

Real-Time Process Algebra with Stochastic Delays

J. Markovski* and E.P. de Vink

Formal Methods Group, Technische Universiteit Eindhoven
Den Dolech 2, 5612 AZ Eindhoven, The Netherlands
j.markovski@tue.nl, evink@win.tue.nl

Abstract

A real-time process algebra is presented that features stochastic delays governed by general distributions. In a setting of weak choice, dependent and independent alternative and parallel composition are distinguished. This enables an expansion law for the parallel operator, as well as modular process definitions. The interplay of real-time, stochastic delays and immediate actions is illustrated by a modeling of the $G/G/1/\infty$ queue.

1. Introduction

Stochastic process algebras emerged as a powerful tool for both qualitative and quantitative analysis of processes. Early stochastic process algebras typically employed exponentially distributed stochastic delays. Because of its memoryless property, usage of the exponential distribution greatly simplifies the treatment of the parallel composition. Prominent Markovian process algebras include TIPP, EMPA, PEPA and the algebra of IMC. The first three associate exponential rates with actions, whereas the latter explicitly distinguishes between actions and rates.

Despite the great success, exponential delays turned out not to be sufficient for a number of modeling purposes, such as for protocols for downloading or media streaming. Consequently, several stochastic process algebras with general distributions were proposed, like SPADES, IGSMP and NMSPA [11, 6, 17]. SPADES uses residual lifetime semantics and has clocks to model stochastic delays. Each clock samples from a general distribution. Sets of clocks guard actions that become enabled after all the clocks have expired. The semantics is given in terms of stochastic automata [10]. IGSMP uses clocks with an associated expiration time

distribution to record spent lifetimes. After a clock expiration the other active clocks are redistributed according to the time that passed. The semantics involves generalized semi-Markov processes extended with actions. To that end, the alternative composition is modeled as a probabilistic choice between differently distributed clocks. NMSPA exploits random variables for the guidance of stochastic delays. Also here, expiration of a stochastic delay induces redistribution of the other variables. The semantics is given in terms of transition systems. The alternative composition is defined over an arbitrary number of summands in order to achieve maximal progress for internal actions. The alternative composition of delays that exhibit the same duration followed by an internal action, is treated as an inherent probabilistic choice. Other stochastic process algebras that we mention here are the stochastic π -calculus and TIPP [20, 14]. More can be found in the review [7].

Typically, Markovian process algebras do not extend real-time process algebras because of the exponential distribution. Generally distributed stochastic delays are usually tackled by clocks. See, e.g., [15, 1]. For SPADES, a structural translation from stochastic automata to timed automata with deadlines is given in [9]. It has been shown that this preserves timed traces. Also, there is a translation from IGSMP into pure real-time models termed Interactive Timed Automata, see [6]. In [2] a proposal of extending timed LOTOS is made by exploiting stochastic timers.

The main goal of our paper is to deal with stochastic time as it is done in real-time process algebras [3, 19], aiming at a conservative extension of real-time process algebras with stochastic delays. Building on our previous work [18], we deal with a semantics that exploits spent-time and avoids explicit clocks in a setting with immediate actions, deadlock and termination. We model stochastic delays as timed delays guided by (finite or countably infinite) discrete random variables, as we wish to distinguish between actions and stochastic

*Corresponding author. Supported by Bsik-project BRICKS AFM 3.2

delays, similar as in IMC [13]. The alternative composition implements weak choice between immediate actions and passage of time along the lines of real-time process algebras in the style of [3]. Relying on this, the parallel operator can be treated by standard means. We give the semantics in terms of stochastic transition systems. In comparison to other stochastic process algebras, our approach is closest to NMSPA. However, we define alternative composition on two processes rather than on arbitrary sums and, in our setting, passage of time makes no choice in case both summands can delay simultaneously. We introduce a so-called dependent alternative composition that guarantees that stochastic delays guided by the same variable will always exhibit the same duration. This way, improving on [18], we are able to obtain an expansion law for the parallel operator, a result that was absent previously. Again, via an embedding of transition systems, the proposed stochastic process algebra can be shown to extend real-time process algebra. In our present work, we focus on discrete stochastic delays, mainly because they almost effortlessly model real-time delays as degenerated discrete random variables. Also, as a technical convenience, they allow two different delays to actually exhibit the same duration, a property not shared by continuous distributions.

The rest of this paper is organized as follows: Section 2 provides mathematical preliminaries. Section 3 introduces a basic stochastic process algebra with alternative composition and prefixes with stochastic delay. Section 4 deals with stochastic transition systems, stochastic bisimulation and α -conversion. Section 5 discusses the parallel operator and standard auxiliary operators. Section 6 revisits the embedding of real-time and the modeling in the present setting of an $G/G/1/\infty$ queue. Section 7 wraps up with concluding remarks.

Acknowledgement Many thanks to Jos Baeten for various discussions on the subject.

2. Preliminaries

We write \mathbb{R}^+ for $\{t \in \mathbb{R} \mid t \geq 0\}$. We use discrete random variables to represent durations of stochastic delays. So, we only consider distribution functions F such that $F(t) = 0$ for $t \leq 0$. We denote the set of such distribution functions by \mathcal{F} and the set of the corresponding random variables by \mathcal{V} . We use X, Y and Z to range over \mathcal{V} and F_X, F_Y and F_Z for their respective distribution functions. If $P(X = t) = 1$, then we say that the random variable X is degenerated or Dirac. We denote such a random variable by D_t .

The support set of a random variable X is denoted

by $\text{supp}(X) = \{t \in \mathbb{R}^+ \mid P(X = t) > 0\}$, which, by assumption, is finite or countably infinite. We put $\text{supp}(S) = \bigcap_{X \in S} \text{supp}(X)$ for a subset $S \subseteq \mathcal{V}$. By $\bar{F}_X(t)$ we denote the residual probability distribution $1 - F_X(t)$. For $S \subseteq \mathcal{V}$, $y \in \mathbb{R}$ and \diamond either $<$, $>$ or $=$, we write $S \diamond y$ if $X \diamond y$, for all $X \in S$. We denote conditional random variables by $\langle X \mid \text{Event} \rangle$, where $X \in \mathcal{V}$ and Event such that $P(\text{Event}) > 0$.

A *stochastic delay* is a timed delay with duration guided by a random variable. We observe simultaneous passage of time for a number of stochastic delays until one or some of them expire. This phenomenon is referred to as the *race condition* and the process as the *race*. For multiple stochastic delays in a race, different stochastic delays can be observed simultaneously as being the shortest. The shortest duration itself can be different and exhibited by different stochastic delays in different observations. The stochastic delays that have the shortest duration are called the *winners*, the others are called the *losers* of the race.

The probability that a subset $W \subseteq V$ is the winning set of a race performed by the variables of the set V with a duration d is denoted by $\text{RC}_d(W, V)$. It is defined as

$$\text{RC}_d(W, V) = \prod_{X \in W} P(X = d) \cdot \prod_{X \in V \setminus W} \bar{F}_X(d).$$

The probability for $W \subseteq V$ being the winning set is denoted by $\text{RC}(W, V)$ and can be calculated from $\text{RC}(W, V) = \sum_{d \in \text{supp}(W)} \text{RC}_d(W, V)$.

3. Basic Sequential Processes

We give the semantics of processes that have discrete stochastic delays, immediate actions, termination and deadlock, and implement weak choice between actions and passage of time. The behavior of these processes is captured by the process algebra $\text{BSP}^{\text{dst}}(\mathcal{A}, \mathcal{V})$ defined below. Here, \mathcal{A} is the set of actions and \mathcal{V} is the set of random variables that guide the stochastic delays. We postulate a function $\varphi: \mathcal{V} \rightarrow \mathcal{F}$ that assigns to a variable X a (possibly infinite) discrete probability distribution. For renaming purposes, we assume that for every distribution in \mathcal{F} there are countably infinitely many variables mapping to it.

The process algebra has two alternative composition operators that differ in their treatment of stochastic delays governed by the same variables in a race. The so-called *dependent* alternative composition $_+ _$ assigns a single duration to racing delays referring to the same random variable. For the so-called *independent* alternative composition $_ \oplus _$ racing delays with the same random variable may have assigned a different duration. We studied independent alternative composition

in previous work [18]. There, we concluded that because of the independence of the stochastic delays, it is not possible to come up with an expansion law for the parallel composition. For that reason we introduce here, the dependent alternative composition that solves this problem. However, having only a dependent alternative composition is not sufficient for modeling purposes. More precisely, when describing a complex system, typically one needs to combine several instances of the same component. Therefore, the stochastic delays in separate instances that happen to be guided by the same variable are considered to sample different durations in general.

Definition 1 *The signature of BSP^{dst} consists of two constants δ and ϵ , two unary operator schemes $a.$, for $a \in \mathcal{A}$, and $\sigma_X.$, for $X \in \mathcal{V}$, and two binary operators $+.$ and $\oplus.$. The syntax is given by*

$$P ::= \delta \mid \epsilon \mid a.P \mid \sigma_X.P \mid P + P \mid P \oplus P$$

with $a \in \mathcal{A}$ and $X \in \mathcal{V}$. We use p and q to range over BSP^{dst} . We write $\mathcal{C}(\text{BSP}^{\text{dst}})$ for its closed terms.

The signature of BSP^{dst} is adopted from [4]. The constant δ represents immediate deadlock which does not allow passage of time. Similarly, the constant ϵ represents a process that immediately terminates successfully. The action prefix scheme $a.p$ comprises processes that immediately execute the action a and continue to behave as p . The delay prefix scheme $\sigma_X.p$ provides processes that execute a stochastic delay guided by X and continue to behave as p . The interpretation of the dependent alternative composition $+$ relies on the context: a non-deterministic choice is made between actions; a weak choice between actions, successful termination and stochastic delays; a race condition is imposed on stochastic delays such that two stochastic delays guided by the same variable have the same duration. For the independent alternative composition \oplus , unlike for the dependent alternative composition, one has a race condition in which different stochastic delays guided by the same variable can exhibit different durations.

Discrete Stochastic Delays Because of the race condition, one cannot observe the execution of a stochastic delay in isolation. An example of a transition system corresponding to a dependent race between three stochastic delays for the variables X , Y and X again, is depicted in Fig. 1. Each \mapsto -transition represents a stochastic delay. The label shows the winners of the race and their observed duration. Between parentheses are the conditions that enable the transition.

For conciseness, the transitions of the stochastic delay guided by the variable X are represented by

one transition scheme labeled by X and x . The observed winning duration x takes its values from $\text{supp}(\langle X \mid X < Y \rangle)$. Thus, the transition scheme replaces $|\text{supp}(\langle X \mid X < Y \rangle)|$ many different transitions, each executed with its own probability. Note that if $P(X < Y) = 0$ then the transition does not exist and, also, the conditional random variable $\langle X \mid X < Y \rangle$ does not exist. The losing stochastic delays, represented by X' and Y' , are adapted using the winning duration as discussed below.




Figure 1. Dependent race condition

Although the race in Fig. 1 is a race between three stochastic delays, the first and the third one will always have the same duration as they are guided by the same random variable in the same race. Thus, the race is actually governed by only two random variables, viz. X and Y . This is the characteristic of the dependent alternative composition. However, in two different races, stochastic delays guided by same random variable do not necessarily have the same duration. For example, let $X \in \mathcal{V}$ be such that $P(X = 1) = \frac{1}{4}$, $P(X = 2) = \frac{3}{4}$. Then the process $\sigma_X.\sigma_X.\epsilon$ delays 2, 3 or 4 time units with probabilities $\frac{1}{16}$, $\frac{6}{16}$ and $\frac{9}{16}$, respectively, and then terminates successfully.

We resolve the interaction of action transitions and termination versus (positive) stochastic delays by a weak choice, i.e., a non-deterministic choice between immediate actions, termination and passage of time. A typical example of weak choice is depicted in Fig. 2. Although the action a is immediate, the stochastic delay is enabled. After the delay the action is no longer available. As a consequence of the weak choice, the




Figure 2. Weak choice

losers of a race remain and have the option of performing every possible duration in their residual support set.

For the independent race in $\sigma_X.p \oplus \sigma_X.q$, we have the transitions as depicted in Fig. 3. Here, each stochastic delay samples its duration independently. The stochastic delays compete with values from equal but independent distributions. Therefore, the two stochastic delays do not need to last equally long (unless the ran-

dom variable is Dirac).




Figure 3. Independent race condition

The situation becomes slightly more involved when mixing both types of alternative composition. We treat the stochastic delays that participate in an independent race also as independent in a comprising race. An example may clarify this matter. Consider $(\sigma_X \oplus \sigma_Y) + \sigma_X$. If the leftmost stochastic delay guided by X wins the race induced by the independent alternative composition, then it will not necessarily have the same duration as the rightmost stochastic delay guided by X as well, even though they are in a race induced by a dependent alternative composition. This is because the leftmost delay sampled independently in the context of the innermost alternative composition in the first place.

Aging and Environments We next discuss how the amount of time that has passed for the winners of a race influences that of the losers. Consider the example in Fig. 1. Note that $F_{X'}$ is defined only if $P(X > y) > 0$. At least one such $y \in \text{supp}(Y)$ exists if $X > Y$ holds. The *aged* distribution of X , $F_{X'}$ is given by

$$F_{X'}(t) = P(X \leq t \mid X > Y, Y = y) = \frac{F_X(t+y) - F_X(y)}{1 - F_X(y)}.$$

In order to calculate the actual distribution functions in each state, we need the original distribution function and its *age*, i.e., the amount of time that the stochastic delay participated in races that it lost.

Definition 2 *The probability distribution aging function is a partial function $_|_$: $[\mathcal{F} \times \mathbb{R} \rightarrow \mathcal{F}]$ given by*

$$(F|d)(t) = \frac{F(t+d) - F(d)}{1 - F(d)},$$

for $t \geq 0$, provided $F(d) < 1$.

We augment the transition systems with a so-called environment to store the ages of the stochastic delays. Put $\mathbb{R}_\perp = \mathbb{R} \cup \{\perp\}$. An environment is the function $\alpha: \mathcal{V} \rightarrow \mathbb{R}_\perp$. We add the a special symbol \perp to denote that a stochastic delay has no age, i.e., it has not participated in any race up to that moment. By convention $F|\perp = F$ and $x + \perp = x$, for $x \in \mathbb{R}_\perp$. We consider an environment α to be well-defined if, for each $X \in \mathcal{V}$ and $t > 0$, the probability distribution function $F_X(t) = \varphi(X)|\alpha(X)$ is defined. The set of all well-defined environments, ranged over by α, β , is denoted by \mathcal{E} .

4. Stochastic Transition Systems

We introduce stochastic transition systems that include stochastic delays. They are based on the race condition and provide the basis for the semantics of BSP^{dst} . We show how to handle conflicts can occur because of variable names by means of α -conversion.

In view of Fig. 1 and Fig. 3, we want to utilize environments to keep track of residual distributions. Therefore, we must make sure that the environments are well-defined. The function $A(_): \mathcal{C}(\text{BSP}^{\text{dst}}) \rightarrow 2^{\mathcal{V}}$ extracts all stochastic delays from a BSP^{dst} -term and it is given as follows:

$$\begin{aligned} A(\delta) = A(\epsilon) = \emptyset, \quad A(a.p) = A(p), \quad A(\sigma_X.p) = \{X\} \cup A(p) \\ A(p+q) = A(p) \cup A(q), \quad A(p \oplus q) = A(p) \cup A(q). \end{aligned}$$

The labels of the stochastic delay transitions are decorated by sets of winning delays and the parameterized winning duration. However, not all stochastic delays will participate in the same race at the same time.

So, we have to identify only the racing stochastic delays, i.e., the ones that participate in the ongoing race. The function $R(_)$ collects all stochastic delays that are directly connected by the topmost alternative composition operator. These are the only stochastic delays of the term that can have an age different from \perp . Formally, the function $R(_): \mathcal{C}(\text{BSP}^{\text{dst}}) \rightarrow 2^{\mathcal{V}}$ is given by

$$\begin{aligned} R(\delta) = R(\epsilon) = \emptyset, \quad R(a.p) = \emptyset, \quad R(\sigma_X.p) = \{X\} \\ R(p+q) = R(p) \cup R(q), \quad R(p \oplus q) = R(p) \cup R(q). \end{aligned}$$

Now, we are ready to give the notion of a stochastic transition system that deals with aging of distributions. The states of a stochastic transition system consist of a closed term and an environment. The term defines the racing delays; the environment keeps track of the distributions. We write $\langle p, \alpha \rangle \in \mathcal{S}$ where $\mathcal{S} = \mathcal{C}(\text{BSP}^{\text{dst}}) \times \mathcal{E}$ for compactness of notation. As usual (cf. [3]) we provide action transitions \rightarrow , timed transitions \mapsto and termination \downarrow .

Definition 3 *A stochastic transition system is a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{V}, \rightarrow, \mapsto, \downarrow)$ where*

- for each state $\langle p, \alpha \rangle \in \mathcal{S}$, α is a well-defined environment and $\alpha(X) = \perp$, for all $X \in A(p) \setminus R(p)$;
- $\rightarrow \subseteq \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ is the labeled transition relation;
- $\downarrow \subseteq \mathcal{S}$ is the immediate termination predicate;
- $\mapsto \subseteq \mathcal{S} \times (2^{\mathcal{V}} \setminus \emptyset) \times \mathbb{R}^+ \times \mathcal{S}$ is the stochastic delay transition relation.

For $u \in \mathcal{S}$, $T_u = \{ u \xrightarrow{S}_t u' \mid S \subseteq R(p), t \in \mathbb{R}^+, u' \in \mathcal{S} \}$ is the set of stochastic delays of u . We require that, the mapping $P_u: T_u \rightarrow [0, 1]$ given by $P_u(u \xrightarrow{S}_t u')$ is a probability distribution.

We write $\langle p, \alpha \rangle \xrightarrow{a} \langle p', \alpha' \rangle$ if the term p in the environment α does an action transition with the label a to the term p' with updated environment α' . We write $\langle p, \alpha \rangle \xrightarrow{S, t} \langle p', \alpha' \rangle$ to denote that a term p in the environment α allows time to pass for a duration t , transforms to p' with updated environment α' . The observed time t is the result of a race won by the set of stochastic delays that are guided by the set of random variables S . The duration for the winners is determined by the random variable $\langle X | S = \min(\mathbf{R}(p)) \rangle$, for every $X \in S$. The race changes the environment α by incrementing the ages of the losing delays. We write X for $\{X\}$ when clear from the context.

For the stochastic transition system of a specific term $p \in \mathcal{C}(\text{BSP}^{\text{dst}})$ with an initial well-defined environment α , we write $\text{STS}(p, \alpha)$. The requirement for an initial well-defined environment ensures that the environment follows the intuition, i.e., terms that do not participate in a race do not have an age. In general, we are interested in stochastic transition systems where no race has occurred previously. Hence, the initial environment is α_{\perp} , where $\alpha_{\perp}(X) = \perp$, for every $X \in \mathcal{V}$. In this case we will use the shorthand $\text{STS}(p)$.

Structural Operational Semantics We define an auxiliary function to age the losers by incrementing their age by the winning duration (overloading the related earlier notation for aging of probability distributions).

Definition 4 *The function $\alpha|_dL : \mathcal{E} \times \mathbb{R} \times 2^{\mathcal{V}} \rightarrow \mathcal{E}$ is defined, for $\alpha \in \mathcal{E}$ and a set of losers $L \subseteq \mathcal{V}$ of a race with winning duration d , by $\alpha|_dL = \alpha(X) + d$ if $X \in L$, and $\alpha|_dL = \alpha(X)$ otherwise.*

The structural operational semantics of stochastic transition systems is given in Table 1. We only give the rules for the left operand. We put $\circ \in \{+, \oplus\}$ for the common rules.

$$\begin{array}{l}
1 \langle \epsilon, \alpha \rangle \downarrow \quad 2 \frac{\langle p, \alpha \rangle \downarrow}{\langle p \circ q, \alpha \rangle \downarrow} \quad 4 \langle a.p, \alpha \rangle \xrightarrow{a} \langle p, \alpha \rangle \\
5 \langle \sigma_X.p, \alpha \rangle \xrightarrow{\{X\}, x} \langle p, \alpha|_x\{X\} \rangle, \quad x \in \text{supp}(X) \\
6 \frac{\langle p, \alpha \rangle \xrightarrow{a} \langle p', \alpha' \rangle}{\langle p \circ q, \alpha \rangle \xrightarrow{a} \langle p', \alpha' \rangle} \quad 8 \frac{\langle p, \alpha \rangle \xrightarrow{S, s} \langle p', \alpha' \rangle, \langle q, \alpha \rangle \not\xrightarrow{S}}{\langle p \circ q, \alpha \rangle \xrightarrow{S, s} \langle p', \alpha' \rangle} \\
10 \frac{\langle p, \alpha \rangle \xrightarrow{S, s} \langle p', \alpha' \rangle, \langle q, \alpha \rangle \xrightarrow{T, t} \langle q', \alpha'' \rangle, S \cap T = \emptyset, s < t}{\langle p + q, \alpha \rangle \xrightarrow{S, s} \langle p' + q, \alpha'|_s\mathbf{R}(q) \rangle} \\
10' \frac{\langle p, \alpha \rangle \xrightarrow{S, s} \langle p', \alpha' \rangle, \langle q, \alpha \rangle \xrightarrow{T, t} \langle q', \alpha'' \rangle, s < t}{\langle p \oplus q, \alpha \rangle \xrightarrow{S, s} \langle p' \oplus q, \alpha'|_s\mathbf{R}(q) \rangle} \\
12 \frac{\langle p, \alpha \rangle \xrightarrow{S, d} \langle p', \alpha' \rangle, \langle q, \alpha \rangle \xrightarrow{T, d} \langle q', \alpha'' \rangle}{\langle p \circ q, \alpha \rangle \xrightarrow{S \cup T, d} \langle p' \circ q', \alpha|_d\mathbf{R}(p \circ q) \rangle}
\end{array}$$

Table 1. SOS for BSP^{dst}

Rules 1, 2, and 4 are the standard rules for termina-

tion and action prefix. Rule 5 states that stochastic delay transitions ($\sigma_X.p$) allow passage of time distributed as $\phi(X)|\alpha(X)$. The non-deterministic choice made by action transitions from the first summand is shown by Rule 6. In case the first summand does performs a stochastic delay as in Rule 8, a weak choice is enabled between action transitions and passage of time, where passage of time disables the action transitions of the second summand. Rule 10 describes the race in case the first summand wins the race. Note that the set of winning and the set of losing delays cannot contain the delays guided by the same variables, which is enforced by the condition $S \cap T = \emptyset$. This condition imposes the desired property that delays simultaneously guided by the same random variable in the same race always observe the same duration. The racing delays of the losing summand $\mathbf{R}(q)$ are aged by the winning duration s by applying the aging function on the environment of the winner α' in which the losers of the winning summand are already aged. Note that since the second summand can perform a winning duration $t > s$, the aging of its racing delays is allowed. Rule 12 states that if both summands have stochastic delays with the same winning duration, then the joint race is won by the union of the winners of the both summands. Note that in the special case where multiple stochastic delays are guided by the same random variable in the same race, we consider all of them to have the same duration and join their stochastic delay transitions into one. Thus, there is no multiplicity of transitions. The new environment is obtained by aging all racing delays of both summands in the original environment.

The rules 3, 7, 9 and 11 for the second summand are analogous to 2, 6, 8 and 10, respectively. Note the difference for $+$ and \oplus in the rules 10 and 10'. The rule 10' shows that the race of the independent alternative composition has no restriction that delays guided by the same variable must exhibit the same duration.

The behaviour given by the structural operational semantics in Table 1 uniquely defines the probabilistic behavior of a stochastic transition system. First, there are no multiple equal transitions to the same state and each stochastic delay transition is uniquely defined by the winning set and the duration. From a straightforward inspection of the rules, we observe that the rules change the environment such that only the past and active racing delays are aged. Next, because of Rules 10, 11 and 12, we conclude that racing delays are allowed to perform all possible stochastic delay transitions. Thus, the probability space defined by the distribution functions of the racing stochastic delays is properly defined.

However, the rules result to conflicting behaviour when multiple stochastic delays that (1) do not partic-

ipate in the same race or (2) participate in the same race enabled by the independent alternative composition are guided by a variable with the same name. This is due to clashes in the environment. Each random variable can obtain only one age that is remembered when the stochastic delay expires. We exploit α -conversion to overcome this problem.

α -conversion For a technical underpinning of the renaming of the variables, we define a relation $\simeq_\alpha \subseteq \mathcal{S} \times \mathcal{S}$. As an example, $\langle \sigma_X.\sigma_X.\epsilon, \alpha \rangle$ is congruent to all of $\langle \sigma_X.\sigma_Y.\epsilon, \beta \rangle$, $\langle \sigma_Y.\sigma_X.\epsilon, \beta \rangle$ and $\langle \sigma_Y.\sigma_Y.\epsilon, \beta \rangle$ as long as $\varphi(X) = \varphi(Y)$ and $\alpha(X) = \beta(X) = \beta(Y)$. However, stochastic delays that (1) are guided by the same random variable (2) in the context of a dependent alternative composition must be renamed simultaneously to preserve equivalent stochastic behaviour. For example, $\langle \sigma_X.\epsilon + \sigma_X.\epsilon, \alpha \rangle$ is not congruent to $\langle \sigma_X.\epsilon + \sigma_Y.\epsilon, \alpha \rangle$, unless X and Y happen to have the same degenerated distribution, i.e., unless $\varphi(X) = \varphi(Y)$ are Dirac and $\alpha(X) = \alpha(Y)$. In case of the independent alternative composition, $\langle \sigma_X.\epsilon \oplus \sigma_X.\epsilon, \alpha \rangle$ is congruent to $\langle \sigma_X.\epsilon \oplus \sigma_Y.\epsilon, \alpha \rangle$, $\langle \sigma_Y.\epsilon \oplus \sigma_X.\epsilon, \alpha \rangle$ as well as $\langle \sigma_Y.\epsilon \oplus \sigma_Y.\epsilon, \alpha \rangle$ provided that $\varphi(X) = \varphi(Y)$ and $\alpha(X) = \alpha(Y)$.

For technical convenience, we define the notion of a ‘maximal distinct representation’ in which all stochastic delays have unique names (modulo permutations of \mathcal{V}), except for the ones under the same dependent alternative composition. For example, $\langle \sigma_X.(\sigma_X.\sigma_X.\epsilon + \sigma_X.\sigma_X.\delta) \oplus (\sigma_X + \sigma_X), \alpha \rangle$ has $\langle \sigma_X.(\sigma_Y.\sigma_U.\epsilon + \sigma_Y.\sigma_Z.\delta) \oplus (\sigma_V + \sigma_V), \alpha \rangle$ as a maximal distinct representation, as long as the variables have the same distribution.

We define an auxiliary relation cf_r on BSP^{dst} for a permutation r of the dependent racing delays in $\text{D}(p)$. The function $\text{D}(_)$, for $p \in \text{BSP}^{\text{dst}}$, is defined by

$$\begin{aligned} \text{D}(\delta) &= \emptyset, \text{D}(\epsilon) = \emptyset, \text{D}(a.p) = \emptyset, \text{D}(\sigma_X.p) = \{X\} \\ \text{D}(p + q) &= \text{D}(p) \cup \text{D}(q), \text{D}(p \oplus q) = \emptyset \end{aligned}$$

whereas the predicate cf_r , for a bijective $r: \mathcal{V} \rightarrow \mathcal{V}$, is given by

$$\begin{aligned} \text{cf}_r(\delta, \delta), \text{cf}_r(\epsilon, \epsilon), \text{cf}_r(a.p, a.p'), \text{cf}_r(p \oplus q, p' \oplus q'), \\ \text{cf}_r(\sigma_X.p, \sigma_Y.p'), \quad \text{if } r(X) = Y, \\ \text{cf}_r(p + q, p' + q'), \quad \text{if } \text{cf}_r(p, p'), \text{cf}_r(q, q'). \end{aligned}$$

As a final technical aid we need the relation mdr on \mathcal{S} . The relation holds if the first state is a maximal distinct representation of the second, taking delays into account:

$$\begin{aligned} \text{mdr}(\langle \delta, \alpha \rangle, \langle \delta, \alpha \rangle) \quad \text{mdr}(\langle \epsilon, \alpha \rangle, \langle \epsilon, \alpha \rangle) \\ \text{mdr}(\langle a.p', \alpha' \rangle, \langle a.p, \alpha \rangle) \text{ if } \text{mdr}(\langle p', \alpha' \rangle, \langle p, \alpha \rangle) \\ \text{mdr}(\langle \sigma_Y.p', \alpha' \rangle, \langle \sigma_X.p, \alpha \rangle) \text{ if } Y \notin A(p') \wedge \varphi(X) = \varphi(Y) \\ \wedge \alpha(X) = \alpha'(Y) \wedge \text{mdr}(\langle p', \alpha' \rangle, \langle p, \alpha \rangle) \end{aligned}$$

$$\begin{aligned} \text{mdr}(\langle p' + q', \alpha' \rangle, \langle p + q, \alpha \rangle) \\ \text{if } \text{cf}_r(p + q, p' + q') \wedge r(\text{D}(p + q)) = \text{D}(p' + q') \\ \wedge \text{mdr}(\langle p', \alpha' \rangle, \langle p, \alpha \rangle) \wedge \text{mdr}(\langle q', \alpha' \rangle, \langle q, \alpha \rangle) \\ \text{mdr}(\langle p' \oplus q', \alpha' \rangle, \langle p \oplus q, \alpha \rangle) \text{ if } A(p') \cap A(q') = \emptyset \\ \wedge \text{mdr}(\langle p', \alpha' \rangle, \langle p, \alpha \rangle) \wedge \text{mdr}(\langle q', \alpha' \rangle, \langle q, \alpha \rangle). \end{aligned}$$

With all the machinery in place, α -conversion becomes easy. Two states can be α -converted if they have the same maximal distinct representations.

Definition 5 *Two states $u, v \in \mathcal{S}$ are α -convertible, notation $u \simeq_\alpha v$, if $\{u' \in \mathcal{S} \mid \text{mdr}(u', u)\} = \{v' \in \mathcal{S} \mid \text{mdr}(v', v)\}$.*

Intuitively, the definition states that the renaming of variables is allowed as long as the variables do not appear in the same race. As a consequence, α -conversion does not alter the stochastic behavior of the stochastic transition systems and \simeq_α is a congruence. This can be proven rigorously by structural induction and case analysis for every rule of the operational semantics and is omitted here.

Definition 6 *A conflict-free stochastic transition system of $\langle p, \alpha \rangle$ is STS($\langle p', \alpha' \rangle$), if $\text{mdr}(\langle p', \alpha' \rangle, \langle p, \alpha \rangle)$.*

We overload the notation $\text{STS}(\langle p, \alpha \rangle)$ for the conflict-free stochastic transition system of $\langle p, \alpha \rangle$. In Fig. 4 we give an example of a stochastic transition system of $p \equiv (\sigma_X.\sigma_X.\epsilon + \sigma_X.a.\epsilon) \oplus \sigma_X.\epsilon$ to illustrate the rules of the structural operational semantics. We choose $p' \equiv (\sigma_X.\sigma_Y.\epsilon + \sigma_X.a.\epsilon) \oplus \sigma_Z.\epsilon$, where $\text{mdr}(p', p)$, as a conflict-free term of p . For clarity we give only the relevant part of the environment and the relevant durations of the stochastic delays. Although we show all possible transitions, not all of them are necessarily enabled because of the race condition (e.g., if $P(X = Z) = 0$, the leftmost top transition does not exist).

Stochastic Bisimulation In defining a suitable process equivalence for stochastic transition systems, we follow the standard approach [16, 10]. We require the bisimulation to be an equivalence, such that every two states from the same class (1) perform the same labeled transitions, (2) perform subsequent stochastic delay transitions to every other class with the same duration and the same accumulative probability, and (3) have the same termination options.

Definition 7 *Let $R \subseteq \mathcal{S} \times \mathcal{S}$ be an equivalence relation and $C \in \mathcal{S}/R$ an arbitrary class. The accumulative probability of doing stochastic delay transitions from a state $u \in \mathcal{S}$ to an equivalence class C with duration d is given as $P_{\text{acc}}(u, C, d) = \sum_{u' \in C} P_u(u \xrightarrow{S, d} u')$. Then R is a stochastic bisimulation if the following conditions hold for all $(u, v) \in R$ and $a \in \mathcal{A}$:*

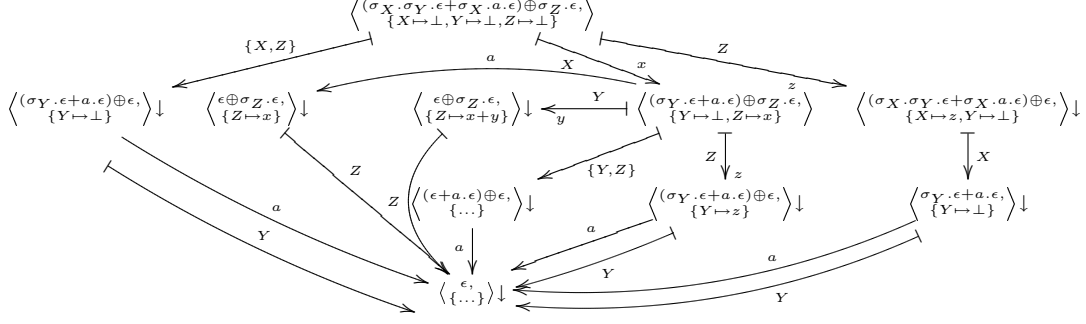


Figure 4. Example stochastic transition system

1. If $u \xrightarrow{a} u'$, then there exists $v' \in \mathcal{S}$, such that $v \xrightarrow{a} v'$ and $(u', v') \in R$.
2. $P_{\text{acc}}(u, d, C) = P_{\text{acc}}(v, d, C)$ for all $d \geq 0$ and $C \in \mathcal{S}/R$.
3. If $u \downarrow$ then $v \downarrow$.

5. Basic Communicating Processes

We add an ACP-style parallel composition operator to BSP^{dst} and obtain the algebra $\text{BCP}^{\text{dst}}(\mathcal{A}, \mathcal{V}, \gamma)$ of Basic Communication Processes with discrete stochastic time, where γ is the ACP-style communication function. As the parallel composition allows both for interleaving and communication of immediate actions, in the present setting it should also cater for interleaving and synchronization of stochastic delays. Similarly to BSP^{dst} , we introduce two types of parallel operators: (1) dependent $- \parallel -$, which enforces stochastic delays in the same race guided by the same random variable to always exhibit the same duration, and (2) independent $- \otimes -$, which treats such delays as equally distributed, but independent stochastic delays. As in real-time process algebras, we merge the delays in case the processes perform stochastic delays of different duration. We synchronize the processes in case their stochastic delays obtain the same value. Immediate actions always take precedence over passage of time in the parallel composition, but do not disable any stochastic delays.

Also, we introduce the standard auxiliary operators left merge $- \ll -$, synchronization $- | -$, encapsulation $\partial_H(-)$ and maximal progress $\theta_H(-)$, for $H \subseteq \mathcal{A}$.

Definition 8 *The signature of BCP^{dst} contains two constants δ and ϵ , four unary operator schemes $a \cdot$, for $a \in \mathcal{A}$, $\sigma_X \cdot$, for $X \in \mathcal{V}$, $\partial_H(-)$ and $\theta_H(-)$, for $H \subseteq \mathcal{A}$, and six binary operators $- \parallel -$, $- | -$, $- \ll -$,*

$- \otimes -$, $- + -$, $- \oplus -$ (in order of precedence). The syntax of BCP^{dst} is given as follows:

$$P ::= \delta \mid \epsilon \mid a.P \mid \sigma_X.P \mid \partial_H(P) \mid \theta_H(P) \mid P + P \mid P \oplus P \mid P \parallel P \mid P \otimes P \mid P \ll P \mid (P \mid P),$$

where $a \in \mathcal{A}$, $X \in \mathcal{V}$ and $H \subseteq \mathcal{A}$.

The encapsulation scheme $\partial_H(p)$ comprises processes for which the actions in H are hidden. The maximal progress operator scheme $\theta_H(p)$ enables actions to execute as soon as they become available by disabling the weak choice. The parallel compositions $p \parallel q$ and $p \otimes q$ impose a race condition in the same way as the dependent and the independent alternative composition, respectively. The auxiliary operators $p \ll q$ and $p \mid q$ impose dependent race conditions as they are used in the axiomatization of dependent parallel composition \parallel , an issue not elaborated further in this paper.

We extend the definition of a stochastic transition system by putting $\mathcal{S} = \mathcal{C}(\text{BCP}^{\text{dst}}) \times \mathcal{E}$. The definitions of $A(-)$ and $R(-)$ are extended straightforwardly to apply to the new operators $\diamond \in \{\parallel, \otimes, \ll, \mid\}$ as follows:

$$\begin{aligned} A(p \diamond q) &= A(p) \cup A(q) & A(\partial_H(p)) &= A(\theta_H(p)) = A(p) \\ R(p \diamond q) &= R(p) \cup R(q) & R(\partial_H(p)) &= A(\theta_H(p)) = R(p) \end{aligned}$$

We give the operational semantics of the additional operators in Table 2. For the sake of compactness of notation we put $\circ \in \{\parallel, \otimes\}$ for the common rules.

We briefly discuss the new rules. Rule 13 states that the parallel composition has the termination option when both operands have a termination option. Rule 14 enables interleaving of actions and rule 16 allows for synchronization of actions defined by γ . Rule 18 allows for successful termination of the right if the left operand can do a stochastic delay. Rule 20 enables the race condition, similar to the Rule 10 for the alternative composition. Rule 22 enables simultaneous passage of time for the left and right operand which allows synchronization of stochastic delays that

13	$\frac{\langle p, \alpha \rangle \downarrow, \langle q, \alpha \rangle \downarrow}{\langle p \circ q, \alpha \rangle \downarrow}$	14	$\frac{\langle p, \alpha \rangle \xrightarrow{a} \langle p', \alpha \rangle}{\langle p \circ q, \alpha \rangle \xrightarrow{a} \langle p' \circ q, \alpha \rangle}$
16	$\frac{\langle p, \alpha \rangle \xrightarrow{a} \langle p', \alpha \rangle, \langle q, \alpha \rangle \xrightarrow{b} \langle q', \alpha \rangle, \gamma(a, b) = c}{\langle p \circ q, \alpha \rangle \xrightarrow{c} \langle p' \circ q', \alpha \rangle}$		
18	$\frac{\langle p, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle, \langle q, \alpha \rangle \not\rightarrow, \langle q, \alpha \rangle \downarrow}{\langle p \circ q, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle}$		
20	$\frac{\langle p, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle, \langle q, \alpha \rangle \xrightarrow{T} \langle q', \alpha'' \rangle, S \cap T = \emptyset, s < t}{\langle p \parallel q, \alpha \rangle \xrightarrow{S} \langle p' \parallel q, \alpha' _s R(q) \rangle}$		
20'	$\frac{\langle p, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle, \langle q, \alpha \rangle \xrightarrow{T} \langle q', \alpha'' \rangle, s < t}{\langle p \otimes q, \alpha \rangle \xrightarrow{S} \langle p' \otimes q, \alpha' _s R(q) \rangle}$		
22	$\frac{\langle p, \alpha \rangle \xrightarrow{S}_d \langle p', \alpha \rangle, \langle q, \alpha \rangle \xrightarrow{T}_d \langle q', \alpha'' \rangle}{\langle p \circ q, \alpha \rangle \xrightarrow{S \cup T}_d \langle p' \circ q', \alpha _d R(p \circ q) \rangle}$		
23	$\frac{\langle p, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle}{\langle \partial_H(p), \alpha \rangle \xrightarrow{S} \langle \partial_H(p'), \alpha \rangle}$	24	$\frac{\langle p, \alpha \rangle \xrightarrow{a} \langle p', \alpha \rangle}{\langle p \parallel q, \alpha \rangle \xrightarrow{a} \langle p' \parallel q, \alpha \rangle}$
25	$\frac{\langle p, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle, \langle q, \alpha \rangle \not\rightarrow, \langle q, \alpha \rangle \downarrow}{\langle p \parallel q, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle}$		
26	$\frac{\langle p, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle, \langle q, \alpha \rangle \xrightarrow{T} \langle q', \alpha'' \rangle, S \cap T = \emptyset, s < t}{\langle p \parallel q, \alpha \rangle \xrightarrow{S} \langle p' \parallel q, \alpha' _s R(q) \rangle}$		
27	$\frac{\langle p, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle, \langle q, \alpha \rangle \xrightarrow{T} \langle q', \alpha'' \rangle, S \cap T = \emptyset, s > t}{\langle p \parallel q, \alpha \rangle \xrightarrow{T} \langle p' \parallel q, \alpha' _t R(p) \rangle}$		
28	$\frac{\langle p, \alpha \rangle \downarrow, \langle q, \alpha \rangle \downarrow}{\langle p q, \alpha \rangle \downarrow}$		
29	$\frac{\langle p, \alpha \rangle \xrightarrow{a} \langle p', \alpha \rangle, \langle q, \alpha \rangle \xrightarrow{b} \langle q', \alpha \rangle, \gamma(a, b) = c}{\langle p q, \alpha \rangle \xrightarrow{c} \langle p' q', \alpha \rangle}$		
30	$\frac{\langle p, \alpha \rangle \xrightarrow{S}_d \langle p', \alpha \rangle, \langle q, \alpha \rangle \xrightarrow{T}_d \langle q', \alpha'' \rangle}{\langle p q, \alpha \rangle \xrightarrow{S \cup T}_d \langle p' q', \alpha _d R(p q) \rangle}$		
31	$\frac{\langle p, \alpha \rangle \downarrow}{\langle \partial_H(p), \alpha \rangle \downarrow}$	32	$\frac{\langle p, \alpha \rangle \xrightarrow{a} \langle p', \alpha \rangle, a \notin H}{\langle \partial_H(p), \alpha \rangle \xrightarrow{a} \langle \partial_H(p'), \alpha \rangle}$
33	$\frac{\langle p, \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle, \langle p, \alpha \rangle \not\rightarrow_a, a \in H}{\langle \theta_H(p), \alpha \rangle \xrightarrow{S} \langle p', \alpha \rangle}$		

Table 2. SOS for BCP^{dst}

exhibit the same duration. Rules 20' and 21' describe the behavior of the independent alternative composition that differs from the dependent one. The rest of the rules define the behavior of the auxiliary operators and we will not embark on their detailed explanation.

Bisimulation remains unaltered. Similar to the alternative compositions, the stochastic transition systems of the parallel compositions may exhibit conflicting behavior. We straightforwardly extend $D(-)$, $\text{cf}_r(-)$ and $\text{mdr}(-, -)$, for $\diamond \in \{\parallel, |, |\}$, \otimes and ∂_H as follows:

$$\begin{aligned}
D(p \diamond q) &= D(p) \cup D(q) & D(p \otimes q) &= \emptyset \\
D(\partial_H(p)) &= D(\partial_H(p)) = D(p) \\
\text{cf}_r(p \diamond q, p' \diamond q') &\text{ if } \text{cf}_r(p, p') \wedge \text{cf}_r(q, q') \\
\text{cf}_r(\partial_H(p)) &\text{ if } \text{cf}_r(p) & \text{cf}_r(\theta_H(p)) &\text{ if } \text{cf}_r(p) \\
&& \text{cf}_r(p \otimes q, p' \otimes q'),
\end{aligned}$$

$$\begin{aligned}
&\text{mdr}(\langle p' \diamond q', \alpha' \rangle, \langle p \diamond q, \alpha \rangle) \\
&\text{ if } \text{cf}_r(p \diamond q, p' \diamond q') \wedge r(D(p \diamond q)) = D(p' \diamond q') \wedge \\
&\quad \text{mdr}(\langle p', \alpha' \rangle, \langle p, \alpha \rangle) \wedge \text{mdr}(\langle q', \alpha' \rangle, \langle q, \alpha \rangle) \\
&\text{mdr}(\langle p' \otimes q', \alpha' \rangle, \langle p \otimes q, \alpha \rangle) \\
&\text{ if } A(p') \cap A(q') = \emptyset \wedge \\
&\quad \text{mdr}(\langle p', \alpha' \rangle, \langle p, \alpha \rangle) \wedge \text{mdr}(\langle q', \alpha' \rangle, \langle q, \alpha \rangle) \\
&\text{mdr}(\langle \partial_H(p'), \alpha' \rangle, \langle \partial_H(p), \alpha \rangle) \text{ if } \text{mdr}(\langle p', \alpha' \rangle, \langle p, \alpha \rangle) \\
&\text{mdr}(\langle \theta_H(p'), \alpha' \rangle, \langle \theta_H(p), \alpha \rangle) \text{ if } \text{mdr}(\langle p', \alpha' \rangle, \langle p, \alpha \rangle).
\end{aligned}$$

In our previous work [18], we concluded that the effect of the winning delays for the losing ones in the presence of weak choice and α -conversion prevented the postulation of a standard expansion law involving the independent alternative composition. However, for the dependent parallel composition, the dependent alternative composition makes it relatively straightforward to obtain this.

6. Modeling $G/G/1/\infty$

Real-time delays can be expressed in our stochastic process algebra by means of degenerated random variables (cf. [18]). We capture (discrete) real-time delays by putting $\sigma^t \equiv \sigma_{D_t}$. This embedding of real-time supports delayable actions: $\underline{a} \equiv a + \sum_{i=1}^{\infty} \sigma^i . a$, for $a \in \mathcal{A}$.

Next, we model the $G/G/1/\infty$ queue, a natural example for a stochastic setting. We assume an arrival rate distributed by F and service time of finite distribution G , say $|\text{supp}(G)| = n$. The queue is given by

$$\begin{aligned}
A &= \sigma_X . s_1 . A \\
Q_0 &= r_1 . Q_1 \\
Q_{k+1} &= r_1 . Q_{k+2} + s_2 . Q_k \quad (k \geq 0) \\
S &= r_2 . \sigma_Y . s_3 . S,
\end{aligned}$$

where X and Y have distributions F and G .

The recursive equation for A models the arrival process to offer a job to the queue after a stochastic delay distributed by F . The queue is modeled as usual by the equations Q_k , $k \in \mathbb{N}$. Note the use of delayable actions, as the queue is always able to receive a new job or to offer a job that is already queued. The server is modeled by the equation for S . It is always ready to accept a job when it is idle or while processing a job with a work-time distributed by G . The specification of the $G/G/1/\infty$ queue itself is given by

$$\theta_I(\partial_H((A \parallel Q_0) \parallel S)),$$

where $\gamma(s_i, r_i) = c_i$ defines the communication, $I = \{c_1, c_2, s_3\}$ enables the instantaneous communication and $H = \{s_1, r_1, s_2, r_2\}$ encapsulates unmatched actions. Note that \parallel is used, rather than \otimes , as there are no stochastic delays for the same variable. The operator \otimes would have been mandatory, if the queue




Figure 5. $G/G/1/\infty$ queue

contained more than one server or more than one arrival process, as in the $G/G/3/\infty$ queue for example, that is specified by $\theta_I(\partial_H((A \parallel Q_0) \parallel ((S \otimes S) \otimes S)))$.

In Fig. 5 we give the stochastic transition system of the $G/G/1/\infty$ queue as specified above. The first index of every recursive variable represents the number of jobs in the queue, whereas the second index indicates the number of jobs that are enqueued since processing of the last job has started. This is required as we need to keep track of the spent time of the delays. By assumption, there are at most n arrivals during the processing of a job. Note that we have $\varphi(X_j) = F$ and $\varphi(X_{ij}) = \varphi(Y_{ij}) = G$. The solution of the above specification of the queue is given by

$$Q_{00} = \langle \sigma_{X_0} \cdot c_1 \cdot \theta_I(\partial_H((A \parallel Q_1) \parallel S)), \{X_0 \mapsto \perp\} \rangle$$

$$Q_{0j} = \langle \sigma_{X_j} \cdot c_1 \cdot Q_{10}, \{X_j \mapsto y_{0,j-1}\} \rangle$$

$$Q_{i0} = \langle c_2 \cdot \theta_I(\partial_H((A \parallel Q_{i-1}) \parallel \sigma_{Y_{i-1,0}} \cdot s_3 \cdot S)), \{X_{i-1,0} \mapsto \perp\} \rangle$$

$$Q_{ij} = \langle c_2 \cdot \theta_I(\partial_H((A \parallel Q_{i-1}) \parallel \sigma_{Y_{i-1,j}} \cdot s_3 \cdot S)), \{X_{i-1,j} \mapsto y_{i,j-1}\} \rangle$$

$$S_{i-1,0} = \langle \sigma_{X_{i-1,0}} \cdot c_1 \cdot \theta_I(\partial_H((A \parallel Q_{i-1,1}) \parallel \sigma_{Y_{i-1,1}} \cdot s_3 \cdot S + \sigma_{Y_{i-1,0}} \cdot s_3 \cdot Q_{i-1})), \{X_{i-1,0} \mapsto \perp\} \rangle$$

$$S_{i-1,j} = \langle \sigma_{X_{i-1,j}} \cdot c_1 \cdot \theta_I(\partial_H((A \parallel Q_{i-1}) \parallel \sigma_{Y_{i-1,j}} \cdot s_3 \cdot S + \sigma_{Y_{i-1,j-1}} \cdot s_3 \cdot Q_{i-1,j})), \{X_{i-1,j} \mapsto y_{i,j-1}, Y_{i-1,j} \mapsto \sum_{k=0}^{j-1} x_{i-k,k}\} \rangle.$$

The extra states arise from the explicit aging of the stochastic delays. A typical modeling of a $G/G/1/\infty$ queue by a using generalized semi-Markov process [12] is given in Fig. 6, where a is an event of an arrival

job that resets a clock with distribution F and s is an event of a processed job that resets a clock with distribution G .




Figure 6. $G/G/1/\infty$ queue (revisited)

We note that our model of the $G/G/1/\infty$ queue of Fig. 5 collapses to the generalized semi-Markov process given in Fig. 6, as all states Q_{ij} and $S_{i-1,j}$, for $j \in \{1, \dots, n\}$, can be joined together in a state Q_i in the presence of clocks with residual lifetime semantics. In Q_i the clocks for s_3 and c_1 are reset, whereas c_2 vanishes as an immediate action that connects Q_{ij} and $S_{i-1,j}$. There is the additional requirement that s_3 and c_1 cannot happen together, i.e., $P(X = Y) = 0$, because of the unique occurrences of these events. We conclude that the explicit bookkeeping of spent lifetimes yields an infinite stochastic transition system if the underlying distribution function is of infinite support. However, the transition systems have a repetitive layered structure as the one in Fig. 5. This structure can be exploited for model checking purposes, as in the case, for example, of quasi birth-death processes (see [21]).

7. Conclusions and Future Work

We have extended a previous version of a stochastic process algebra with dependent alternative composition to cater for the parallel composition. We revis-

ited a notion of stochastic bisimulation and extended it to apply to subsequent stochastic delays. We showed how clashes of stochastic delays can be dealt with using α -conversion. We exploited the embedding of real-time into our algebra to compose delayable actions out of timed delays and immediate actions. With the delayable actions we modeled the $G/G/1/\infty$ queue, showing the interaction between real-time and stochastic time. Because of the explicit treatment of spent lifetimes of delays we obtained a model that has more detail than the other clock-based approaches.

As future work, we continue our axiomatization effort. Presently, we do have a sound theory that we expect to be ground-complete. However, proving this requires an enhancement of standard techniques. By construction, the theory conservatively extends real-time process algebra [3]. However, the mix of real-time and stochastic delays may require the addition of an explicit probabilistic choice. Furthermore, we tend to introduce the abstraction operator that produces silent transitions, also called τ -steps, and the notion of branching or weak bisimulation in that setting. For verification purposes, it will be advantageous to rely on an observational congruence. However, obtaining a congruence result may require a substantial effort [5, 6, 8]. Also, we plan to extend the current setting with continuous stochastic time. Afterwards, we will consider case studies, in particular verification of Internet protocols, as a successful modeling of real-time delays paves the way for a convenient specification of time-outs. Accurate performance specification is feasible with general distributions, like, for example, the heavy-tail distributions that model media information flow.

References

- [1] M. Ajmone Marsan, G. Balbo, A. Bobbio, G. Chiola, G. Conte, and A. Cumani. The effect of execution policies on the semantics and analysis of stochastic Petri nets. *IEEE Transactions on Software Engineering*, 15(7):832–846, 1989.
- [2] M. Ajmone Marsan, A. Bianco, L. Ciminiera, R. Sisto, and A. Valenzano. A LOTOS extension for the performance analysis of distributed systems. *IEEE/ACM Transactions on Networking*, 2(2):151–165, 1994.
- [3] J. Baeten and C. Middelburg. *Process Algebra with Timing*. Monographs in Theoretical Computer Science. Springer, 2002.
- [4] J. Baeten and M. Reniers. Timed process algebra (with a focus on explicit termination and relative timing). In M. Bernardo and F. Corradini, editors, *Proc. SFM 2004*, pages 59–97. LNCS 3185, 2004.
- [5] E. Bandini and R. Segala. Axiomatizations for probabilistic bisimulation. In F. Orejas, P. Spirakis, and J. van Leeuwen, editors, *Proc. ICALP 2001*, pages 370–381. LNCS 2076, 2001.
- [6] M. Bravetti. *Specification and Analysis of Stochastic Real-time Systems*. PhD thesis, Università di Bologna, 2002.
- [7] M. Bravetti and P. D’Argenio. Tutte le algebre insieme – concepts, discussions and relations of stochastic process algebras with general distributions. In C. Baier et al., editor, *Validation of Stochastic Systems*, pages 44–88. LNCS 2925, 2004.
- [8] S. Cattani, R. Segala, M. Kwiatkowska, and G. Norman. Stochastic transition systems for continuous state spaces and non-determinism. In V. Sassone, editor, *Proc. FoSSaCS*, pages 125–139. LNCS 3441, 2005.
- [9] P. D’Argenio. From stochastic automata to timed automata: Abstracting probability in a compositional manner. In M. Fiore and D. Fridlender, editors, *Proc. WAIT 2003*, Buenos Aires, 2003.
- [10] P. D’Argenio and J.-P. Katoen. A theory of stochastic systems, part I: Stochastic automata. *Information and Computation*, 203(1):1–38, 2005.
- [11] P. D’Argenio and J.-P. Katoen. A theory of stochastic systems, part II: Process algebra. *Information and Computation*, 203(1):39–74, 2005.
- [12] P. Glynn. A GSMP formalism for discrete event systems. *Proceedings of the IEEE*, 77(1):14–23, 1989.
- [13] H. Hermanns. *Interactive Markov Chains and the Quest For Quantified Quantity*. LNCS 2428, 2002.
- [14] H. Hermanns, V. Mertsiotakis, and M. Rettelbach. Performance analysis of distributed systems using TIPP. In R.J. Pooley et al., editor, *Proc. UKPEW’94*, pages 131–144. University of Edinburgh, 1994.
- [15] J.-P. Katoen and P. R. D’Argenio. General distributions in process algebra. In H. Brinksma et al., editor, *Lectures on Formal Methods and Performance Analysis*, pages 375–429. LNCS 2090, 2001.
- [16] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, 1991.
- [17] N. Lopez and M. Nunez. NMSPA: A non-Markovian model for stochastic processes. In T.-H. Lai, editor, *Proc. ICDS 2000*, pages 33–40. IEEE, 2000.
- [18] J. Markovski and E.P. de Vink. Embedding real-time in stochastic process algebras. In A. Horvath and M. Telek, editors, *Proc. EPEW 2006*, pages 47–62. LNCS 4054, 2006.
- [19] X. Nicollin and J. Sifakis. An overview and synthesis of timed process algebras. In J.W. de Bakker et al., editor, *Real-Time: Theory in Practice*, pages 526–548. LNCS 600, 1992.
- [20] C. Priami. Stochastic π -calculus with general distributions. In M. Ribaud, editor, *Proc. PAPM 1996*, pages 41–57, Torino, 1996.
- [21] A. Remke, B. Haverkort, and L. Cloth. Model checking infinite-state Markov chains. In N. Halbwachs and L. Zuck, editors, *Proc. TACAS 2005*, pages 237–252. LNCS 3440, 2005.