma', \rho' \rangle$  denote the i/o scheduler generated by  $\langle \sigma_D, \rho_D \rangle$  and  $\langle \sigma_C, \rho_C \rangle$  and write  $\mathbb{L}'$  for  $\text{tr}(Q_{\sigma', \rho'})$ .

Now, let  $I$  denote  $I_{\mathcal{A} \parallel \mathcal{C}} = I_{\mathcal{D} \parallel \mathcal{C}}$  and  $O$  denote  $O_{\mathcal{A} \parallel \mathcal{C}} = O_{\mathcal{D} \parallel \mathcal{C}}$ . Applying Lemma 35, we have for all  $\alpha \in (I \cup O)^{<\omega}$ ,  $\mathbb{L}(\alpha) = \mathbb{L}_A(\pi_A(\alpha)) \cdot \mathbb{L}_C(\pi_C(\alpha))$ . Since  $\mathcal{A}$  and  $\mathcal{D}$  have the same external signature, we know that  $\pi_A(\alpha) = \pi_D(\alpha)$ . Moreover, by the choice of  $\langle \sigma_D, \rho_D \rangle$ , we have  $\mathbb{L}_A = \mathbb{L}_D$ . Hence  $\mathbb{L}_A(\pi_A(\alpha)) = \mathbb{L}_D(\pi_D(\alpha))$ .

Applying Lemma 35 again, we have

$$\mathbb{L}(\alpha) = \mathbb{L}_A(\pi_A(\alpha)) \cdot \mathbb{L}_C(\pi_C(\alpha)) = \mathbb{L}_D(\pi_D(\alpha)) \cdot \mathbb{L}_C(\pi_C(\alpha)) = \mathbb{L}'(\alpha).$$

This proves that  $\mathbb{L} = \mathbb{L}' \in \text{ExtBeh}(\mathcal{D} \parallel \mathcal{C})$ .  $\square$

## 8 Centralized Scheduling with Arbiters

Our switched PIOA framework implements a distributed scheduling scheme: components rely on a token structure to avoid conflicts and scheduling decisions are always made by the (unique) active component. Some may argue that such a scheduling scheme does not realistically represent situations such as asynchronous message passing via an unpredictable network. In response, we outline a setting in which a designated component takes on the role of an *arbiter*, which is responsible for all global scheduling decisions in the system. In other words, we use our switched PIOA framework to recreate a centralized interpretation of component scheduling. The obvious advantage is that our external behavior semantics is compositional and hence we can freely replace components with others that are behaviorally equivalent.

First, we fix a nonempty, finite index set  $\mathcal{I}$  and assume that the universal set  $CAct$  of control actions is  $\bigcup_{i \in \mathcal{I}} \{\mathbf{go}_i, \mathbf{done}_i\}$ . We restrict our attention to controllable automata, defined as follows.

**Definition 36** *Let  $A$  be a switched PIOA and let  $i \in \mathcal{I}$  be given. We say that  $A$  is controllable for  $i$  provided:*

- (1)  $A$  is initially inactive;
- (2)  $CI_A = \{\mathbf{go}_i\}$  and  $CO_A = \{\mathbf{done}_i\}$ .

In other words,  $A$  has a limited control interface,  $\{\mathbf{go}_i, \mathbf{done}_i\}$ , and must wait for an activation signal at the beginning of each execution. Aside from these restrictions,  $A$  is free to communicate with other components (not necessarily the arbiter) via synchronization of basic actions.

Various requirements can be placed on the i/o schedulers for  $A$ . For example, we may require that  $A$  performs at most one locally controlled action during each activation. Or  $A$  may take a finite number of internal steps, possibly followed by a visible action. The latter can be seen as a *fairness* condition, so that no one component is allowed to retain the activity token indefinitely.

To compose a set of (pairwise compatible) controllable automata, we use an arbiter automaton, which models uncertainties in the parallel environment.

**Definition 37** *Let  $X \subseteq BAct$  be given. An arbiter for  $\langle \mathcal{I}, X \rangle$  is a switched PIOA  $Arb$  satisfying the following:*

- (1)  $I_{Arb} = \{\mathbf{done}_i \mid i \in \mathcal{I}\} \cup X$  and  $O_{Arb} = \{\mathbf{go}_i \mid i \in \mathcal{I}\}$ ;
- (2)  $\mathbf{active}_{Arb}(s_{Arb}^0) = 1$ .

Such an arbiter manages the flow of the activity token among components, so that token exchange does not take place directly between components. This is depicted in Figure 10 below.




Fig. 10. Arbitrated Composition

Different notions of parallel composition can be obtained by varying the choice of local i/o schedulers as well as arbiters. A simple example is the parameterized composition operator (cf. Section 1.2), which can be implemented with

- local i/o schedulers that always return control after one locally controlled move and

- an arbiter that schedules  $\text{go}_A$  with probability  $p$  and  $\text{go}_B$  with probability  $1 - p$ .

More complex examples can be obtained by varying the parameter  $X$  in Definition 37. This determines the observational power of the arbiter, i.e., the amount of information which can be used by the arbiter to make scheduling decisions. Such flexibility can be very useful when we wish to limit scheduling freedom in order to improve performance of algorithms. For example, the *write-oblivious* adversary model of [Cha96] requires that random outcomes cannot be used by adversaries until they are read by at least one process. This can be modeled by simply excluding all write-related actions from the set  $X$ .

## 9 Conclusions and Future Work

We have presented the switched PIOA framework, which is designed for the purpose of modeling and analyzing stochastic systems. This framework accommodates for both non-deterministic and probabilistic choices within components, and the associated notion of parallel composition is based on asynchronous communication under a distributed scheduling scheme. We define a trace-style semantics for this framework and prove it is compositional.

Throughout our development, a main focus is the notion of scheduling, i.e., the mechanism with which non-deterministic choices are eliminated. Since the choices between parallel components are often considered non-deterministic, scheduling directly affects semantic behaviors of composite systems. However, in our experience with the literature, scheduling mechanisms are often just mentioned in passing, without due justification. Therefore, we provide a summary of some common scheduling schemes and try to compare them against our distributed scheduling scheme.

Compared to earlier versions [CLSV04a] and [CLSV04b], the current paper presents several technical improvements. First of all, we introduce a new formulation of PIOAs, applying i/o distinction to reactive and generative system types. Moreover, we have modified some of the defining axioms for switched PIOAs, simplifying the definition of external behavior. Finally, we provide a more flexible mechanism for reasoning with systems with open inputs. In particular, the notions of execution trees and likelihood assignments are directly defined for open components, without reference to closing contexts. This allows us to eliminate some of the cumbersome proofs involving renaming and hiding.

As for future research, we see much potential in the proposal of arbiters and controllable automata. We believe it can serve as a theoretical foundation in

many application areas, including distributed consensus and process coordination. In particular, we would like to explore possibilities in modeling noisy scheduling [Asp00], as well as quantum-based and priority-based scheduling [AM99].

We are also interested in adapting the testing scenario of [SV03] to switched PIOAs. Since our semantics focuses on externally visible behavior, we expect to be able to derive a characterization based on frequencies of external observations.

## References

- [AB04] Y. Aumann and M.A. Bender. Efficient low-contention asynchronous consensus with the value-oblivious adversary scheduler. *Distributed Computing*, 2004. Accepted in 2004.
- [Agg94] S. Aggarwal. Time optimal self-stabilizing spanning tree algorithms. Master's thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, May 1994. Available as Technical Report MIT/LCS/TR-632.
- [AM99] J. H. Anderson and M. Moir. Wait-free synchronization in multi-programmed systems: integrating priority-based and quantum-based scheduling. In *Proceedings of the Eighteenth Annual ACM Symposium on Principles of Distributed Computing*, pages 123–132, 1999.
- [Asp00] J. Aspnes. Fast deterministic consensus in a noisy environment. In *Proceedings of the Nineteenth Annual ACM Symposium on Principles of Distributed Computing*, pages 299–309, 2000.
- [Asp03] J. Aspnes. Randomized protocols for asynchronous consensus. *Distributed Computing*, 16(2-3):165–175, 2003.
- [BDEP02] R. Blute, J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for labeled markov processes. *Information and Computaton*, 179(2):163–193, 2002. Special Issue LICS'97.
- [BPW04] M. Backes, B. Pfitzmann, and M. Waidner. Secure asynchronous reactive systems. Cryptology ePrint Archive Report 2004/082, 2004.
- [Can01] R. Canetti. Universally composable security: a new paradigm for cryptographic protocols. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computing*, pages 136–145, 2001.
- [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 *Proceedings of the 15th Annual ACM Symposium on Principles of Distributed Computing*, pages 166–175, 1996.
- [CLSV04a] L. Cheung, N.A. Lynch, R. Segala, and F.W. Vaandrager. Switched probabilistic i/o automata. In *Proceedings First International Colloquium on Theoretical Aspects of Computing (ICTAC2004)*, Lecture Notes in Computer Science. Springer-Verlag, 2004.
- [CLSV04b] L. Cheung, N.A. Lynch, R. Segala, and F.W. Vaandrager. Switched probabilistic i/o automata. Technical Report NIII-R0437, University of Nijmegen, September 2004.
- [dAHJ01] L. de Alfaro, T.A. Henzinger, and R. Jhala. Compositional methods for probabilistic systems. In K.G. Larsen and M. Nielsen, editors, *Proceedings CONCUR 01*, volume 2154 of *Lecture Notes in Computer Science*, pages 351–365. Springer-Verlag, 2001.
- [DGRV00] M.C.A. Devillers, W.O.D. Griffioen, J.M.T. Romijn, and F.W. Vaandrager. Verification of a leader election protocol - formal methods applied to ieee 1394. *Formal Methods in System Design*, 16(3), 2000.
- [DHK98] P. D’Argenio, H. Hermanns, and J.-P. Katoen. On generative parallel composition. In *Proceedings PROBMIV’98*, volume 22 of *ENTCS*, pages 105–122, 1998.
- [Hav98] B.R. Haverkort. *Performance of Computer Communication Systems*. John Wiley & Sons, Ltd, 1998.
- [JLY01] B. Josson, K.G. Larsen, and W. Yi. *Handbook of Process Algebras*, chapter Probabilistic extensions of process algebras. Elsevier, 2001.
- [KS76] J.G. Kennedy and J.L. Snell. *Finite Markov Chains*. Springer-Verlag, New York, 1976.
- [LS91] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 91:1–28, 1991.
- [LSS94] N.A. Lynch, I. Saias, and R. Segala. Proving time bounds for randomized distributed algorithms. In *Proceedings of the 13th Annual 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 R. Amadio and D. Lugiez, editors, *Proceedings 14th International Conference on Concurrency Theory (CONCUR 2003)*, volume 2761 of *Lecture Notes in Computer Science*, 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.

- [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.
- [SAGG<sup>+</sup>93] J. Søgaard-Andersen, S. Garland, J. Guttag, N. Lynch, and A. Pogosyants. Computer-assisted simulation proofs. In *Proceedings of the 4th Conference on Computer Aided Verification, CAV'93*, pages 305–319, 1993.
- [SdV04] A. Sokolova and E.P. de Vink. Probabilistic automata: system types, parallel composition and comparison. In *Validation of Stochastic Systems*, volume 2925 of *Lecture Notes in Computer Science*, 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.
- [Ste94] W.J. Stewart. *Introduction to the Numerical Solutions of Markov Chains*. Princeton University Press, 1994.
- [SV99] M.I.A. Stoelinga and F.W. Vaandrager. Root contention in iee 1394. In J.-P. Katoen, editor, *Proceedings 5th International AMAST Workshop on Formal Methods for Real-Time and Probabilistic Systems*, volume 1601 of *Lecture Notes in Computer Science*, pages 53–74. Springer-Verlag, 1999.
- [SV03] M.I.A. Stoelinga and F.W. Vaandrager. A testing scenario for probabilistic automata. In *Proceedings 30 ICALP*, volume 2719 of *Lecture Notes in Computer Science*, pages 407–418. Springer-Verlag, 2003.
- [vGSS95] R.J. van Glabbeek, S.A. Smolka, and B. Steffen. Reactive, generative, and stratified models of probabilistic processes. *Information and Computation*, 121:59–80, 1995.
- [WSS94] S.-H. Wu, S.A. Smolka, and E.W. Stark. Composition and behaviors of probabilistic i/o automata. In B. Jonsson and J. Parrow, editors, *Proceedings CONCUR 94*, volume 836 of *Lecture Notes in Computer Science*, pages 513–528. Springer-Verlag, 1994.
- [YL92] W. Yi and K.G. Larsen. Testing probabilistic and nondeterministic processes. *Testing and Verification*, 12:47–61, 1992. Protocol Specification.