Input/Output Stochastic Automata with Urgency: Confluence and Weak Determinism

Pedro R. D'Argenio[1,2,3]([✉]) and Raúl E. Monti[1,2]

[1] Universidad Nacional de Córdoba, FAMAF, Córdoba, Argentina
{dargenio,rmonti}@famaf.unc.edu.ar
[2] CONICET, Córdoba, Argentina
[3] Saarland University, Department of Computer Science, Saarbrücken, Germany

Abstract. In a previous work, we introduced an input/output variant of stochastic automata (IOSA) that, once the model is closed (i.e., all synchronizations are resolved), the resulting automaton is fully stochastic, that is, it does not contain non-deterministic choices. However, such variant is not sufficiently versatile for compositional modelling. In this article, we extend IOSA with urgent actions. This extension greatly increases the modularization of the models, allowing to take better advantage on compositionality than its predecessor. However, this extension introduces non-determinism even in closed models. We first show that confluent models are weakly deterministic in the sense that, regardless the resolution of the non-determinism, the stochastic behaviour is the same. In addition, we provide sufficient conditions to ensure that a network of interacting IOSAs is confluent without the need to analyse the larger composed IOSA.

1 Introduction

The advantages of compositional modelling complex systems can hardly be overestimated. On the one hand, compositional modelling facilitates systematic design, allowing the designer to focus on the construction of small models for the components whose operational behavior is mostly well understood, and on the synchronization between the components, which are in general quite evident. On the other hand, it facilitates the interchange of components in a model, enables compositional analysis, and helps on attacking the state explosion problem.

In particular we focus on modelling of stochastic system for dependability and performance analysis, and aim to general models that require more than the usual negative exponential distribution. Indeed, phenomena such as timeouts in communication protocols, hard deadlines in real-time systems, human response times or the variability of the delay of sound and video frames (so-called jitter)

in modern multi-media communication systems are typically described by non-memoryless distributions such as uniform, log-normal, or Weibull distributions.

The analysis of this type of model quite often can only be performed through discrete event simulation [22]. However, simulation requires that the model under study is fully stochastic, that is, they should not contain non-deterministic choices. Unfortunately, compositional modelling languages such as stochastic process algebras with general distributions (see [5] and references therein) and Modest [4,18,19], were designed so that the non-determinism arises naturally as the result of composition.

Based on stochastic automata [10–12] and probabilistic I/O automata [26], we introduced input/output stochastic automata (IOSA) [13]. IOSAs were designed so that parallel composition works naturally and, moreover, the system becomes fully stochastic –not containing non-determinism– when closed, i.e., when all interactions are resolved and no input is left available in the model. IOSA splits the set of actions into inputs and outputs and let them behave in a reactive and generative manner respectively [17]. Thus, inputs are passive and their occurrence depends only on their interaction with outputs. Instead, occurrence of outputs are governed by the expiration of a timer which is set according to a given random variable. In addition, and not to block the occurrence of outputs, IOSAs are required to be input enabled.

We have used IOSA as input language of the rare event simulation tool FIG [6,7] and have experienced the limitations of the language, in particular when transcribing models originally given in terms of variants of dynamic fault trees (DFT) with repairs [24]. To illustrate the problem, suppose the simple digital system of Fig. 1.
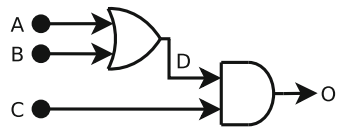

Fig. 1. A simple digital system.

We would like to measure the average time that the output O is 1 given that we know the distributions of the times in which the values on inputs A, B, and C change from 0 to 1 and vice-versa. The natural modelling of such system is to define 5 IOSA modules, three of them modelling the behaviour of the input signals and the other two modelling the OR and AND gates. Then we compose and synchronize the 5 modules properly. The main problem is that, while the dynamic behaviour of the input signal modules are governed by stochastically timed actions, the dynamic behavior of the gates are instantaneous and thus, for instance the output D of the OR gate, may change immediately after the arrival of signals A or B. Similar situations arise when modeling the behaviour of DFT under complex gates like priority AND, Spares or Repair boxes. As a consequence, we observe that the introduction of urgent actions will allow for a direct and simple compositional modelling of situations like the one recently described. Also, it is worth to notice that the need for instantaneous but causally dependent synchronization have been observed in many other timed modelling languages, notably, in Uppaal, with the introduction of committed locations, urgent locations and urgent synchronization [2,3].

Based on IMC [20] and, particularly, on I/O-IMC [9], in this article we extended IOSA with urgent actions (Sect. 2). Urgent actions are also partitioned in input and output actions and, though inputs behave reactively and passively as before, urgent outputs are executed instantaneously as soon as the enabling state is reached. We also give semantics to IOSA with urgent actions (from now on, we simply call it IOSA) in terms of NLMP [14,25] (Sect. 3), and define its parallel composition (Sect. 4).

The problem is that urgent actions on IOSA introduce non-determinism. Fortunately, non-determinism is limited to urgent actions and, in many occasions, it is introduced by confluent urgent output actions as a result of a parallel composition. Such non-determinism turns to be spurious in the sense that it does not change the stochastic behaviour of the model. In this paper, we characterize confluence on IOSAs (Sect. 5), define the concept of weak determinism, and show that a confluent closed IOSA is weakly deterministic (Sect. 6). Notably, a weakly deterministic IOSA is amenable to discrete event simulation. Milner [23] has provided a proof that confluence preserves weak determinism but it is confined to a discrete non-probabilistic setting. A similar proof has been used by Crouzen [9] on I/O-IMC but, though the model is stochastic, the proof is limited to discrete non-probabilistic transitions. Contrarily, our proof has to deal with continuous probabilities (since urgent action may sample on continuous random variables), hence making use of the solid measure theoretical approach. In particular, we address the complications of defining a particular form of weak transition on a setting that is normally elusive.

Based on the work of Crouzen [9] for I/O-IMC, in Sect. 7, we provide sufficient conditions to ensure that a closed IOSA is confluent and hence, weakly deterministic. If the IOSA is the result of composing several smaller IOSAs, the verification of the conditions is performed by inspecting the components rather than the resulting composed IOSA.

2   Input/Output Stochastic Automata with Urgency

Stochastic automata [10,11] use continuous random variables (called *clocks*) to observe the passage of time and control the occurrence of events. These variables are set to a value according to their associated probability distribution, and, as time evolves, they count down at the same rate. When a clock reaches zero, it may trigger some action. This allows the modelling of systems where events occur at random continuous time steps.

Following ideas from [26], IOSAs restrict Stochastic Automata by splitting actions into input and output actions which will act in a reactive and generative way respectively [17]. This splitting reflects the fact that input actions are considered to be controlled externally, while output actions are locally controlled.

Therefore, we consider the system to be input enabled. Moreover, output actions could be stochastically controlled or instantaneous. In the first case, output actions are controlled by the expiration of a single clock while in the second case the output actions take place as soon as the enabling state is reached.

We called these instantaneous actions *urgent*. A set of restrictions over IOSA will ensure that, almost surely, no two non-urgent outputs are enabled at the same time.

Definition 1. *An* input/output stochastic automaton with urgency *(IOSA) is a structure* $(\mathcal{S}, \mathcal{A}, \mathcal{C}, \rightarrow, C_0, s_0)$, *where* $\mathcal{S}$ *is a (denumerable) set of states,* $\mathcal{A}$ *is a (denumerable) set of labels partitioned into disjoint sets of* input *labels* $\mathcal{A}^{\mathsf{i}}$ *and* output *labels* $\mathcal{A}^{\mathsf{o}}$, *from which a subset* $\mathcal{A}^{\mathsf{u}} \subseteq \mathcal{A}$ *is marked as* urgent. *We consider the distinguished* silent *urgent action* $\tau \in \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\mathsf{o}}$ *which is not amenable to synchronization.* $\mathcal{C}$ *is a (finite) set of clocks such that each* $x \in \mathcal{C}$ *has an associated continuous probability measure* μ_x *on* $\mathbb{R}$ *s.t.* $\mu_x(\mathbb{R}_{>0}) = 1$, $\rightarrow \subseteq \mathcal{S} \times \mathcal{C} \times \mathcal{A} \times \mathcal{C} \times \mathcal{S}$ *is a transition function,* C_0 *is the set of clocks that are initialized in the initial state, and* $s_0 \in \mathcal{S}$ *is the initial state.*

In addition, an IOSA with urgency should satisfy the following constraints:

(a) *If* $s \xrightarrow{C,a,C'} s'$ *and* $a \in \mathcal{A}^{\mathsf{i}} \cup \mathcal{A}^{\mathsf{u}}$, *then* $C = \varnothing$.
(b) *If* $s \xrightarrow{C,a,C'} s'$ *and* $a \in \mathcal{A}^{\mathsf{o}} \setminus \mathcal{A}^{\mathsf{u}}$, *then* C *is a singleton set.*
(c) *If* $s \xrightarrow{\{x\},a_1,C_1} s_1$ *and* $s \xrightarrow{\{x\},a_2,C_2} s_2$ *then* $a_1 = a_2$, $C_1 = C_2$ *and* $s_1 = s_2$.
(d) *For every* $a \in \mathcal{A}^{\mathsf{i}}$ *and state* s, *there exists a transition* $s \xrightarrow{\varnothing,a,C} s'$.
(e) *For every* $a \in \mathcal{A}^{\mathsf{i}}$, *if* $s \xrightarrow{\varnothing,a,C_1'} s_1$ *and* $s \xrightarrow{\varnothing,a,C_2'} s_2$, $C_1' = C_2'$ *and* $s_1 = s_2$.
(f) *There exists a function* active $: \mathcal{S} \to 2^{\mathcal{C}}$ *such that: (i)* active$(s_0) \subseteq C_0$, *(ii)* enabling$(s) \subseteq$ active(s), *(iii) if* s *is stable,* active$(s) =$ enabling(s), *and (iv) if* $t \xrightarrow{C,a,C'} s$ *then* active$(s) \subseteq ($active$(t) \setminus C) \cup C'$.

where enabling$(s) = \{y \mid s \xrightarrow{\{y\},\_,\_} \_\}$, *and* s *is* stable, *denoted* st(s), *if there is no* $a \in \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\mathsf{o}}$ *such that* $s \xrightarrow{\varnothing,a,\_} \_$. *(_ indicates the existential quantification of a parameter.)*

The occurrence of an output transition is controlled by the expiration of clocks. If $a \in \mathcal{A}^{\mathsf{o}}$, $s \xrightarrow{C,a,C'} s'$ indicates that there is a transition from state s to state s' that can be taken only when all clocks in C have expired and, when taken, it triggers action a and sets all clocks in C' to a value sampled from their associated probability distribution. Notice that if $C = \varnothing$ (which means $a \in \mathcal{A}^{\mathsf{o}} \cap \mathcal{A}^{\mathsf{u}}$) $s \xrightarrow{C,a,C'} s'$ is immediately triggered. Instead, if $a \in \mathcal{A}^{\mathsf{i}}$, $s \xrightarrow{\varnothing,a,C'} s'$ is only intended to take place if an external output synchronizes with it, which means, in terms of an open system semantics, that it may take place at any possible time.

Restrictions (a) to (e) ensure that any *closed* IOSA without urgent actions is deterministic [13]. An IOSA is closed if all its synchronizations have been resolved, that is, the IOSA resulting from a composition does not have input actions ($\mathcal{A}^{\mathsf{i}} = \varnothing$). Restriction (a) is two-folded: on the one hand, it specifies that output actions must occur as soon as the enabling state is reached, on the other hand, since input actions are reactive and their time occurrence can only depend on the interaction with an output, no clock can control their enabling.

Restriction (b) specifies that the occurrence of a non-urgent output is locally controlled by a single clock. Restriction (c) ensures that two different non-urgent output actions leaving the same state are always controlled by different clocks (otherwise it would introduce non-determinism). Restriction (d) ensures input enabling. Restriction (e) determines that IOSAs are input deterministic. Therefore, the same input action in the same state can not jump to different states, nor set different clocks. Finally, (f) guarantees that clocks enabling some output transition have not expired before, that is, they have not been used before by another output transition (without being reset in between) nor inadvertently reached zero. This is done by ensuring the existence of a function "active" that, at each state, collects clocks that are required to be active (i.e. that have been set but not yet expired). Notice that enabling clocks are required to be active (conditions (f)(ii) and (f)(iii)). Also note that every clock that is active in a state is allowed to remain active in a successor state as long as it has not been used, and clocks that have just been set may become active in the successor state (condition (f)(iv)).

Note that since clocks are set by sampling from a continuous random variable, the probability that the values of two different clocks are equal is 0. This fact along with restriction (c) and (f) guarantee that almost never two different non-urgent output transitions are enabled at the same time.

Example 1. Figure 2 depicts three simple examples of IOSAs. Although IOSAs are input enabled, we have omitted self loops of input enabling transitions for the sake of readability. In the figure, we represent output actions suffixed by '!' and by '!!' when they are urgent, and input actions suffixed by '?' and by '??' when they are urgent.

Fig. 2. Examples of IOSAs.
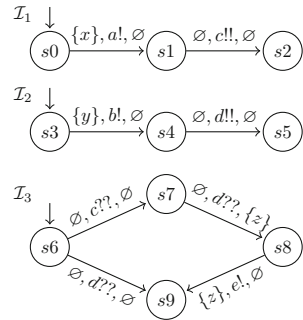
3 Semantics of IOSA

The semantics of IOSA is defined in terms of non-deterministic labeled Markov processes (NLMP) [14,25] which extends LMP [15] with *internal* non-determinism.

The foundations of NLMP is strongly rooted in measure theory, hence we recall first some basic definitions. Given a set S and a collection Σ of subsets of S, we call Σ a *σ-algebra* iff $S \in \Sigma$ and Σ is closed under complement and denumerable union. We call the pair (S, Σ) a *measurable space*. Let $\mathscr{B}(S)$ denote the Borel σ-algebra on the topology S. A function $\mu : \Sigma \to [0, 1]$ is a *probability measure* if (i) $\mu(\bigcup_{i \in \mathbb{N}} Q_i) = \sum_{i \in \mathbb{N}} \mu(Q_i)$ for all countable family of pairwise disjoint measurable sets $\{Q_i\}_{i \in \mathbb{N}} \subseteq \Sigma$, and (ii) $\mu(S) = 1$. In particular, for $s \in S$, δ_s denotes the Dirac measure so that $\delta_s(\{s\}) = 1$. Let $\Delta(S)$ denote the set of all probability measures over (S, Σ). Let (S_1, Σ_1) and (S_2, Σ_2) be two measurable spaces. A function $f : S_1 \to S_2$ is said to be *measurable* if for all

$Q_2 \in \Sigma_2$, $f^{-1}(Q_2) \in \Sigma_1$. There is a standard construction to endow $\Delta(S)$ with a σ-algebra [16] as follows: $\Delta(\Sigma)$ is defined as the smallest σ-algebra containing the sets $\Delta^q(Q) \doteq \{\mu \mid \mu(Q) \geq q\}$, with $Q \in \Sigma$ and $q \in [0,1]$. Finally, we define the *hit σ-algebra* $H(\Delta(\Sigma))$ as the minimal σ-algebra containing all sets $H_\xi = \{\zeta \in \Delta(\Sigma) \mid \zeta \cap \xi \neq \varnothing\}$ with $\xi \in \Delta(\Sigma)$.

A *non-deterministic labeled Markov process* (NLMP for short) is a structure $(\mathbf{S}, \Sigma, \{\mathcal{T}_a \mid a \in \mathcal{L}\})$ where Σ is a σ-algebra on the set of states $\mathbf{S}$, and for each label $a \in \mathcal{L}$ we have that $\mathcal{T}_a : \mathbf{S} \to \Delta(\Sigma)$ is measurable from Σ to $H(\Delta(\Sigma))$.

The formal semantics of an IOSA is defined by a NLMP with two classes of transitions: one that encodes the discrete steps and contains all the probabilistic information introduced by the sampling of clocks, and another describing the time steps, that only records the passage of time synchronously decreasing the value of all clocks. For simplicity, we assume that the set of clocks has a total order and their current values follow the same order in a vector.

Definition 2. *Given an IOSA* $\mathcal{I} = (\mathcal{S}, \mathcal{A}, \mathcal{C}, \to, C_0, s_0)$ *with* $\mathcal{C} = \{x_1, \dots, x_N\}$, *its semantics is defined by the NLMP* $\mathcal{P}(\mathcal{I}) = (\mathbf{S}, \mathscr{B}(\mathbf{S}), \{\mathcal{T}_a \mid a \in \mathcal{L}\})$ *where*

- $\mathbf{S} = (\mathcal{S} \cup \{\mathsf{init}\}) \times \mathbb{R}^N$, $\mathcal{L} = \mathcal{A} \cup \mathbb{R}_{>0} \cup \{\mathsf{init}\}$, *with* $\mathsf{init} \notin \mathcal{S} \cup \mathcal{A} \cup \mathbb{R}_{>0}$
- $\mathcal{T}_{\mathsf{init}}(\mathsf{init}, \vec{v}) = \{\delta_{s_0} \times \prod_{i=1}^N \mu_{x_i}\}$,
- $\mathcal{T}_a(s, \vec{v}) = \{\mu_{C',s'}^{\vec{v}} \mid s \xrightarrow{C,a,C'} s', \bigwedge_{x_i \in C} \vec{v}(i) \leq 0\}$, *for all* $a \in \mathcal{A}$, *where* $\mu_{C',s'}^{\vec{v}} = \delta_{s'} \times \prod_{i=1}^N \overline{\mu}_{x_i}$ *with* $\overline{\mu}_{x_i} = \mu_{x_i}$ *if* $x_i \in C'$ *and* $\overline{\mu}_{x_i} = \delta_{\vec{v}(i)}$ *otherwise, and*
- $\mathcal{T}_d(s, \vec{v}) = \{\delta_s \times \prod_{i=1}^N \delta_{\vec{v}(i)-d}\}$ *if there is no urgent* $b \in \mathcal{A}^\circ \cap \mathcal{A}^{\mathsf{u}}$ *for which* $s \xrightarrow{-,b,-} \_$ *and* $0 < d \leq \min\{\vec{v}(i) \mid \exists_{a \in \mathcal{A}^\circ}, C' \subseteq \mathcal{C}, s' \in \mathcal{S} : s \xrightarrow{\{x_i\},a,C'} s'\}$, *and* $\mathcal{T}_d(s, \vec{v}) = \varnothing$ *otherwise, for all* $d \in \mathbb{R}_{\geq 0}$.

The state space is the product space of the states of the IOSA with all possible clock valuations. A distinguished initial state init is added to encode the random initialization of all clocks (it would be sufficient to initialize clocks in C_0 but we decided for this simplification). Such encoding is done by transition $\mathcal{T}_{\mathsf{init}}$. The state space is structured with the usual Borel σ-algebra. The discrete step is encoded by $\mathcal{T}_a$, with $a \in \mathcal{A}$. Notice that, at state $(s, \vec{v})$, the transition $s \xrightarrow{C,a,C'} s'$ will only take place if $\bigwedge_{x_i \in C} \vec{v}(i) \leq 0$, that is, if the current values of all clocks in C are not positive. For the particular case of the input or urgent actions this will always be true. The next actual state would be determined randomly as follows: the symbolic state will be s' (this corresponds to $\delta_{s'}$ in $\mu_{C',s'}^{\vec{v}} = \delta_{s'} \times \prod_{i=1}^N \overline{\mu}_{x_i}$), any clock not in C' preserves the current value (hence $\overline{\mu}_{x_i} = \delta_{\vec{v}(i)}$ if $x_i \notin C'$), and any clock in C' is set randomly according to its respective associated distribution (hence $\overline{\mu}_{x_i} = \mu_{x_i}$ if $x_i \in C'$). The time step is encoded by $\mathcal{T}_d(s, \vec{v})$ with $d \in \mathbb{R}_{\geq 0}$. It can only take place at d units of time if there is no output transition enabled at the current state within the next d time units (this is verified by condition $0 < d \leq \min\{\vec{v}(i) \mid \exists_{a \in \mathcal{A}^\circ}, C' \subseteq \mathcal{C}, s' \in \mathcal{S} :$ $s \xrightarrow{\{x_i\},a,C'} s'\}$). In this case, the system remains in the same symbolic state (this corresponds to δ_s in $\delta_{(s,\vec{v})}^{-d} = \delta_s \times \prod_{i=1}^N \delta_{\vec{v}(i)-d}$), and all clock values are

Table 1. Parallel composition on IOSA

$$\frac{s_1 \xrightarrow{C,a,C'}_1 s_1'}{s_1\|s_2 \xrightarrow{C,a,C'} s_1'\|s_2} \; a \in (\mathcal{A}_1 \backslash \mathcal{A}_2) \cup \{\tau\} \;\; (\text{R1}) \qquad \frac{s_2 \xrightarrow{C,a,C'}_2 s_2'}{s_1\|s_2 \xrightarrow{C,a,C'} s_1\|s_2'} \; a \in (\mathcal{A}_2 \backslash \mathcal{A}_1) \cup \{\tau\} \;\; (\text{R2})$$

$$\frac{s_1 \xrightarrow{C_1,a,C_1'}_1 s_1' \quad s_2 \xrightarrow{C_2,a,C_2'}_2 s_2'}{s_1\|s_2 \xrightarrow{C_1 \cup C_2, a, C_1' \cup C_2'} s_1'\|s_2'} \; a \in (\mathcal{A}_1 \cap \mathcal{A}_2) \backslash \{\tau\} \;\; (\text{R3})$$

decreased by d units of time (represented by $\delta_{\vec{v}(i)-d}$ in the same formula). Note the difference from the timed transitions semantics of pure IOSA [13]. This is due to the maximal progress assumption, which forces to take urgent transition as soon as they get enabled. We encode this by not allowing to make time transitions in presence of urgent actions, i.e. we check that there is no urgent $b \in \mathcal{A}^o \cap \mathcal{A}^u$ for which $s \xrightarrow{\_,b,\_} \_$. (Notice that b may be τ.) Otherwise, $\mathcal{T}_d(s,\vec{v}) = \varnothing$. Instead, notice the *patient* nature of a state $(s,\vec{v})$ that has no output enabled. That is, $\mathcal{T}_d(s,\vec{v}) = \{\delta_s \times \prod_{i=1}^{N} \delta_{\vec{v}(i)-d}\}$ for all $d > 0$ whenever there is no output action $b \in \mathcal{A}^o$ such that $s \xrightarrow{\_,b,\_} \_$.

In a similar way to [13], it is possible to show that $\mathcal{P}(\mathcal{I})$ is indeed a NLMP, i.e. that $\mathcal{T}_a$ maps into measurable sets in $\Delta(\mathscr{B}(\mathbf{S}))$, and that $\mathcal{T}_a$ is a measurable function for every $a \in \mathcal{L}$.

4   Parallel Composition

In this section, we define parallel composition of IOSAs. Since outputs are intended to be autonomous (or locally controlled), we do not allow synchronization between them. Besides, we need to avoid name clashes on the clocks, so that the intended behavior of each component is preserved and moreover, to ensure that the resulting composed automaton is indeed an IOSA. Furthermore, synchronizing IOSAs should agree on urgent actions in order to ensure their immediate occurrence. Thus we require to compose only *compatible* IOSAs.

Definition 3. *Two IOSAs $\mathcal{I}_1$ and $\mathcal{I}_2$ are* compatible *if they do not share synchronizable output actions nor clocks, i.e. $\mathcal{A}_1^o \cap \mathcal{A}_2^o \subseteq \{\tau\}$ and $\mathcal{C}_1 \cap \mathcal{C}_2 = \varnothing$ and, moreover, they agree on urgent actions, i.e. $\mathcal{A}_1 \cap \mathcal{A}_2^u = \mathcal{A}_2 \cap \mathcal{A}_1^u$.*

Definition 4. *Given two compatible IOSAs $\mathcal{I}_1$ and $\mathcal{I}_2$, the parallel composition $\mathcal{I}_1\|\mathcal{I}_2$ is a new IOSA $(\mathcal{S}_1 \times \mathcal{S}_2, \mathcal{A}, \mathcal{C}, \rightarrow, C_0, s_0^1\|s_0^2)$ where (i) $\mathcal{A}^o = \mathcal{A}_1^o \cup \mathcal{A}_2^o$ (ii) $\mathcal{A}^i = (\mathcal{A}_1^i \cup \mathcal{A}_2^i) \backslash \mathcal{A}^o$ (iii) $\mathcal{A}^u = \mathcal{A}_1^u \cup \mathcal{A}_2^u$ (iv) $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2$ (v) $C_0 = C_0^1 \cup C_0^2$ and $\rightarrow$ is defined by rules in Table 1 where we write $s\|t$ instead of (s,t).*

Definition 4 does not ensure *a priori* that the resulting structure satisfies conditions (a)–(f) in Definition 1. This is only guaranteed by the following proposition.
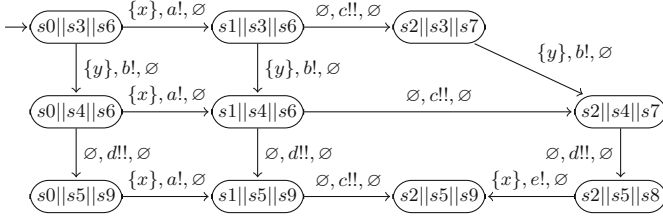
Fig. 3. IOSA resulting from the composition $\mathcal{I}_1||\mathcal{I}_2||\mathcal{I}_3$ of IOSAs in Fig. 2.

Proposition 1. *Let $\mathcal{I}_1$ and $\mathcal{I}_2$ be two compatible IOSAs. Then $\mathcal{I}_1||\mathcal{I}_2$ is indeed an IOSA.*

Example 2. The result of composing $\mathcal{I}_1||\mathcal{I}_2||\mathcal{I}_3$ from Example 1 is depicted in Fig. 3.

Larsen and Skou's probabilistic bisimulation [21] has been extended to NLMPs in [14]. It can be shown that the bisimulation equivalence is a congruence for parallel composition of IOSA. In fact, this has already been shown for IOSA without urgency in [13] and since the characteristics of urgency do not play any role in the proof over there, the result immediately extends to our setting. So we report the theorem and invite the reader to read the proof in [13].

Theorem 1. *Let $\sim$ denote the bisimulation equivalence relation on NLMPs [14] properly lifted to IOSA [13], and let $\mathcal{I}_1$, $\mathcal{I}_1'$, $\mathcal{I}_2$, $\mathcal{I}_2'$ be IOSAs such that $\mathcal{I}_1 \sim \mathcal{I}_1'$ $\mathcal{I}_2 \sim \mathcal{I}_2'$. Then, $\mathcal{I}_1||\mathcal{I}_2 \sim \mathcal{I}_1'||\mathcal{I}_2'$.*

5   Confluence

Confluence, as studied by Milner [23], is related to a form of weak determinism: two silent transitions taking place on an interleaving manner do not alter the behaviour of the process regardless of which happens first. In particular, we will eventually assume that urgent actions in a closed IOSA are silent as they do not delay the execution. Thus we focus on confluence of urgent actions only. The notion of confluence is depicted in Fig. 4 and formally defined as follows.
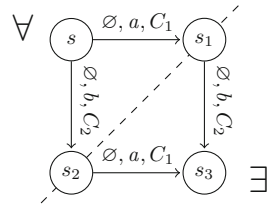
Fig. 4. Confluence in IOSA.

Definition 5. *An IOSA $\mathcal{I}$ is confluent with respect to actions $a, b \in \mathcal{A}^{\mathsf{u}}$ if, for every state $s \in \mathcal{S}$ and transitions $s \xrightarrow{\varnothing, a, C_1} s_1$ and $s \xrightarrow{\varnothing, b, C_2} s_2$, there exists a state $s_3 \in \mathcal{S}$ such that $s_1 \xrightarrow{\varnothing, b, C_2} s_3$ and $s_2 \xrightarrow{\varnothing, a, C_1} s_3$. $\mathcal{I}$ is confluent if it is confluent with respect to every pair of urgent actions.*

Note that we are asking that the two actions converge in a single state, which is stronger than Milner's strong confluence, where convergence takes place on bisimilar but potentially different states.
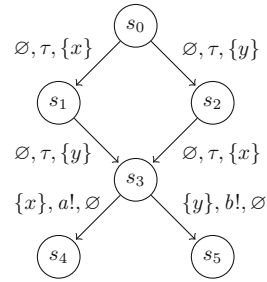
Confluence is preserved by parallel composition:

Proposition 2. *If both $\mathcal{I}_1$ and $\mathcal{I}_2$ are confluent w.r.t. actions $a, b \in \mathcal{A}^{\mathsf{u}}$, then so is $\mathcal{I}_1 || \mathcal{I}_2$. Therefore, if $\mathcal{I}_1$ and $\mathcal{I}_2$ are confluent, $\mathcal{I}_1 || \mathcal{I}_2$ is also confluent.*

However, parallel composition may turn non-confluent components into a confluent composed system.

By looking at the IOSA in Fig. 5, one can notice that the non-determinism introduced by confluent urgent output actions is spurious in the sense that it does not change the stochastic behaviour of the model after the output urgent actions have been abstracted. Indeed, since time does not progress, it is the same to sample first clock x and then clock y passing through state s_1, or first y and then x passing through s_2, or even sampling both clocks simultaneously through a transition $s_1 \xrightarrow{\varnothing, \tau, \{x,y\}} s_3$. In any of the cases, the stochastic resolution of the execution of a or b in the stable state s_3 is the same. This could be generalized to any number of confluent transitions.

Thus, it will be convenient to use term rewriting techniques to collect all clocks that are active in the convergent stable state and have been activated through a path of urgent actions. Therefore, we recall some basic notions of rewriting systems. An *abstract reduction system* [1] is a pair $(E, \rightarrowtail)$, where the reduction $\rightarrowtail$ is a binary relation over the set E, i.e. $\rightarrowtail \subseteq E \times E$. We write $a \rightarrowtail b$ for $(a, b) \in \rightarrowtail$. We also write $a \overset{*}{\rightarrowtail} b$ to denote that there is a path $a_0 \rightarrowtail a_1 \ldots \rightarrowtail a_n$ with $n \geq 0$, $a_0 = a$ and $a_n = b$. An element $a \in E$ is in *normal form* if there is no b such that $a \rightarrowtail b$. We say that b is a normal form of a if $a \overset{*}{\rightarrowtail} b$ and b is in normal form. A reduction system $(E, \rightarrowtail)$ is *confluent* if

Fig. 5. Confluence is weakly deterministic

for all $a, b, c \in E$ $a \overset{*}{\leftarrowtail} c \overset{*}{\rightarrowtail} b$ implies $a \overset{*}{\rightarrowtail} d \overset{*}{\leftarrowtail} b$ for some $d \in E$. This notion of confluence is implied by the following statement: for all $a, b, c \in E$, $a \leftarrowtail c \rightarrowtail b$ implies that either $a \rightarrowtail d \leftarrowtail b$ for some $d \in E$, or $a = b$. A reduction system is *normalizing* if every element has a normal form, and it is *terminating* if there is no infinite chain $a_0 \rightarrowtail a_1 \rightarrowtail \cdots$. A terminating reduction system is also normalizing. In a confluent reduction system every element has at most one normal form. If in addition it is also normalizing, then the normal form is unique.

We now define the abstract reduction system introduced by the urgent transitions of an IOSA.

Definition 6. *Given an IOSA $\mathcal{I} = (\mathcal{S}, \mathcal{A}, \mathcal{C}, \rightarrow_{\mathcal{I}}, C_0, s_0)$, define the abstract reduction system $\mathcal{U}_I$ as $(\mathcal{S} \times \mathcal{P}(\mathcal{C}) \times \mathbb{N}_0, \rightarrowtail)$ where $(s, C, n) \rightarrowtail (s', C \cup C', n+1)$ if and only if there exists $a \in \mathcal{A}^{\mathsf{u}}$ such that $s \xrightarrow{\varnothing, a, C'} s'$.*

An IOSA is *non-Zeno* if there is no loop of urgent actions. The following result can be straightforwardly proven.

Proposition 3. *Let the IOSA $\mathcal{I}$ be closed and confluent. Then $\mathcal{U}_{\mathcal{I}}$ is confluent, and hence every element has at most one normal form. Moreover, an element (s, C, n) is in normal form iff s is stable in $\mathcal{I}$. If in addition $\mathcal{I}$ is non-Zeno, $\mathcal{U}_{\mathcal{I}}$ is also terminating and hence every element has a unique normal form.*

6   Weak Determinism

As already shown in Fig. 5, the non-determinism introduced by confluence is spurious. In this section, we show that closed confluent IOSAs behave deterministically in the sense that the stochastic behaviour of the model is the same, regardless the way in which non-determinism is resolved. Thus, we say that a closed IOSA is *weakly deterministic* if (i) almost surely at most one discrete non-urgent transition is enabled at every time point, (ii) the election over enabled urgent transitions does not affect the non urgent-behavior of the model, and (iii) no non-urgent output and urgent output are enabled simultaneously. To avoid referring explicitly to time in (i), we say instead that a closed IOSA is weakly deterministic if it almost never reaches a state in which two different non-urgent discrete transitions are enabled. Moreover, to ensure (ii), we define the following weak transition.

For this definition and the rest of the section we will assume that the IOSA is closed and all its urgent actions have been abstracted, that is, all actions in $\mathcal{A}^{\mathsf{u}}$ have been renamed to τ.

Definition 7. *For a non stable state s, and $v \in \mathbb{R}^N$, we define $(s, \vec{v}) \overset{C}{\Longrightarrow}_n \mu$ inductively by the following rules:*

$$
(T1) \quad \frac{\begin{array}{c} s \xrightarrow{\varnothing, \tau, C} s' \\ \mathrm{st}(s') \end{array}}{(s, \vec{v}) \overset{C}{\Longrightarrow}_1 \mu_{C, s'}^{\vec{v}}} \qquad\qquad (T2) \quad \frac{\begin{array}{c} s \xrightarrow{\varnothing, \tau, C'} s' \\ \forall \vec{v}' \in \mathbb{R}^N : \exists C'', \mu' : (s', \vec{v}') \overset{C''}{\Longrightarrow}_n \mu' \end{array}}{(s, \vec{v}) \overset{C' \cup C''}{\Longrightarrow}_{n+1} \hat{\mu}}
$$

where $\mu_{C,s}^{\vec{v}}$ is defined as in Definition 2 and $\hat{\mu} = \int_{\mathcal{S} \times \mathbb{R}^N} f_n^{C''} d\mu_{C', s'}^{\vec{v}}$, with $f_n^{C''}(t, \vec{w}) = \nu$, if $(t, \vec{w}) \overset{C''}{\Longrightarrow}_n \nu$, and $f_n^{C''}(t, \vec{w}) = \mathbf{0}$ otherwise. We define the weak transition $(s, \vec{v}) \Rightarrow \mu$ if $(s, \vec{v}) \overset{C}{\Longrightarrow}_n \mu$ for some $n \geq 1$ and $C \subseteq \mathcal{C}$.

As given above, there is no guarantee that $\overset{C}{\Longrightarrow}_n$ is well defined. In particular, there is no guarantee that $f_n^{C''}$ is a well defined measurable function. We postpone this to Lemma 1 below.

With this definition, we can introduce the concept of weak determinism:

Definition 8. *A closed IOSA $\mathcal{I}$ is* weakly deterministic *if $\Rightarrow$ is well defined in $\mathcal{I}$ and, in $P(\mathcal{I})$, any state $(s, v) \in \boldsymbol{S}$ that satisfies one of the following conditions is almost never reached from any $(\mathsf{init}, v_0) \in \boldsymbol{S}$: (a) s is stable and*

$\cup_{a\in\mathcal{A}\cup\{\text{init}\}}\mathcal{T}_a(s,v)$ *contains at least two different probability measures, (b) s is not stable, $(s,v) \Rightarrow \mu$, $(s,v) \Rightarrow \mu'$ and $\mu \neq \mu'$, or (c) s is not stable and $(s,v) \xrightarrow{a} \mu$ for some $a \in \mathcal{A}^\circ \setminus \mathcal{A}^u$.*

By "almost never" we mean that the measure of the set of all paths leading to any measurable set in $\mathscr{B}(\mathbf{S})$ containing only states satisfying (a), (b), or (c) is zero. Thus, Definition 8 states that, in a weakly deterministic IOSA, a situation in which a non urgent output action is enabled with another output action, being it urgent (case (c)) or non urgent (case (a)), or in which sequences of urgent transitions lead to different stable situations (case (b)), is almost never reached.

For the previous definition to make sense we need that $\mathcal{P}(\mathcal{I})$ satisfies *time additivity, time determinism*, and *maximal progress* [27]. This is stated in the following theorem whose proof follows as in [13, Theorem 16].

Theorem 2. *Let $\mathcal{I}$ be an IOSA $\mathcal{I}$. Its semantics $\mathcal{P}(\mathcal{I})$ satisfies, for all $(s,\vec{v}) \in \mathbf{S}$, $a \in \mathcal{A}^\circ$ and $d, d' \in \mathbb{R}_{>0}$, (i) $\mathcal{T}_a(s,\vec{v}) \neq \varnothing \Rightarrow \mathcal{T}_d(s,\vec{v}) = \varnothing$ (maximal progress), (ii) $\mu, \mu' \in \mathcal{T}_d(s,\vec{v}) \Rightarrow \mu = \mu'$ (time determinism), and (iii) $\delta^{-d}_{(s,\vec{v})}\in\mathcal{T}_d(s,\vec{v}) \wedge \delta^{-d'}_{(s,\vec{v}-d)}\in\mathcal{T}_{d'}(s,\vec{v} - d) \iff \delta^{-(d+d')}_{(s,\vec{v})}\in\mathcal{T}_{d+d'}(s,\vec{v})$ (time additivity).*

The next lemma states that, under the hypothesis that the IOSA is closed and confluent, $\overset{C}{\Longrightarrow}_n$ is well defined. Simultaneously, we prove that $\overset{C}{\Longrightarrow}_n$ is deterministic.

Lemma 1. *Let $\mathcal{I}$ be a closed and confluent IOSA. Then, for all $n \geq 1$, the following holds:*

1. *If $(s,\vec{v}) \overset{C}{\Longrightarrow}_n \mu$ then there is a stable state s' such that (i) $\mu = \mu^{\vec{v}}_{C,s'}$, (ii) $(s,C',m) \overset{*}{\rightarrowtail} (s',C'\cup C,m+n)$ for all $C' \subseteq C$ and $m \geq 0$, and (iii) if $(s,\vec{v}') \overset{C'}{\Longrightarrow}_n \mu'$ then $C' = C$ and moreover, if $\vec{v}' = \vec{v}$, also $\mu' = \mu$; and*
2. *f^C_n is a measurable function.*

The proof of the preceding lemma uses induction on n to prove item 1 and 2 simultaneously. It makes use of the previous results on rewriting systems in conjunction with measure theoretical tools such as Fubini's theorem to deal with Lebesgue integrals on product spaces. All these tools make the proof that confluence preserves weak determinism radically different from those of Milner [23] and Crouzen [9].

The following corollary follows by items 1.(ii) and 1.(iii) of Lemma 1.

Corollary 1. *Let $\mathcal{I}$ be a closed and confluent IOSA. Then, for all $(s,\vec{v})$, if $(s,\vec{v}) \Rightarrow \mu_1$ and $(s,\vec{v}) \Rightarrow \mu_2$, $\mu_1 = \mu_2$.*

This corollary already shows that closed and confluent IOSAs satisfy part (b) of Definition 8. In general, we can state:

Theorem 3. *Every closed confluent IOSA is weakly deterministic.*

The rest of the section is devoted to discuss the proof of this theorem. From now on, we work with the closed confluent IOSA $\mathcal{I} = (\mathcal{S}, \mathcal{C}, \mathcal{A}, \rightarrow, s_0, C_0)$, with $|\mathcal{C}| = N$, and its semantics $\mathcal{P}(\mathcal{I}) = (\mathbf{S}, \mathscr{B}(\mathbf{S}), \{\mathcal{T}_a \mid a \in \mathcal{L}\})$.

The idea of the proof of Theorem 3 is to show that the property that all active clocks have non-negative values and they are different from each other is almost surely an invariant of $\mathcal{I}$, and that at most one non-urgent transition is enabled in every state satisfying such invariant. Furthermore, we want to show that, for unstable states, active clocks have strictly positive values, which implies that non-urgent transitions are never enabled in these states. Formally, the invariant is the set

$$
\begin{aligned}
\mathsf{Inv} = \quad & \{(s, \vec{v}) \mid \mathrm{st}\,(s) \text{ and } \forall x_i, x_j \in \mathrm{active}(s) : i \neq j \Rightarrow \vec{v}(i) \neq \vec{v}(j) \wedge \vec{v}(i) \geq 0\} \\
& \cup \{(s, \vec{v}) \mid \neg\mathrm{st}\,(s) \text{ and } \forall x_i, x_j \in \mathrm{active}(s) : i \neq j \Rightarrow \vec{v}(i) \neq \vec{v}(j)\vec{v}(i) > 0\} \\
& \cup (\{\mathsf{init}\} \times \mathbb{R}^N)
\end{aligned}
\tag{1}
$$

with active as in Definition 1. Note that its complement is:

$$
\begin{aligned}
\mathsf{Inv}^c = \quad & \{(s, \vec{v}) \mid \exists x_i, x_j \in \mathrm{active}(s) : i \neq j \wedge \vec{v}(i) = \vec{v}(j)\} \\
& \cup \{(s, \vec{v}) \mid \mathrm{st}\,(s) \text{ and } \exists x_i \in \mathrm{active}(s) : \vec{v}(i) < 0\} \\
& \cup \{(s, \vec{v}) \mid \neg\mathrm{st}\,(s) \text{ and } \exists x_i \in \mathrm{active}(s) : \vec{v}(i) \leq 0\}
\end{aligned}
\tag{2}
$$

It is not difficult to show that Inv^c is measurable and, in consequence, so is Inv. The following lemma states that Inv^c is almost never reached in one step from a state satisfying the invariant.

Lemma 2. *If $(s, \vec{v}) \in \mathsf{Inv}$, $a \in \mathcal{L}$, and $\mu \in \mathcal{T}_a(s, \vec{v})$, then $\mu(\mathsf{Inv}^c) = 0$.*

From this lemma we have the following corollary.

Corollary 2. *The set Inv^c is almost never reachable in $\mathcal{P}(\mathcal{I})$.*

The proof of the corollary requires the definitions related to schedulers and measures on paths in NLMPs (see [25, Chap. 7] for a formal definition of scheduler and probability measures on paths in NLMPs.) We omit the proof of the corollary since it eventually boils down to an inductive application of Lemma 2.

The next lemma states that any stable state in the invariant Inv has at most one discrete transition enabled. Its proof is the same as that of [13, Lemma 20].

Lemma 3. *For all $(s, \vec{v}) \in \mathsf{Inv}$ with s stable or $s = \mathsf{init}$, the set $\bigcup_{a \in \mathcal{A} \cup \{\mathsf{init}\}} \mathcal{T}_a(s, \vec{v})$ is either a singleton set or the empty set.*

The next lemma states that any unstable state in the invariant Inv can only produce urgent actions.

Lemma 4. *For every state $(s, \vec{v}) \in \mathsf{Inv}$, if $\neg\mathrm{st}\,(s)$ and $(s, \vec{v}) \xrightarrow{a} \mu$, then $a \in \mathcal{A}^{\mathsf{u}}$.*

Proof. First recall that $\mathcal{I}$ is closed; hence $\mathcal{A}^{\mathsf{i}} = \varnothing$. If $(s, \vec{v}) \in \mathsf{Inv}$ and $\neg\mathsf{st}(s)$ then $\vec{v}_i > 0$ for all $x_i \in \mathsf{enabling}(s) \subseteq \mathsf{active}(s)$. Therefore, by Definition 2, $\mathcal{T}_a(s, \vec{v}) = \varnothing$ if $a \in \mathcal{A}^{\mathsf{o}} \setminus \mathcal{A}^{\mathsf{u}}$. Furthermore, for any $d \in \mathbb{R}_{>0}$, $\mathcal{T}_d(s, \vec{v}) = \varnothing$ since s is not stable and hence $s \xrightarrow{-,b,-} \_$ for some $b \in \mathcal{A}^{\mathsf{o}} \cap \mathcal{A}^{\mathsf{u}}$.   □

Finally, Theorem 3 is a consequence of Lemma 3, Lemma 4, Corollary 2, and Corollary 1.

7   Sufficient Conditions for Weak Determinism

Figure 3 shows an example in which the composed IOSA is weakly deterministic despite that some of its components are not confluent. The potential nondeterminism introduced in state $s_1\|s_4\|s_6$ is never reached since urgent actions at states $s_0\|s_4\|s_6$ and $s_1\|s_3\|s_6$ prevent the execution of non urgent actions leading to such state. We say that state $s_1\|s_4\|s_6$ is not *potentially reachable*. The concept of potentially reachable can be defined as follows.

Definition 9. *Given an IOSA $\mathcal{I}$, a state s is* potentially reachable *if there is a path $s_0 \xrightarrow{-,a_0,-} s_1 \ldots, s_{n-1} \xrightarrow{-,a_{n-1},-} s_n = s$ from the initial state, with $n \geq 0$, such that for all $0 \leq i < n$, if $s_i \xrightarrow{-,b,-} \_$ for some $b \in \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\mathsf{o}}$ then $a_i \in \mathcal{A}^{\mathsf{u}}$. In such case we call the path* plausible.

Notice that none of the paths leading to $s_1\|s_4\|s_6$ in Fig. 3 are plausible. Also, notice that an IOSA is bisimilar to the same IOSA when its set of states is restricted to only potentially reachable states.

Proposition 4. *Let $\mathcal{I}$ be a closed IOSA with set of states $\mathcal{S}$ and let $\overline{\mathcal{I}}$ be the same IOSA as $\mathcal{I}$ restricted to the set of states $\overline{\mathcal{S}} = \{s \in \mathcal{S} \mid$ is potentially reachable in $\mathcal{I}\}$. Then $\mathcal{I} \sim \overline{\mathcal{I}}$.*

Although we have not formally introduced bisimulation, it should be clear that both semantics are bisimilar through the identity relation since a transition $s \xrightarrow{\{x\},a,C} s'$ with s unstable does not introduce any concrete transition. (Recall the IOSA is closed so there is no input action on $\mathcal{I}$.)

For a state in a composed IOSA to be potentially reachable, necessarily each of the component states has to be potentially reachable in its respective component IOSA.

Lemma 5. *If a state $s_1\|\cdots\|s_n$ is potentially reachable in $\mathcal{I}_1\|\cdots\|\mathcal{I}_n$ then s_i is potentially reachable in $\mathcal{I}_i$ for all $i = 1, \ldots, N$.*

By Theorem 3, it suffices to check whether a closed IOSA is confluent to ensure that it is weakly deterministic. In this section, and following ideas introduced in [9], we build on a theory that allows us to ensure that a closed composed IOSA is confluent in a compositional manner, even when its components may not be confluent. Theorem 5 provides the sufficient conditions to guarantee that

the composed IOSA is confluent. Because of Proposition 2, it suffices to check whether two urgent actions that are not confluent in a single component are potentially reached. Since potential reachability depends on the composition, the idea is to overapproximate by inspecting the components. The rest of the section builds on concepts that are essential to construct such overapproximation.

Let $\text{uen}(s) = \{a \in \mathcal{A}^{\mathsf{u}} \mid s \xrightarrow{\,\text{-},a,\text{-}\,} \_\}$ be the set of urgent actions enabled in a state s. We say that a set B of output urgent actions is spontaneously enabled by a non-urgent action b if b is potentially reached and it transitions to a state enabling all actions in B.

Definition 10. *A set $B \subseteq \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\mathsf{o}}$ is* spontaneously enabled *by $a \in \mathcal{A} \setminus \mathcal{A}^{\mathsf{u}}$ in $\mathcal{I}$, if either $B = \varnothing$ or there are potentially reachable states s and s' such that s is stable, $s \xrightarrow{\,\text{-},a,\text{-}\,} s'$, and $B \subseteq \text{uen}(s')$. B is* maximal *if for any B' spontaneously enabled by b in $\mathcal{I}$ such that $B \subseteq B'$, $B = B'$.*

A set that is spontaneously enabled in a composed IOSA, can be constructed as the union of spontaneously enabled sets in each of the components as stated by the following proposition. Therefore, spontaneously enabled sets in a composed IOSA can be overapproximated by unions of spontaneously enabled sets of its components.

Proposition 5. *Let B be spontaneously enabled by action a in $\mathcal{I}_1 || \dots || \mathcal{I}_n$. Then, there are $B_1, \dots, B_n$ such that each B_i is spontaneously enabled by a in $\mathcal{I}_i$, and $B = \bigcup_{i=1}^{n} B_i$. If in addition B is maximal, there are $B_1, \dots, B_n$ such that each B_i is maximal spontaneously enabled by a in $\mathcal{I}_i$, and $B \subseteq \bigcup_{i=1}^{n} B_i$.*

Proof. We only prove it for $\mathcal{I}_1 || \mathcal{I}_2$. The generalization to any n follows easily. Let $\bar{B}_i = B \cap \mathcal{A}_i$ for $i = 1, 2$ and note that $B = \bar{B}_1 \cup \bar{B}_2$. We show that $\bar{B}_1$ is spontaneously enabled by a in $\mathcal{I}_1$. The case of $\bar{B}_2$ follows similarly. Since B is spontaneously enabled by a in $\mathcal{I}_1 || \mathcal{I}_2$, there exist potentially reachable states $s_1 || s_2$ and $s'_1 || s'_2$, such that $s_1 || s_2$ is stable, $s_1 || s_2 \xrightarrow{\,\text{-},a,\text{-}\,} s'_1 || s'_2$, and $B \subseteq \text{uen}(s'_1 || s'_2)$. First notice that $\bar{B}_1 \subseteq \text{uen}(s_1)$. Also, suppose $\bar{B}_1 \neq \varnothing$, otherwise $\bar{B}_1$ is spontaneously enabled by a trivially. Consider first the case that $a \in \mathcal{A}_2 \setminus \mathcal{A}_1$. By (R2), $s_1 = s'_1$, but, since there is some $b \in \bar{B}_1$, $s_1 \xrightarrow{\,\text{-},b,\text{-}\,} \_$ and hence $s_1 || s_2 \xrightarrow{\,\text{-},b,\text{-}\,} \_$ rendering $s_1 || s_2$ unstable, which is a contradiction. So $a \in \mathcal{A}_1$ and $s_1 \xrightarrow{\,\text{-},a,\text{-}\,} s'_1$. By Lemma 5, s_1 and s'_1 are potentially reachable and, necessarily, s_1 is stable (otherwise $s_1 || s_2$ has to be unstable as shown before). Therefore $\bar{B}_1$ is spontaneously enabled by a in $\mathcal{I}_1$. The second part of the proposition is immediate from the first part. $\qquad\square$

Spontaneously enabled sets refer to sets of urgent output actions that are enabled after some steps of execution. Urgent output actions can also be enabled at the initial state.

Definition 11. *A set $B \subseteq \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\mathsf{o}}$ is* initial *in an IOSA $\mathcal{I}$ if $B \subseteq \text{uen}(s_0)$, with s_0 being the initial state of $\mathcal{I}$. B is* maximal *if $B = \text{uen}(s_0) \cap \mathcal{A}^{\mathsf{o}}$.*

An initial set of a composed IOSA can be constructed as the union of initial sets of its components. In particular the maximal initial set is the union of all the maximal sets of its components. The proof follows directly from the definition of parallel composition taking into consideration that IOSAs are input enabled.

Proposition 6. *Let B be initial in $\mathcal{I} = (\mathcal{I}_1|| \ldots ||\mathcal{I}_n)$. Then, there are $B_1, \ldots, B_2$, with B_i initial of $\mathcal{I}_i$, $1 \le i \le n$ and $B = \bigcup_{i=1}^{n} B_i$. Moreover, $\mathrm{uen}(s_0) \cap \mathcal{A}_{\mathcal{I}}^{\circ} = \bigcup_{i=1}^{n} \mathrm{uen}(s_i^0) \cap \mathcal{A}_i^{\circ}$.*

We say that an urgent action triggers an urgent output action if the first one enables the occurrence of the second one, which was not enabled before.

Definition 12. *Let $a \in \mathcal{A}^{\mathsf{u}}$ and $b \in \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\circ}$. a triggers b in an IOSA $\mathcal{I}$ if there are potentially reachable states s_1, s_2, and s_3 such that $s_1 \xrightarrow{-,a,-} s_2 \xrightarrow{-,b,-} s_3$ and, if $a \ne b$, $b \notin \mathrm{uen}(s_1)$.*

Notice that, for the particular case in which $a = b$, $b \notin \mathrm{uen}(s)$ is not required. The following proposition states that if one action triggers another one in a composed IOSA, then the same triggering occurs in a particular component.

Proposition 7. *Let $a \in \mathcal{A}^{\mathsf{u}}$ and $b \in \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\circ}$ such that a triggers b in $\mathcal{I}_1|| \ldots ||\mathcal{I}_n$. Then there is a component $\mathcal{I}_i$ such that $b \in \mathcal{A}_i^{\circ}$ and a triggers b in $\mathcal{I}_i$.*

Proof. We only prove it for $\mathcal{I}_1||\mathcal{I}_2$. The generalization to any n follows easily. Because $b \in \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\circ}$ necessarily $b \in \mathcal{A}_1^{\circ}$ or $b \in \mathcal{A}_2^{\circ}$. W.l.o.g. suppose $b \in \mathcal{A}_1^{\circ}$. Since a triggers b in $\mathcal{I}_1||\mathcal{I}_1$, $s_1||s_2 \xrightarrow{-,a,-} s_1'||s_2' \xrightarrow{-,b,-} s_1''||s_2''$ with $s_1||s_2$, $s_1'||s_2'$, and $s_1''||s_2''$ being potentially reachable.

Suppose first that $a \ne b$. Then $b \notin \mathrm{uen}(s_1||s_2)$. Recall that, by Lemma 5, s_1, s_1', and s_1'' are potentially reachable in $\mathcal{I}_1$. Since $b \in \mathcal{A}_1^{\circ}$, $s_1' \xrightarrow{-,b,-} s_1''$. Suppose $a \in \mathcal{A}_2 \backslash \mathcal{A}_1$. Then, necessarily, $s_1 = s_1'$ which gives $b \in \mathrm{uen}(s_1) \cap \mathcal{A}^{\circ} \subseteq \mathrm{uen}(s_1||s_2)$, yielding a contradiction. Thus, necessarily $a \in \mathcal{A}_1^{\mathsf{u}}$ and hence $s_1 \xrightarrow{-,a,-} s_1'$, by the definition of parallel composition. It remains to show that $b \notin \mathrm{uen}(s_1)$, but this is immediate since $\mathrm{uen}(s_1) \cap \mathcal{A}^{\circ} \subseteq \mathrm{uen}(s_1||s_2)$ and $b \notin \mathrm{uen}(s_1||s_2)$. Thus a triggers b in $\mathcal{I}_1$ in this case. If instead $a = b$, by the definition of parallel composition we immediately have that $s_1 \xrightarrow{-,b,-} s_1' \xrightarrow{-,b,-} s_1''$, proving thus the proposition. □

Proposition 7 tells us that the triggering relation of a composed IOSA can be overapproximated by the union of the triggering relations of its components. Thus we define:

Definition 13. *The approximate triggering relation of $\mathcal{I}_1|| \ldots ||\mathcal{I}_n$ is defined by $\rightsquigarrow = \bigcup_{i=1}^{n} \{(a, b) \mid a \text{ triggers } b \text{ in } \mathcal{I}_i\}$. Its reflexive transitive closure $\rightsquigarrow^*$ is called approximate indirect triggering relation.*

The next definition characterizes all sets of urgent output actions that are simultaneously enabled in any potentially reachable state of a given IOSA.

Definition 14. *A set $B \subseteq \mathcal{A}^u \cap \mathcal{A}^\circ$ is an* enabled set *in an IOSA $\mathcal{I}$ if there is a potentially reachable state s such that $B \subseteq \text{uen}(s)$. If $a \in B$, we say that a is* enabled *in s. Let $\text{ES}_\mathcal{I}$ be the set of all enabled sets in $\mathcal{I}$.*

If an urgent output action is enabled in a potentially reachable state of a IOSA, then it is either initial, spontaneously enabled, or triggered by some action.

Theorem 4. *Let $b \in \mathcal{A}^u \cap \mathcal{A}^\circ$ be enabled in some potentially reachable state of the IOSA $\mathcal{I}$. Then there is a set B with $b \in B$ that is either initial or spontaneously enabled by some action $a \in \mathcal{A}^u$, or b is triggered by some action $a \in \mathcal{A}^\circ \setminus \mathcal{A}^u$.*

Proof. Let s be potentially reachable in $\mathcal{I}$ such that $b \in \text{uen}(s) \cap \mathcal{A}^\circ$. We prove the theorem for b by induction on the plausible path σ leading to s. If $|\sigma| = 0$, then $\sigma = s$ and s is the initial state. Then the set $\text{uen}(s) \cap \mathcal{A}^\circ$ is initial and we are done in this case. If $|\sigma| > 0$, then $\sigma = \sigma' \cdot (s' \xrightarrow{\; \cdot, a, \cdot \;} s)$ for some s', a, and plausible σ'. If $a \in \mathcal{A} \setminus \mathcal{A}^u$ then s' is stable (since σ is plausible) and thus $\text{uen}(s) \cap \mathcal{A}^\circ$ is spontaneously enabled by a. If instead $a \in \mathcal{A}^u$, two possibilities arise. If $b \notin \text{uen}(s')$, then b is triggered by a. If $b \in \text{uen}(s')$, the conditions are satisfied by induction since $|\sigma'| = |\sigma| - 1$. $\qquad\square$

The next definition is auxiliary to prove the main theorem of this section. It constructs a graph from a closed and composed IOSA whose vertices are sets of urgent output actions. It has the property that, if there is a path from one vertex to another, all actions in the second vertex are approximately indirectly triggered by actions in the first vertex (Lemma 7). This will allow to show that any set of simultaneously enabled urgent output actions is approximately indirectly triggered by initial actions or spontaneously enabled sets (Lemma 8).

Definition 15. *Let $\mathcal{I} = (\mathcal{I}_1 || \ldots || \mathcal{I}_n)$ be a closed IOSA. The* enabled graph *of $\mathcal{I}$ is defined by the labelled graph $\text{EG}_\mathcal{I} = (V, E)$, where $V \subseteq 2^{\mathcal{A}^\circ \cap \mathcal{A}^u}$ and $E \subseteq V \times (\mathcal{A}^u \cap \mathcal{A}^\circ) \times V$, with $V = \bigcup_{k \geq 0} V_k$ and $E = \bigcup_{k \geq 0} E_k$, and, for all $k \in \mathbb{N}$, V_k and E_k are inductively defined by*

$$V_0 = \bigcup_{a \in \mathcal{A}} \{\textstyle\bigcup_{i=1}^n B_i \mid \forall 1 \leq i \leq n :$$
$$B_i \text{ is spontaneously enabled by } a \text{ and maximal in } \mathcal{I}_i\}$$
$$\cup \{\textstyle\bigcup_{i=1}^n \text{uen}(s_i^0) \cap \mathcal{A}_i^\circ \mid \forall 1 \leq i \leq n : s_i^0 \text{ is the initial state in } \mathcal{I}_i\}$$
$$E_k = \{(v, a, (v \setminus \{a\}) \cup \{b \mid a \rightsquigarrow b\}) \mid v \in V_k, a \in v\}$$
$$V_{k+1} = \{v' \mid v \in V_i, (v, v') \in E_k, v' \notin \textstyle\bigcup_{j=0}^k V_j\}$$

Notice that V_0 contains the maximal initial set of $\mathcal{I}$ and an overapproximation of all its maximal spontaneously enabled sets. Notice also that, by construction, there is a path from any vertex in V to some vertex in V_0.

The set closure of V in $\text{EG}_\mathcal{I}$, defined by $\overline{\text{ES}}_\mathcal{I} = \{B \mid B \subseteq v, v \in V\}$, turns out to be an overapproximation of the actual set $\text{ES}_\mathcal{I}$ of all enabled sets in $\mathcal{I}$.

Lemma 6. *For any closed IOSA $\mathcal{I} = (\mathcal{I}_1 || \cdots || \mathcal{I}_n)$, $\mathsf{ES}_{\mathcal{I}} \subseteq \overline{\mathsf{ES}}_{\mathcal{I}}$.*

Proof. Let $B \in \mathsf{ES}_{\mathcal{I}}$. We proceed by induction on the length of the plausible path σ that leads to the state s s.t. $B \subseteq \mathrm{uen}(s)$. If $|\sigma| = 0$ then s is the initial state and thus B is initial in $\mathcal{I}$. Thus, by Definition 11, Proposition 6, and Definition 15, $B \subseteq (\mathrm{uen}(s_0) \cap \mathcal{A}^{\circ}_{\mathcal{I}}) = (\bigcup_{i=1}^{n} \mathrm{uen}(s_i^0) \cap \mathcal{A}_i^{\circ}) \in V_0 \subseteq \overline{\mathsf{ES}}_{\mathcal{I}}$. As a consequence $B \in \overline{\mathsf{ES}}_{\mathcal{I}}$.

If $|\sigma| > 0$ then $\sigma = \sigma' \cdot (s' \xrightarrow{\text{-},a,\text{-}} s)$, for some s', a, and plausible σ'. If $a \in \mathcal{A} \backslash \mathcal{A}^{\mathsf{u}}$ then s' is stable (since σ is plausible) and thus B is spontaneously enabled by a. By Proposition 5, there are $B_1, \ldots, B_n$ such that each B_i is spontaneously enabled by a and maximal in $\mathcal{I}_i$, and $B \subseteq \bigcup_{i=1}^{n} B_i$. Since $\bigcup_{i=1}^{n} B_i \in V_0 \subseteq \overline{\mathsf{ES}}_{\mathcal{I}}$, then $B \in \overline{\mathsf{ES}}_{\mathcal{I}}$. If instead $a \in \mathcal{A}^{\mathsf{u}}$, let $B' = \{a\} \cup (B \cap \mathrm{uen}(s'))$. Notice that $B' \subseteq \mathrm{uen}(s') \cap \mathcal{A}^{\circ}$. Since s' is the last state on σ' and $|\sigma'| = |\sigma| - 1$, $B' \in \overline{\mathsf{ES}}_{\mathcal{I}}$ by induction. Hence, there is a vertex $v' \in V$ in $\mathsf{EG}_{\mathcal{I}}$ such that $B' \subseteq v$ and, by Definition 15, $v' \in V_k$ for some $k \geq 0$. Let $v = (v' \backslash \{a\}) \cup \{b \mid a \leadsto b\}$, then $(v', a, v) \in E_k$ and hence $v \in V_{k+1}$. We show that $B \subseteq v$. Let $b \in B$. If $b = a$, then $a \in \mathrm{uen}(s) \cap \mathcal{A}^{\circ}$ and hence a triggers a in $\mathcal{I}$. By Proposition 7, $a \leadsto a$ which implies $a \in v$. Suppose, instead, that $b \neq a$. If $b \in \mathrm{uen}(s')$, then $b \in B' \backslash \{a\} \subseteq v' \backslash \{a\} \subseteq v$. If $b \notin \mathrm{uen}(s')$, then a triggers b in $\mathcal{I}$, and by Proposition 7, $a \leadsto b$ which implies $b \in v$. This proves $B \subseteq v \in \overline{\mathsf{ES}}_{\mathcal{I}}$ and hence $B \in \overline{\mathsf{ES}}_{\mathcal{I}}$. $\qquad\square$

The next lemma states that if there is a path from a vertex of $\mathsf{EG}_{\mathcal{I}}$ to another vertex, every action in the second vertex is approximately indirectly triggered by some action in the first vertex.

Lemma 7. *Let $\mathcal{I}$ be a closed IOSA, let $v, v' \in V$ be vertices of $\mathsf{EG}_{\mathcal{I}}$ and let ρ be a path following E from v to v'. Then for every $b \in v'$ there is an action $a \in v$ such that $a \leadsto^* b$.*

Proof. We proceed by induction in the length of ρ. If $|\rho| = 0$ then $v = v'$ and the lemma holds since $\leadsto^*$ is reflexive. If $|\rho| > 0$, there is a path ρ', $v'' \in V$, and $c \in \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\circ}$ such that $\rho = \rho' \cdot (v'', c, v')$. By induction, for every action $d \in v''$ there is some $a \in v$ such that $a \leadsto^* d$. Because of the definition of E in Definition 15, either $b \in v''$ or $c \leadsto b$ and $c \in v''$. The first case follows by induction. In the second case, also by induction, $a \leadsto^* c$ for some $a \in v$ and hence $a \leadsto^* b$. $\qquad\square$

The next lemma states that every enabled set B in a composed IOSA is either approximately triggered by a set of initial actions of the components of the IOSA or by a subset of the union of spontaneously enabled sets in each component where such sets are spontaneously enabled by the same event.

Lemma 8. *Let $\mathcal{I} = (\mathcal{I}_1 || \ldots || \mathcal{I}_n)$ be a closed IOSA and let $\{b_1, \ldots, b_m\} \subseteq \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\circ}$ be enabled in $\mathcal{I}$. Then, there are (not necessarily different) $a_1, \ldots, a_m$ such that $a_j \leadsto^* b_j$, for all $1 \leq j \leq m$, and either (i) $\{a_1, \ldots, a_m\} \subseteq \bigcup_{i=1}^{n} \mathrm{uen}(s_i^0) \cap \mathcal{A}_i^{\circ}$, or (ii) there exists $e \in \mathcal{A}$ and (possibly empty) sets $B_1, \ldots, B_n$ spontaneously enabled by e in $\mathcal{I}_1, \ldots, \mathcal{I}_n$ respectively, such that $\{a_1, \ldots, a_m\} \subseteq \bigcup_{i=1}^{n} B_i$.*

Proof. Because of Lemma 6 there is a vertex v of $\mathsf{EG}_{\mathcal{I}}$ such that $\{b_1, \ldots, b_n\} \subseteq v$. Because of the inductive construction of E and V, there is a path from some $v' \in V_0$ to v in $\mathsf{EG}_{\mathcal{I}}$. From Lemma 7, for each $1 \leq j \leq m$, there is an $a_j \in v'$ such that $a_j \rightsquigarrow^* b_j$. Because $v' \in V_0$, then either $v' = \bigcup_{i=1}^{n} \mathrm{uen}(s_i^0) \cap \mathcal{A}_i^{\circ}$ or there is some $e \in \mathcal{A}$ such that $v' = \bigcup_{i=1}^{n} B_i$ with B_i spontaneously enabled by e in $\mathcal{I}_i$ $\square$

The following theorem is the main result of this section and provides sufficient conditions to guarantee that a closed composed IOSA is confluent or, as stated in the theorem, necessary conditions for the IOSA to be non-confluent.

Theorem 5. *Let $\mathcal{I} = (\mathcal{I}_1 || \cdots || \mathcal{I}_n)$ be a closed IOSA. If $\mathcal{I}$ potentially reaches a non-confluent state then there are actions $a, b \in \mathcal{A}^{\mathsf{u}} \cap \mathcal{A}^{\circ}$ such that some $\mathcal{I}_i$ is not confluent w.r.t. a and b, and there are c and d such that $c \rightsquigarrow^* a$, $d \rightsquigarrow^* b$, and, either (i) c and d are initial actions in any component, or (ii) there is some $e \in \mathcal{A}$ and (possibly empty) sets $B_1, \ldots, B_n$ spontaneously enabled by e in $\mathcal{I}_1, \ldots, \mathcal{I}_n$ respectively, such that $c, d \in \bigcup_{i=1}^{n} B_i$.*

Proof. Suppose $\mathcal{I}$ potentially reaches a non confluent state s. Then there are necessarily $a, b \in \mathrm{uen}(s)$ that show it and hence $\mathcal{I}$ is not confluent w.r.t. a and b. By Proposition 2, there is necessarily a component $\mathcal{I}_i$ that is not confluent w.r.t. a and b. Since $\{a, b\}$ is an enabled set in $\mathcal{I}$, the rest of the theorem follows by Lemma 8. $\square$

Because of Proposition 4 and Theorem 3, if all potentially reachable states in a closed IOSA $\mathcal{I}$ are confluent, then $\mathcal{I}$ is weakly deterministic. Thus, if no pair of actions satisfying conditions in Theorem 5 are found in $\mathcal{I}$, then $\mathcal{I}$ is weakly deterministic.

Notice that the IOSA $\mathcal{I} = \mathcal{I}_1 || \mathcal{I}_2 || \mathcal{I}_3$ of Example 2 (see also Figs. 2 and 3) is an example that does not meet the conditions of Theorem 5, and hence detected as confluent. c and d are the only potential non-confluent actions, which is noticed in state s_6 of $\mathcal{I}_3$. The approximate indirect triggering relation can be calculated to $\rightsquigarrow^* = \{(c, c), (d, d)\}$. Also, $\{c\}$ is spontaneously enabled by a in $\mathcal{I}_1$ and $\{d\}$ is spontaneously enabled by b in $\mathcal{I}_2$. Since both sets are spontaneously enabled by *different* actions and c and d are not initial, the set $\{c, d\}$ does not appear in V_0 of $\mathsf{EG}_{\mathcal{I}}$ which would be required to meet the conditions of the theorem.

Conditions in Theorem 5 are not suffi-cient and confluent IOSAs may satisfy them. Consider the IOSAs in Fig. 6. $\mathcal{I}_1 || \mathcal{I}_2 || \mathcal{I}_3$ is a closed IOSA with a single state and no outgoing transition. Hence, it is confluent. However, $\mathcal{I}_3$ is not confluent w.r.t. b and c, $\rightsquigarrow^* = \{(b, b), (c, c)\}$, $B_1 = \{b\}$ is sponta-neously enabled by a in $\mathcal{I}_1$, and $B_2 = \{c\}$ is spontaneously enabled by a in $\mathcal{I}_2$. Hence $b, c \in \bigcup_{i=1}^{n} B_i$, thus meeting the conditions of Theorem 5.
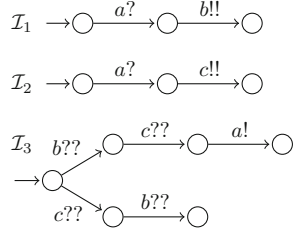
Fig. 6. $\mathcal{I}_1 || \mathcal{I}_2 || \mathcal{I}_3$ meets conditions in Theorem 5

8   Concluding Remarks

In this article, we have extended IOSA as introduced in [13] with urgent actions. Though such extension introduces non-determinism even if the IOSA is closed, it does so in a limited manner. We were able to characterize when a IOSA is weakly deterministic, which is an important concept since weakly deterministic IOSAs are amenable to discrete event simulation. In particular, we showed that closed and confluent IOSAs are weakly deterministic and provided conditions to check compositionally if a closed IOSA is confluent. Open IOSAs are naturally non-deterministic due to input enabledness: at any moment of time either two different inputs may be enabled or an input is enabled jointly with a possible passage of time. Thus, the property of non-determinism can only be possible in closed IOSAs. However, Theorem 5 relates open IOSAs to the concept of weak determinism by providing sufficient properties on open IOSAs whose composition leads to a closed weakly deterministic IOSA. In addition, we notice that languages like Modest [4, 18, 19], that have been designed for compositional modelling of complex timed and stochastic systems, embrace the concept of non-determinism as a fundamental property. Thus, ensuring weak determinism on Modest models using compositional tools like Theorem 5 will require significant limitations that may easily boil down to reduce it to IOSA. Notwithstanding this observation, we remark that some translation between IOSA and Modest is possible through Jani [8].

Finally, we remark that, though not discussed in this paper, the conditions provided by Theorem 5, can be verified in polynomial time respect to the size of the components and the number of actions.

References

1. Baader, F., Nipkow, T.: Term Rewriting and All That. Cambridge University Press, Cambridge (1998). https://doi.org/10.1017/cbo9781139172752
2. Behrmann, G., David, A., Larsen, K.G.: A tutorial on Uppaal. In: Bernardo, M., Corradini, F. (eds.) SFM-RT 2004. LNCS, vol. 3185, pp. 200–236. Springer, Heidelber (2004). https://doi.org/10.1007/978-3-540-30080-9_7
3. Bengtsson, J., et al.: Verification of an audio protocol with bus collision using Uppaal. In: Alur, R., Henzinger, T.A. (eds.) CAV 1996. LNCS, vol. 1102, pp. 244–256. Springer, Heidelberg (1996). https://doi.org/10.1007/3-540-61474-5_73
4. Bohnenkamp, H.C., D'Argenio, P.R., Hermanns, H., Katoen, J.: MODEST: a compositional modeling formalism for hard and softly timed systems. IEEE Trans. Softw. Eng. **32**(10), 812–830 (2006). https://doi.org/10.1109/tse.2006.104
5. Bravetti, M., D'Argenio, P.R.: Tutte le algebre insieme: concepts, discussions and relations of stochastic process algebras with general distributions. In: Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.-P., Siegle, M. (eds.) Validation of Stochastic Systems. LNCS, vol. 2925, pp. 44–88. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24611-4_2
6. Budde, C.E.: Automation of importance splitting techniques for rare event simulation. Ph.D. thesis, Universidad Nacional de Córdoba (2017)

7. Budde, C.E., D'Argenio, P.R., Monti, R.E.: Compositional construction of importance functions in fully automated importance splitting. In: Puliafito, A., Trivedi, K.S., Tuffin, B., Scarpa, M., Machida, F., Alonso, J. (eds.) Proceedings of 10th EAI International Conference on Performance Evaluation Methodologies and Tools, VALUETOOLS 2016, October 2016, Taormina. ICST (2017). https://doi.org/10.4108/eai.25-10-2016.2266501

8. Budde, C.E., Dehnert, C., Hahn, E.M., Hartmanns, A., Junges, S., Turrini, A.: JANI: quantitative model and tool interaction. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10206, pp. 151–168. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54580-5_9

9. Crouzen, P.: Modularity and determinism in compositional markov models. Ph.D. thesis, Universität des Saarlandes, Saarbrücken (2014)

10. D'Argenio, P.R.: Algebras and automata for timed and stochastic systems. Ph.D. thesis, Universiteit Twente (1999)

11. D'Argenio, P.R., Katoen, J.P.: A theory of stochastic systems, part I: Stochastic automata. Inf. Comput. **203**(1), 1–38 (2005). https://doi.org/10.1016/j.ic.2005.07.001

12. D'Argenio, P.R., Katoen, J., Brinksma, E.: An algebraic approach to the specification of stochastic systems (extended abstract). In: Gries, D., de Roever, W.P. (eds.) PROCOMET 1998. IFIP Conference Proceedings, vol. 125, pp. 126–147. Chapman & Hall, Boca Raton (1998). https://doi.org/10.1007/978-0-387-35358-6_12

13. D'Argenio, P.R., Lee, M.D., Monti, R.E.: Input/Output stochastic automata. In: Fränzle, M., Markey, N. (eds.) FORMATS 2016. LNCS, vol. 9884, pp. 53–68. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44878-7_4

14. D'Argenio, P.R., Sánchez Terraf, P., Wolovick, N.: Bisimulations for nondeterministic labelled Markov processes. Math. Struct. Comput. Sci. **22**(1), 43–68 (2012). https://doi.org/10.1017/s0960129511000454

15. Desharnais, J., Edalat, A., Panangaden, P.: Bisimulation for labelled Markov processes. Inf. Comput. **179**(2), 163–193 (2002). https://doi.org/10.1006/inco.2001.2962

16. Giry, M.: A categorical approach to probability theory. In: Banaschewski, B. (ed.) Categorical Aspects of Topology and Analysis. LNM, vol. 915, pp. 68–85. Springer, Heidelberg (1982). https://doi.org/10.1007/BFb0092872

17. van Glabbeek, R.J., Smolka, S.A., Steffen, B.: Reactive, generative and stratified models of probabilistic processes. Inf. Comput. **121**(1), 59–80 (1995). https://doi.org/10.1006/inco.1995.1123

18. Hahn, E.M., Hartmanns, A., Hermanns, H., Katoen, J.: A compositional modelling and analysis framework for stochastic hybrid systems. Form. Methods Syst. Des. **43**(2), 191–232 (2013). https://doi.org/10.1007/s10703-012-0167-z

19. Hartmanns, A.: On the analysis of stochastic timed systems. Ph.D. thesis, Saarlandes University, Saarbrücken (2015). http://scidok.sulb.uni-saarland.de/volltexte/2015/6054/

20. Hermanns, H.: Interactive Markov Chains: And the Quest for Quantified Quality. LNCS, vol. 2428. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45804-2

21. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. Inf. Comput. **94**(1), 1–28 (1991). https://doi.org/10.1016/0890-5401(91)90030-6

22. Law, A.M., Kelton, W.D.: Simulation Modeling and Analysis, 3rd edn. McGraw-Hill Higher Education, New York City (1999)

23. Milner, R.: Communication and Concurrency. Prentice-Hall, Englewood Cliffs (1989)
24. Ruijters, E., Stoelinga, M.: Fault tree analysis: a survey of the state-of-the-art in modeling, analysis and tools. Comput. Sci. Rev. **15**, 29–62 (2015). https://doi.org/10.1016/j.cosrev.2015.03.001
25. Wolovick, N.: Continuous probability and nondeterminism in labeled transition systems. Ph.D. thesis, Universidad Nacional de Córdoba, Argentina (2012)
26. Wu, S., Smolka, S.A., Stark, E.W.: Composition and behaviors of probabilistic I/O automata. Theor. Comput. Sci. **176**(1–2), 1–38 (1997). https://doi.org/10.1016/S0304-3975(97)00056-X
27. Wang, Y.: Real-time behaviour of asynchronous agents. In: Baeten, J.C.M., Klop, J.W. (eds.) CONCUR 1990. LNCS, vol. 458, pp. 502–520. Springer, Heidelberg (1990). https://doi.org/10.1007/BFb0039080