Compositional Analysis of Expected Delays in Networks of Probabilistic I/O Automata*

Eugene W. Stark, Scott A. Smolka
Department of Computer Science
State University of New York at Stony Brook
Stony Brook, NY 11794 USA

December 5, 1997

Abstract

In previous work, we defined a notion of *probabilistic I/O automata* (PIOA), and we showed that certain functions, which we called "probabilistic behavior maps," constitute a compositional semantics for PIOAs that is fully abstract with respect to a notion of testing equivalence. We also observed that information about *completion probability* and *expected completion time* for a "closed PIOA" can be extracted from its behavior map.

In the present paper, we greatly extend and refine our previous results, thereby obtaining a practical method for computing completion probabilities and expected completion times. Our method is *compositional*, in the sense that it can be applied to a system of PIOAs one component at a time, without ever calculating the global state space of the system. The method is based on symbolic calculations with vectors and matrices of rational functions, and it draws upon a theory of *observables*, which are mappings from *delayed traces* to real numbers that generalize "formal power series" from algebra and combinatorics. We define *rational observables* to be those satisfying certain conditions, among which is the condition that a certain vector space of "derivatives" be finite-dimensional. Central to the theory is a notion of *representation* for an observable, which generalizes the notion "linear representation" for formal power series. We prove that the representable observables coincide with the rational ones; this generalizes to observables a result of Carlyle and Paz equating the recognizable series with those whose "syntactic right ideal" has finite codimension. We also present a minimization algorithm for representations of observables that generalizes a result of Schützenberger for formal power series. The minimization algorithm is applied in our analysis method to limit combinatorial explosion that would otherwise occur.

*Research supported in part by AFOSR Grant F49620-96-1-0087. E-mail addresses: `stark@cs.sunysb.edu` `sas@cs.sunysb.edu`,

1 Introduction

In our previous paper [WSS97], we defined the class of *probabilistic I/O automata* (PIOA), which are a model for distributed or concurrent systems that incorporates a notion of probabilistic choice. The basic intuition underlying the model is the following: the time a PIOA spends in a state before performing its next action is described by an exponentially distributed random variable whose parameter (the so-called *delay parameter*) depends on the state. Under an independence assumption, a simple *composition rule* can be given for producing, given a collection of interacting PIOAs, a single "composite" PIOA representing the entire system.

We also showed how to associate with a PIOA a *probabilistic behavior map*, which in a sense represents the externally observable aspects of the behavior of the PIOA. We showed that behavior map semantics is compositional, in the sense that the behavior map associated with a composite PIOA is uniquely determined by the behavior maps associated with the components. We further showed that, for PIOAs satisfying a certain "delay restriction" concerning their internal actions, behavior map semantics is also fully abstract with respect to a behavioral equivalence based on a notion of probabilistic testing.

As a byproduct of the way of way probability is represented in the PIOA model, it is meaningful to consider certain aspects of timing for PIOA executions. In [Wu96], it is noted that the expected time for a PIOA to complete a specified finite sequence of actions (called a *trace*) can be extracted from the probabilistic behavior map associated with that automaton, and this idea was applied there to analyze some examples.

Certain limitations inherent in our previous work restricted its applicability as a method for analyzing expected completion times in a practical setting. A major problem was that our theory only supported "one trace at a time" analysis: given a PIOA A and a finite trace, the expected time for A to complete an execution having that trace could be determined, but the theory did not provide any useful method by which to specify an infinite set of traces and to determine the expected time for A to complete some execution having one of the traces in that set. The latter problem, rather than the former, is the type of timing analysis that is more often encountered in practice. Another problem was that timing analysis could not be performed on a system of PIOAs "one component at a time"; essentially, a full description of the global state space system had to be constructed and the timing information extracted from that. Any "non-compositional" analysis method that requires the construction of the global state space of a system will in general only be able to handle very small systems, due to the exponential growth of the state space as the number of components increases.

In this paper, we present a new theory and associated analysis methods that overcome the limitations inherent in our previous work. One important part of our new theory is a revised definition of probabilistic behavior map which does not have the "trace at a time" limitation of our previous version. Our new definition makes use of a new notion of *delayed trace*, which generalizes to PIOAs the standard notion of the trace of an execution of an automaton, so that certain probabilistic scheduling information is represented along with the sequence of actions. An *observable* is defined to be a function from delayed traces to real numbers. The *behavior* of a PIOA is defined to be a *transformation of observables*; that is, a mapping from observables to observables. Our revised definition of PIOA behavior admits a much simpler compositionality result (Theorem 1) than the previous version. In particular, we show that the behavior of the composition of "compatible" probabilistic I/O automata is given by the ordinary function composition of the corresponding behaviors.

We show (Lemma 3) that information about *completion probability* and *expected completion time* for a "closed" PIOA A can be obtained by applying its "empty alphabet behavior" $\mathcal{B}_{\emptyset}^{A}$ to appropriate observables. In particular, given a set T of finite action sequences, pairwise incomparable with respect to the prefix relation, one can define an observable Π_T, such that the value of $\mathcal{B}_{\emptyset}^{A}\Pi_T$ on a delayed trace (0) having no actions, is the probability of the set of executions of A whose delayed traces lie in the upward closure of T with respect to the prefix relation on delayed traces. We also define the "expected completion time" for A with respect to T to be the expected time for A to complete some execution having a delayed trace that "just reaches" the set T, and we show that, for a particular observable Ω_T, this time is given by the value of $\mathcal{B}_{\emptyset}^{A}\Omega_T$ on the delayed trace (0).

The above results suggest a naive approach to computing completion probability and expected completion time for a system of PIOAs: construct the composite PIOA representing the entire system, then evaluate the result of applying the behavior map for that system to the observables Π_T and Ω_T and then to the empty delayed trace (0). There are two problems with this approach: (1) the evaluation of the probabilistic behavior map involves computing the value of a summation over a very large (and potentially infinite) set of executions for the system; (2) the method requires the construction of the full global state space for the system. We can avoid problem (1) by observing that the desired summation can be obtained, without enumerating executions, by solving a system of linear equations that can be constructed from the state space of the automaton. In fact, the resulting method works quite well for small systems. However, the size (number of variables) of the system of equations grows linearly with the number of global states, hence this method will use too much space to be useful for large systems.

It turns out that we can do much better than the naive approach described in the previous paragraph. In particular, we can compute the result of applying the behavior map for a system to a specific observable like Π_T or Ω_T without enumerating executions, by working compositionally, "a component at a time," in such a way that the global state space is never constructed. This method is based on the realization that the observables Π_T and Ω_T can be represented in a certain way by a by a kind of automata, having states in a finite-dimensional vector space over the reals, that execute on delayed traces. We call such observables *representable*. We also show (Theorems 3 and 4) that the class of representable observables is closed under the application of PIOA behaviors, and that the result of applying a PIOA behavior to a representable observable can be effectively computed in terms of a construction on representations. Although this construction is a kind of "product construction," which produces an output representation whose size depends on the product of the size of the input representation and the number of states in the PIOA, we can mitigate the blow-up in size by applying a minimization algorithm to the result. We present a minimization algorithm (Theorem 5) that, given the representation of an observable as input, outputs a representation that in a sense has minimum size over all representations of the same observable.

Our theory of observables and their representations can be seen as a generalization of work by Carlyle and Paz [CP71], Schützenberger [Sch61a, Sch61b], and others (see [BR84] for references), on formal power series and linear representations. In particular, our "observables" generalize "formal power series," our "representations" generalize the "linear representations" for formal power series, and our "representable observables" generalize "recognizable series." We define a class of *rational observables*, which are those for which an associated space of *derivatives* is a finite-dimensional vector space, and we show (Theorem 2) that an observable is rational if and only if it is representable. This in a sense generalizes to observables a result of Carlyle and Paz [CP71], which equates the recognizable series with those whose "syntactic right ideal" has finite codimension. Our minimization algorithm for representations of observables corresponds to a result of Schützenberger [Sch61a, Sch61b] for formal power series. The novel aspects of our work are: (1) the introduction of "delayed traces" as a generalization of "words over a finite alphabet", and "observables" as a generalization of "formal power series"; (2) the recognition that "transformations of observables" yield a compositional semantics for PIOAs that is expressive enough to permit the treatment of expected termination time; (3) extension of the theory of "linear representations of formal power series" to a theory of "representable observables"; and (4) use of the theory of representable observables as a basis for deriving compositional algorithms for the analysis of PIOAs. Though closed PIOA's are examples of continuous-time semi-Markov processes [**?**], and as such have a variety of well-developed analysis techniques applicable to them, we are not aware of such techniques that do not have as a prerequisite the construction of a global system description such as a transition matrix or flowgraph.

In other related work, Campos et al. in [CCM97] present BDD-based algorithms that determine the exact bounds on the delay between two specified events and the number of occurrences of another event in all such intervals. Segala et al. [LSS94, PS95] have developed a method for the analysis of the expected time complexity of randomized distributed algorithms. The method consists of proving auxiliary probabilistic time bound statements of the form $U \text{---}\{t,p\} \rightarrow U'$, which mean that whenever the algorithm begins in a state in a set U, it will reach a state in set U' within time t with probability at least p. Finally, a number of "stochastically timed" process algebras and Petri net formalisms have been proposed for the performance analysis of concurrent systems, including [MBC84, GHR93, Hil96, Pri96, BDG98]. In the case of process algebra, these approaches are sometimes referred to as "compositional" in the sense that a composite stochastic system can be specified algebraically in terms of its components.

The remainder of this paper is organized as follows: Section 1 is devoted to the basic definitions and theory of PIOAs and probabilistic behavior maps. Section 2 treats rational observables and their representations. Section 3 considers a simple example of the use of the techniques.

2 Probabilistic I/O Automata and Their Behaviors

2.1 Probabilistic I/O Automata

In this section, we recall the basic definitions from [WSS97], to which the reader is referred for additional details and discussion. We give here simplified versions of the definitions, which are equivalent to those of [WSS97] in the case of *finite* PIOAs, which are all that we consider in the present paper.

A *finite probabilistic I/O automaton* is a quadruple $A = (Q, q^{\mathrm{I}}, E, \Delta, \mu, \delta)$, where

- Q is a finite set of *states*.
- $q^{\mathrm{I}} \in Q$ is a distinguished *start state*.
- E is a finite set of *actions*, partitioned into disjoint sets of *input*, *output*, and *internal* actions, which are denoted by E^{in}, E^{out}, and E^{int}, respectively. The set $E^{\mathrm{loc}} = E^{\mathrm{out}} \cup E^{\mathrm{int}}$ of output and internal actions is called the set of *locally controlled* actions, and the set $E^{\mathrm{ext}} = E^{\mathrm{in}} \cup E^{\mathrm{out}}$ is called the set of *external* actions.
- $\Delta \subseteq Q \times E \times Q$ is the *transition relation*, which satisfies the following *input-enabledness* condition: for any state $q \in Q$ and input action $e \in E^{\mathrm{in}}$, there exists a state $r \in Q$ such that $(q, e, r) \in \Delta$.
- $\mu : (Q \times E \times Q) \to [0, 1]$ is the *transition probability* function, which is required to satisfy the following conditions:
  1. $\mu(q, e, r) > 0$ iff $(q, e, r) \in \Delta$.
  2. $\sum_{r \in Q} \mu(q, e, r) = 1$, for all $q \in Q$ and all $e \in E^{\mathrm{in}}$.
  3. For all $q \in Q$, if there exist $e \in E^{\mathrm{loc}}$ and $r \in Q$ such that $(q, e, r) \in \Delta$, then $\sum_{r \in Q} \sum_{e \in E^{\mathrm{loc}}} \mu(q, e, r) = 1$,
- $\delta : Q \to [0, \infty)$ is the *state delay* function, which is required to satisfy the following condition: for all $q \in Q$, we have $\delta(q) > 0$ if and only if there exist $e \in E^{\mathrm{loc}}$ and $r \in Q$ such that $(q, e, r) \in \Delta$.

A *finite execution fragment* for a probabilistic I/O automaton A is an alternating sequence of states and actions of the form

$$q_0 \xrightarrow{e_0} q_1 \xrightarrow{e_1} \ldots \xrightarrow{e_{n-1}} q_n,$$

such that for each k with $0 \leq k < n$, one of the following two conditions holds:

1. $e_k \in E$ and $(q_k, e_k, q_{k+1}) \in \Delta$.

2. $e_k \notin E$ and $q_{k+1} = q_k$.

An execution fragment with $q_0 = q^{\mathrm{I}}$ (the distinguished start state) is called an *execution.*

In case (1) above, we say that action e_k is a *native action* of A, and that the triple (q_k, e_k, q_{k+1}) is a *native transition* of A. In case (2), we say that e_k is a *non-native action* of A and that (q_k, e_k, q_{k+1}) is a *non-native transition* of A. We often use the notation $q_k \xrightarrow{e_k} q_{k+1}$ to assert the disjunction of conditions (1) and (2) above. We adopt a convention whereby μ can be applied to triples (q, e, r), where e is a non-native action of A, by defining $\mu(q, e, q) = 1$ and $\mu(q, e, r) = 0$ for all other $r \in Q$. We use the terms *native execution fragment* and *native execution* to refer to an execution fragment or execution of A in which only native actions appear.

If σ denotes an execution fragment as above, then we will use $\sigma(k)$ to denote the state q_k, for $0 \leq k \leq n$, and we will use $\sigma(k, k+1)$ to denote the action e_k, for $0 \leq k < n$. We use the term *trace* to refer to a sequence of actions. If σ is an execution fragment as above, then the *trace of* σ, denoted $\mathrm{tr}(\sigma)$, is the sequence of actions $e_0 e_1 \ldots e_{n-1}$ appearing in σ.

2.2 Composition

A finite collection $\{A_i : i \in I\}$ of probabilistic I/O automata, where $A_i = (Q_i, q_i^{\mathrm{I}}, E_i, \Delta_i, \mu_i, \delta_i)$, is called *compatible* if for all $i, j \in I$, $i \neq j$, we have $E_i^{\mathrm{out}} \cap E_j^{\mathrm{out}} = \emptyset$ and $E_i^{\mathrm{int}} \cap E_j = \emptyset$. We define the *composition* $\|_{i \in I} A_i$ of a finite compatible collection to be the probabilistic I/O automaton $(Q, q^{\mathrm{I}}, E, \Delta, \mu, \delta)$, defined as follows:

- $Q = \|_{i \in I} Q_i$.
- $q^{\mathrm{I}} = \langle q_i^{\mathrm{I}} : i \in I \rangle$.
- $E = \bigcup_{i \in I} E_i$, where
$$E^{\mathrm{out}} = \bigcup_{i \in I} E_i^{\mathrm{out}} \qquad E^{\mathrm{int}} = \bigcup_{i \in I} E_i^{\mathrm{int}} \qquad E^{\mathrm{in}} = (\bigcup_{i \in I} E_i^{\mathrm{in}}) \setminus E^{\mathrm{out}}.$$
- Δ is the set of all $(\langle q_i : i \in I \rangle, e, \langle r_i : i \in I \rangle)$ such that for all $i \in I$, if $e \in E_i$, then $(q_i, e, r_i) \in \Delta_i$, otherwise $r_i = q_i$.
- $\delta(\langle q_i : i \in I \rangle) = \sum_{i \in I} \delta_i(q_i)$.
- If $e \in E^{\mathrm{in}}$, then
$$\mu(\langle q_i : i \in I \rangle, e, \langle r_i : i \in I \rangle) = \prod_{i \in I} \mu_i(q_i, e, r_i).$$
If $e \in E_k^{\mathrm{loc}}$ for some k, then
$$\mu(\langle q_i : i \in I \rangle, e, \langle r_i : i \in I \rangle) = \frac{\delta_k(q_k)}{\sum_{i \in I} \delta_i(q_i)} \prod_{i \in I} \mu_i(q_i, e, r_i).$$

We use the notation $A \| B$ to denote the composition $\|\{A, B\}$ of a compatible 2-element set of PIOAs.

As discussed in [WSS97], the definitions of μ and δ above reflect the intuition we wish to capture concerning the execution of a system of PIOAs. Upon arrival in a state, a component PIOA chooses randomly the length of time it will spend in that state before executing its next "locally controlled" transition. The random choice is made according to an exponential distribution described by the delay parameter associated with that state, and it is made independently of the other PIOAs in the system. The definitions of μ and δ for the composite system express the idea that the various component PIOAs are in a race to see which of them will execute the next locally controlled action. This competition will be won by the component that has chosen the smallest delay time, and the probability that any given component will win the competition is given by the ratio of the local delay parameter for that component over the sum of the local delay parameters for all components. The time the system remains in a particular global state before executing the next locally controlled action is the minimum of the times that each component spends in its respective local state. This time is governed by an exponential distribution, whose parameter is the sum of the parameters of the distributions for each of the components.

2.3 Probability Space

In [WSS97], we showed how a closed PIOA A (one with no input actions) induces a probability space over the set of all its executions. If σ is a finite execution of A, of the form:

$$q_0 \xrightarrow{e_0} q_1 \xrightarrow{e_1} \dots \xrightarrow{e_{n-1}} q_n,$$

then let $[\sigma]$ denote the set of all executions of A having σ as a prefix. The measurable sets of executions are those generated by declaring each set $[\sigma]$ to be measurable and closing up under countable unions and complement. The probability measure pr_A is the extension, to the full class of measurable sets, of the mapping that assigns to each set of the form $[\sigma]$ the quantity

$$\mathrm{p}_A(\sigma) = \prod_{k=0}^{n-1} \mu(q_k, e_k, q_{k+1}).$$

Note the difference between pr_A, which is a countably additive set function defined on the class of measurable sets of executions, and p_A, which is a function from finite executions to real numbers. In this paper, we shall also be interested in the related function $\mathrm{w}_A(\sigma)$ on finite executions, defined by:

$$\mathrm{w}_A(\sigma) = \left(\prod_{\{k : e_k \in E_A^{\mathrm{loc}}\}} \delta_A(q_k) \right) \mathrm{p}_A(\sigma).$$

We call $\mathrm{w}_A(\sigma)$ the *weight* of the execution σ. Although the definition of w_A may at first seem somewhat *ad hoc*, it turns out that w_A behaves in a more convenient fashion than $\mathrm{p}_A(\sigma)$

when considering the composition of PIOAs. In particular, w_A has the useful property stated in Lemma 1 below.

2.4 Delayed Traces, Observables, and Behaviors

Let E be a set of actions. A (finite) *delayed trace* α *over* E consists of an alternating sequence of the form:

$$d_0 \xrightarrow{e_0} d_1 \xrightarrow{e_1} \dots \xrightarrow{e_{n-1}} d_n,$$

where the d_k are nonnegative real numbers and the the e_k are actions in E. The sequence $e_0, e_1, \dots e_{n-1}$ is called the *trace of* α, and we sometimes denote it by $\mathrm{tr}(\alpha)$. The sequence $d_0, d_1, \dots, d_n$ is called the *sequence of delay parameters* of α. We often use the notation $\alpha(k)$ to denote d_k, and the notation $\alpha(k, k+1)$ to denote e_k. The number n is called the *length* of α, and we denote it by $|\alpha|$.

We use $\mathrm{DTraces}(E)$ to denote the set of all delayed traces over E. We also use the notation $(d)_E$, or just d, when E is clear from the context, to denote the *empty delayed trace* in $\mathrm{DTraces}(E)$, consisting of the single delay parameter d and no actions.

Suppose $\alpha \in \mathrm{DTraces}(E)$. If $E \subseteq E'$, then a delayed trace $\alpha' \in \mathrm{DTraces}(E')$ is a *refinement* of α, and we write $\alpha' \triangleright \alpha$, if there exists a monotone injection

$$\phi : \{j : 0 \leq j \leq |\alpha|\} \to \{k : 0 \leq k \leq |\alpha'|\}$$

such that $\phi(0) = 0$ and such that the following conditions hold:

1. $\alpha'(\phi(j) - 1, \phi(j)) = \alpha(j-1, j)$ for $0 < j \leq |\alpha|$.
2. $\alpha'(k-1, k) \in E' \setminus E$ for all k outside the image of ϕ.
3. $\alpha'(k) = \begin{cases} \alpha(j), & \text{for } 0 \leq j < |\alpha| \text{ and } \phi(j) \leq k < \phi(j+1), \\ \alpha(|\alpha|), & \text{for } \phi(|\alpha|) \leq k \leq |\alpha'|. \end{cases}$

Figure 1 (a) depicts graphically the refinement relationship between α' and α.

Suppose A is a PIOA. If α is a delayed trace over E, then an execution σ of A is *conformant with* α, and we write $\sigma \propto \alpha$, if there exists a monotone injection

$$\phi : \{j : 0 \leq j \leq |\alpha|\} \to \{k : 0 \leq k \leq |\sigma|\}$$

such that $\phi(0) = 0$ and such that the following conditions hold:

1. $\sigma(\phi(j) - 1, \phi(j)) = \alpha(j-1, j)$ for $0 < j \leq |\alpha|$.
2. $\sigma(k-1, k) \in E_A \setminus E$, for all k outside the image of ϕ.

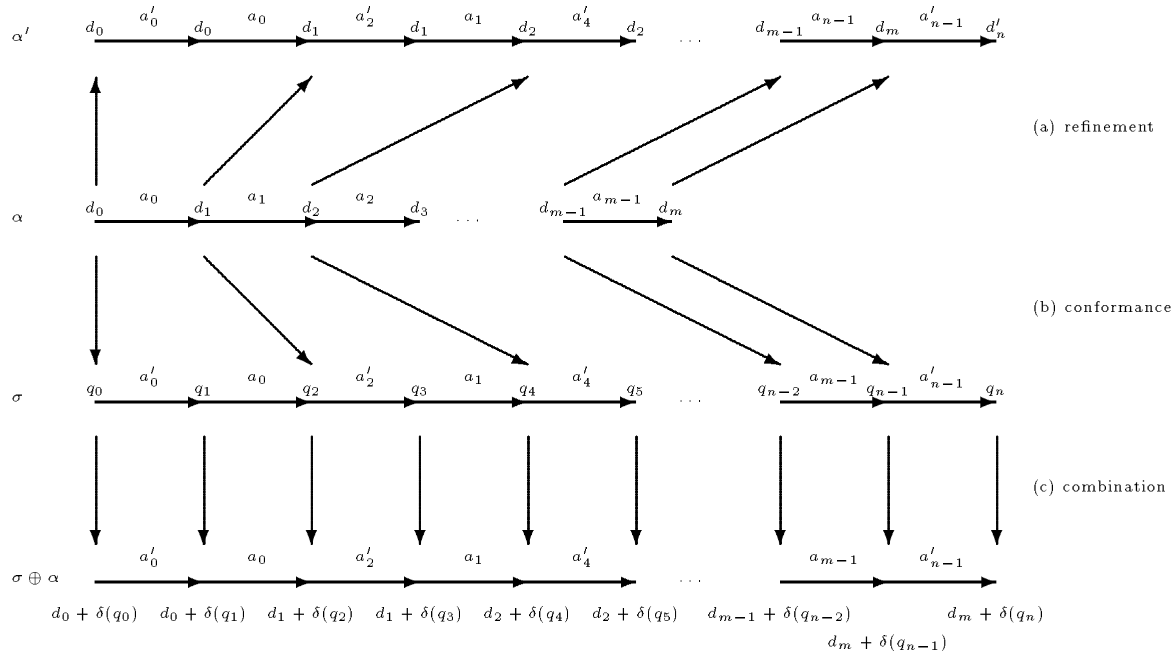

Figure 1: Refinement, Conformance, and Combination

It is easy to check that if $\sigma \propto \alpha$, then there is exactly one monotone injection ϕ that satisfies these conditions. Figure 1 (b) depicts graphically the conformance relationship between σ and α.

Suppose $\alpha \in \mathrm{DTraces}(E)$ and $\sigma \propto \alpha$, with ϕ the corresponding monotone injection. Then the *combination* of σ and α is the delayed trace $\sigma \oplus \alpha \in \mathrm{DTraces}(E \cup E_A)$ of the form:

$$d_0 \xrightarrow{e_0} d_1 \xrightarrow{e_1} \dots \xrightarrow{e_{n-1}} d_n,$$

where

1. $e_k = \sigma(k, k+1)$ for $0 \leq k < |\sigma|$.
2. $d_k = \begin{cases} \delta_A(\sigma(k)) + \alpha(j), & \text{for } 0 \leq j < |\alpha| \text{ and } \phi(j) \leq k < \phi(j+1), \\ \delta_A(\sigma(k)) + \alpha(|\alpha|), & \text{for } \phi(|\alpha|) \leq k \leq |\sigma|. \end{cases}$

Figure 1 (c) depicts graphically the result of combining σ and α.

An *observable* over a set of actions E is a mapping:

$$\Phi : \mathrm{DTraces}(E) \to \mathcal{R}.$$

If A is a PIOA and E is a set of actions, then the *E-behavior* of A is the *transformation of observables*:

$$\mathcal{B}_E^A : (\mathrm{DTraces}(E \cup E_A) \to \mathcal{R}) \to (\mathrm{DTraces}(E) \to \mathcal{R})$$

defined by:

$$\mathcal{B}_E^A \Phi\alpha = \sum_{\sigma \propto \alpha} \Phi(\sigma \oplus \alpha) \mathrm{w}_A(\sigma).$$

In general, $\mathcal{B}_E^A \Phi\alpha$ will not be defined for all Φ and α, because the defining summation above need not converge.

2.5 Compositionality

We now prove a compositionality result that shows how the behavior $\mathcal{B}_E^{A|B}$ for a composite PIOA $A|B$ can be derived from the component behaviors $\mathcal{B}_E^A$ and $\mathcal{B}_{E \cup E_A}^B$ for the PIOAs A and B, respectively.

We first establish a technical lemma.

Lemma 1 *Suppose A and B are compatible PIOAs. Then, given a delayed trace α, the set of all executions σ of $A|B$ such that $\sigma \propto \alpha$, is in bijective correspondence with the set of all pairs of executions (σ^A, σ^B), where σ^A is an execution of A such that $\sigma^A \propto \alpha$, and σ^B is an execution of B such that $\sigma^B \propto (\sigma^A \oplus \alpha)$. Moreover, whenever σ corresponds under the bijection to the pair (σ^A, σ^B) we have:*

$$\sigma \oplus \alpha = \sigma^B \oplus (\sigma^A \oplus \alpha) \qquad \text{and} \qquad \mathrm{w}_{A|B}(\sigma) = \mathrm{w}_A(\sigma^A)\, \mathrm{w}_B(\sigma^B).$$

Proof – The bijection is given by the mapping that takes an execution σ of $A|B$, of the form:

$$(q_0^A, q_0^B) \xrightarrow{e_0} (q_1^A, q_1^B) \xrightarrow{e_1} \dots \xrightarrow{e_{n-1}} (q_n^A, q_n^B),$$

to the pair (σ^A, σ^B), where σ^B is the following execution of B:

$$q_0^B \xrightarrow{e_0} q_1^B \xrightarrow{e_1} \dots \xrightarrow{e_{n-1}} q_n^B,$$

and σ^A is the execution

$$q_{k_0}^A \xrightarrow{e_{k_0}} q_{k_1}^A \xrightarrow{e_{k_1}} \dots \xrightarrow{e_{k_{m-1}}} q_{k_m}^A \xrightarrow{e_{k_m}} q_n^A,$$

where $k_0 < k_1 < \dots < k_m$ is the sequence of all indices k with $0 \le k < n$ for which either $e_k \in E_A$ or else $e_k \notin E_B$.

To describe the inverse mapping, suppose (σ^A, σ^B) satisfies $\sigma^A \propto \alpha$ and $\sigma^B \propto (\sigma^A \oplus \alpha)$. Let

$$\phi : \{j : 0 \le j \le |\sigma^A \oplus \alpha|\} \to \{k : 0 \le k \le |\sigma^B|\}$$

be the monotone injection that exists due to the relationship $\sigma^B \propto (\sigma^A \oplus \alpha)$, and for each k with $0 \le k < |\sigma^B|$, let j_k be the greatest j such that $\phi(j) \le k$. Then it is easy to see that there is an execution σ of $A|B$ uniquely defined by the conditions $|\sigma| = |\sigma^B|$, $\sigma(k, k+1) = \sigma^B(k, k+1)$ for $0 \le k < |\sigma|$, and

$$\sigma(k) = \begin{cases} (\sigma^A(j), \sigma^B(k)), & \text{for } 0 \le j < |\alpha| \text{ and } \phi(j) \le k < \phi(j+1), \\ (\sigma^A(|\sigma^A|), \sigma^B(k)), & \text{for } \phi(|\sigma^A|) \le k \le |\sigma|, \end{cases}$$

that this execution correponds to the pair (σ^A, σ^B) under the map defined above, and satisfies the relationship:

$$\sigma \oplus \alpha = \sigma^B \oplus (\sigma^A \oplus \alpha).$$

To verify the identity

$$\mathrm{w}_{A|B}(\sigma) = \mathrm{w}_A(\sigma^A)\ \mathrm{w}_B(\sigma^B)$$

for σ corresponding to (σ^A, σ^B), suppose σ has the form

$$(q_0^A, q_0^B) \xrightarrow{e_0} (q_1^A, q_1^B) \xrightarrow{e_1} (q_2^A, q_2^B) \ldots \xrightarrow{e_{n-1}} (q_n^A, q_n^B).$$

We compute as follows, using the definition of $A\|B$:

$$\begin{aligned}
\mathrm{w}_{A\|B}(\sigma) &= \left(\prod_{\{k : e_k \in E^{\mathrm{loc}}_{A\|B}\}} \delta_{A\|B}(q_k^A, q_k^B) \right) \left(\prod_{k=0}^{n-1} \mu_{A\|B}((q_k^A, q_k^B), e_k, (q_{k+1}^A, q_{k+1}^B)) \right) \\
&= \left(\prod_{\{k : e_k \in E^{\mathrm{loc}}_A \cup E^{\mathrm{loc}}_B\}} \delta_A(q_k^A) + \delta_B(q_k^B) \right) \\
&\qquad \left(\prod_{k=0}^{n-1} \mu_A(q_k^A, e_k, q_{k+1}^A) \right) \left(\prod_{k=0}^{n-1} \mu_B(q_k^B, e_k, q_{k+1}^B) \right) \\
&\qquad \left(\prod_{\{k : e_k \in E^{\mathrm{loc}}_A\}} \frac{\delta_A(q_k^A)}{\delta_A(q_k^A) + \delta_B(q_k^B)} \right) \left(\prod_{\{k : e_k \in E^{\mathrm{loc}}_B\}} \frac{\delta_B(q_k^B)}{\delta_A(q_k^A) + \delta_B(q_k^B)} \right) \\
&= \left(\prod_{k=0}^{n-1} \mu_A(q_k^A, e_k, q_{k+1}^A) \right) \left(\prod_{k=0}^{n-1} \mu_B(q_k^B, e_k, q_{k+1}^B) \right) \\
&\qquad \left(\prod_{\{k : e_k \in E^{\mathrm{loc}}_A\}} \delta_A(q_k^A) \right) \left(\prod_{\{k : e_k \in E^{\mathrm{loc}}_B\}} \delta_B(q_k^B) \right) \\
&= \mathrm{w}_A(\sigma^A)\ \mathrm{w}_B(\sigma^B).
\end{aligned}$$

In the last step, we have made use of the fact that

$$\left(\prod_{k=0}^{n-1} \mu_A(q_k^A, e_k, q_{k+1}^A) \right) \left(\prod_{\{k : e_k \in E^{\mathrm{loc}}_A\}} \delta_A(q_k^A) \right) = \mathrm{w}_A(\sigma^A).$$

Even though σ^A, in general, will be shorter than σ^B, the fact that $\sigma^B \propto (\sigma^A \oplus \alpha)$ means that any actions in σ^B that are not also in σ^A will be in $E_B \setminus E_A$. Thus, if e_k is such an action, corresponding to the transition $(q_k^A, q_k^B) \xrightarrow{e_k} (q_{k+1}^A, q_{k+1}^B)$ in σ, then by the defining conditions for σ, the transition $q_k^A \xrightarrow{e_k} q_{k+1}^A$ of A is a nonnative transition of A. Since μ_A by definition has value one on nonnative transitions, it follows that we may drop the e_k term from the first product. Since we cannot have $e_k \in E_A^{\mathrm{loc}}$ if e_k is nonnative for A, the second product does not contain any e_k term. The stated fact follows immediately. ∎

Theorem 1 *Suppose A and B are compatible PIOAs, and E is a set of actions. Then*

$$\mathcal{B}_E^{A\|B} = \mathcal{B}_E^A \circ \mathcal{B}_{E\cup E_A}^B.$$

Proof – We compute, using the definitions of $\mathcal{B}_E^A$, $\mathcal{B}_{E\cup E_A}^B$, and $\mathcal{B}_E^{A\|B}$:

$$\begin{aligned}
(\mathcal{B}_E^A \circ \mathcal{B}_{E\cup E_A}^B)\Phi\alpha &= (\mathcal{B}_E^A(\mathcal{B}_{E\cup E_A}^B\Phi))\alpha \\
&= \sum_{\sigma^A\propto\alpha} (\mathcal{B}_{E\cup E_A}^B\Phi(\sigma^A \oplus \alpha))\ \mathrm{w}_A(\sigma^A) \\
&= \sum_{\sigma^A\propto\alpha} \quad \sum_{\sigma^B\propto(\sigma^A\oplus\alpha)} \Phi(\sigma^B \oplus (\sigma^A \oplus \alpha))\ \mathrm{w}_B(\sigma^B)\ \mathrm{w}_A(\sigma^A) \\
&= \sum_{\sigma\propto\alpha} \Phi(\sigma \oplus \alpha)\ \mathrm{w}_{A\|B}(\sigma),
\end{aligned}$$

where we have used Lemma 1 in the last step to replace the double summation by a single one. But the last expression above is precisely $\mathcal{B}_E^{A\|B}\Phi\alpha$. ∎

The following corollary is worth noting. It expresses a commutativity property of behavior composition that derives from the more obvious commutativity property $A\|B = B\|A$ of PIOA composition.

Corollary 2 *Suppose A and B are compatible PIOAs. Then*

$$\mathcal{B}_E^A \circ \mathcal{B}_{E\cup E_A}^B = \mathcal{B}_E^B \circ \mathcal{B}_{E\cup E_B}^A.$$

The above definitions and compositionality result are a generalization and simplification of those in [WSS97]. For a fixed action sequence u, the definition given in [WSS97] for the probabilistic behavior map $\mathcal{E}_u^A$ was as follows:

$$\mathcal{E}_u^A[g(\mathbf{D})] = \sum_{\mathbf{d}} g(\mathbf{d}) p_u^A(\mathbf{d}),$$

where the index of summation $\mathbf{d}$ ranges over $\mathcal{R}^{|u|+1}$, where $g : \mathcal{R}^{|u|+1} \to \mathcal{R}$, and where $p_u^A(\mathbf{d})$ denotes the summation of $\mathrm{p}_A(\sigma)$ over all executions σ with action sequence u and sequence of delay parameters $\mathbf{d}$. We now recognize that u and $\mathbf{d}$ are best regarded as two attributes of a single, more general entity (a delayed trace), that g should be correspondingly generalized (to an observable Φ), and that $\mathcal{E}^A$ is properly regarded as a transformation $\mathcal{B}_E^A$ of observables. In addition, we see that the correct place for the summing of delay parameters to appear is in the definition of $\mathcal{B}_E^A$, rather than in the compositionality law, and that the ugly "correction factor" $h(\mathbf{d}^A, \mathbf{d}^B)$ appearing in the compositionality law proved in [WSS97] can be made to disappear if the definition of $\mathcal{B}_E^A$ uses the "weight" $\mathrm{w}_A(\sigma)$, rather than the probability $\mathrm{p}_A(\sigma)$.

2.6 Completion Probability and Expected Completion Time

We are interested in calculating the expected time taken for a PIOA A to perform a finite execution having an action sequence in a specified set, which in general will be infinite. To avoid ambiguity surrounding executions that "complete" multiple times in the sense of having more than one prefix with an action sequence lying in the specified set, we restrict our attention to sets of action sequences that are pairwise incomparable with respect to the prefix relation.

Formally, we define a *target set* to be a set T of finite sequences of actions that is pairwise incomparable with respect to the prefix relation. We write $T\uparrow$ to denote the upward-closure of T with respect to prefix.

Suppose $T \subseteq E^*$ is a target set. The *characteristic observable* of T is the map $\chi_T : \mathrm{DTraces}(E) \to \mathcal{R}$ defined as follows:

$$\chi_T(\alpha) = \begin{cases} 1, & \text{if } \mathrm{tr}(\alpha) \in T \\ 0, & \text{otherwise.} \end{cases}$$

Two other observables will be of interest to us. The *probability observable* is the map $\Pi : \mathrm{DTraces}(E) \to \mathcal{R}$ defined by:

$$\Pi(\alpha) = \prod_{k=0}^{|\alpha|-1} \frac{1}{\alpha(k)}.$$

The *completion time observable* is the mapping $\Omega : \mathrm{DTraces}(E) \to \mathcal{R}$ defined by:

$$\Omega(\alpha) = \left(\sum_{k=0}^{|\alpha|-1} \frac{1}{\alpha(k)}\right)\left(\prod_{k=0}^{|\alpha|-1} \frac{1}{\alpha(k)}\right).$$

If T is a target set, then we define:

$$\Pi_T(\alpha) = \Pi(\alpha)\chi_T(\alpha) \qquad\qquad \Omega_T(\alpha) = \Omega(\alpha)\chi_T(\alpha).$$

If A is a closed PIOA, and T is a target set, then the *completion probability* $\mathrm{pr_c}(A, T)$ for A with respect to T is the quantity:

$$\mathrm{pr_c}(A, T) = \mathrm{pr}_A\{\sigma : \sigma \text{ native},\ \mathrm{tr}(\sigma) \in T\uparrow\}.$$

We say that A *almost certainly completes* T, if $\mathrm{pr_c}(A, T) = 1$.

Lemma 3 *Suppose A is a closed PIOA, and T is a target set. Then*

$$\mathrm{pr_c}(A, T) = \mathcal{B}_\emptyset^A \Pi_T\ (0),$$

where (0) denotes the delayed trace with no actions and zero as its sole delay parameter.

Proof – We compute:

$$
\begin{aligned}
\mathcal{B}^A_\emptyset \Pi_T\ (0) &= \sum_{\sigma \propto (0)} \Pi_T(\sigma \oplus 0)\ \mathrm{w}_A(\sigma) \\
&= \sum_{\substack{\sigma\ \propto\ (0) \\ \mathrm{tr}(\sigma)\ \in\ T}} \left(\prod_{k=0}^{|\sigma|-1} \frac{1}{\delta_A(\sigma(k))} \right) \mathrm{w}_A(\sigma) \\
&= \sum_{\substack{\sigma\ \propto\ (0) \\ \mathrm{tr}(\sigma)\ \in\ T}} \mathrm{p}_A(\sigma) \\
&= \mathrm{pr}_A\{\sigma : \sigma \text{ native},\ \mathrm{tr}(\sigma) \in T\uparrow\} \\
&= \mathrm{pr}_\mathrm{c}(A, T),
\end{aligned}
$$

where in the first line we have used the definition of $\mathcal{B}^A$, in the second line we have used the definition of Π_T and the fact that if $\sigma \propto (0)$, then $\mathrm{tr}(\sigma \oplus (0)) = \mathrm{tr}(\sigma)$, and in the third line we have used the fact that, for A closed, if $\sigma \propto (0)$, then every action in σ is in E^{loc}, hence

$$
\left(\prod_{k=0}^{|\sigma|-1} \frac{1}{\delta_A(\sigma(k))} \right) \mathrm{w}_A(\sigma) = \mathrm{p}_A(\sigma),
$$

and in the fourth line we have used the definition of the probability measure pr_A and the fact that if $\sigma \propto (0)$, then σ is a native execution of A. ∎

If A is a closed PIOA, and T is a target set such that A almost certainly completes T, then the *expected completion time* $\exp_\mathrm{c}(A, T)$ for A with respect to T is the quantity:

$$
\exp_\mathrm{c}(A, T) = \sum_{\substack{\sigma \text{ native} \\ \mathrm{tr}(\sigma)\ \in\ T}} \left(\sum_{k=0}^{|\sigma|-1} \frac{1}{\delta_A(\sigma(k))} \right) \mathrm{p}_A(\sigma).
$$

Lemma 4 *Suppose A is a closed PIOA, and T is a target set. Then*

$$
\exp_\mathrm{c}(A, T) = \mathcal{B}^A_\emptyset \Omega_T\ (0).
$$

Proof – We compute:

$$
\begin{aligned}
\mathcal{B}^A_\emptyset \Omega_T\ (0) &= \sum_{\sigma \propto (0)} \Omega_T(\sigma \oplus (0))\ \mathrm{w}_A(\sigma) \\
&= \sum_{\substack{\sigma \text{ native} \\ \mathrm{tr}(\sigma)\ \in\ T}} \left(\sum_{k=0}^{|\sigma|-1} \frac{1}{\delta_A(\sigma(k))} \right) \left(\prod_{k=0}^{|\sigma|-1} \frac{1}{\delta_A(\sigma(k))} \right) \mathrm{w}_A(\sigma)
\end{aligned}
$$

$$= \sum_{\substack{\sigma \text{ native} \\ \mathrm{tr}(\sigma) \in T}} \left(\sum_{k=0}^{|\sigma|-1} \frac{1}{\delta_A(\sigma(k))} \right) \mathrm{p}_A(\sigma)$$

$$= \exp_c(A, T),$$

using similar reasoning to that of Lemma 3. ∎

3 Computing Expected Completion Time Compositionally

In this section, we develop the theory of "rational observables," and show how this theory, together with that of the previous section, can be used to obtain a compositional method for computing expected completion time.

3.1 Rational Observables and Their Representations

Let $\mathrm{Obs}(E)$ denote the set of all observables $\Phi : \mathrm{DTraces}(E) \to \mathcal{R}$. Then $\mathrm{Obs}(E)$ is a a vector space under the usual pointwise addition and scalar multiplication:

$$(\Phi + \Phi')(\alpha) = \Phi(\alpha) + \Phi'(\alpha). \qquad (a\Phi)(\alpha) = a(\Phi(\alpha))$$

Suppose $\Phi : \mathrm{DTraces}(E) \to \mathcal{R}$ is an observable. If $d \in \mathcal{R}$ and $a \in E$, then the *derivative* of Φ by d and a is the observable $\Phi_{d \xrightarrow{a}}$ defined by:

$$\Phi_{d \xrightarrow{a}}(\alpha) = \Phi(d \xrightarrow{a} \alpha),$$

where if α is the delayed trace:

$$d_0 \xrightarrow{e_0} d_1 \xrightarrow{e_1} \dots \xrightarrow{e_{n-1}} d_n,$$

then $d \xrightarrow{a} \alpha$ denotes the delayed trace:

$$d \xrightarrow{a} d_0 \xrightarrow{e_0} d_1 \xrightarrow{e_1} \dots \xrightarrow{e_{n-1}} d_n.$$

Lemma 5 *For all $d \in \mathcal{R}$ and $a \in E$, the mapping taking $\Phi \in \mathrm{Obs}(E)$ to its derivative $\Phi_{d \xrightarrow{a}} \in \mathrm{Obs}(E)$ is a linear transformation on $\mathrm{Obs}(E)$.*

Proof – Simply observe that, for a fixed $d \in \mathcal{R}$ and $a \in E$, for all Φ and Φ' in $\mathrm{Obs}(E)$, and c and c' in $\mathcal{R}$, we have:

$$\begin{aligned} (c\Phi + c'\Phi')_{d \xrightarrow{a}}(\alpha) &= (c\Phi + c'\Phi')(d \xrightarrow{a} \alpha) \\ &= c\Phi(d \xrightarrow{a} \alpha) + c'\Phi'(d \xrightarrow{a} \alpha) \\ &= c\Phi_{d \xrightarrow{a}}(\alpha) + c'\Phi'_{d \xrightarrow{a}}(\alpha). \end{aligned}$$

∎

If S is an arbitrary subset of $\mathrm{Obs}(E)$, then define

$$\mathcal{D}S = \{\Phi_{d \xrightarrow{a}} : \Phi \in S, d \in \mathcal{R}, a \in E\},$$

and let $\mathcal{D}^*S$ denote the least subspace of $\mathrm{Obs}(E)$ containing S and satisfying $\mathcal{D}(\mathcal{D}^*S) \subseteq \mathcal{D}^*S$.

Define an observable $\Phi \in \mathrm{Obs}(E)$ to be *rational* if the following three conditions hold:

1. The space $\mathcal{D}^*\Phi$ is a finite dimensional subspace of $\mathrm{Obs}(E)$.
2. For all $\Psi \in \mathcal{D}^*\Phi$, the quantity $\Psi(d)$ (the value of Ψ on the delayed trace of length zero with single delay parameter d) is a rational function of d.
3. For all $\Psi \in \mathcal{D}^*\Phi$, all $a \in E$, and all linear maps $L : \mathcal{D}^*\Phi \to \mathcal{R}$, the quantity $\Psi_{d \xrightarrow{a}} L$ is a rational function of d (note that we denote the application of a linear transformation by writing it to the right of its argument).

Define the *dimension* of a rational observable Φ to be the dimension of $\mathcal{D}^*\Phi$.

Lemma 6 *If $\Phi \in \mathrm{Obs}(E)$ is rational, then every element of $\mathcal{D}^*\Phi$ is rational.*

Proof – By construction, every element of $\mathcal{D}^*\Phi$ can be expressed as a (finite) linear combination of derivatives of Φ. It is easy to check that a linear combination of rational observables is again rational, so it remains to be shown that if Ψ is a rational observable, then $\Psi_{d \xrightarrow{a}}$ is also rational, for all $d \in \mathcal{R}$ and $a \in E$. We simply verify that conditions (1)-(3) hold for $\Psi_{d \xrightarrow{a}}$. For condition (1), observe that clearly we have $\mathcal{D}^*\Psi_{d \xrightarrow{a}} \subseteq \mathcal{D}^*\Psi$, hence the dimension of $\mathcal{D}^*\Psi_{d \xrightarrow{a}}$ can be no greater than that of $\mathcal{D}^*\Psi$, which is finite. For condition (2), note that the map "evaluation at (d')," which takes Ψ to $\Psi(d')$, is linear. Thus, since Ψ is rational, the quantity $\Psi_{d \xrightarrow{a}}(d')$ is a rational function of d'. Finally, for condition (3), recall that the map taking Ψ to its derivative $\Psi_{d \xrightarrow{a}}$ is linear. Thus, if $L : \mathcal{D}^*\Phi \to \mathcal{R}$ is linear, then so is the map taking Ψ to $\Psi_{d \xrightarrow{a}} L$. Since Ψ is rational, the quantity $\Psi_{d \xrightarrow{a}} L$ is a rational function of d. ∎

Let $\mathrm{Rat}(x)$ denote the set of all real-valued rational functions of a single real parameter x. For n a nonnegative integer, an *n-dimensional representation* of an observable $\Phi \in \mathrm{Obs}(E)$ consists of

- An n-dimensional row vector C with entries in $\mathcal{R}$,
- An n-dimensional column vector $D(x)$ with entries in $\mathrm{Rat}(x)$,
- For each $a \in E$, an $n \times n$ matrix $M_a(x)$, with entries in $\mathrm{Rat}(x)$,

such that for all delayed traces $\alpha \in \mathrm{DTraces}(E)$, the quantity $\Phi(\alpha)$ is given by the formula:

$$\Phi(\alpha) = C\left(\prod_{k=0}^{|\alpha|-1} M_{\alpha(k,k+1)}(\alpha(k))\right) D(\alpha(|\alpha|)),$$

An observable $\Phi \in \mathrm{Obs}(E)$ is called *representable* if there exists a n-dimensional representation of Φ, for some n.

A representation is essentially a kind of automaton that computes a function on delayed traces (*i.e.* an observable). The states of the automaton are n-dimensional row vectors of real numbers, with the vector C serving as the initial state. If the automaton is in state X, and the initial portion of the input is $d\xrightarrow{a}$, then the automaton multiplies the current state vector by the matrix $M_a(d)$, and advances the input pointer. Upon reaching the end of the input, if the current state is X and the single remaining delay parameter is d, then the row vector X is multiplied by the column vector $D(d)$, to obtain a scalar, which becomes the output produced by the automaton.

Lemma 7 *A triple $(C, D(x), \{M_a(x) : a \in E\})$ is an n-dimensional representation of an observable $\Phi \in \mathrm{Obs}(E)$ if and only if there exists a linear transformation $R : \mathcal{R}^n \to \mathrm{Obs}(E)$ such that the following conditions hold:*

1. *$CR = \Phi$.*
2. *$XD(d) = (XR)(d)$, for all $X \in \mathcal{R}^n$, and all $d \in \mathcal{R}$.*
3. *$XM_a(d)R = (XR)_{d\xrightarrow{a}}$, for all $X \in \mathcal{R}^n$, all $d \in \mathcal{R}$, and all $a \in E$.*

Proof – Suppose $(C, D(x), \{M_a(x) : a \in E\})$ is an n-dimensional representation of Φ. Define the map $R : \mathcal{R}^n \to \mathrm{Obs}(E)$ so that for each $X \in \mathcal{R}^n$, the observable $XR \in \mathrm{Obs}(E)$ is defined by the condition of having $(X, D(x), \{M_a(x) : a \in E\})$ as a representation. Observe that if X and X' are in $\mathcal{R}^n$, then for any delayed trace $\alpha \in \mathrm{DTraces}(E)$ we have:

$$\begin{aligned}
&(cX + c'X')\left(\prod_{k=0}^{|\alpha|-1} M_{\alpha(k,k+1)}(\alpha(k))\right) D(\alpha(|\alpha|)) \\
&\quad = c\left\{X\left(\prod_{k=0}^{|\alpha|-1} M_{\alpha(k,k+1)}(\alpha(k))\right) D(\alpha(|\alpha|))\right\} \\
&\qquad\qquad + c'\left\{X'\left(\prod_{k=0}^{|\alpha|-1} M_{\alpha(k,k+1)}(\alpha(k))\right) D(\alpha(|\alpha|))\right\},
\end{aligned}$$

thus showing $(cX + c'X')R = c(XR) + c'(X'R)$; that is, R is linear. Clearly, $CR = \Phi$, showing that condition (1) holds. Also, $(XR)(d) = XD(d)$ is immediate by the fact that, by construction, $(X, D(x), \{M_a(x) : a \in E\})$ is a representation of XR, hence condition (2) holds. To see that condition (3) holds as well, observe that

$$\begin{aligned}((XM_a(d))R)(\alpha) &= XM_a(d)\left(\prod_{k=0}^{|\alpha|-1} M_{\alpha(k,k+1)}(\alpha(k))\right) D(\alpha(|\alpha|)) \\ &= (XR)(d \xrightarrow{a} \alpha) \\ &= (XR)_{d \xrightarrow{a}}(\alpha).\end{aligned}$$

Conversely, suppose we are given $(C, D(x), \{M_a(x) : a \in E\})$ and Φ, such that there exists a linear transformation $R : \mathcal{R}^n \to \mathrm{Obs}(E)$ satisfying conditions (1)-(3). Then a straightforward induction on $|\alpha|$, using conditions (1)-(3), establishes that, for all delayed traces $\alpha \in \mathrm{Obs}(E)$, we have:

$$\Phi(\alpha) = C\left(\prod_{k=0}^{l-1} M_{\alpha(k,k+1)}(\alpha(k))\right) D(\alpha(|\alpha|)),$$

thus showing that $(C, D(x), \{M_a(x) : a \in E\})$ is an n-dimensional representation of Φ. ∎

In the previous proof, we saw that, to any n-dimensional representation $(C, D(x), \{M_a(x) : a \in E\})$ of an observable $\Phi \in \mathrm{Obs}(E)$ there corresponds in a natural way a linear transformation $R : \mathcal{R}^n \to \mathrm{Obs}(E)$, which is defined to take $X \in \mathcal{R}^n$ to the observable XR having $(X, D(x), \{M_a(x) : a \in E\})$ as a representation. In the sequel, we shall refer to the map R as the linear transformation *associated with* the representation $(C, D(x), \{M_a(x) : a \in E\})$.

Theorem 2 *An observable $\Phi \in \mathrm{Obs}(E)$ is rational if and only if it is representable. Moreover, if an observable Φ is representable, then it has a representation whose dimension is equal to the dimension of Φ, and this dimension is the minimum possible among representations of Φ.*

Proof – We first show that representable observables are rational. Suppose Φ has an n-dimensional representation $(C, D(x), \{M_a(x) : a \in E\})$. Let

$$R : \mathcal{R}_n \to \mathrm{Obs}(E),$$

be the associated linear transformation, then the image of R contains the space $\mathcal{D}^*\Phi$. Since the image of a finite dimensional vector space under a linear transformation is finite dimensional, this shows that $\mathcal{D}^*\Phi$ is a finite dimensional subspace S of $\mathrm{Obs}(E)$, thus establishing condition (1) of the definition of a rational observable. To prove condition (2) of that definition, simply observe that, for each $X \in \mathcal{R}_n$, the value $(XR)(d)$ is given by the formula

$$XD(d),$$

which is clearly a rational function of d. To prove condition (3) in the definition of a rational observable, let $\Psi \in \mathcal{D}^*\Phi$, $a \in E$, and $L : \mathcal{D}^*\Phi \to \mathcal{R}$ be given. Then $\Psi = XR$ for some $X \in \mathcal{R}^n$, and also $\Psi_{d \xrightarrow{a}} = XM_a(d)R$. Thus, we have

$$\Psi_{d \xrightarrow{a}} L = XM_a(d)RL.$$

Since the right-hand side is a rational function of d (as can easily be seen by noting that application of the linear map RL corresponds to multiplication on the right by a column vector), condition (3) follows.

Conversely, suppose $\Phi \in \mathrm{Obs}(E)$ is rational. We show Φ is representable by constructing an n-dimensional representation of Φ, where n is the dimension of $\mathcal{D}^*\Phi$.

Let $\mathcal{B} = \{\Psi_1, \Psi_2, \ldots, \Psi_n\}$ be a basis for $\mathcal{D}^*\Phi$. Let $C \in \mathcal{R}^n$ be the row vector of coordinates of Φ with respect to the basis $\mathcal{B}$. For each $x \in \mathcal{R}$, the map "evaluate at x," which takes each $\Psi \in S$ to its value $\Psi(x)$ on the delayed trace of length zero with delay parameter x, is a linear functional on $\mathcal{D}^*\Phi$. Since by Lemma 6, every observable in $\mathcal{D}^*\Phi$, including the Ψ_i, is rational, it follows that the quantity $\Psi_i(x)$ is a rational function $r_i \in \mathrm{Rat}(x)$ for $1 \leq i \leq n$. Let $D(x)$ be the column vector having r_i as its ith entry, for $1 \leq i \leq n$.

For $x \in \mathcal{R}$ and $a \in E$, let $M_a(x)$ be the matrix, with respect to the basis $\mathcal{B}$, of the linear transformation on S taking each $\Psi \in S$ to its derivative $\Psi_{x \xrightarrow{a}}$. Since the mapping taking an element of S to its Ψ_j-coordinate is linear, and since every element of $\mathcal{D}^*\Phi$ is rational, it follows by condition (3) in the definition of an observable that the Ψ_j-coordinate of $(\Psi_i)_{x \xrightarrow{a}}$ is a rational function $r_{ij} \in \mathrm{Rat}(x)$. Let $M_a(x)$ be the matrix whose entries are the r_{ij}, for $1 \leq i \leq n$, $1 \leq j \leq n$.

We claim that the triple $(C, D(x), \{M_a(x) : a \in E\})$ is an n-dimensional representation of Φ. But this is clear, because the map $R : \mathcal{R}^n \to \mathrm{Obs}(E)$ taking $X \in \mathcal{R}^n$ to the observable Φ having $(X, D(x), \{M_a(x) : a \in E\})$ as a representation is clearly a linear transformation satisfying the conditions of Lemma 7.

Finally, note that the representation of Φ constructed above has dimension n which is equal to the dimension of $\mathcal{D}^*\Phi$. Moreover, this dimension is minimal among representations of Φ, because given any m-dimensional representation of Φ the image of the associated linear transformation $R : \mathcal{R}^m \to \mathrm{Obs}(E)$ can have dimension at most m. If $m < n$, this image cannot contain $\mathcal{D}^*\Phi$. ∎

3.2 Examples of Rational Observables

In this section we show that certain observables of interest are representable, hence rational.

The *completion probability* observable Π on $\mathrm{DTraces}(E)$ is defined by:

$$\Pi(\alpha) = \prod_{k=0}^{|\alpha|-1} \frac{1}{\alpha(k)}.$$

Recall from previous sections that this observable is related to the expected probability for a PIOA to complete a target set.

Lemma 8 *The completion probability observable* Π *has a* 1*-dimensional representation*

$$(C, D(x), \{M_a(x) : a \in E\}),$$

where

$$C = (1) \qquad D(x) = (1) \qquad M_a(x) = (1/x).$$

Proof – Obvious from the definition. ∎

The *completion time* observable Ω on DTraces(E) is defined by:

$$\Omega(\alpha) = \left(\prod_{k=0}^{|\alpha|-1} \frac{1}{\alpha(k)}\right)\left(\sum_{k=0}^{|\alpha|-1} \frac{1}{\alpha(k)}\right).$$

Recall from previous sections that the completion time observable is related to the expected completion time for a PIOA with respect to a target set.

Lemma 9 *The completion time observable* Ω *has the* 2*-dimensional representation*

$$(C, D(x), \{M_a(x) : a \in E\}),$$

where

$$C = \begin{pmatrix} 0 & 1 \end{pmatrix} \qquad D(x) = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \qquad M_a(x) = \begin{pmatrix} 1/x & 0 \\ 1/x^2 & 1/x \end{pmatrix}.$$

Proof – For a delayed trace α of the form:

$$d_0 \xrightarrow{a_0} d_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{m-1}} d_m,$$

where $m > 0$, let α' denote the delayed trace

$$d_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{m-1}} d_m.$$

We may then write:

$$\begin{aligned}
\Omega(\alpha) &= \left(\prod_{k=0}^{m-1} \frac{1}{d_k}\right)\left(\sum_{k=0}^{m-1} \frac{1}{d_k}\right) \\
&= \frac{1}{d_0}\left(\prod_{k=1}^{m-1} \frac{1}{d_k}\right)\left(\frac{1}{d_0} + \sum_{k=1}^{m-1} \frac{1}{d_k}\right) \\
&= \frac{1}{d_0^2}\Pi(\alpha') + \frac{1}{d_0}\Omega(\alpha')
\end{aligned}$$

It follows from this computation that Ω and Π satisfy the system of "differential equations":

$$\begin{aligned} \Pi_{d \xrightarrow{a}} &= \frac{1}{d}\Pi \\ \Omega_{d \xrightarrow{a}} &= \frac{1}{d^2}\Pi + \frac{1}{d}\Omega. \end{aligned}$$

Thus, the set $\mathcal{B} = \begin{pmatrix} \Pi & \Omega \end{pmatrix}$ is a basis for $\mathcal{D}^*\Omega$. Note further that C is the row vector of coefficients of Ω with respect to this basis, $D(x)$ is the column vector of coefficients of the "evaluate at x" functional on S with respect to the dual basis $\mathcal{B}^*$, and $M_a(x)$ is the matrix, with respect to $\mathcal{B}$ of the linear transformation "derivative with respect to $x \xrightarrow{a}$" on S. It follows that $(C, D, \{M_a : a \in E\})$ is a representation of Ω of minimal dimension. ∎

Lemma 10 *Suppose T is a target set which is also a regular subset of E^*. Then the observables Π_T and Ω_T are rational.*

Proof – Since T is regular, it is accepted by a DFA $\mathcal{M}$. Let $\{q_1, q_2, \ldots, q_m\}$ be an enumeration of the states of $\mathcal{M}$, with q_1 the start state. We also assume that q_m is the unique accept state of $\mathcal{M}$. Since T is pairwise incomparable under prefix, this can always be arranged, by manipulating $\mathcal{M}$ if necessary.

We first construct an m-dimensional representation of Π_T. Define

$$C = \begin{pmatrix} 1 & 0 & \ldots & 0 \end{pmatrix} \qquad\qquad D(x) = \begin{pmatrix} 0 \\ 0 \\ \ldots \\ 1 \end{pmatrix}$$

For each $a \in E$, define the matrix M_a as follows:

$$(M_a)_{ij}(x) = \begin{cases} 1/x, & \text{if } q_i \xrightarrow{a} q_j \text{ in } \mathcal{M}, \\ 0, & \text{otherwise.} \end{cases}$$

To prove that this is indeed a representation of Π_T, we claim that for all delayed traces α of the form:

$$d_0 \xrightarrow{a_0} d_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{l-1}} d_l,$$

the jth component of the row vector

$$C \left(\prod_{k=0}^{l-1} M_{\alpha(k,k+1)}(\alpha(k)) \right)$$

is equal to $\prod_{k=0}^{l-1} \frac{1}{d_k}$, if the input string $\alpha(0,1)\alpha(1,2)\ldots\alpha(l-1,l)$ takes $\mathcal{M}$ from state q_1 to state q_j, otherwise 0. This can be shown by a straightforward induction on l. It follows from this that for all delayed traces α of the above form, the value

$$C \left(\prod_{k=0}^{l-1} M_{\alpha(k,k+1)}(\alpha(k)) \right) D(\alpha(l))$$

is equal to $\prod_{k=0}^{l-1} \frac{1}{d_k}$, if the input string $\alpha(0,1)\alpha(1,2)\ldots\alpha(l-1,l)$ is accepted by $\mathcal{M}$, otherwise 0. But this is precisely the value of $\Pi_T(\alpha)$.

We next construct a $2m$-dimensional representation of Ω_T. The idea is the same as that for Π_T, except that we construct C, $D(x)$, and the $M_a(x)$ in 2-dimensional blocks. Define

$$C = \begin{pmatrix} 0 & 1 & 0 & 0 & \ldots & 0 & 0 \end{pmatrix} \qquad D(x) = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \ldots \\ 1 \\ 0 \end{pmatrix}$$

For each $a \in E$, let $M_a(x)$ be the $m \times m$ matrix of 2×2 blocks defined as follows:

$$(M_a)_{ij}(x) = \begin{cases} \begin{pmatrix} 1/x & 0 \\ 1/x^2 & 1/x \end{pmatrix}, & \text{if } q_i \xrightarrow{a} q_j \text{ in } \mathcal{M}, \\ \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, & \text{otherwise.} \end{cases}$$

We claim that for all delayed traces α of the form:

$$d_0 \xrightarrow{a_0} d_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{l-1}} d_l,$$

the jth 2-dimensional block of the row vector

$$C \left(\prod_{k=0}^{l-1} M_{\alpha(k,k+1)}(\alpha(k)) \right)$$

is equal to

$$\begin{pmatrix} 0 & 1 \end{pmatrix} \prod_{k=0}^{l-1} \begin{pmatrix} 1/d_k & 0 \\ 1/d_k^2 & 1/d_k \end{pmatrix}$$

if the input string $\alpha(0,1)\alpha(1,2)\ldots\alpha(l-1,l)$ takes $\mathcal{M}$ from state q_1 to state q_j, otherwise $\begin{pmatrix} 0 & 0 \end{pmatrix}$. Again, this can be shown by a straightforward induction on l. It follows from this that for all delayed traces α of the above form, the value

$$C \left(\prod_{k=0}^{l-1} M_{\alpha(k,k+1)}(\alpha(k)) \right) D(\alpha(l))$$

is equal to $\Omega(\alpha)$, if the input string $\alpha(0,1)\alpha(1,2)\ldots\alpha(l-1,l)$ is accepted by $\mathcal{M}$, otherwise 0. But this is precisely the definition of $\Omega_T(\alpha)$. ∎

3.3 Rational Observables and PIOA Behaviors

In this section we prove that the class of rational observables is closed under the application of PIOA behaviors. Moreover, a representation of $\mathcal{B}_E^A(\Phi)$ can be effectively computed from a representation of Φ.

We first consider the case of the behavior $\mathcal{B}_E^A$ of a PIOA A, where $E_A \subseteq E$.

Theorem 3 *Suppose A is a PIOA. If Φ is a rational observable in $\mathrm{Obs}(E)$, where $E_A \subseteq E$, then $\mathcal{B}_E^A\Phi$ is also a rational observable in $\mathrm{Obs}(E)$. Moreover, a representation of $\mathcal{B}_E^A\Phi$ can be effectively computed from a representation of Φ.*

Proof – Suppose $\Phi \in \mathrm{Obs}(E)$ is a rational observable, where $E_A \subseteq E$, and let

$$(C, D(x), \{M_a(x) : a \in E\})$$

be an n-dimensional representation of Φ. Suppose the PIOA A has m states. We show how to construct an mn-dimensional representation

$$(C', D'(x), \{M'_a(x) : a \in E\})$$

of $\mathcal{B}_E^A\Phi$.

The idea is similar to that of Lemma 10. Let $q_1, q_2, \ldots, q_m$, be an enumeration of the states of A, with q_1 the distinguished start state. Let C' be the mn-dimensional row vector consisting of n-dimensional blocks as follows:

$$C' = \begin{pmatrix} C & 0 & \ldots & 0 \end{pmatrix}.$$

Let D' be the mn-dimensional column vector consisting of n-dimensional blocks as follows:

$$D'(x) = \begin{pmatrix} D(x + \delta_A(q_1)) \\ D(x + \delta_A(q_2)) \\ \ldots \\ D(x + \delta_A(q_m)) \end{pmatrix}$$

For $a \in E$, let M'_a be the $mn \times mn$ matrix consisting of $n \times n$ blocks $(M'_a)_{ij}$ defined by:

$$(M'_a)_{ij}(x) = \mu_{ij}\nu_i M_a(x + \delta_A(q_i)),$$

where

$$\mu_{ij} = \mu_A(q_i, a, q_j)$$

and

$$\nu_i = \begin{cases} \delta_A(q_i), & \text{if } a \in E_A^{\mathrm{loc}}, \\ 1, & \text{otherwise.} \end{cases}$$

The following is the basic correctness property for the above representation.

Claim: *For all delayed traces α in* $\mathrm{DTraces}(E)$ *of the form:*

$$d_0 \xrightarrow{a_0} d_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{l-1}} d_l,$$

the jth n-dimensional block of the row vector:

$$C' \left(\prod_{k=0}^{l-1} M'_{\alpha(k,k+1)}(\alpha(k)) \right)$$

is equal to the following sum:

$$\sum_{\sigma \in \mathrm{Exec}_A(\alpha, q_j)} C \left(\prod_{k=0}^{l-1} M_{\sigma(k,k+1)}(\alpha(k) + \delta_A(\sigma(k))) \right) \mathrm{w}_A(\sigma),$$

where $\mathrm{Exec}_A(\alpha, q_j)$ *denotes the set of all executions of A of the form:*

$$r_0 \xrightarrow{a_0} r_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{l-1}} r_l,$$

with $r_0 = q_1$ and $r_l = q_j$.

To see that the theorem follows from the claim, note that, in view of the definition of D', the quantity

$$C' \left(\prod_{k=0}^{l-1} M'_{\alpha(k,k+1)}(\alpha(k)) \right) D'(\alpha(l))$$

is equal to

$$\sum_{j=1}^{m} \sum_{\sigma \in \mathrm{Exec}_A(\alpha, q_j)} C \left(\prod_{k=0}^{l-1} M_{\sigma(k,k+1)}(\alpha(k) + \delta_A(\sigma(k))) \right) D(\alpha(l) + \delta_A(\sigma(l)))\, \mathrm{w}_A(\sigma),$$

which in turn is equal to

$$\sum_{\sigma \propto \alpha} \Phi(\sigma \oplus \alpha) \mathrm{w}_A(\sigma),$$

or, more simply,

$$\mathcal{B}_E^A \Phi \alpha.$$

To prove the claim, we proceed by induction on l. The basis case $l = 0$ simply asserts that

$$(C')_j = \begin{cases} C, & \text{if } j = 1, \\ 0, & \text{otherwise,} \end{cases}$$

which is true by definition of C'.

Suppose now the result has been established for $l \geq 0$ and consider the case of $l+1$. Given delayed trace α of the form:

$$d_0 \xrightarrow{a_0} d_1 \xrightarrow{a_1} \ldots \xrightarrow{a_l} d_{l+1},$$

let α' denote the prefix of length l:

$$d_0 \xrightarrow{a_0} d_1 \xrightarrow{a_1} \dots \xrightarrow{a_{l-1}} d_l,$$

We observe:

$$C' \left(\prod_{k=0}^{l} M'_{\alpha(k,k+1)}(\alpha(k)) \right) = C' \left(\prod_{k=0}^{l-1} M'_{\alpha(k,k+1)}(\alpha(k)) \right) M'_{\alpha(l,l+1)}(\alpha(l)).$$

Applying the induction hypothesis and using the definition of $M'_{\alpha(l,l+1)}(\alpha(l))$ we have:

$$\begin{aligned}(C')_j &= \sum_{i=1}^{m} \left\{ \sum_{\sigma \in \mathrm{Exec}_A(\alpha', q_i)} C \left(\prod_{k=0}^{l-1} M_{\sigma(k,k+1)}(\alpha(k) + \delta_A(\sigma(k))) \right) \mathrm{w}_A(\sigma) \right\} \\ &\qquad \cdot M_{\alpha(l,l+1)}(\alpha(l) + \delta_A(\sigma(l))) \mu_{aij} \nu_{ai}.\end{aligned}$$

But this is easily seen to be equal to:

$$\sum_{\sigma \in \mathrm{Exec}_A(\alpha, q_j)} C \left(\prod_{k=0}^{l} M_{\sigma(k,k+1)}(\alpha(k) + \delta_A(\sigma(k))) \right) \mathrm{w}_A(\sigma),$$

as required. ∎

We now show how to extend the previous result to the case of $\mathcal{B}_E^A$, where we do not necessarily have $E_A \subseteq E$. If $E \subseteq E'$, then define the map $[\text{ - }]_E : \mathrm{Obs}(E') \to \mathrm{Obs}(E)$ by:

$$[\Psi]_E(\alpha) = \sum_{\alpha' \rhd \alpha} \Psi(\alpha').$$

Note that the sum on the right need not converge, in general, so that $[\Psi]_E$ will be defined only for certain $\Psi \in \mathrm{Obs}(E')$.

The following result states that the E-behavior of A is determined by the $(E \cup E_A)$-behavior of A.

Lemma 11 *Suppose A is a PIOA. Then for all sets of actions E, for all observables $\Phi \in \mathrm{Obs}(E \cup E_A)$, and for all delayed traces $\alpha \in \mathrm{DTraces}(E)$ we have:*

$$\mathcal{B}_E^A \Phi \alpha = \sum_{\alpha' \rhd \alpha} \mathcal{B}_{E \cup E_A}^A \Phi \alpha' = [\mathcal{B}_{E \cup E_A}^A \Phi]_E(\alpha).$$

Proof – Write out the definition and observe that for an execution σ of A we have $\sigma \propto \alpha$ if and only if there exists a unique $\alpha' \in \mathrm{DTraces}(E \cup E_A)$ such that $\alpha' \rhd \alpha$ and $\sigma \propto \alpha'$. ∎

In view of Lemma 11, the map $[\text{ - }]_E$ allows us to reduce the problem of computing a representation of $\mathcal{B}_E^A \Phi$ to that of computing, given a representation of an observable $\Psi \in \mathrm{Obs}(E \cup E_A)$, a representation of the observable $[\Psi]_E \in \mathrm{Obs}(E)$. This is the content of Theorem 4 below.

Lemma 12 *Suppose $E \subseteq E'$. Suppose further that S is a linear subspace of* $\mathrm{Obs}(E')$ *such that:*

1. *$\mathcal{D}S \subseteq S$.*

2. *$[\Psi]_E$ is defined for all $\Psi \in S$.*

Then the following relations are satisfied for all $\Psi \in S$, all $d \in \mathcal{R}$, and all $a \in E$:

$$([\Psi]_E)_{d \xrightarrow{a}} = [\Psi_{d \xrightarrow{a}}]_E + \sum_{a' \in E' \setminus E} ([\Psi_{d \xrightarrow{a'}}]_E)_{d \xrightarrow{a}}$$

$$[\Psi]_E(d) = \Psi(d) + \sum_{a' \in E' \setminus E} [\Psi_{d \xrightarrow{a'}}]_E(d).$$

Proof – To prove the first relation, observe that if $\alpha' \rhd (d \xrightarrow{a} \alpha)$, then either

1. α' is $d \xrightarrow{a} \alpha''$, where $\alpha'' \rhd \alpha$, or

2. α' is $d \xrightarrow{a'} \alpha''$, where $\alpha'' \rhd (d \xrightarrow{a} \alpha)$.

Thus, assuming $[\Psi]_E$ is defined, we may write:

$$\begin{aligned}
([\Psi]_E)_{d \xrightarrow{a}} \alpha &= \sum_{\alpha'' \rhd \alpha} \Psi(d \xrightarrow{a} \alpha'') + \sum_{a' \in E' \setminus E} \sum_{\alpha'' \rhd (d \xrightarrow{a} \alpha)} \Psi(d \xrightarrow{a'} \alpha'') \\
&= [\Psi_{d \xrightarrow{a}}]_E \alpha + \sum_{a' \in E' \setminus E} ([\Psi_{d \xrightarrow{a'}}]_E)_{d \xrightarrow{a}} \alpha,
\end{aligned}$$

from which the first relation follows.

To prove the second relation, observe that if $\alpha' \rhd d$, then either

1. α' is d, or

2. α' is $d \xrightarrow{a'} \alpha''$, where $\alpha'' \rhd d$.

Thus, we may write:

$$\begin{aligned}
[\Psi]_E(d) &= \Psi(d) + \sum_{a' \in E' \setminus E} \sum_{\alpha'' \rhd d} \Psi(d \xrightarrow{a'} \alpha'') \\
&= \Psi(d) + \sum_{a' \in E' \setminus E} [\Psi_{d \xrightarrow{a'}}]_E(d),
\end{aligned}$$

which is the second relation. ∎

Theorem 4 *Suppose* $(C', D', \{M'_{a'} : a' \in E'\})$ *is a representation of an observable* $\Phi' \in \mathrm{Obs}(E')$*, and suppose* $E \subseteq E'$*, and suppose* $[\text{ - }]_E$ *is well defined on* $\mathcal{D}^*\Phi'$*. Suppose further that the power series:*

$$I + \hat{M}(x) + \hat{M}^2(x) + \ldots$$

converges (componentwise) for all nonnegative $x \in \mathcal{R}$*, where we define*

$$\hat{M}(x) = \sum_{a' \in E' \setminus E} M'_{a'}(x).$$

Then the matrix $I - \hat{M}(x)$ *is nonsingular for all nonnegative* $x \in \mathcal{R}$*, and an* n*-dimensional representation of* $[\Phi']_E$ *is given by*

$$(C', \;\; (I - \hat{M}(x))^{-1} D'(x), \;\; \{(I - \hat{M}(x))^{-1} M'_a(x) : a \in E\}).$$

Proof – First, note that it is easy to show that the componentwise convergence of the indicated power series implies that $I - \hat{M}(x)$ is nonsingular, and its inverse is the sum of the series.

Now, let $R : \mathcal{R}^n \to \mathrm{Obs}(E)$ be defined by:

$$XR = [XR']_E,$$

where $R' : \mathcal{R}^n \to \mathrm{Obs}(E')$ is the linear transformation associated with the given representation of Φ'. Clearly, the linearity of R follows from that of R' and the fact that $[\text{ - }]_E$ is well-defined, hence linear by the form of its definition, on $\mathcal{D}^*\Phi'$. We claim that R satisfies the three conditions of Lemma 7 with respect to the data $(C, D(x), \{M_a(x) : a \in E\})$, thus the latter is a representation of $[\Phi']_E$.

1. That $CR = [\Phi']_E$ is immediate by construction.

2. From Lemma 12 (1), the following relation holds for all Ψ in $\mathcal{D}^*\Phi'$, all $d \in \mathcal{R}$, and all $a \in E$:

   $$([\Psi]_E)_{d \xrightarrow{a}} = [\Psi_{d \xrightarrow{a}}]_E + \sum_{a' \in E' \setminus E} ([\Psi_{d \xrightarrow{a'}}]_E)_{d \xrightarrow{a}}.$$

   Thus, using the fact that $(C', D'(x), \{M'_{a'}(x) : a' \in E'\})$ is a representation of Φ', for all $Y \in \mathcal{R}^n$, all $d \in \mathcal{R}$, and all $a \in E$ we have:

   $$(YR)_{d \xrightarrow{a}} = YM'_a(d)R + \sum_{a' \in E' \setminus E} (YM'_{a'}(d)R)_{d \xrightarrow{a}}.$$

   Rearranging terms using linearity and using the definition of $\hat{M}(x)$, we have

   $$(Y(I - \hat{M}(d))R)_{d \xrightarrow{a}} = YM'_a(d)R.$$

Since this holds for all $Y \in \mathcal{R}^n$, it certainly holds for $Y = X(I - \hat{M}(d))^{-1}$, hence

$$\begin{aligned}(XR)_{d \xrightarrow{a}} &= (X(I - \hat{M}(d))^{-1} M'_a(d) R \\ &= X M_a(d) R\end{aligned}$$

holds for all $X \in \mathcal{R}^n$, all $d \in \mathcal{R}$, and all $a \in E$, as required.

3. From Lemma 12 (2), the following relation holds for all Ψ in $\mathcal{D}^*\Phi'$ and all $d \in \mathcal{R}$:

$$[\Psi]_E(d) = \Psi(d) + \sum_{a' \in E' \setminus E} [\Psi_{d \xrightarrow{a'}}]_E(d).$$

Thus, using the fact that $(C', D'(x), \{M'_{a'}(x) : a' \in E'\})$ is a representation of Φ', for all $Y \in \mathcal{R}^n$ and all $d \in \mathcal{R}$, we have:

$$(YR)(d) = YD'(d) + \sum_{a' \in E' \setminus E} (YM'_{a'}(d)R)(d).$$

Rearranging terms using linearity and using the definition of $\hat{M}(x)$, we have

$$(Y(I - \hat{M}(d))R)(d) = YD'(d).$$

Since this holds for all $Y \in \mathcal{R}^n$, it certainly holds for $Y = X(I - \hat{M}(d))^{-1}$, hence

$$(XR)(d) = X(I - \hat{M}(d))^{-1} D'(d)$$

holds for all $X \in \mathcal{R}^n$ and all $d \in \mathcal{R}$, as required.

∎

We now obtain sufficient conditions for the technical hypotheses of the previous theorem to hold.

Lemma 13 *Suppose M is an $n \times n$ matrix over an arbitrary field. If the matrix $I - M$ is nonsingular, then its inverse $(I - M)^{-1}$ is the unique solution X to the equation:*

$$X = I + MX.$$

Proof – Since

$$(I - M)(I - M)^{-1} = I,$$

we have

$$(I - M)^{-1} - M(I - M)^{-1} = I,$$

hence

$$(I - M)^{-1} = I + M(I - M)^{-1},$$

so that $(I - M)^{-1}$ solves $X = I + MX$.

Suppose X and X' are both solutions to the above equation. Then

$$X - X' = M(X - X'),$$

hence

$$(I - M)(X - X') = 0;$$

Multiplying both sides by $(I - M)^{-1}$, we have

$$X - X' = 0,$$

hence $X = X'$. ∎

Lemma 14 *Suppose M is an $n \times n$ matrix over the reals such that*

1. *$M \geq 0$ componentwise.*
2. *The matrix $I - M$ is nonsingular, and its inverse $(I - M)^{-1}$ satisfies $(I - M)^{-1} \geq 0$ componentwise.*

Then the power series:

$$I + M + M^2 + \ldots$$

converges componentwise to $(I - M)^{-1}$.

Proof – An inductive argument, using the fact that $M \geq 0$ and that $(I - M)^{-1} \geq 0$ and satisfies $X = I + MX$, shows that:

$$0 \leq I \leq I + M \leq I + M + M^2 \leq \ldots \leq (I - M)^{-1}$$

holds componentwise. This shows that the power series is bounded componentwise, hence converges componentwise. Clearly then the sum is a solution to $X = I + MX$. But by Lemma 13, $(I - M)^{-1}$ is the unique solution, hence it is equal to the sum. ∎

Corollary 15 *Suppose $(C', D'(x), \{M'_{a'} : a' \in E'\})$ is a representation of $\Phi' \in \mathrm{Obs}(E')$, and suppose $E \subseteq E'$. Let $\hat{M}(x) = \sum_{a' \in E' \setminus E} M'_{a'}(x)$. Suppose*

1. *$M_{a'}(x) \geq 0$ and $D'(x) \geq 0$ componentwise for all $a' \in E'$ and all nonnegative $x \in \mathcal{R}$.*
2. *For all $x \in \mathcal{R}$, the matrix $I - \hat{M}(x)$ is nonsingular, and its inverse satisfies $(I - \hat{M}(x))^{-1} \geq 0$ componentwise.*

Then the power series:

$$I + \hat{M}(x) + \hat{M}^2(x) + \ldots$$

converges componentwise for all nonnegative $x \in \mathcal{R}$, and $[\Psi]_E$ is defined for all $\Psi \in \mathcal{D}^* \Phi'$.

Proof – By Lemma 14, the hypotheses imply the componentwise convergence of the indicated power series.

It remains to be shown that $[\Psi]_E$ is defined for all $\Psi \in \mathcal{D}^*\Phi'$. Now, if $\Psi \in \mathcal{D}^*\Phi'$, then $\Psi = XR'$ for some $X \in \mathcal{R}^n$, where $R' : \mathcal{R}^n \to \mathrm{Obs}(E')$ is the linear transformation associated with the given representation of Φ'. We show that $[\Psi]_E$ is is defined in case $X \geq 0$ holds componentwise; the case of arbitrary Ψ in $\mathcal{D}^*\Phi'$ follows easily from this by linearity, using the fact that an arbitrary $X \in \mathcal{R}^n$ can be written uniquely as $X = X^+ - X^-$, where $X^+ \geq 0$ and $X^- \geq 0$ hold componentwise.

Suppose, then, that $X \geq 0$ holds componentwise. Let $\alpha \in \mathrm{DTraces}(E)$ be arbitrary, and suppose $\alpha' \in \mathrm{DTraces}(E')$ satisfies $\alpha' \rhd \alpha$, with corresponding monotone injection ϕ. Suppose $|\alpha'| = n$. Then the fact that $(X, D'(x), \{M'_{a'} : a' \in E'\})$ is a representation for Ψ gives us the following formula for $\Psi(\alpha')$:

$$\begin{aligned}
\Psi(\alpha') \quad = \quad & X \left(\prod_{k=\phi(0)}^{\phi(1)-1} M'_{\alpha'(k)}(d_0) \right) M'_{a_0}(d_0) \\
& \left(\prod_{k=\phi(1)}^{\phi(2)-1} M'_{\alpha'(k)}(d_1) \right) M'_{a_1}(d_1) \\
& \quad \ldots \left(\prod_{k=\phi(n-2)}^{\phi(n-1)-1} M'_{\alpha'(k)}(d_{n-1}) \right) M'_{a_{n-1}}(d_{n-1}) \\
& \qquad \left(\prod_{k=\phi(n-1)}^{n} M'_{\alpha'(k)}(d_n) \right) D'(d_n).
\end{aligned}$$

As all terms in the above formula are nonnegative, summing the above formula over all $\alpha' \rhd \alpha$ we obtain:

$$\begin{aligned}
\sum_{\alpha' \rhd \alpha} \Psi(\alpha') \quad \leq \quad & X(I - \hat{M}(d_0))^{-1} M'_{a_0}(d_0) \\
& (I - \hat{M}(d_1))^{-1} M'_{a_1}(d_1) \\
& \quad \ldots (I - \hat{M}(d_2))^{-1} M'_{a_{n-1}}(d_{n-1}) \\
& \qquad (I - \hat{M}(d_n))^{-1} D'(d_n).
\end{aligned}$$

In particular, the above summation converges. But this summation is the definition of $[\Psi]_E(\alpha)$. ∎

3.4 Example

We now apply the method of the previous section to calculate the completion probability and expected completion time for a simple example.

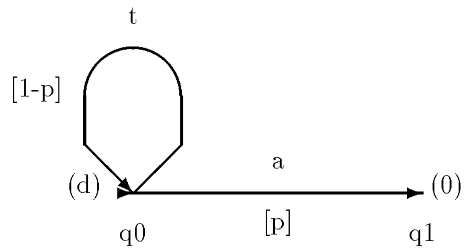

Figure 2: Example PIOA for Calculation of Expected Completion Time

Let A be the PIOA with $E_A = \{t, a\}$, where t is an internal action and a is an output action, with $Q_A = \{q_0, q_1\}$, where q_0 is the start state, with $\mu_A(q_0, a, q_1) = p$, $\mu_A(q_0, t, q_0) = 1 - p$, and $\mu_A(q_i, a', q_j) = 0$ for all other cases, and with $\delta_A(q_0) = d$ and $\delta_A(q_1) = 0$ (see Figure 2). Let T be the target set consisting of the single string a.

We wish to calculate the completion probability for A with respect to T, and the expected completion time for A with respect to T. Note that for this simple example, it is possible to carry out by hand a non-compositional method for determining these quantities. In particular, the expected completion time x satisfies the following linear equation:

$$x = 1/d + (1-p)x + p \cdot 0,$$

which expresses the expected completion time from state q_0 as the sum of the expected "dwell time" in state q_0, plus the sum of the expected delays from the successor states q_0 and q_1 of q_0, weighted respectively by the probability of transitions to these successor states. Solving this equation for x yields the result:

$$x = 1/dp$$

In a similar fashion, the completion probability can be shown to be 1.

We now apply the theory of the preceding sections to provide an alternative calculation of the same quantities. We wish to emphasize that, although the calculations are more involved for this simple example, the real advantage of our method over the "equation-solving" method will be realized on very large systems having many components. For these cases, the non-compositional equation-solving method will yield an unmanageably large system of equations to be solved, whereas our method can be applied one component at a time, potentially avoiding thereby a similar explosion in the size of the data to be stored (assuming, of course that we systematically apply the minimization algorithm to be presented in Section 3.5).

Following the theory of the preceding sections, the completion probability for the above example is given by:

$$\mathcal{B}_{\emptyset}^{A} \Pi_T(0)$$

and the expected completion time is given by:

$$\mathcal{B}_{\emptyset}^{A} \Omega_T(0)$$

To calculate the completion probability, we first note that the observable Ω_T has the 4-dimensional representation given by:

$$C_{\Omega_T} = \begin{pmatrix} 0 & 1 & 0 & 0 \end{pmatrix} \qquad D(x) = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}$$

$$M_a(x) = \begin{pmatrix} 0 & 0 & 1/x & 0 \\ 0 & 0 & 1/x^2 & 1/x \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \qquad M_t(x) = \begin{pmatrix} 1/x & 0 & 0 & 0 \\ 1/x^2 & 1/x & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

For Π_T, we simply replace C_{Ω_T} by

$$C_{\Pi_T} = \begin{pmatrix} 1 & 0 & 0 & 0 \end{pmatrix}.$$

This does not give a representation of minimal dimension, but this is not required.

From the above, using Theorem 3, we can compute an 8-dimensional representation for $\mathcal{B}^A_{E_A}\Omega_T$:

$$C_{\Omega_T} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \qquad D(x) = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}$$

$$M_a(x) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & \frac{dp}{x+d} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{dp}{(x+d)^2} & \frac{dp}{x+d} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$M_t(x) = \begin{pmatrix} \frac{d(1-p)}{x+d} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{d(1-p)}{(x+d)^2} & \frac{d(1-p)}{x+d} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

For $\mathcal{B}^A_{E_A}\Pi_T$, we replace C_{Ω_T} by

$$C_{\Pi_T} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Next, we wish to compute representations for

$$\mathcal{B}^A_\emptyset \Pi_T = [\mathcal{B}^A_{E_A}\Pi_T]_\emptyset.$$

and

$$\mathcal{B}^A_\emptyset \Omega_T = [\mathcal{B}^A_{E_A}\Omega_T]_\emptyset;$$

We first compute the matrix $I - \hat{M}(x) = (I - M_a(x) - M_t(x))$:

$$\begin{pmatrix} \frac{x+dp}{x+d} & 0 & 0 & 0 & 0 & 0 & -\frac{dp}{x+d} & 0 \\ -\frac{d(1-p)}{(x+d)^2} & \frac{x+dp}{x+d} & 0 & 0 & 0 & 0 & -\frac{dp}{(x+d)^2} & -\frac{dp}{x+d} \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

Using a computer algebra system (GNU Emacs Calc 2.02) to invert this matrix symbolically, we obtain:

$$\begin{pmatrix} \frac{x+d}{x+dp} & 0 & 0 & 0 & 0 & 0 & \frac{dp}{x+dp} & 0 \\ \frac{d(1-p)}{(x+dp)^2} & \frac{x+d}{x+dp} & 0 & 0 & 0 & 0 & \frac{dp}{(x+dp)^2} & \frac{dp}{x+dp} \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

From this, we can calculate the completion probability.

$$C_{\Pi_T}(I - \hat{M}(x))^{-1}D(x) = \frac{dp}{x + dp}.$$

Evaluating at $x = 0$ and assuming $dp \neq 0$ yields the result: 1.

We can also compute the expected completion time.

$$C_{\Omega_T}(I - \hat{M}(x))^{-1}D(x) = \frac{dp}{(x + dp)^2}.$$

Evaluating at $x = 0$ gives:

$$\frac{1}{dp}.$$

3.5 Minimization of Representations

In this section, we present an algorithm that, given a representation $(C, D, \{M_a : a \in E\})$ for an observable Φ, computes a representation $(C', D', \{M'_a : a \in E\})$ for Φ which is of minimal dimension.

We first obtain necessary and sufficient conditions for a representation to be minimal.

Lemma 16 *Suppose $(C, D(x), \{M_a(x) : a \in E\})$ is an n-dimensional representation of an observable $\Phi \in \mathrm{Obs}(E)$. Then this representation is minimal if and only if the associated linear transformation*

$$R : \mathcal{R}^n \to \mathrm{Obs}(E)$$

is an isomorphism from $\mathcal{R}^n$ to $\mathcal{D}^*\Phi$.

Proof – If R is an isomorphism from $\mathcal{R}^n$ to $\mathcal{D}^*\Phi$, then the dimension of $\mathcal{D}^*\Phi$ must be n, thus showing that the representation is minimal.

Conversely, if the representation is minimal, then $\mathcal{D}^*\Phi$ must have dimension n. Since $\mathcal{D}^*\Phi$ is contained in the image of R, it follows that the image of R has dimension at least n. But the image of R can have dimension no more than n (the dimension of its domain, $\mathcal{R}^n$), hence the image of R has dimension exactly n. Thus, R is injective and $\mathcal{D}^*\Phi$ coincides with the image of R, showing that R is an isomorphism from $\mathcal{R}^n$ to $\mathcal{D}^*\Phi$. ∎

Lemma 17 *Suppose $(C, D(x), \{M_a(x) : a \in E\})$ is an n-dimensional representation of an observable $\Phi \in \mathrm{Obs}(E)$. Then this representation is minimal if and only if neither of the following two (infinite) systems of equations has any nontrivial solutions:*

1. *The set of all equations of the form*

$$X(\prod_{k=0}^{l-1} M_{a_k}(d_k))D(d_l) = 0$$

in the unknown row vector X, where l ranges over all nonnegative integers, the a_k range over all elements of E, and the d_k range over all nonnegative reals.

2. *The set of all equations of the form*

$$C(\prod_{k=0}^{l-1} M_{a_k}(d_k))Y = 0$$

in the unknown column vector Y, where l ranges over all nonnegative integers, the a_k range over all elements of E, and the d_k range over all nonnegative reals.

Proof – We first show that a nontrivial solution either to system (1) or system (2) would imply that the representation is non-minimal. Let $R : \mathcal{R}^n \to \mathrm{Obs}(E)$ be the linear transformation associated with the representation.

If system (1) has a nontrivial solution X, then that would imply that the observable XR is identically zero. But then $(cX)R = c(XR) = 0 = XR$ for all $c \in \mathcal{R}$, showing that in this case the mapping R cannot be injective. It follows by Lemma 16 that the representation cannot be minimal.

If system (2) has a nontrivial solution Y, then that means Y is orthogonal to the subspace of $\mathcal{R}^n$ spanned by all vectors of the form $C(\prod_{k=0}^{l-1} M_{a_k}(d_k))$. But then the dimension of this subspace must be strictly less than n. Since $\mathcal{D}^*\Phi$ is contained in this subspace, it follows that the dimension of Φ is strictly less than n, so that an n-dimensional representation cannot be minimal.

Conversely, suppose that the representation is non-minimal. Then by Lemma 16 the mapping R is not an isomorphism of $\mathcal{R}^n$ to $\mathcal{D}^*\Phi$. Then either R is not injective, or else the image of R contains an observable that is not in $\mathcal{D}^*\Phi$. If R is not injective, then $XR = X'R$ for some distinct X and X' in $\mathcal{R}^n$. This implies $(X - X')R$ is the identically zero observable, which then implies that $X - X'$ is a nontrivial solution to system (1). Suppose now that the image of R contains an observable Ψ that is not in $\mathcal{D}^*\Phi$. Then Ψ is independent of all observables XR, where X can be expressed in the form

$$X = C(\prod_{k=0}^{l-1} M_{a_k}(d_k)),$$

so that the space spanned by such observables has dimension strictly less than n. Since we have already shown that R is injective, it follows that the subspace of $\mathcal{R}^n$ spanned by the set of all X that can be expressed in the above form also has dimension strictly less than n, hence is a proper subspace of $\mathcal{R}^n$. Therefore, there exists a nontrivial $Y \in \mathcal{R}^n$ which is orthogonal to this subspace. Such a Y is a nontrivial solution to the system (2). ∎

Lemma 18 *There exist algorithms for:*

1. *Computing a basis for the subspace of all n-dimensional row vectors $X \in \mathcal{R}^n$ that satisfy a system of identities of the form:*

$$XD_j(x) = 0, \qquad \text{all nonnegative } x \in \mathcal{R}, 1 \le j \le m$$

where $D_1(x), D_2(x), \ldots, D_m(x)$ are given n-dimensional column vectors with entries in $\mathrm{Rat}(x)$.

2. *Computing a basis for the subspace of all n-dimensional column vectors $Y \in \mathcal{R}^n$ that satisfy a system of identities of the form:*

$$C_j(x)Y = 0, \qquad \text{all nonnegative } x \in \mathcal{R}, 1 \leq j \leq m$$

where $C_1(x), C_2(x), \ldots, C_m(x)$ are given n-dimensional row vectors with entries in $\mathrm{Rat}(x)$.

Proof – We only prove (1), the proof of (2) is analogous. Suppose

$$D_j(x) = \begin{pmatrix} d_{j1}(x) \\ d_{j2}(x) \\ \ldots \\ d_{jn}(x) \end{pmatrix}.$$

A row vector

$$X = \begin{pmatrix} x_1 & x_2 & \ldots & x_n \end{pmatrix}$$

satisfies the jth identity if and only if the rational function

$$x_1 d_{j1}(x) + x_2 d_{j2}(x) + \ldots + x_n d_{jn}(x)$$

is identically zero. By expressing each $d_{jk}(x)$ as a quotient of polynomials in x, combining the above terms over a common denominator, reducing by cancellation of factors common between the numerator and the denominator, and then equating the coefficient of each power of x in the numerator to zero, we obtain a homogeneous system of linear equations in the unknowns $x_1, x_2, \ldots, x_n$, with the property that a row vector

$$X = \begin{pmatrix} x_1 & x_2 & \ldots & x_n \end{pmatrix}$$

solves this system if and only if it is a solution to the jth identity. Solving simultaneously the m systems of linear equations obtained in this way yields a basis for the space of all row vectors X that simultaneously satisfy the original identities. ∎

Lemma 19 *There exist algorithms that:*

1. *Given $D(x)$ and $\{M_a(x) : a \in E\}$, compute a basis for the solution space of the system of all equations of the form:*

$$X(\prod_{k=0}^{l-1} M_{a_k}(d_k))D(d_k) = 0, \qquad l \geq 0, a_k \in E, d_k \in \mathcal{R}, d_k \geq 0$$

2. *Given C and $\{M_a(x) : a \in E\}$, compute a basis for the solution space of the system of all equations of the form:*

$$C(\prod_{k=0}^{l-1} M_{a_k}(d_k))Y = 0. \qquad l \geq 0, a_k \in E, d_k \in \mathcal{R}, d_k \geq 0$$

Proof – (1) The algorithm is based on the observation that if S denotes the solution space, then $S = \cap_{k=0}^{\infty} S_k$, where

$$S_0 = \{X \in \mathcal{R}^n : XD(x) = 0, \text{ all } x \in \mathcal{R}\},$$

and

$$S_{k+1} = S_k \cap \{X \in \mathcal{R}^n : XM_a(x) \in S_k, \text{ all } a \in E, x \in \mathcal{R}\}.$$

The algorithm works by successively computing a basis for each S_k, until a point is reached where S_{k+1} has the same dimension as S_k, which implies that $S_{k+1} = S_k = S$.

We now give the full description of the algorithm. First, apply Lemma 19 to obtain a basis

$$\mathcal{B}_0 = \{B_{01}, B_{02}, \ldots B_{0n_0}\}$$

for the subspace S_0 of $\mathcal{R}^n$. Then, repeat the following step until no further reduction in dimension is achieved:

- Given a linearly independent set $\mathcal{B}_k = \{B_{k1}, B_{k2}, \ldots, B_{kn_k}\}$, spanning a subspace S_k of $\mathcal{R}^n$, solve the system of simultaneous linear equations:

  $$B_{kj}Y = 0, \qquad (1 \leq j \leq n_k)$$

  to obtain a basis

  $$\mathcal{C}_k = \{C_{k1}, C_{k2}, \ldots, C_{km_k}\}$$

  for the orthogonal complement $S_k^{\perp}$ of S_k. Then, apply Lemma 19 to solve the system consisting of the equations (1) and identities (2) below:

  1. $$XC_{kj} = 0, \qquad (1 \leq j \leq m_k)$$

  2. $$XM_a(x)C_{kj} = 0, \qquad (1 \leq j \leq m_k, a \in E, x \in \mathcal{R})$$

  to obtain a basis

  $$\mathcal{B}_{k+1} = \{B_{k+1,1}, B_{k+1,2}, \ldots, B_{k+1,n_{k+1}}\}$$

  for a subspace S_{k+1} of $\mathcal{R}^n$. Observe that S_{k+1} is a subspace of S_k, due to the presence of the equations (1).

We claim that when the above step yields no further reduction in dimension (*i.e.* $n_{k+1} = n_k$), then the resulting set $\mathcal{B}_k$ is a basis for the space S of solutions X to the set of all equations of the form:

$$X(\prod_{k=0}^{l-1} M_{a_k}(d_k))D(d_k) = 0.$$

For, clearly $S = \cap_k S_k$. Moreover, when $S_{k+1} = S_k$, then $S_{k'} = S_k$ for all $k' \geq k$, and hence $S_k = S$.

(2) is proved in an entirely analogous fashion, using the observation that if S denotes the solution space, then $S = \cap_{k=0}^{\infty} S_k$, where

$$S_0 = \{Y \in \mathcal{R}^n : CY = 0\},$$

and

$$S_{k+1} = S_k \cap \{Y \in \mathcal{R}^n : M_a(x)Y \in S_k, \text{ all } a \in E, x \in \mathcal{R}\}.$$

∎

Lemma 20 *Suppose* $(C, D(x), \{M_a(x) : a \in E\})$ *is an* n*-dimensional representation of an observable* $\Phi \in \mathrm{Obs}(E)$*. Suppose the system of all equations of the form:*

$$X(\prod_{k=0}^{l-1} M_{a_k}(d_k))D(d_k) = 0.$$

has a nontrivial solution X*. Let* m *be the dimension of the solution space. Then* Φ *has an* $(n-m)$*-dimensional representation* $(C', D'(x), \{M'_a(x) : a \in E\})$*, which is effectively computable from the given* n*-dimensional representation.*

Proof – Using Lemma 19, we can compute a basis $\mathcal{B} = \{B_1, B_2, \ldots, B_m\}$ (of row vectors) for the solution space S of the above system of equations and a basis $\mathcal{C} = \{C_1, C_2, \ldots, C_{n-m}\}$ for the orthogonal complement $S^\perp$ of S. Assume that the basis $\mathcal{C}$ is orthonormal, which can be ensured using the Gram-Schmidt procedure. Let $P_{S^\perp}$ be the $(n \times (n-m))$-dimensional matrix of the projection of $\mathcal{R}^n$ (row vectors) to $S^\perp$, with respect to the basis $\mathcal{C}$ for $S^\perp$ and the natural basis for $\mathcal{R}^n$. Explicitly, for $1 \leq i \leq n-m$, the ith column of the matrix $P_{S^\perp}$ contains the components of the basis vector C_i. Let $Q_{S^\perp} = P^{\mathrm{t}}_{S^\perp}$, which is the $((n-m) \times n)$-dimensional matrix of the embedding of $S^\perp$ in $\mathcal{R}^n$, with respect to the basis $\mathcal{C}$ for $S^\perp$ and the natural basis for $\mathcal{R}^n$.

Before proceeding to exhibit the $(n-m)$-dimensional representation for Φ, we first observe that the following relationship holds for all $a \in E$:

$$P_{S^\perp} Q_{S^\perp} M_a(x) P_{S^\perp} = M_a(x) P_{S^\perp}.$$

This is because of the fact that, for an arbitrary vector X, the vector $X - X P_{S^\perp} Q_{S^\perp}$ is the orthogonal projection of X onto the subspace S of $\mathcal{R}^n$. From the defining condition for S, it is easy to see that S is invariant under $M_a(x)$; that is, if $X \in S$, then also $XM_a(x) \in S$. Thus the vector

$$XM_a(x) - XP_{S^\perp}Q_{S^\perp}M_a(x)$$

is also in S, hence

$$XM_a(x)P_{S^\perp} - XP_{S^\perp}Q_{S^\perp}M_a(x)P_{S^\perp} = 0,$$

that is to say,

$$XM_a(x)P_{S^\perp} = XP_{S^\perp}Q_{S^\perp}M_a(x)P_{S^\perp}.$$

Since X was arbitrary, the relation

$$M_a(x)P_{S^\perp} = P_{S^\perp}Q_{S^\perp}M_a(x)P_{S^\perp}$$

follows. Similar reasoning establishes $D(x) = P_{S^\perp}Q_{S^\perp}D(x)$.

We now define $C' = CP_{S^\perp}$, $M'_a(x) = Q_{S^\perp}M_a(x)P_{S^\perp}$, and $D'(x) = Q_{S^\perp}D(x)$. We claim that that $(C', D'(x), \{M'_a(x) : a \in E\})$ is also a representation of Φ. To show this, it suffices to show that

$$C\left(\prod_{k=0}^{l-1} M_{a_k}(d_k)\right)P_{S^\perp} \quad = \quad C'\left(\prod_{k=0}^{l-1} M'_{a_k}(d_k)\right) \tag{1}$$

for all delayed traces

$$d0 \xrightarrow{a_0} d_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{l-1}} d_l.$$

For then, since $D(x) = P_{S^\perp}Q_{S^\perp}D(x)$, it follows that

$$\begin{aligned}
C'\left(\prod_{k=0}^{l-1} M'_{a_k}(d_k)\right)D'(d_l) &= C'\left(\prod_{k=0}^{l-1} M'_{a_k}(d_k)\right)Q_{S^\perp}D(d_l) \\
&= C\left(\prod_{k=0}^{l-1} M_{a_k}(d_k)\right)P_{S^\perp}Q_{S^\perp}D(d_l) \\
&= C\left(\prod_{k=0}^{l-1} M_{a_k}(d_k)\right)D(d_l).
\end{aligned}$$

To prove the stated equations, we proceed by induction on l. If $l = 0$, then

$$CP_{S^\perp} = C'$$

is simply the definition of C'. Suppose now that the stated equations hold for all delayed traces of length l or less, and consider a delayed trace:

$$d_0 \xrightarrow{a_0} d_1 \xrightarrow{a_1} \ldots \xrightarrow{a_l} d_{l+1}.$$

Then, using the induction hypothesis and the definition of $M'_{a_l}(d_l)$, we have

$$\begin{aligned} C'(\prod_{k=0}^{l} M'_{a_k}(d_k)) &= C'(\prod_{k=0}^{l-1} M'_{a_k}(d_k))M'_{a_l}(d_l) \\ &= C(\prod_{k=0}^{l-1} M_{a_k}(d_k))P_{S^\perp}Q_{S^\perp}M_{a_l}(d_l)P_{S^\perp}. \end{aligned}$$

But, as we have observed, $P_{S^\perp}Q_{S^\perp}M_{a_l}(d_l)P_{S^\perp} = M_{a_l}(d_l)P_{S^\perp}$, hence the last term above is equal to:

$$C(\prod_{k=0}^{l} M_{a_k}(d_k))P_{S^\perp},$$

as required. ∎

Lemma 21 *Suppose $(C, D(x), \{M_a(x) : a \in E\})$ is an n-dimensional representation of an observable $\Phi \in \mathrm{Obs}(E)$. Suppose the system of all equations of the form:*

$$C(\prod_{k=0}^{l-1} M_{a_k}(d_k))Y = 0$$

has a nontrivial solution Y. Let m be the dimension of the solution space. Then Φ has an $(n-m)$-dimensional representation $(C', D'(x), \{M'_a(x) : a \in E\})$, which is effectively computable from the given n-dimensional representation.

Proof – Using Lemma 19, we can compute a basis $\mathcal{C} = \{C_1, C_2, \ldots, C_m\}$ (of column vectors) for the solution space S of the above system of equations and a basis $\mathcal{B} = \{B_1, B_2, \ldots, B_{n-m}\}$ for the orthogonal complement $S^\perp$ of S. Assume that the basis $\mathcal{B}$ is orthonormal, which can be ensured using the Gram-Schmidt procedure. Let $P_{S^\perp}$ be the $((n-m) \times n)$-dimensional matrix of the projection of $\mathcal{R}^n$ (column vectors) to $S^\perp$, with respect to the basis $\mathcal{B}$ and the natural basis for $\mathcal{R}^n$. Explicitly, for $1 \leq i \leq n-m$, the ith row of the matrix $P_{S^\perp}$ consists of the components of the basis vector B_i. Let $Q_{S^\perp} = (P_{S^\perp})^{\mathrm{t}}$, the $(n \times (n-m))$-dimensional matrix of the embedding of $S^\perp$ in $\mathcal{R}^n$, with respect to the basis $\mathcal{B}$ and the natural basis for $\mathcal{R}^n$.

We now define $C' = CQ_{S^\perp}$, $M'_a(x) = P_{S^\perp}M_a(x)Q_{S^\perp}$, and $D'(x) = P_{S^\perp}D(x)$. We claim that that $(C', D'(x), \{M'_a(x) : a \in E\})$ is also a representation of Φ. The proof is essentially a "time-reversed" version of the proof of Lemma 20. We use the invariance of S under left multiplication by $M_a(x)$ to establish the relations:

$$P_{S^\perp}M_a(x) = P_{S^\perp}M_a(x)Q_{S^\perp}P_{S^\perp},$$

$$C = CQ_{S^\perp}P_{S^\perp},$$

then use these relations to prove that:

$$P_{S^\perp}\left(\prod_{k=0}^{l-1} M_{a_k}(d_k)\right) D(d_l) \;=\; \left(\prod_{k=0}^{l-1} M'_{a_k}(d_k)\right) D'(d_l) \tag{2}$$

for all delayed traces

$$d_0 \xrightarrow{a_0} d_1 \xrightarrow{a_1} \dots \xrightarrow{a_{l-1}} d_l.$$

The stated equations follow from this. ∎

Theorem 5 *There exists an algorithm that, given an n-dimensional representation*

$$(C, D(x), \{M_a(x) : a \in E\})$$

of an observable $\Phi \in \mathrm{Obs}(E)$, outputs an m-dimensional representation $(C', D'(x), \{M'_a(x) : a \in E\})$ of Φ, which is minimal.

Proof –

1. Determine the space of solutions X to the system of all equations of the form:

$$X(\prod_{k=0}^{l-1} M_{a_k}(d_k))D(d_k) = 0.$$

   If this space is nontrivial, use Lemma 20 to produce an n'-dimensional representation of Φ, where $n' < n$.

2. Determine the space of solutions Y to the system of all equations of the form:

$$C(\prod_{k=0}^{l-1} M_{a_k}(d_k))Y = 0$$

   If this space is nontrivial, use Lemma 21 to produce an n''-dimensional representation of Φ, where $n'' < n'$.

Let

$$(C'', D''(x), \{M''_a(x) : a \in E\})$$

denote the n''-dimensional representation resulting from step (2). Letting $P_{S^\perp}$ and $Q_{S^\perp}$ denote the projection and embedding matrices constructed in performing step (2), and using the property $P_{S^\perp} M'_a(x) Q_{S^\perp} P_{S^\perp} = P_{S^\perp} M'_a(x)$, for all $a_k \in E$ and all nonnegative $d_k \in \mathcal{R}^{n''}$ we have:

$$\begin{aligned} X(\prod_{k=0}^{l-1} M''_{a_k}(d_k))D''(d_k) &= X(\prod_{k=0}^{l-1} P_{S^\perp} M'_{a_k}(d_k) Q_{S^\perp}) P_{S^\perp} D'(d_k) \\ &= X P_{S^\perp}(\prod_{k=0}^{l-1} M'_{a_k}(d_k)) D'(d_k). \end{aligned}$$

Thus, if the system of equations:

$$X(\prod_{k=0}^{l-1} M''_{a_k}(d_k))D''(d_k) = 0,$$

had a nontrivial solution X, then $XP_{S^\perp}$ would satisfy:

$$(XP_{S^\perp})(\prod_{k=0}^{l-1} M'_{a_k}(d_k))D'(d_k) = 0.$$

Since the mapping from $\mathcal{R}^{n''}$ to $\mathcal{R}^{n'}$ defined by multiplication on the right by $P_{S^\perp}$ is injective (recall that multiplication on the left by $P_{S^\perp}$ is a projection), it follows that $XP_{S^\perp} \neq 0$ if and only if $X \neq 0$. Since step (1) guarantees that we cannot have any $XP_{S^\perp} \neq 0$ satisfying the above equations, it follows that step (2) does not introduce any additional possibility of nontrivial solutions X to the system of the form (1). Since the n''-dimensional representation resulting from step (2) thus satisfies the conditions of Lemma 16, it is minimal. ∎

3.6 Example

To illustrate the minimization algorithm, we now apply it to the 8-dimensional representation of the observable $\mathcal{B}^A_{E_A}\Omega_T$ obtained in Section 3.4, to obtain a minimal, 3-dimensional representation for this same observable.

Recall the 8-dimensional representation from Section 3.4:

$$C_{\Omega_T} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \qquad D(x) = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}$$

$$M_a(x) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & \frac{dp}{x+d} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{dp}{(x+d)^2} & \frac{dp}{x+d} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$M_t(x) = \begin{pmatrix} \frac{d(1-p)}{x+d} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{d(1-p)}{(x+d)^2} & \frac{d(1-p)}{x+d} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

We first consider the system of equations:

$$X(\prod_{k=0}^{l-1} M_{a_k}(d_k))D(d_k) = 0.$$

We begin by observing that the set:

$$\mathcal{D}_0 = \{D_1\}$$

where

$$D_1 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}$$

is a basis for the orthogonal complement $S_0^{\perp}$ of

$$S_0 = \{X \in \mathcal{R}^8 : XD(x) = 0,\ \text{all } x \in \mathcal{R}\}.$$

To obtain S_1, we solve simultaneously:

$$XD_1 = 0, \qquad XM_a(x)D_1 = 0, \qquad XM_t(x)D_1;$$

that is,

$$X \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} = 0, \qquad X \begin{pmatrix} \frac{dp}{x+d} \\ \frac{dp}{(x+d)^2} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} = 0, \qquad X \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} = 0.$$

Obviously, the third equation is vacuous. In view of the fact that the two components of the vector in the second equation are independent rational functions, a basis for the orthogonal complement of the space S_1 is as follows:

$$\mathcal{D}_1 = \{D_1, D_2, D_3\},$$

where

$$D_2 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \qquad D_3 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

We now proceed to solve simultaneously the system of identities:

$$XD_1 = 0, \qquad XD_2 = 0, \qquad XD_3 = 0,$$

$$XM_a(x)D_1 = 0, \qquad XM_a(x)D_2 = 0, \qquad XM_a(x)D_3 = 0,$$
$$XM_t(x)D_1 = 0, \qquad XM_t(x)D_2 = 0, \qquad XM_t(x)D_3 = 0.$$

The equations involving D_1 are the same as we had at the previous stage. The first of the $M_t(x)$ equations is vacuous, as are the second two of the $M_a(x)$ equations. The remaining new equations are:

$$X \begin{pmatrix} \frac{d(1-p)}{x+d} \\ \frac{d(1-p)}{(x+d)^2} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} = 0, \qquad X \begin{pmatrix} 0 \\ \frac{d(1-p)}{x+d} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} = 0.$$

It is easily seen that the solution space of these equations is the same as the space S_1, so that a fixed point S has been reached, which is the (5-dimensional) space of solutions of the original system. A basis for the complement of $S^\perp$ is $\mathcal{D}_1$, which after orthonormalization gives:

$$D_1' = \begin{pmatrix} 0 \\ 0 \\ \frac{\sqrt{2}}{2} \\ 0 \\ 0 \\ 0 \\ \frac{\sqrt{2}}{2} \\ 0 \end{pmatrix}, \qquad D_2' = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \qquad D_3' = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

The projection and embedding matrices $P_{S^\perp}$ and $Q_{S^\perp}$ are as follows:

$$P_{S^\perp} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ \frac{\sqrt{2}}{2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ \frac{\sqrt{2}}{2} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \qquad Q_{S^\perp} = \begin{pmatrix} 0 & 0 & \frac{\sqrt{2}}{2} & 0 & 0 & 0 & \frac{\sqrt{2}}{2} & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

We now compute a new, three-dimensional representation:

$$C' = CP = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix} \qquad D' = QD = \begin{pmatrix} \sqrt{2} \\ 0 \\ 0 \end{pmatrix}$$

$$M_a'(x) = QM_a(x)P = \begin{pmatrix} 0 & 0 & 0 \\ \frac{\sqrt{2}}{2}\frac{dp}{x+d} & 0 & 0 \\ \frac{\sqrt{2}}{2}\frac{dp}{(x+d)^2} & 0 & 0 \end{pmatrix} \qquad M_t'(x) = QM_t(x)P = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \frac{d(1-p)}{x+d} & 0 \\ 0 & \frac{d(1-p)}{(x+d)^2} & \frac{d(1-p)}{x+d} \end{pmatrix}.$$

If we attempt further reduction on the above representation, by solving the system:

$$C'\left(\prod_{k=0}^{l-1} M_{a_k}'(d_k)\right)Y = 0,$$

we find that there are no nontrivial solutions. Thus, we have found a minimal representation for the originally given observable.

References

[BDG98] M. Bernardo, L. Donatiello, and R. Gorrieri. A formal approach to the integration of performance aspects in the modeling and analysis of concurrent systems. *Information and Computation*, 1998. To appear.

[BR84] J. Berstel and C. Reutenauer. *Rational Series and Their Languages*, volume 12 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, 1984.

[CCM97] S. Campos, E. M. Clarke, and M. Minea. Symbolic techniques for formally verifying industrial systems. *Science of Computer Programming*, 29:79–88, 1997.

[CP71] J. W. Carlyle and A. Paz. Realizations by stochastic finite automata. *J. Comput. System Sci.*, 5:26–40, 1971.

[GHR93] N. Götz, U. Herzog, and M. Rettelbach. Multiprocessor and distributed system design: the integration of functional specification and performance analysis using stochastic process algebras. In *Proceedings of PERFORMANCE '93*, pages LNCS 729:121–146. Springer-Verlag, 1993.

[Hil96] J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.

[How71] Ronald A. Howard. *Dynamic Probabilistic Systems, Volume II: Semi-Markov and Decision Processes*, volume 2 of *Series in Decision and Control*. John Wiley & Sons, 1971.

[LSS94] N. Lynch, I. Saias, and R. Segala. Proving time bounds for randomized distributed algorithms. In *Proceedings of the 13th ACM PODC*, Los Angeles, CA, August 1994.

[MBC84] M. Ajmone Marsan, G. Balbo, and G. Conte. A class of generalized stochastic Petri nets for the performance evaluation of multiprocessor systems. *ACM Trans. on Computer Systems*, 2:143–172, 1984.

[Pri96] C. Priami. Stochastic π-calculus with general distributions. In M. Ribaudo, editor, *Proceedings of the Fourth Workshop on Process Algebras and Performance Modelling*, Torino, Italy, 1996.

[PS95] A. Pogosyants and R. Segala. Formal verification of time properties for randomized distributed algorithms. In *Proceedings of the 14th ACM PODC*, Ottawa, Ontario, Canada, August 1995.

[Sch61a] M. P. Schützenberger. On a special class of recurrent events. *Annals Math. Stat.*, 32:1201–1213, 1961.

[Sch61b] M. P. Schützenberger. On the definition of a family of automata. *Information and Control*, 4:245–270, 1961.

[WSS97] S.-H. Wu, S. A. Smolka, and E. W. Stark. Composition and behaviors of probabilistic I/O automata. *Theoretical Computer Science*, 176(1-2):1–38, 1997.

[Wu96] S.-H. Wu. *Composition and Behaviors of Probabilistic Input/Output Automata.* PhD thesis, SUNY at Stony Brook, 1996.