URL: http://www.elsevier.nl/locate/entcs/volume22.html 25 pages

On Generative Parallel Composition [1]

Pedro R. D'Argenio [a] Holger Hermanns [a] Joost-Pieter Katoen [b]

[a] *Department of Computer Science, University of Twente*
P.O. Box 217, 7500 AE Enschede, The Netherlands
{dargenio, hermanns}@cs.utwente.nl

[b] *Lehrstuhl für Informatik 7*
Friedrich-Alexander-Universität Erlangen-Nürnberg
Martensstraße 3, 91058 Erlangen, Germany
katoen@informatik.uni-erlangen.de

Abstract

A major reason for studying probabilistic processes is to establish a link between a formal model for describing functional system behaviour and a stochastic process. Compositionality is an essential ingredient for specifying systems. Parallel composition in a probabilistic setting is complicated since it gives rise to non-determinism, for instance due to interleaving of independent autonomous activities. This paper presents a detailed study of the resolution of non-determinism in an asynchronous generative setting. Based on the intuition behind the synchronous probabilistic calculus PCCS we formulate two criteria that an asynchronous parallel composition should fulfill. We provide novel probabilistic variants of parallel composition for CCS and CSP and show that these operators satisfy these general criteria, opposed to most existing proposals. Probabilistic bisimulation is shown to be a congruence for these operators and their expansion is addressed.

Key words: bisimulation; bundle transition systems; CCS; CSP; probabilistic process algebra; PCCS; semantics

1 Introduction

In the last decade the study of probabilistic processes using formal methods has received significant attention. A major reason for studying probabilistic processes is to establish a link between a formal model for describing functional system behaviour and a stochastic process. In the setting of process algebras relations with several stochastic models have been established, such

[1] Supported by the NWO/SION project 612-33-006 and the System Validation Centre/CTIT.

as continuous-time and discrete-time Markov chains and generalised semi-Markov processes. So-called *probabilistic process algebras* incorporate a probabilistic choice operator $+_p$ such that in $P +_p Q$ process P is selected with probability p and Q with $1-p$. The underlying semantic model, a labelled transition system equipped with discrete probabilities, can be viewed as a discrete-time Markov chain.

In order to calculate performance measures it is essential that non-determinism is absent or resolved. Although for several reasons non-determinism is of significant importance for the specification of reactive systems, it under-specifies the quantities with which certain alternative computations can appear. A stochastic process therefore does not exhibit non-determinism. Basically two approaches have been pursued to overcome this different treatment of non-determinism. In the *alternating* approach [9] both non-deterministic and probabilistic transitions are allowed. The outgoing transitions of a state are either all probabilistic or all non-deterministic. For performance analysis the present non-determinism is resolved using schedulers [19]. In the *generative* approach non-determinism is ruled out by means of a probability distribution that assigns a probability to each possible action. Since non-determinism is absent, a generative probabilistic transition system is easily converted into a discrete-time Markov chain. (Reactive and stratified [8] and simple and fully probabilistic transition systems [16] are variants or combinations of these two approaches.) This paper considers the generative setting.

Compositional specification of generative probabilistic transition systems is, however, to be treated carefully. Parallel composition in a generative setting is complicated since it gives rise to non-determinism, for instance due to interleaving of independent autonomous activities. To overcome these problems one typically resorts to some synchronous parallel composition in which one avoids to make a scheduling decision of independent processes, since all components must proceed in a 'lock-step' fashion. This is the approach of PCCS, the well-established probabilistic variant of Milner's synchronous version of CCS [5]. Since we do not want to stay in such a strict synchronous setting we take a different route.

In this paper we consider asynchronous generative processes and discuss the resolution of non-determinism in this setting. Based on the intuitions behind PCCS we formulate criteria (with respect to a congruence relation) that an asynchronous parallel composition should fulfill. For most existing generative parallel composition operators we show that they do not fulfill the criteria, and thus have a rather weak connection to the PCCS-approach. A notably exception is parallel composition in probabilistic ACP [2]. We argue that this calculus cannot be reduced to an appropriate asynchronous probabilistic CCS- or CSP-algebra, in contrast to the non-probabilistic case. Therefore, we provide probabilistic variants of parallel composition for CCS [12] and (to complete the picture) CSP [10,14] and show that these novel operators satisfy the criteria. The resulting calculi can be considered as asynchronous variants of PCCS. Probabilistic bisimulation is shown to be a congruence for these operators and the expansion law for asynchronous probabilistic CCS and CSP is addressed. We show that re-normalisation, a phenomenon that appears when certain alternatives become impossible, has to redistribute the probability mass that a deadlock can appear. Unlike several other approaches for asynchronous generative processes [13,6] this is identical to the interpretation of re-normalisation in PCCS (for restriction).

The organisation of this paper is as follows. Section 2 introduces generative probabilistic transition systems, PCCS, and discusses intuitively how the PCCS approach is expected to be transferred to asynchronous CCS and CSP. The basic ingredients of this discussion, so-called *bundle transition systems*, are formalised in Section 3 and two criteria for parallel composition, *respectfulness* and *stochasticity*, are defined. Section 4 makes our ideas for asynchronous probabilistic CCS and CSP concrete and presents our technical results. Section 5 discusses several existing generative parallel composition operators using the criteria introduced in Section 3. Finally, Section 6 concludes the paper. Proofs of the most important results are provided in the appendix.

2 Motivation

In this section we discuss the conditions that an appropriate definition of the parallel composition of generative probabilistic transition systems should satisfy. To do so, we first review the concepts of generative probabilistic transition systems and the definition of parallel composition in a synchronous way, i.e. à la SCCS. Based on these concepts, we informally discuss how a parallel composition à la CCS or CSP should look like in a generative setting.

2.1 A synchronous calculus for generative probabilistic systems

A generative probabilistic model

A discrete probability space is a structure (Ω, Pr) where Ω is a discrete sample space and Pr a probability measure on 2^{Ω}. Let $\mathsf{Prob}(H)$, for some universe H, be the set of discrete probability spaces with $\Omega \subseteq H$. The following definition is basically adopted from [8] phrased in a style that fits our purpose.

Definition 2.1 *$\mathcal{G} = (\Sigma, \mathbf{A}, \mathbf{I}, T)$ is a* generative probabilistic transition system *(GPTS for short) with Σ, a set of* states, *$\mathbf{A}$, a set of* actions, *$\mathbf{I}$, a set of* indices, *and $T : \Sigma \nrightarrow \mathsf{Prob}(\,(\mathbf{A} \times \mathbf{I} \times \Sigma) \cup \{\varnothing\}\,)$, a* probabilistic transition function, *such that the following condition is satisfied: if $T(s) = (\Omega_s, \mathsf{Pr}_s)$ then $(a, i, t), (b, i, u) \in \Omega_s \implies a = b \wedge t = u$.*

Here $\nrightarrow$ denotes a partial function. The constraint requires that each element in the sample space of $T(s)$ is uniquely identifiable through the index.

If $T(s)$ is defined we denote by Ω_s its sample space and by Pr_s its probability measure, i.e. $T(s) = (\Omega_s, \mathsf{Pr}_s)$; if $T(s)$ is not defined we say that s is an *endpoint*. Let $\mathsf{Pr}_s(a, i, s')$ denote $\mathsf{Pr}_s(\{(a, i, s')\})$ and $s \xrightarrow{a,p}_i s'$ denote $\mathsf{Pr}_s(a, i, s') = p$. The purpose of the index i in the probabilistic transition $\longrightarrow_i$ is to distinguish occurrences of the same probabilistic transition, and is standard in a probabilistic setting [8]. A GPTS $\mathcal{G}$ is called *stochastic* if for all states s we have $\varnothing \notin \Omega_s$; otherwise it is called *sub-stochastic*. If $\varnothing \in \Omega_s$ then $\mathsf{Pr}_s(\varnothing)$ can be considered as the probability to deadlock.

Synchronous probabilistic CCS

The reference language for our discussion is PCCS, the well-accepted probabilistic variant of synchronous CCS introduced by [5]. In PCCS atomic actions form a commutative semi-group $(\mathbf{A}, \cdot)$ generated from the set of basic actions $\{ a, b, c, \ldots\}$. Thus all elements of $\mathbf{A}$ are of the form a or $\alpha \cdot \beta$ with $\alpha, \beta \in \mathbf{A}$. The atomic action $\alpha \cdot \beta$ can be considered as the simultaneous (unordered) occurrence of actions α and β. Let X be a process variable, $A \subseteq \mathbf{A}$, and $f : A \to A$. The syntax of PCCS is

$$P ::= X \;\Big|\; \sum_{i\in I} [p_i]a_i.P \;\Big|\; P \times P \;\Big|\; P\backslash A \;\Big|\; P[f] \mid \mathbf{fix}\ X.P$$

such that $\sum_{i\in I} p_i = 1$, $p_i \in (0, 1]$ and I is a finite set of indices.

The term $\sum[p_i]a_i.P_i$ offers a probabilistic choice among the prefixes $a_i.P_i$. It performs a_i with probability p_i and then behaves like P_i. To be more precise, we should say at least probability p_i since there might be identical summands with distinct indices. (Action-prefix and probabilistic choice have been separated originally in PCCS [5].) For $I = \varnothing$ let $\sum[p_i]a_i.P_i = \mathbf{0}$, the process that cannot perform any action. $P \times Q$ represents synchronous parallel composition, and $P\backslash A$ a process that behaves like P except that actions in A are disallowed. (This operator is the dual of restriction in PCCS.) The term $P[f]$ denotes a process that behaves like P except that actions are renamed according to f. $\mathbf{fix}\ X.P$ defines a recursive process X by P, that possibly contains occurrences of X.

The operational semantics of PCCS is given in Table 1. Here,

$$\nu(P, A) =_{df} 1 - \sum_i \{\!|p \mid P \xrightarrow{a,p}_i P',\ a \in A|\!\}.$$

The inference rules determine a mapping of PCCS terms onto GPTSs. The rules for most operators are self-explanatory. The rule for $P\backslash A$ uses the function ν for normalisation of probabilities. In the definition of ν, $\{\!|\,|\!\}$ denotes a multi-set. The role of ν is extensively discussed below. Since $\backslash A$ and $\times$ will become important for the definition of our probabilistic calculi later on, we discuss these operators more extensively.

$$\sum_{i\in I}[p_i]a_i.P_i \xrightarrow{a_j,p_j}_j P_j \quad (j\in I) \qquad\qquad \frac{P \xrightarrow{a,p}_i P'}{P[f] \xrightarrow{f(a),p}_i P'[f]}$$

$$\frac{P \xrightarrow{a,p}_i P' \quad Q \xrightarrow{b,q}_j Q'}{P\times Q \xrightarrow{ab,pq}_{(i,j)} P'\times Q'} \qquad\qquad \frac{P[\mathbf{fix}\ X.P/X] \xrightarrow{a,p}_i P'}{\mathbf{fix}\ X.P \xrightarrow{a,p}_i P'}$$

$$\frac{P \xrightarrow{a,p}_i P'}{P\backslash A \xrightarrow{a,\frac{p}{\nu(P,A)}}_i P'\backslash A} \quad (a\notin A)$$

Table 1
Operational semantics of PCCS

Restriction
Consider $P = [\frac{1}{6}]a.\mathbf{0} + [\frac{1}{2}]b.\mathbf{0} + [\frac{1}{3}]c.\mathbf{0}$, a process that can either perform action a, b or c with probability $\frac{1}{6}$, $\frac{1}{2}$ and $\frac{1}{3}$, respectively. The corresponding GPTS of P is depicted in Figure 1(a). For convenience we omit transition indices. Consider the transition $P \xrightarrow{c,\frac{1}{3}} \mathbf{0}$. The value $\frac{1}{3}$ denotes the probability that

(a) (b)

Fig. 1. GPTSs for (a) $P = [\frac{1}{6}]a.\mathbf{0} + [\frac{1}{2}]b.\mathbf{0} + [\frac{1}{3}]c.\mathbf{0}$ and (b) $P\backslash c$

P *intends to* perform action c. We deliberately say "intends to perform" rather than "performs": when P is considered in a context that is not able to participate in c, action c is prohibited even if P intends to perform it. In such a case the probability to perform c is 0, and its (local) probability $\frac{1}{3}$ needs to be redistributed among the remaining possible actions. This principle is applied when P is considered in the context of $\backslash A$ where A contains c. For instance, consider $P\backslash c$. In principle, the behaviour of $P\backslash c$ is determined by the *conditional* probabilities of the following three situations:

(i) P performs a, provided that P does not perform c.

(ii) P performs b, provided that P does not perform c.

(iii) P performs c, provided that P does not perform c.

Thus, the probabilities in $P\backslash c$ are conditioned to the fact that P does not perform c, and clearly, the third option has probability 0. Accordingly, we obtain for the probability of performing a:

$$\mathsf{Pr}_{P\backslash c}(a, \mathbf{0}\backslash c) = \mathsf{Pr}_{P}(a, \mathbf{0} \mid \neg(c, \mathbf{0})) = \frac{1}{4}$$

where we have omitted transition indices for convenience. Similarly, we obtain that the probability of performing b (c) is $\frac{3}{4}$ (resp. 0). The resulting GPTS for $P\backslash c$ is depicted in Figure 1(b). The probability of not performing c, $1-\frac{1}{3}$, is the *normalisation factor* $\nu(P, \{c\})$. In general, the normalisation factor $\nu(P, A)$ denotes the probability that P does not perform actions in A[2]. This interpretation of normalisation will be adopted for asynchronous probabilistic CSP later on.

The principle of normalisation can intuitively be explained as follows: a process probabilistically selects one of its alternatives repeatedly, until the selected transition can actually be taken. In case of process $P\backslash c$, it means that if the outcome of the experiment is (the prohibited) action c, then a subsequent experiment is carried out, until c is not selected. Accordingly, the resulting probability with which, for instance, action a happens is given by $\frac{1}{6}$ (choose a in the first experiment) plus $\frac{1}{3} \cdot \frac{1}{6}$ (first select c and then a) plus $\left(\frac{1}{3}\right)^2 \cdot \frac{1}{6}$ (select c twice, and then a), and so on. So,

$$\mathsf{Pr}_{P\backslash c}(a, \mathbf{0}\backslash c) = \frac{1}{6} \cdot \sum_{k=0}^{\infty} \left(\frac{1}{3}\right)^k = \frac{1}{6} \cdot \frac{1}{1-\frac{1}{3}}$$

which indeed equals $\frac{1}{4}$, the result obtained above by applying normalisation.

Synchronous parallel composition

Let $Q = [\frac{1}{2}]a.\mathbf{0} + [\frac{1}{2}]b.\mathbf{0}$ and $R = [\frac{1}{3}]\overline{a}.\mathbf{0} + [\frac{2}{3}]c.\mathbf{0}$, and consider the construction of $Q \times R$. Since Q and R intend to perform two actions each, four possible scenarios result: Q performs a and R performs $\overline{a}$, Q performs a and R performs action c, etcetera. The probabilities of the transitions of $Q \times R$ are simply determined by the *product* of the probabilities of the constituents. This is based on the fact that probabilistic choices of Q and R are stocastically independent. The GPTSs of Q, R and $Q \times R$ are depicted in Figure 2(a), (b) and (c), respectively.

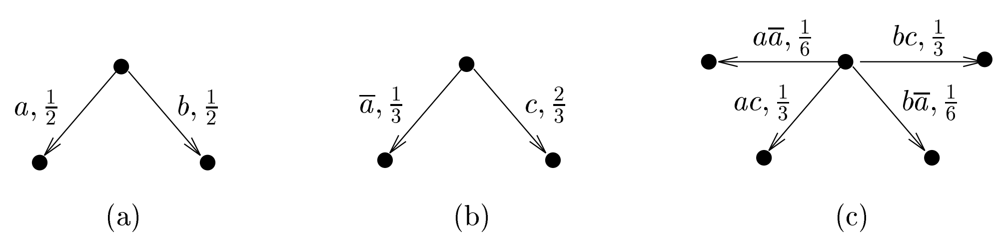

Fig. 2. GPTSs of (a) $Q = [\frac{1}{2}]a.\mathbf{0} + [\frac{1}{2}]b.\mathbf{0}$, (b) $R = [\frac{1}{3}]\overline{a}.\mathbf{0} + [\frac{2}{3}]c.\mathbf{0}$ and (c) $Q \times R$

[2] Notice that $\nu(P, A)$ is only used in Table 1 for restriction with the precondition that P can perform some action $a \notin A$. This guarantees that $\nu(P, A) > 0$ for all used cases.

2.2 Asynchronous probabilistic parallel composition

Basically, two different kinds of parallel composition have been defined in process algebra: one à la CCS [12], in which any action can be performed independently by each process, and besides, the processes can synchronise if they are allowed, and the other, à la CSP [10,14], in which actions that are intended to synchronise are listed in a synchronisation set and can only be performed synchronously, and the other actions are performed always independently. In the following, we investigate how these operators should look like in generative PTSs following the line of thought of PCCS discussed above.

CCS

For (non-probabilistic) CCS, parallel composition, denoted by $|$, is defined by the following inference rules:

$$\frac{P \xrightarrow{a} P'}{P|Q \xrightarrow{a} P'|Q} \qquad \frac{P \xrightarrow{a} P'}{P|Q \xrightarrow{a} P|Q'} \qquad \frac{P \xrightarrow{a} P' \quad Q \xrightarrow{\overline{a}} Q'}{P|Q \xrightarrow{\tau} P'|Q'} \quad (a \neq \tau)$$

Here τ denotes a distinguished action that models internal activity. Notice that processes are not forced to synchronise; they can equally well autonomously perform actions that could be synchronised.

In order to motivate our ideas concerning a probabilistic version of $|$, consider the processes $Q = [\frac{1}{2}]a.\mathbf{0} + [\frac{1}{2}]b.\mathbf{0}$ and $R = [\frac{1}{3}]\overline{a}.\mathbf{0} + [\frac{2}{3}]c.\mathbf{0}$ of Figure 2(a) and (b), respectively. Like for the synchronous case, four different scenarios for $Q|R$ may arise: Q performs a and R performs $\overline{a}$, Q performs a and R action c, etcetera. The probabilities for these scenarios are simply determined by the product of the probabilities of the involved actions in Q and R, in analogy to the synchronous case. The difference with the synchronous case, however, is that actions are executed asynchronously. That is, the occurrence of e.g. a and c does no longer constitute a single atomic action (but two). As a result there are different ways in which a given scenario occurs. For instance, the scenario that Q performs a and R does c can be obtained — through interleaving — by first performing a followed by c, or in the reverse order. The probabilities of these sub-scenarios are unspecified (i.e., they are non-deterministic); we only know that together they have a probability $\frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3}$.

Due to the nature of CCS parallel composition the scenario "Q performs a and R does $\overline{a}$" can be established in three ways: the two possible ways of interleaving a and $\overline{a}$ and the possibility of synchronising a and $\overline{a}$, yielding τ. Once more, the probabilities of the individual sub-scenarios are unknown; together they have probability $\frac{1}{2} \cdot \frac{1}{3} = \frac{1}{6}$. Figure 3 depicts the transition system that results if we apply a similar reasoning to all possible scenarios. In the picture we have grouped with a small connecting line the different transitions that constitute a single scenario. The attached probabilities are associated to these "bundles" of transitions.

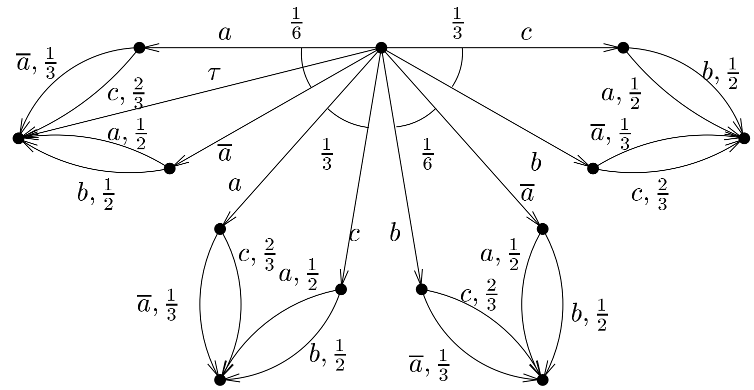

Fig. 3. The "bundle" probabilistic transition system for $Q|R$

This example suggests that an appropriate probabilistic version of CCS parallel composition should preserve the bundle probabilities. If this is not the case, then probabilities of the autonomous moves are in some way weighted, which is not intended. So,

> *an appropriate probabilistic CCS parallel composition should only quantify the unresolved non-determinism and nothing else.*

The principle that lies behind the bundle construction is analog to the intuition behind restriction (and normalisation). That is, a process probabilistically selects one of its alternatives. If this alternative is not executed for some reason, either because the environment is not willing to participate, the action is prohibited (in case of restriction), or an autonomous move is selected rather than a potentially possible synchronisation (in case of CCS), then the process carries out a next experiment, until the process can actually perform a transition (if any). This scheme has also been applied to Figure 3. Consider, for instance, the leftmost bundle in this figure. In this case Q has selected a and R has selected $\bar{a}$ for execution. Suppose that a happens. Although process R intends to perform $\bar{a}$, this is prevented since Q autonomously performs a instead of proceeding synchronously together with R. According to the above principle — which directly has been adopted from the treatment of restriction in PCCS — R now carries out another experiment after the occurrence of a. Hence, it again has the choice between $\bar{a}$ and c.

CSP

Unlike parallel composition in CCS, actions that can be synchronised are forced to synchronise in CSP; those actions cannot autonomously be performed. The set A of synchronising actions is a parameter of parallel composition $||$. Its semantics is defined by the following inference rules [14]:

$$\frac{P \xrightarrow{a} P'}{P \,||_A\, Q \xrightarrow{a} P' \,||_A\, Q}\ (a \notin A) \qquad \frac{P \xrightarrow{a} P'}{Q \,||_A\, P \xrightarrow{a} Q \,||_A\, P'}\ (a \notin A)$$

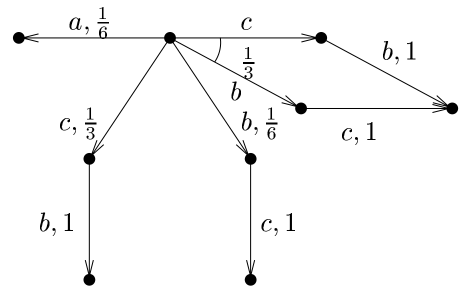

Fig. 4. The "bundle" probabilistic transition system of $Q \,||_{\{a\}}\, R'$

$$\frac{P \xrightarrow{a} P' \quad Q \xrightarrow{a} Q'}{P \,||_A\, Q \xrightarrow{a} P' \,||_A\, Q'} \; (a \in A)$$

Consider our example $Q \,||_{\{a\}} R'$ with $R' = [\frac{1}{3}]a.\mathbf{0} + [\frac{2}{3}]c.\mathbf{0}$ and Q as before. The four scenarios for this CSP-term are analog to those for the synchronous case (and CCS). However, due to the different synchronisation policy, the occurrence of some actions may be prohibited. For instance, if Q intends to perform a and R' wants to perform c, action a cannot occur since its occurrence requires participation of R'. Instead, action c can be performed autonomously with (in this case) probability $\frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3}$. The thus resulting PTS for $Q \,||_{\{a\}} R'$ is depicted in Figure 4 where we used the bundle notation introduced before. Like for the probabilistic variant of CCS parallel composition we conclude that

> *an appropriate probabilistic version of* $||_A$ *should only schedule the present non-determinism and nothing else.*

Due to the different synchronisation policy in CSP, the difference with CCS is twofold. Since synchronisation actions cannot be performed autonomously, the scenario that Q intends to perform a and R' wants to do a gives rise to a single case. There is no distinction to be made whether a synchronisation or an individual move takes place. This simplifies the definition of the probabilistic variant of $||_A$. On the other hand, however, CSP parallel composition may give rise to normalisation of probabilities. (This should not surprise the reader, since restriction can be described using CSP-style parallel composition [7]: $P \backslash A$ can, for instance, be encoded as $P \,||_A\, \mathbf{0}$.) This occurs, for instance, in our example above. Consider the term $Q \,||_{\{a\}}\, \mathbf{0}$ that is reached after R' performs c (with probability $\frac{1}{3}$). Now, Q may, with probability $\frac{1}{2}$, choose in favour of a, but synchronisation on a is permanently impossible. So, one might decide that a deadlock occurs with probability $\frac{1}{2}$. Inspired by the treatment of restriction in PCCS, we redistribute the probability mass of deadlocking among the remaining possibilities. This is depicted in Figure 4, where b occurs (after c) with probability 1. As we will show later on, this normalisation complicates the definition of the probabilistic variant of $||_A$.

3 Bundle probabilistic transition systems

In this section we formalise the kind of probabilistic transition systems discussed informally just above and we define the notion of appropriate parallel composition on this model. For set S let $\wp_{\mathrm{fin}}(S)$ denote the set of finite subsets of S.

Definition 3.1 *A bundle probabilistic transition system (BPTS, for short) is a quadruple* $\mathcal{B} = (\Sigma, \mathbf{A}, \mathbf{I}, T)$ *with* Σ*, a set of* states*,* $\mathbf{A}$*, a set of* actions*,* $\mathbf{I}$*, a set of* indices*, and* $T : \Sigma \rightarrow\!\!\!\!\rightarrow \mathsf{Prob}(\wp_{\mathrm{fin}}(\mathbf{A} \times \mathbf{I} \times \Sigma))$*, a* probabilistic transition function*, such that the following conditions are satisfied: if* $T(s) = (\Omega_s, \mathsf{Pr}_s)$ *then*

(i) $(a, i, t), (b, i, u) \in \bigcup \Omega_s \implies a = b \wedge t = u$*, and*

(ii) $\forall B, B' \in \Omega_s.\ B \cap B' \neq \varnothing \implies B = B'$*.*

The first constraint requires that indices uniquely determine transitions, in analogy to the constraint on GPTS (Def. 2.1). The second constraint requires that elements of Ω are pairwise disjoint.

Each probabilistic transition in a BPTS is a "bundle" of non-deterministic transitions as depicted previously. So, a certain set of non-deterministic alternatives is chosen with a certain probability. From this point of view, a BPTS is the converse of the simple model of [16] where probability distributions on successor states can be chosen non-deterministically. Both models are (action-labelled) simplified cases of the probabilistic finite-state programs of [15].

A BPTS is isomorphic to a GPTS, if all bundles are singletons (or empty), i.e. if for all states s the sample space Ω_s of $T(s)$ satisfies: $B \in \Omega_s$ implies $|B| \leq 1$. Such a BPTS is called *generative*. Let $T(s) = (\Omega_s, \mathsf{Pr}_s)$. For $B \in \Omega_s$ we abbreviate $\mathsf{Pr}_s(\{B\})$ by $\mathsf{Pr}_s(B)$ and let $s \xrightarrow{p} B$ denote that $\mathsf{Pr}_s(B) = p$. $\rightarrow$ is called a *bundle*. If for all states s the sample space Ω_s does not contain $\varnothing$ we call the BPTS *stochastic*, otherwise, it is called *sub-stochastic*. If $\varnothing \in \Omega_s$, the value $\mathsf{Pr}_s(\varnothing)$ can be considered as a deadlock or termination probability.

Parallel composition of BPTSs

Although BPTSs are an interesting model in themselves, it is not our intention to develop a complete theory around BPTSs in this paper but just to give the necessary tools to understand what is an appropriate definition of a probabilistic parallel composition. To do so, we first define a general parallel composition, denoted $\otimes$, on BPTSs. It constructs a full product of the involved BPTSs where transitions may always happen independently or synchronously (even if they are unequally labelled[3]). The general parallel composition cannot resolve the introduced non-determinism, so, the bundles

[3] This is similar to parallel composition in probabilistic ACP [2] for which the synchronisation function $\gamma : \mathbf{A} \times \mathbf{A} \rightarrow \mathbf{A}$ has to be total.

executed independently for each process are joined in a new bundle.

Definition 3.2 *Let* $(\mathbf{A}, \cdot)$ *be a commutative semi-group, and* $\mathcal{B}_1 = (\Sigma_1, \mathbf{A}, \mathbf{I}_1, T_1)$*,* $\mathcal{B}_2 = (\Sigma_2, \mathbf{A}, \mathbf{I}_2, T_2)$ *be two BPTSs defined over* $\mathbf{A}$*. The* general parallel composition *of* $\mathcal{B}_1$ *and* $\mathcal{B}_2$*, denoted* $\mathcal{B}_1 \otimes \mathcal{B}_2$*, is defined by* $(\Sigma_1 \times \Sigma_2, \mathbf{A}, \mathbf{I}, T)$ *where*

- $\mathbf{I} =_{df} \{(i,j], [i,j), (i,j) \mid i \in \mathbf{I}_1 \cup \{0\}, j \in \mathbf{I}_2 \cup \{0\}\}$*, and*
- T *is defined according to the rule*

$$\frac{s \xrightarrow{p} B_s \quad t \xrightarrow{q} B_t}{s \otimes t \xrightarrow{pq} B_s \otimes B_t}$$

with $s \otimes t$ *denoting* $(s,t) \in \Sigma_1 \times \Sigma_2$ *and* $B_s \otimes B_t$ *defined by*

$$\{(a, (i,j], s' \otimes t) \mid (a,i,s') \in B_s \wedge (\, (b,j,t') \in B_t \vee (B_t = \varnothing \wedge j = 0)\,)\}$$

$$\cup\ \{(b, [i,j), s \otimes t') \mid (\, (a,i,s') \in B_s \vee (B_s = \varnothing \wedge i = 0)\,) \wedge (b,j,t') \in B_t\}$$

$$\cup\ \{(ab, (i,j), s' \otimes t') \mid (a,i,s') \in B_s \wedge (b,j,t') \in B_t\}.$$

In the index of the transition relation, the parentheses indicate whether the left or right process moves (performs the action) and the square brackets indicate if the process remains passive. The fact that indices uniquely determine the individual transitions ensures that transitions are still uniquely determined in $\bigcup \Omega_{s \otimes t}$, and moreover, that elements in $\Omega_{s \otimes t}$ are pairwise disjoint. Moreover, we recall that we are dealing with discrete probability spaces, and hence our definition of T induces a unique probability measure. So, $\mathcal{B}_1 \otimes \mathcal{B}_2$ is indeed a BPTS. Remark that $\otimes$ does neither rule out any possible transition nor resolves any possibly introduced non-determinism.

Normalisation of BPTSs

Sometimes we are only interested in dealing with stochastic BPTSs (or GPTSs). Some operations may map a stochastic BPTS into a sub-stochastic BPTS. An example of this situation is, in fact, the restriction operation that we have discussed before. There, some transitions are pruned and the lost probability must be redistributed by means of normalisation. The process of normalisation is defined for BPTSs as follows.

Definition 3.3 *Let* $\mathcal{B} = (\Sigma, \mathbf{A}, \mathbf{I}, T)$ *be a BPTS. The* normalisation *of* $\mathcal{B}$ *is the stochastic BPTS* $\mathcal{N}(\mathcal{B}) =_{df} (\Sigma, \mathbf{A}, \mathbf{I}, T')$ *where, for all* $s \in \Sigma$*,* $T'(s)$ *is defined if and only if* $T(s)$ *is defined and* $\mathsf{Pr}_s(\varnothing) < 1$*. In such a case* $T'(s) =_{df} (\Omega_s - \{\varnothing\}, \mathsf{Pr}'_s)$ *with, for all* $B \in \Omega_s - \{\varnothing\}$*,*

$$\mathsf{Pr}'_s(B) =_{df} \textbf{ if } \varnothing \notin \Omega_s \textbf{ then } \mathsf{Pr}_s(B) \textbf{ else } \frac{\mathsf{Pr}_s(B)}{1 - \mathsf{Pr}_s(\varnothing)}.$$

It is straightforward to check that $\mathcal{N}(\mathcal{B})$ is indeed a BPTS. Since $\varnothing$ is not contained in any of the sample spaces of T', it follows that $\mathcal{N}(\mathcal{B})$ is stochastic;

i.e. there is no deadlocking possibility. The probability mass of deadlocking, $\mathsf{Pr}_s(\varnothing)$, is redistributed over the remaining bundles if appropriate.

Resolving non-determinism

To resolve non-determinism in a BPTS we introduce a simplified (and restricted) variant of adversary [19,16] that we call *determinisation.*

Definition 3.4 *A* determinisation *is a function* $\mathcal{D} : \mathsf{Prob}(\wp_{\text{fin}}(H)) \to \mathsf{Prob}(H \cup \{\varnothing\})$ *such that, if* $\mathcal{D}(\Omega, \mathsf{Pr}) = (\Omega', \mathsf{Pr}')$ *then* $\Omega' \subseteq \bigcup \Omega \cup \{\varnothing\}$ *and*

(i) $\sum_{\gamma \in B \cap \Omega'} \mathsf{Pr}'(\gamma) = \mathsf{Pr}(B)$*, provided* $B \in \Omega$ *and* $B \cap \Omega' \neq \varnothing$*, and*

(ii) $\mathsf{Pr}'(\varnothing) = \sum\{\mathsf{Pr}(B) \mid B \in \Omega, B \cap \Omega' = \varnothing\}$.

We call $\mathcal{D}$ a *determinisation* because it resolves non-determinism in bundles. Given a BPTS $\mathcal{B} = (\Sigma, \mathbf{A}, \mathbf{I}, T)$, its determinisation according to $\mathcal{D}$ is the GPTS $\mathcal{D}(\mathcal{B}) = (\Sigma, \mathbf{A}, \mathbf{I}, \mathcal{D} \circ T)$, where $\circ$ denotes ordinary function composition. The first constraint requires that the bundle probability in the BPTS $\mathcal{B}$ is equal to the sum of the probabilities of each element of that bundle in the determinised GPTS $\mathcal{D}(\mathcal{B})$. The second constraint determines that the probability of a deadlock in $\mathcal{D}(\mathcal{B})$ is the cumulated probability of having a bundle B that is eliminated by $\mathcal{D}$, that is, for which $B \cap \Omega' = \varnothing$.

Respectful and stochastic

Using parallel composition ($\otimes$), normalisation ($\mathcal{N}$) and determinisation ($\mathcal{D}$) we now formalise two general properties, called *respectfulness* and *stochasticity*, for probabilistic parallel composition. Let $\mathcal{G}_1 = (\Sigma_1, \mathbf{A}, \mathbf{I}_1, T_1)$ and $\mathcal{G}_2 = (\Sigma_2, \mathbf{A}, \mathbf{I}_2, T_2)$ be two GPTSs defined over $\mathbf{A}$, and let par be a parallel composition operator on GPTSs. (For the sake of generality, we define these concepts on the level of GPTSs, although the GPTSs that we consider are obtained from a probabilistic process algebra.)

Definition 3.5 *Operator* par *is* respectful *(for* $\sim$*) if* $\mathcal{G}_1 \;\mathsf{par}\; \mathcal{G}_2 \sim \mathcal{D}(\mathcal{G}_1 \otimes \mathcal{G}_2)$ *for some determinisation* $\mathcal{D}$*, semi-group* $(\mathbf{A}, \cdot)$ *and congruence* $\sim$*.*

A respectful parallel composition respects the bundle probabilities that are determined using $\otimes$. That is, the bundle probabilities are obtained by summation of the probabilities of the individual transitions in the bundle. A parallel composition operator is called stochastic if the bundle probabilities are respected after normalisation. Formally,

Definition 3.6 *For stochastic* $\mathcal{G}_1$ *and* $\mathcal{G}_2$*,* par *is* stochastic *(for* $\sim$*) if* $\mathcal{G}_1 \mathsf{par} \mathcal{G}_2 \sim \mathcal{N} \circ \mathcal{D}(\mathcal{G}_1 \otimes \mathcal{G}_2)$ *for some determinisation* $\mathcal{D}$*, semi-group* $(\mathbf{A}, \cdot)$ *and congruence* $\sim$*.*

In the probabilistic setting we often instantiate $\sim$ by probabilistic bisimulation [11]. This equivalence notion is defined for GPTSs as follows.

Definition 3.7 *For a GPTS $\mathcal{G} = (\Sigma, \mathbf{A}, \mathbf{I}, T)$ let the function $\mu : \Sigma \times \mathbf{A} \times \wp(\Sigma) \to [0,1]$ be defined by $\mu(s,a,C) =_{df} \sum_i \{\!|p \mid s \xrightarrow{a,p}_i s', s' \in C|\!\}$. An equivalence relation R on Σ is a* probabilistic bisimulation *if $s_1 \, R \, s_2$ implies for all $C \in \Sigma/R$ and $a \in \mathbf{A}$ that*

$$\mu(s_1, a, C) = \mu(s_2, a, C).$$

States s_1 and s_2 are (probabilistically) bisimilar*, denoted $s \sim_p t$, if there exists a probabilistic bisimulation R with $s_1 \, R \, s_2$. GPTSs $\mathcal{G}_1$ and $\mathcal{G}_2$ are* bisimilar*, notation $\mathcal{G}_1 \sim_p \mathcal{G}_2$, if their respective initial states are bisimilar on the disjoint union of $\mathcal{G}_1$ and $\mathcal{G}_2$.*

To illustrate the introduced concepts we provide the following (expected) result for PCCS:

Theorem 3.8 *$\times$ is respectful and stochastic for $\sim_p$.*

This result can be explained as follows. Let $(\mathbf{A}, .)$ be the same as for $\times$. In order to characterise determinisation $\mathcal{D}$ we first observe that for any two GPTSs composed according to $\otimes$, bundles are either: (1) complete, (2) incomplete, or (3) empty. A complete bundle is of the form

$$\{(a, (i,j], s' \otimes t), (b, [i,j), s \otimes t'), (ab, (i,j), s' \otimes t')\}.$$

Bundles $\{(a, (i,0], s' \otimes t)\}$ and $\{(b, [0,j), s \otimes t')\}$ are incomplete bundles. Let $\mathcal{D}(\Omega, \mathsf{Pr}) = (\Omega', \mathsf{Pr}')$ be defined by

- $\Omega' = \{(ab, (i,j), s \otimes t) \mid (ab, (i,j), s \otimes t) \in \bigcup \Omega\}$
  $\cup \ \{\varnothing \mid \exists B \in \Omega.\ B \text{ is not complete}\}$,
- $\mathsf{Pr}'(\gamma) = \mathsf{Pr}(B) \iff \gamma \in B$, and
- $\mathsf{Pr}'(\varnothing) = \sum \{\mathsf{Pr}(B) \mid B \text{ is not complete}\}$

By conditions imposed on BPTSs (see Definition 3.1) it follows that Pr' is a well-defined probability measure. It directly follows that $\mathcal{D}$ is a determinisation and that the equivalence closure of the relation

$$\{(s \times t, s \otimes t) \mid s \in \Sigma_1, t \in \Sigma_2\}$$

is a probabilistic bisimulation between $\mathcal{G}_1 \times \mathcal{G}_2$ and $\mathcal{D}(\mathcal{G}_1 \otimes \mathcal{G}_2)$, where every $s \otimes t$ is a state in the GPTS $\mathcal{D}(\mathcal{G}_1 \otimes \mathcal{G}_2)$. If s and t come from stochastic GPTSs, the same relation is a probabilistic bisimulation where every $s \otimes t$ is a state in the GPTS $\mathcal{N}(\mathcal{D}(\mathcal{G}_1 \otimes \mathcal{G}_2))$.

4 Asynchronous probabilistic CCS and CSP

In this section we introduce two composition operators that naturally correspond to CCS- respectively CSP-style parallel composition. Since we intend to avoid the synchrony assumption of PCCS we call them asynchronous composition operators.

Asynchronous probabilistic CCS

This language is obtained from PCCS by replacing the synchronous composition $\times$ by an operator ${}^{\theta}|^{\sigma}$ with two parameters $\sigma, \theta \in (0,1)$. Both parameters are conditional probabilities and can be considered as the relevant information for an adversary (or determinisation) to resolve the non-determinism that arises by putting two processes in parallel. The two probabilistic parameters σ and θ in the term $P\,{}^{\theta}|^{\sigma}\,Q$ are interpreted as follows. σ denotes the probability that P performs an autonomous action, given that both P and Q do not want to synchronise, and θ denotes the probability that some autonomous action occurs given that a synchronisation is possible. In other words, if a synchronisation is possible, it will take place with probability $1-\theta$. The formal semantics of APCCS is defined by the least relation satisfying the inference rules in Table 1, where the rules in Table 2 replace the rule for synchronous composition. Here we use $P \xrightarrow{a,p}_i$ as an abbreviation of $\exists P'.\, P \xrightarrow{a,p}_i P'$ and 'P *endpoint*' as an abbreviation of $\forall a,p,i.\, \neg(P \xrightarrow{a,p}_i\,)$. Note that σ does not play any role in the inference rule for synchronisation (last rule), while θ is irrelevant for the case in which a synchronisation cannot take place (first two rules). It is not difficult to check that GPTSs are closed under ${}^{\theta}|^{\sigma}$. For

$$\frac{P \xrightarrow{a,p}_i P' \qquad Q \xrightarrow{b,q}_j}{P\,{}^{\theta}|^{\sigma}\,Q \xrightarrow{a,pq\sigma}_{(i,j]} P'\,{}^{\theta}|^{\sigma}\,Q} \quad (b \neq \overline{a} \;\vee\; a = \tau)$$

$$\frac{Q \xrightarrow{a,p}_i \qquad P \xrightarrow{b,q}_j P'}{Q\,{}^{\theta}|^{\sigma}\,P \xrightarrow{b,pq(1-\sigma)}_{[i,j)} Q\,{}^{\theta}|^{\sigma}\,P'} \quad (b \neq \overline{a} \;\vee\; a = \tau)$$

$$\frac{P \xrightarrow{a,p}_i P' \qquad Q \text{ endpoint}}{P\,{}^{\theta}|^{\sigma}\,Q \xrightarrow{a,p}_{(i,0]} P'\,{}^{\theta}|^{\sigma}\,Q} \qquad \frac{Q \text{ endpoint} \qquad P \xrightarrow{a,q}_j P'}{Q\,{}^{\theta}|^{\sigma}\,P \xrightarrow{a,q}_{[0,j)} Q\,{}^{\theta}|^{\sigma}\,P'}$$

$$\frac{P \xrightarrow{a,p}_i P' \qquad Q \xrightarrow{\overline{a},q}_j}{P\,{}^{\theta}|^{\sigma}\,Q \xrightarrow{a,pq\theta\sigma}_{(i,j]} P'\,{}^{\theta}|^{\sigma}\,Q} \;(a \neq \tau) \qquad \frac{Q \xrightarrow{a,p}_i \qquad P \xrightarrow{\overline{a},q}_j P'}{Q\,{}^{\theta}|^{\sigma}\,P \xrightarrow{\overline{a},pq\theta(1-\sigma)}_{[i,j)} Q\,{}^{\theta}|^{\sigma}\,P'} \;(a \neq \tau)$$

$$\frac{P \xrightarrow{a,p}_i P' \qquad Q \xrightarrow{\overline{a},q}_j Q'}{P\,{}^{\theta}|^{\sigma}\,Q \xrightarrow{\tau,pq(1-\theta)}_{(i,j)} P'\,{}^{\theta}|^{\sigma}\,Q'} \quad (a \neq \tau)$$

Table 2
Operational semantics of APCCS parallel composition

APCCS we have the following technical results. The proofs of these facts can be found in the appendix.

Theorem 4.1 *$\sim_p$ is a congruence with respect to ${}^{\theta}|^{\sigma}$.*

Theorem 4.2 *${}^{\theta}|^{\sigma}$ is respectful and stochastic for $\sim_p$.*

Respectfulness can be seen intuitively as follows. In a bundle without synchronisation (e.g. the three bundles of cardinality two in Figure 3) one branch is assigned probability $pq\sigma$ and the other $pq(1-\sigma)$, together yielding the bundle probability pq. In case of a bundle with synchronisation these two probabilities are multiplied with θ, while the synchronisation itself gets probability $pq(1-\theta)$. Also in this case the probabilities sum up to the bundle probability pq. The APCCS parallel composition of stochastic processes is also a stochastic process, i.e., the composition of processes without deadlocks (or empty bundles) yields another process without deadlocks. Since normalisation does not have any effect on stochastic GPTSs, it also follows that ${}^{\theta}|^{\sigma}$ is stochastic. The use of σ and θ is reflected in the following expansion law.

Theorem 4.3 *Let $P = \sum_i [p_i]a_i.P_i$ and $Q = \sum_j [q_j]b_j.Q_j$ such that P, Q differ from $\mathbf{0}$. Then $P\,{}^{\theta}|^{\sigma}\,Q$ equals*

$$
\begin{aligned}
&\sum_{i,j}\{\,[r]\,\tau.(P_i\,{}^{\theta}|^{\sigma}\,Q_j) \mid \overline{a_i} = b_j \neq \tau, r = p_iq_j(1-\theta)\,\}\\
+&\sum_{i,j}\{\,[r]\,a_i.(P_i\,{}^{\theta}|^{\sigma}\,Q) \mid \overline{a_i} = b_j \neq \tau, r = p_iq_j\theta\sigma\,\}\\
+&\sum_{i,j}\{\,[r]\,b_j.(P\,{}^{\theta}|^{\sigma}\,Q_j) \mid \overline{a_i} = b_j \neq \tau, r = p_iq_j\theta(1-\sigma)\,\}\\
+&\sum_{i,j}\{\,[r]\,a_i.(P_i\,{}^{\theta}|^{\sigma}\,Q) \mid \overline{a_i} \neq b_j \vee a_i = \tau, r = p_iq_j\sigma\,\}\\
+&\sum_{i,j}\{\,[r]\,b_j.(P\,{}^{\theta}|^{\sigma}\,Q_j) \mid \overline{a_i} \neq b_j \vee a_i = \tau, r = p_iq_j(1-\sigma)\,\}.
\end{aligned}
$$

If Q equals $\mathbf{0}$ we have $P\,{}^{\theta}|^{\sigma}\,Q$ equals P. Similarly for P equals $\mathbf{0}$. Notice that the probability θ only is of importance if a synchronisation is possible, that is if $\overline{a_i} = b_j$. The probability σ plays a role only if a process performs an autonomous action and is irrelevant in case of synchronisation (first summation).

Asynchronous probabilistic CSP

We introduce an operator denoted $\|_A^{\sigma}$ with two parameters, probability $\sigma \in (0,1)$ and synchronisation set A. For $P\,\|_A^{\sigma}\,Q$, parameter σ denotes the probability that P performs an autonomous action, given that both P and Q have decided not to synchronise. (Notice that σ has the same interpretation for APCCS.) One probabilistic parameter suffices in the case of CSP, since the only non-determinism that has to be resolved is the one occurring if both processes autonomously decide to perform actions not in A. This is exactly the purpose of parameter σ. The semantics of APCSP is given by the least relation satisfying the rules in Table 1, where the rules in Table 3 replace the rule for $\times$.

$$\frac{P \xrightarrow{b,p}_i P' \qquad Q \xrightarrow{c,q}_j}{P \parallel_A^\sigma Q \xrightarrow{b,\frac{pq\sigma}{\nu(P,Q,A)}}_{(i,j]} P' \parallel_A^\sigma Q}\ (b,c \notin A) \qquad \frac{Q \xrightarrow{b,p}_i \qquad P \xrightarrow{c,q}_j P'}{Q \parallel_A^\sigma P \xrightarrow{c,\frac{pq(1-\sigma)}{\nu(Q,P,A)}}_{[i,j)} Q \parallel_A^\sigma P'}\ (b,c \notin A)$$

$$\frac{P \xrightarrow{b,p}_i s_1' \qquad Q \xrightarrow{a,q}_j}{P \parallel_A^\sigma Q \xrightarrow{b,\frac{pq}{\nu(P,Q,A)}}_{(i,j]} P' \parallel_A^\sigma Q}\ \begin{matrix}(a \in A,\\ b \notin A)\end{matrix} \qquad \frac{Q \xrightarrow{a,p}_i \qquad P \xrightarrow{b,q}_j P'}{Q \parallel_A^\sigma P \xrightarrow{b,\frac{pq}{\nu(Q,P,A)}}_{[i,j)} Q \parallel_A^\sigma P'}\ \begin{matrix}(a \in A,\\ b \notin A)\end{matrix}$$

$$\frac{P \xrightarrow{b,p}_i s_1' \qquad Q\ \mathit{endpoint}}{P \parallel_A^\sigma Q \xrightarrow{b,\frac{p}{\nu'(P,A)}}_{(i,0]} P' \parallel_A^\sigma Q}\ (b \notin A) \qquad \frac{Q\ \mathit{endpoint} \qquad P \xrightarrow{b,q}_j P'}{Q \parallel_A^\sigma P \xrightarrow{b,\frac{q}{\nu'(P,A)}}_{[0,j)} Q \parallel_A^\sigma P'}\ (b \notin A)$$

$$\frac{P \xrightarrow{a,p}_i s_1' \qquad Q \xrightarrow{a,q}_j s_2'}{P \parallel_A^\sigma Q \xrightarrow{a,\frac{pq}{\nu(P,Q,A)}}_{(i,j)} P' \parallel_A^\sigma Q'}\ (a \in A) \qquad \nu'(P,A) =_{df} 1 - \sum_i \{\!|p \mid P \xrightarrow{a,p}_i,\ a \in A|\!\}$$

$$\nu(P,Q,A) =_{df} 1 - \sum_{i,j} \{\!|pq \mid P \xrightarrow{a,p}_i,\ Q \xrightarrow{b,q}_j,\ a,b \in A,\ a \neq b|\!\}$$

Table 3
Operational semantics of APCSP parallel composition

Notice that σ appears only in the rules in the first row, where it is used as a weight for an autonomous move of P (and $(1-\sigma)$ for Q). In all inference rules, each transition probability is normalised by some factor ν, or ν'. (These factors can never equal 0; e.g. ν' is only used if P can perform an action not in A, which guarantees that $\nu' \neq 0$ if it appears as a denominator.) These factors redistribute the probability mass that is due to autonomous decisions of both processes that would otherwise lead to a deadlock. There may be two different reasons for such a situation.

- Redistribution of probability mass is required if one component, say P, autonomously decides in favour of a synchronisation, while Q is an endpoint, i.e. Q cannot move at all. In this case $\nu'(P,A)$ is the probability that P intends to perform a synchronisation (i.e. a deadlock occurs) provided that Q is an endpoint.
- Another source of normalisation is the case that both P and Q decide in favour of a synchronisation, but the labels of these actions do not match. Function $\nu(P,Q,A)$ collects the probability mass of all these mismatching synchronisations.

It is not difficult to check that GPTSs are closed under $\parallel_A^\sigma$. For APCSP we have the following technical results.

Theorem 4.4 *$\sim_p$ is a congruence with respect to $||^{\sigma}_A$.*

Theorem 4.5 *$||^{\sigma}_A$ is stochastic for $\sim_p$.*

$||^{\sigma}_A$ is not respectful for the simple reason that no determinisation takes care of normalisation, whereas the rules for $||^{\sigma}_A$ do. Therefore, bundle probabilities are not respected. As we have discussed in Section 2 the operator $||_A$, and consequently also $||^{\sigma}_A$, can express restriction. Since we want — like in PCCS — that in case of restriction probabilities are redistributed, we perform this normalisation as part of the definition of $||^{\sigma}_A$. (One may argue that from this point of view it is not even desired to consider respectfulness.) Hence the probability of deadlocking is redistributed (using $\mathcal{N}$) after determinisation, and so the probabilities of the newly obtained bundles are respected. In this way $||^{\sigma}_A$ is stochastic.

Theorem 4.6 *Let $P = \sum_i [p_i]a_i.P_i$ and $Q = \sum_j [q_j]b_j.Q_j$ such that P, Q differ from $\mathbf{0}$. Then $P ||^{\sigma}_A Q$ equals*

$$\begin{array}{l}
\sum_{i,j} \left\{ [r]\, a_i.(P_i \,||^{\sigma}_A\, Q_j) \mid a_i = b_j, a_i, b_j \in A, r = \frac{p_i q_j}{\nu(P,Q,A)} \right\} \\
+\sum_{i,j} \left\{ [r]\, a_i.(P_i \,||^{\sigma}_A\, Q) \mid a_i \notin A, b_j \in A, r = \frac{p_i q_j}{\nu(P,Q,A)} \right\} \\
+\sum_{i,j} \left\{ [r]\, b_j.(P \,||^{\sigma}_A\, Q_j) \mid a_i \in A, b_j \notin A, r = \frac{p_i q_j}{\nu(P,Q,A)} \right\} \\
+\sum_{i,j} \left\{ [r]\, a_i.(P_i \,||^{\sigma}_A\, Q) \mid a_i \notin A, b_j \notin A, r = \frac{p_i q_j \sigma}{\nu(P,Q,A)} \right\} \\
+\sum_{i,j} \left\{ [r]\, b_j.(P \,||^{\sigma}_A\, Q_j) \mid a_i \notin A, b_j \notin A, r = \frac{p_i q_j (1-\sigma)}{\nu(P,Q,A)} \right\}.
\end{array}$$

If one of the processes equals $\mathbf{0}$ we obtain, for instance

$$P \,||^{\sigma}_A\, \mathbf{0} = \sum_i \left\{ [r]\, a_i.P_i \mid a_i \notin A, r = \frac{p_i}{\nu'(P,A)} \right\}$$

and similarly for $\mathbf{0} \,||^{\sigma}_A\, Q$. The reader is invited to check that for all P, processes $P \,||^{\sigma}_A\, \mathbf{0}$ and $P \backslash A$ are equivalent (i.e. probabilistic bisimilar).

5 Appropriate parallel compositions

In this section we consider several existing generative probabilistic operators. These operators have been defined for probabilistic variants of process algebras CCS, CSP, ACP and LOTOS. We consider respectfulness and stochasticity of these calculi with respect to probabilistic bisimulation $\sim_p$ (unless stated otherwise). The results of this comparative study are summarised in Table 4, where '$\surd$' indicates that parallel composition in the respective calculus is respectful or stochastic, and '–' indicates that this is not the case.

Language	Respectful	Stochastic
APCCS	√	√
APCSP	–	√
PCCS [5]	√	√
PACP [2]	√	√
PACP$^{\delta}$	√ / –	√ / –
PCSP [17]	√	–
PTPA [6]	–	–
PLOTOS [18]	–	–
PL [13]	–	–

Table 4
Appropriateness of existing generative probabilistic calculi

PTPA, PLOTOS and PL

For the latter three calculi that all have a CSP-like synchronisation, we consider our running example $Q \,||_{\{a\}}\, R'$, see Figure 5 (where only the initial steps are depicted) and Figure 4 for the bundle view. In PTPA, $||$ is not equipped with a probabilistic parameter, and the resulting bundle probabilities are not respected, consider, for instance, the transition labelled a for which one expects $\frac{1}{6}$. PLOTOS and PL contain a variant of $||_A^{\sigma}$ which, however, both result

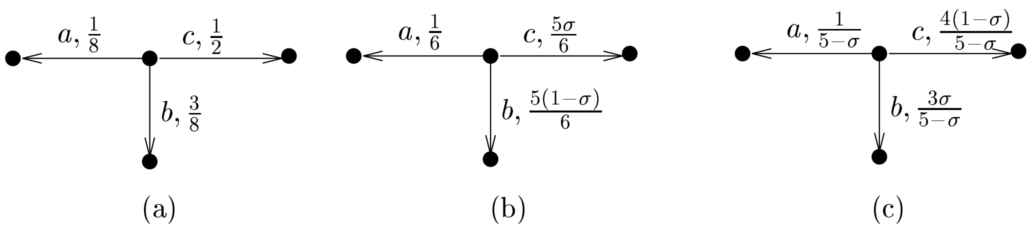

Fig. 5. Parallel composition of Q and R' in (a) PTPA, (b) PLOTOS and (c) PL

in the undesired phenomenon that the probability of an autonomous move (b or c) can be made arbitrarily small. A major source of the inappropriateness of these parallel composition operators is the 'meaningless' normalisation factor.

Probabilistic CSP

In the most simple case, parallel composition in PCSP works just like the PCCS synchronisation where operation $\cdot$ of the semi-group $(\mathbf{A}, \cdot)$ is not completely defined: $aa = a$ and ab is not defined if $a \neq b$. In this way, all mismatching synchronisations introduce some probability of deadlock. For instance, process $Q||R'$ has only one available transition $Q||R' \xrightarrow{a,\frac{1}{6}} \mathbf{0}||\mathbf{0}$ and a deadlock probability of $\frac{5}{6}$. Using a similar reasoning to that of Theorem 3.8, we can state that this parallel composition is respectful (with respect to probabilistic trace equivalence). However, it is not stochastic, since normalisation would remove the deadlock that the operator introduces, whereas PCSP-parallel composition does not. Notice that for APCSP we decided to include normalisation as part of the semantics of $\|_A^\sigma$. As a result, $\|_A^\sigma$ is stochastic, but not respectful (for $\sim_p$); the reverse of PCSP. Since CSP-parallel composition can express restriction, we consider normalisation to be a natural part of $\|_A^\sigma$ like it is for restriction in PCCS. We conjecture that it is impossible to obtain a probabilistic variant of $||_A$ that is both respectful and stochastic (for $\sim_p$).

Probabilistic ACP

Parallel composition in probabilistic ACP (PACP) has two parameters, σ and θ, both in $(0, 1)$. For $P \,||_{\sigma,\theta}\, Q$ a synchronisation between P and Q occurs with probability $1{-}\theta$ and an autonomous action (of either P or Q) with probability θ. Note that θ is unconditioned, as opposed to APCCS. Given that an autonomous move occurs, it comes from P with probability σ, and from Q with probability $1{-}\sigma$. The initial steps of the transition system for $Q \,||_{\sigma,\theta}\, R$ are depicted in Figure 6. Here we assumed that the communication function γ is defined by: $\gamma(a, \overline{a}) = a\overline{a}, \gamma(b, \overline{a}) = b\overline{a}, \gamma(a, c) = ac$ and $\gamma(b, c) = bc$. (The fact that θ is unconditioned introduces the need to define γ as a total function in PACP, as opposed to original ACP, where γ may be partially defined.) The single transition labelled with a is the superposition of the two a-transitions in the bundle view, see Figure 3. Similar for the transitions labelled with $\overline{a}, b$

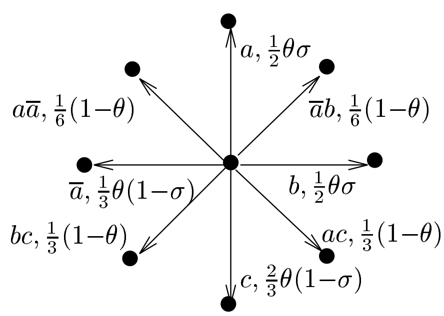

Fig. 6. Parallel composition of Q and R in probabilistic ACP

and c. The probabilities of these transitions can be dispersed in such a way that the bundle probabilities are respected. Since, in addition, PACP-parallel composition does not introduce deadlocks, $||_{\sigma,\theta}$ is also stochastic.

Probabilistic ACP with encapsulation

It is known that for the non-probabilistic case, CCS and CSP parallel composition can be encoded in terms of ACP parallel composition composed with encapsulation (what we have called restriction so far). In fact, it is the usual treatment in ACP to encapsulate processes composed in parallel in order to obtain an adequate specification of the system that is being modelled. As a consequence, given that PACP parallel composition in isolation is stochastic and respectful, it is interesting to investigate whether a combination of parallel composition and encapsulation is stochastic and respectful, as well.

Consider encapsulation in PACP denoted PACP^δ in Table 4 above. By using the encapsulation operator, we can prohibit the execution of the autonomous actions a, $\overline{a}$, b, and c such that only synchronisation actions can be executed. This yields a normalisation with $1-\theta$, the probability that a synchronisation occurs, and the result is Figure 1(c), the PCCS synchronisation. Complementary, prohibiting all synchronisation actions yields an APCSP-view with empty synchronisation set. In both cases the result is respectful and stochastic.

These are two special cases of encapsulation in which each bundle is equally treated. If, however, encapsulation affects bundles in an unequal fashion, it is no longer guaranteed that the bundle probabilities are respected. For instance, allowing only the actions $a, \overline{a}, b, c$ and $a\overline{a}$ (yielding a view similar to APCCS), does not affect the structure of the bundle containing a, $\overline{a}$ and $a\overline{a}$, but affects its probability. Although for some specific choices of σ and θ this might result in a respectful probability assignment, in general this is *not* the case. This differs from our proposals for APCCS and APCSP where normalisation only affects the bundles from which a branch is pruned, and not the others.

It is interesting to note that recently a version of probabilistic ACP has been proposed [1] in which probabilistic and non-deterministic choice co-exist, where parallel composition is based on our bundle concept.

6 Concluding remarks

In this paper we have extensively discussed parallel composition in an asynchronous probabilistic setting. Based on the line of thought in PCCS we formulated two criteria for such parallel composition operators in the context of a congruence relation. These were formalised using the novel notion of *bundle probabilistic transition systems*, transition systems that contain probabilistic 'bundles' of non-deterministic transitions. The basic idea of an appropriate parallel composition operator is that it should leave the bundle probabilities unaffected; only the non-determinism within a bundle should be resolved. This aspect is considered with and without normalisation. We proposed an asynchronous probabilistic variant of CCS and CSP that satisfy this criterion with normalisation (and that preserve probabilistic bisimulation $\sim_p$). Since CCS parallel composition does not introduce deadlocks it also satisfies the criterion (for $\sim_p$) without normalisation. In addition we argued that various existing generative probabilistic calculi do not satisfy these criteria, with the notably exception of probabilistic ACP. Nonetheless, probabilistic ACP with restriction (encapsulation) is, in general, not appropriate (for $\sim_p$).

We like to point out that we have been slightly restrictive in our notion of appropriate parallel composition. Determinisation only operates in a static way, i.e. it only looks at the structure of a bundle probabilistic transition system. In this way, appropriate parallel compositions have to be static operators (with the usual notion of static operator, see [12]). Instead of defining appropriate parallel compositions in terms of determinisations, we could also do it in terms of adversaries [19,16]. This would allow a more dynamic view on the system, since adversaries are typically defined on executions (i.e. runs) of the system. In this setting, parallel compositions that change probabilities or priorities along the execution could also be considered as appropriate. We will report this in the future. In the future we also plan to adopt the notion of scheduling as proposed in this paper in the context of stochastic automata and the syntax of the stochastic process algebra $\spadesuit$ [4].

Acknowledgement

We would like to thank the reviewers for their constructive criticism and for pointing out the relation between BPTSs and the model of Pnueli and Zuck. We also thank Ed Brinksma and Rom Langerak (both of the University of Twente) for fruitful discussions.

A Proof of Theorem 4.1

This appendix illustrates the proof of congruence for probabilistic bisimulation with respect to the operators introduced in Section 5. We restrict ourselves to the case of APCCS parallel composition. The proof strategy for APCSP (Theorem 4.4) follows similar lines.

Let $\sigma, \theta \in (0, 1)$. We show that whenever $P_1 \sim P_2$ and $Q_1 \sim Q_2$ we have that $P_1 \,{}^{\theta}|^{\sigma}\, Q_1 \sim P_2 \,{}^{\theta}|^{\sigma}\, Q_2$. Having fixed σ and θ, we abbreviate $P \,{}^{\theta}|^{\sigma}\, Q$ by (P, Q). So, we are aiming to deduce $(P_1, Q_1) \sim (P_2, Q_2)$. By Definition 3.7 it is sufficient to show $(P_1, Q_1)\ R\ (P_2, Q_2)$ for some probabilistic bisimulation R. To do so, we define R as the reflexive closure of

$$\Big\{((P, Q), (P', Q')) \mid P \sim P', Q \sim Q'\Big\}.$$

Obviously R is an equivalence relation and it satisfies $(P_1, Q_1)\ R\ (P_2, Q_2)$ whenever $P_1 \sim P_2$ and $Q_1 \sim Q_2$. Note that the equivalence classes of R are of the form $C \times C'$ where C and C' are arbitrary equivalence classes of $\sim$. It remains to be shown that R is a probabilistic bisimulation. For this purpose, we fix an equivalence class $C \times C'$ of R and an action $a \in \mathbf{A}$. We have to show that $(P, Q)\ R\ (P', Q')$ implies

$$\mu((P, Q), a, C \times C') = \mu((P', Q'), a, C \times C').$$

Notice that $P \sim P'$ and $Q \sim Q'$ holds by the definition of R. Thus, if Q *endpoint*, this implies Q' *endpoint* and the proof obligation follows from $\mu(P, a, C) = \mu(P', a, C)$ (a consequence of the fact that $\sim$ is itself a probabilistic bisimulation), together with the third pair of rules in Table 2. Since the case of P *endpoint* is symmetric, only the situation where neither P *endpoint* nor Q *endpoint*, remains to be tackled by a detailed analysis of the rules in Table 2. We distinguish the cases $a = \tau$ and $a \neq \tau$, and illustrate the necessary reasoning for $a = \tau$. In this case, the operational rules give us

$$\begin{aligned}\mu\Big((P,Q),\tau,C\times C'\Big) = \quad & \sum_{b\in\mathbf{A}} \mu(P,\tau,C)\ \mu(Q,b,\Sigma)\ \sigma && P \text{ moves}\\ & +\sum_{b\in\mathbf{A}} \mu(P,b,\Sigma)\ \mu(Q,\tau,C')\ (1-\sigma) && Q \text{ moves}\\ & +\sum_{c\neq\tau} \mu(P,c,C)\ \mu(Q,\overline{c},C')\ (1-\theta) && \text{synchronize}\end{aligned}$$

Since C and C' are equivalence classes of $\sim$, we use $\mu(P,a,C'') = \mu(P',a,C'')$ and $\mu(Q,a,C'') = \mu(Q',a,C'')$ (for arbitrary actions a and classes C'') to equate the above right hand side with

$$\begin{aligned}&\sum_{b\in\mathbf{A}} \mu(P',\tau,C)\ \mu(Q',b,\Sigma)\ \ \sigma\\ &+\sum_{b\in\mathbf{A}} \mu(P',b,\Sigma)\ \mu(Q',\tau,C')\ \ (1-\sigma)\\ &+\sum_{c\neq\tau} \mu(P',c,C)\ \mu(Q',\overline{c},C')\ \ (1-\theta)\ = \mu\Big((P',Q'),\tau,C\times C'\Big)\end{aligned}$$

completing the proof for this case. The converse case, $a \neq \tau$ is shown in the same way. It differs with respect to the summands appearing in the above equations, but the proof strategy remains unchanged.

B Proof of Theorem 4.2

To illustrate how proofs of respectfulness and stochasticity are conducted we provide a detailed proof of Theorem 4.2. The proofs of Theorem 4.5 and the statements for probabilistic CSP and probabilistic ACP in Section 5 are constructed in a similar way and are omitted here.

Let the commutative semi-group $(\mathbf{A},.)$ be the comm. group $(\mathbf{A},.,^{-},\tau)$. Define the determinisation function $\mathcal{D}$ by $\mathcal{D}(\Omega,\mathsf{Pr}) = (\Omega',\mathsf{Pr}')$ where the sample space

$$\begin{aligned}\Omega' \ =_{df}\ & \{(a,\mathbf{i},s) \mid (a,\mathbf{i},s)\in \textstyle\bigcup\Omega,\ (\mathbf{i}=(i,j]\vee \mathbf{i}=[i,j))\}\\ \cup\ \ & \{(\tau,(i,j),s) \mid \{(a,(i,j],t),(\overline{a},[i,j),t'),(\tau,(i,j),s)\}\in\Omega\}\\ \cup\ \ & \{\varnothing \mid \varnothing\in\Omega\}\end{aligned}$$

and probability measure Pr' is defined as follows

(*i*) if $B = \varnothing \in \Omega$ then

$$\mathsf{Pr}'(\varnothing) =_{df} \mathsf{Pr}(B),$$

(*ii*) if $B = \{(a, \mathbf{i}, s)\} \in \Omega$, and $\mathbf{i} = (i, 0]$ or $\mathbf{i} = [0, j)$ then

$$\mathsf{Pr}'(a, \mathbf{i}, s) =_{df} \mathsf{Pr}(B),$$

(*iii*) if $B = \{(a, (i, j], s), (b, [i, j), t), (ab, (i, j), r)\} \in \Omega$ and $a \neq \overline{b}$ then

$$\mathsf{Pr}'(a, (i, j], s) =_{df} \sigma \mathsf{Pr}(B), \quad \text{and}$$

$$\mathsf{Pr}'(b, [i, j), t) =_{df} (1 - \sigma)\mathsf{Pr}(B), \quad \text{and}$$

$$\mathsf{Pr}'(ab, (i, j), r) =_{df} 0$$

(*iv*) if $B = \{(a, (i, j], s), (\overline{a}, [i, j), t), (\tau, (i, j), r)\} \in \Omega$ then

$$\mathsf{Pr}'(a, (i, j], s) =_{df} \theta\sigma \mathsf{Pr}(B),$$

$$\mathsf{Pr}'(\overline{a}, [i, j), t) =_{df} \theta(1 - \sigma)\mathsf{Pr}(B), \quad \text{and}$$

$$\mathsf{Pr}'(\tau, (i, j), r) =_{df} (1 - \theta)\mathsf{Pr}(B).$$

Notice that there is no other possible form for B than those considered above.

To state that $\mathcal{D}$ is indeed a determinisation, we must check that Pr' is a probability measure and that it satisfies conditions 1 and 2 in Definition 3.4. First, notice that for all $B \in \Omega$, $B \neq \varnothing$ implies $B \cap \Omega' \neq \varnothing$. Thus, condition 2 follows immediately from item (*i*) above. Satisfaction of condition 1 follows from simple calculations taking into account cases (*ii*), (*iii*), and (*iv*) above. We check now that Pr' is a probability measure. Since

$$\text{(B.1)} \quad \Omega' - \{\varnothing\} = (\bigcup \Omega) \cap \Omega' - \{\varnothing\} = \bigcup_{B \in \Omega - \{\varnothing\}} B \cap \Omega'$$

we can derive that

$$
\begin{aligned}
& \sum_{\gamma \in \Omega'} \mathsf{Pr}'(\gamma) \\
= \ & \{ \text{ calculus } \} \\
& \sum_{\gamma \in \Omega' - \{\varnothing\}} \mathsf{Pr}'(\gamma) \ + \ \mathsf{Pr}'(\varnothing) \\
= \ & \{ \text{ (B.1) and } \Omega \text{ is pairwise disjoint } \} \\
& \sum_{B \in \Omega - \{\varnothing\}} \sum_{\gamma \in B \cap \Omega'} \mathsf{Pr}'(\gamma) \ + \ \mathsf{Pr}'(\varnothing) \\
= \ & \{ \text{ Conditions 1 and 2 } \} \\
& \sum_{B \in \Omega - \{\varnothing\}} \mathsf{Pr}(B) \ + \ \mathsf{Pr}(\varnothing) \\
= \ & \{ \ \mathsf{Pr} \text{ is a probability measure } \} \\
& \sum_{B \in \Omega} \mathsf{Pr}(B) \ = \ 1
\end{aligned}
$$

To prove that ${}^{\sigma}|^{\theta}$ is respectful, we should check that $\mathcal{G}_1 \,{}^{\sigma}|^{\theta}\, \mathcal{G}_2$ and $\mathcal{D}(\mathcal{G}_1 \otimes \mathcal{G}_2)$ are probabilistically bisimilar. To do so, it is sufficient to prove that the reflexive and symmetric closure of the relation

$$\{(s \,{}^{\sigma}|^{\theta}\, t, s \otimes t) \mid s \in \Sigma_1, t \in \Sigma_2\}$$

where $s \otimes t$ indicates a state of $\mathcal{D}(\mathcal{G}_1 \otimes \mathcal{G}_2)$, is an equivalence relation and moreover a probabilistic bisimulation. We leave this last proof obligation to the reader.

Finally, from (i) above, we can conclude that ${}^{\sigma}|^{\theta}$ does not introduce any deadlock which was not present already in the composed processes. Thus, it is not difficult to prove that the same relation above is a probabilistic bisimulation when $s \otimes t$ indicates a state of $\mathcal{N}(\mathcal{D}(\mathcal{G}_1 \otimes \mathcal{G}_2))$, provided $\mathcal{G}_1$ and $\mathcal{G}_2$ are stochastic. This implies that ${}^{\sigma}|^{\theta}$ is also stochastic.

References

[1] S. Andova. Process algebra with probabilistic choice. In J.-P. Katoen, editor, *Formal Methods for Real-Time and Probabilistic Systems*, LNCS 1601, pages 111–129, 1999.

[2] J.C.M. Baeten, J.A. Bergstra, and S.A. Smolka. Axiomatizing probabilistic processes: ACP with generative probabilities. *Inf. and Comp.*, **121**:234–255, 1995.

[3] J.C.M. Baeten and W.P. Weijland. *Process Algebra*. Cambridge University Press, 1990.

[4] P.R. D'Argenio, J.-P. Katoen and E. Brinksma. An algebraic approach to the specification of stochastic systems (extended abstract). In D. Gries and W.-P. de Roever, editors, *Proceedings of the Working Conference on Programming Concepts and Methods*, pages 126–148. Chapman & Hall, 1998.

[5] A. Giacalone, C.-C. Jou, and S.A. Smolka. Algebraic reasoning for probabilistic concurrent systems. In M. Broy and C.B. Jones, editors, *Proceedings of the Working Conference on Programming Concepts and Methods*, pages 443–458. North-Holland, 1990.

[6] C. Gregorio-Rodriguez, L. Llana-Diaz, M. Núñez, and P. Paloa-Gostanza. In M. Bertran and T. Rus, editors, *Transformation-Based Reactive Systems Development*, LNCS 1231, pages 353–367. Springer-Verlag, 1997.

[7] R.J. van Glabbeek. Notes on the methodology of CCS and CSP. *Th. Comp. Sci.*, **177**:329–349, 1997.

[8] R.J. van Glabbeek, S.A. Smolka, and B. Steffen. Reactive, generative, and stratified models of probabilistic processes. *Inf. and Comp.*, **121**:59–80, 1995.

[9] H. Hansson and B. Jonsson. A calculus for communicating systems with time and probabilities. In *Proceedings of 11th IEEE Real-Time Systems Symposium*, pages 278–287. IEEE Comp. Soc. Press, 1990.

[10] C.A.R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.

[11] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inf. and Comp.*, **94**(1): 1–28, 1992.

[12] R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.

[13] M. Núñez and D. de Frutos. Testing semantics for probabilistic LOTOS. In G. von Bochmann, R. Dssouli, and O. Rafiq, editors, *Formal Description Techniques VIII*, pages 365–380. Chapman & Hall, 1995.

[14] E.-R. Olderog and C.A.R. Hoare. Specification-oriented semantics for communicating processes. *Acta Inf.*, **23**:9–66, 1986.

[15] A. Pnueli and L.D. Zuck. Probabilistic verification. *Inf. and Comp.*, **103**:1–29, 1993.

[16] R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. *Nordic J. of Computing*, **2**(2):250–273, 1995.

[17] K. Seidel. Probabilistic communicating processes. *Th. Comp. Sc.*, **152**:219–249, 1995.

[18] R. Sisto, L. Ciminiera, and A. Valenzano. Probabilistic characterization of algebraic protocol specifications. In *Proceedings 12th Int. Conf. on Distributed Computing Systems*, pages 260–268. IEEE Comp. Soc. Press, 1992.

[19] M.Y. Vardi. Automatic verification of probabilistic concurrent finite state programs. In *Proceedings* 26^{th} *Symposium on Foundations of Computer Science*, pages 327–338. IEEE Comp. Soc. Press, 1985.