

TR-QC-3-2014

Stochastically timed predicate-based communication primitives for autonomic computing - Full Paper

Revision: 0.0; Feb 5, 2014

Author(s): Diego Latella (CNR), Michele Loreti (IMT-Lucca), Mieke Massink (CNR),
Valerio Senni (IMT-Lucca)

Publication date: Feb 5, 2014

Funding Scheme: Small or medium scale focused research project (STREP)

Topic: ICT-2011 9.10: FET-Proactive 'Fundamentals of Collective Adaptive Systems' (FOCAS)

Project number: 600708

Coordinator: Jane Hillston (UEDIN)

e-mail: Jane.Hillston@ed.ac.uk

Fax: +44 131 651 1426

| Part. no. | Participant organisation name | Acronym | Country |
|------------|--|---------|-------------|
| 1 (Coord.) | University of Edinburgh | UEDIN | UK |
| 2 | Consiglio Nazionale delle Ricerche – Istituto di Scienza e Tecnologie della Informazione “A. Faedo” | CNR | Italy |
| 3 | Ludwig-Maximilians-Universität München | LMU | Germany |
| 4 | Ecole Polytechnique Fédérale de Lausanne | EPFL | Switzerland |
| 5 | IMT Lucca | IMT | Italy |
| 6 | University of Southampton | SOTON | UK |



Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 2 |
| 2 | StocS: a Stochastic extension of SCEL | 3 |
| 2.1 | Syntax | 3 |
| 2.2 | Informal Timed Semantics - Four Variants | 7 |
| 2.2.1 | Network-oriented Semantics | 8 |
| 2.2.2 | Action-oriented | 8 |
| 2.2.3 | Interaction-oriented | 9 |
| 2.2.4 | Activity-oriented | 9 |
| 3 | Preliminary Definitions for Operational Semantics | 9 |
| 3.1 | Preliminaries | 10 |
| 4 | Action-oriented Operational Semantics | 11 |
| 4.1 | Operational semantics of processes | 12 |
| 4.2 | Operational semantics of components and systems | 13 |
| 5 | Interaction-oriented Operational Semantics | 16 |
| 5.1 | Operational semantics of processes | 18 |
| 5.2 | Operational semantics of components and systems | 19 |
| 6 | Network-oriented Operational Semantics | 22 |
| 6.1 | Explanatory example | 22 |
| 6.2 | Operational semantics of processes | 23 |
| 6.3 | Operational semantics of components and systems | 26 |
| 7 | Activity-oriented Operational Semantics | 28 |
| 7.1 | Operational semantics of processes | 28 |
| 7.2 | Operational semantics of components and systems | 30 |
| 8 | Case Study | 32 |
| 9 | Conclusions and Future Work | 36 |

Abstract

Predicate-based communication allows components of a system to send messages and requests to ensembles of components that are determined at execution time through the evaluation of a predicate, in a multicast fashion. Predicate-based communication can greatly simplify the programming of autonomous and adaptive systems. We present a stochastically timed extension of the Software Component Ensemble Language (SCEL) that was introduced in previous work. Such an extension raises a number of non-trivial design and formal semantics issues with different options as possible solutions at different levels of abstraction. We discuss four of these options. We provide formal semantics and an illustration of the use of the language modeling a variant of a bike sharing system, together with some preliminary analysis of the system performance.

1 Introduction

SCEL (Software Component Ensemble Language) [6], is a kernel language that is equipped with programming abstractions for the specification of system models within the framework of the *autonomic computing* paradigm, and for programming such systems. These abstractions are specifically designed for representing behaviours, knowledge, and aggregations according to specific policies, and to support programming context-awareness, self-awareness, and adaptation. SCEL is parametric with respect to some syntactic categories, namely KNOWLEDGE, POLICIES, TEMPLATES, and ITEMS (TEMPLATES and ITEMS determine the part of KNOWLEDGE to be retrieved/removed or added, respectively).

The main focus of the SCEL language is on supporting the development of autonomous, loosely-coupled, component-based software systems. For this purpose, a number of underlying assumptions are made on the kind of peculiarities of these software systems, among which adaptivity, open-endedness, ensemble-orientedness, high ability of reconfiguration, and support for heterogeneity. Two novel key aspects of SCEL, that distinguish it from other languages, are designed to support these peculiarities: predicate-based communication and the role of the component knowledge-base. Predicate-based communication allows to send messages to *ensembles* of components that are not predetermined at modeling time, but are defined at execution time, depending on how the communication predicate evaluates w.r.t. the receiver interface. The component knowledge-base allows to realize various adaptation patterns, by explicit separation of adaptation data in the spirit of [4], and to model components view on (and awareness of) the environment. In the first three years of the ASCENS project SCEL has been used to specify many scenarios related to the project Case Studies [13, 11, 16, 15]. These specifications witness how SCEL primitives simplify the programming of autonomous and adaptive systems. In these systems, the emerging behavior is realized through the coordination of the components activities.

In this paper we address the problem of enriching SCEL with information about action durations, by providing a stochastic semantics for the language. There exist various frameworks that support the systematic development of stochastic languages, such as [8]. However, the main challenge in developing a stochastic semantics for SCEL is in making appropriate modeling choices, both taking into account the specific application needs and allowing to manage model complexity and size. Our contribution in this work is the proposal of four variants of STOCS, a Markovian extension of a significant fragment of SCEL, that can be used to support quantitative analysis of adaptive systems composed of ensembles of cooperating components. These variants adopt *the same language syntax* of SCEL but denote different underlying stochastic models, having a different level of granularity. The choices we make will be motivated via examples and considerations on the model complexity and size. Of course one could encode these four semantics in an expressive language like Bio-PEPA [5] or HYPE [1], thus obtaining four fragments of Bio-PEPA/HYPE, one for each semantics.

Instead, what we do in this paper is to develop a single language and define four possible semantics for it, allowing to describe models at a different level of detail (by simply modifying labels and relations used to construct the LTS). From a software engineering point of view, this is a more pragmatic and flexible solution.

In summary, STOCS is essentially a *modeling language* which inherits the purpose and focus of SCEL. STOCS extends SCEL by modeling the average time duration of state-permanence and by replacing non-determinism by a probability distribution over outgoing transitions, thus adopting a CTMC-based operational semantics [7]. In the current preliminary phase of the design of STOCS, we deliberately omit to incorporate certain advanced features of SCEL, such as the presence and role of policies. Certain other features, such as the possibility of attaching components a unique identity, are discussed and evaluated for their impact on STOCS (See Section 2). A further issue of discussion is whether STOCS actions should have an atomic semantics or not (See Section 2.2) and the physical meaning of these assumptions.

Finally, an important aspect in a modeling language concerns the need of devising an appropriate syntax to express the environment model. In STOCS and SCEL the only point of contact with the environment is the knowledge base, which contains both internal information and externally-sensed events. In our approach, the knowledge is the most appropriate part of the language to specify environment models.

The outline of the present report is as follows. Section 2 discusses the trade-offs between four stochastic variants of SCEL, followed by the presentation of their formal semantics in Sections 4, 5, 6, and 7, after some preliminary definitions are recalled in Section 3. Section 8 introduces a case study to illustrate various aspects of the use of the design language in the context of a smart bike sharing system. Concluding remarks and lines for possible future research are presented in Section 9.

Note: *The material presented in Section 6 has been originally developed in the context of the EU project ASCENS (nr.257414) and is reported on in [12]. We included it in the present report for completeness and self-containment reasons.*

2 StocS: a Stochastic extension of SCEL

In this section we present the main features of STOCS. We start by illustrating its main syntactic ingredients. Then, we discuss possible choices for its timed semantics.

2.1 Syntax

The syntax of STOCS is presented in Table 1, where the syntactic categories of the language are defined. The basic category defines PROCESSES that are used to specify the order in which ACTIONS can be performed. Sets of processes are used to define the behavior of COMPONENTS, that in turn are used to define SYSTEMS. ACTIONS operate on local or remote knowledge-bases and have a TARGET to determine which other components are involved in the action. As we mentioned in the Introduction, for the sake of simplicity, in this version of STOCS we do not include POLICIES, whereas, like SCEL, STOCS is parametric w.r.t. KNOWLEDGE, TEMPLATES and ITEMS.

We define the following domains for variables and for defining functions signature: \mathbb{A} is the set of attribute names (which include the constant `id` used to indicate the component identifier), \mathbb{V} is the set of values, \mathbb{K} is the set of possible knowledge states, \mathbb{I} is the set of knowledge items, \mathbb{T} is the set of knowledge templates. So, in Table 1, $a \in \mathbb{A}$, $v \in \mathbb{V}$, $K \in \mathbb{K}$, $t \in \mathbb{I}$, $T \in \mathbb{T}$.

| | |
|----------------------|---|
| SYSTEMS: | $S ::= C \mid S \parallel S$ |
| COMPONENTS: | $C ::= I[K, P]$ |
| PROCESSES: | $P ::= \mathbf{nil} \mid a.P \mid P + P \mid P \mid P \mid X \mid A(\bar{p})$ |
| ACTIONS: | $a ::= \mathbf{get}(T)@c \mid \mathbf{qry}(T)@c \mid \mathbf{put}(t)@c$ |
| TARGETS: | $c ::= \mathbf{self} \mid \mathbf{p}$ |
| ENSEMBLE PREDICATES: | $\mathbf{p} ::= tt \mid e \bowtie e \mid \neg \mathbf{p} \mid \mathbf{p} \wedge \mathbf{p}$ with $\bowtie \in \{<, \leq, >, \geq\}$ |
| EXPRESSIONS: | $e ::= v \mid x \mid a \mid \dots$ |

Table 1: STOCS syntax (KNOWLEDGE K , TEMPLATES T , and ITEMS t are parameters)

Example 1 (Items and Templates as Tuples and Patterns). Consider a signature $(\mathcal{V}, \mathcal{F})$ where \mathcal{V} is a set of variables and \mathcal{F} is a set of function symbols with arity (we indicate by f/n a function symbol f with arity n) such that $\langle \rangle / i \in \mathcal{F}$ for $i = 0, 1, 2, \dots$. We denote by $\text{Terms}(\mathcal{V}, \mathcal{F})$ the set of all possible finite terms on the given signature (i.e. the terms with variables, constructed respecting function symbols arities) and by $\text{Terms}(\mathcal{F})$ the set of all possible finite ground terms. A pattern is a term of the form $\langle t_1, \dots, t_n \rangle$, with $t_i \in \text{Terms}(\mathcal{V}, \mathcal{F})$ for $i = 1, \dots, n$. A tuple is a term of the form $\langle t_1, \dots, t_n \rangle$, with $t_i \in \text{Terms}(\mathcal{F})$ for $i = 1, \dots, n$. In this example we have defined the set of Templates \mathbb{T} as the set of patterns and the set of Items \mathbb{I} as the set of tuples.

Systems and components

We let Sys denote the set of systems defined by the Syntax in Table 1 and $S, S_1, \dots, S' \dots$ variables ranging over Sys . A system S consists of an aggregation of COMPONENTS obtained via the (parallel) *composition* operator $- \parallel -$. A component $I[K, P]$ consists of:

1. An *interface*, which is a function I in the set $\mathbb{K} \rightarrow (\mathbb{A} \rightarrow \mathbb{V})$ used for publishing information about the component state in the form of attribute values. In detail, an interface defines a (partial) function from a pair knowledge-base and attribute-name to a domain of values. Among the possible attributes, \mathbf{id} is mandatory and is bound to the name of the component. Component names are not required to be unique, so that replicated service components can be modeled. The *evaluation* of an interface I in a knowledge state K is denoted as $I(K)$. The set of possible interface evaluations is denoted as \mathbb{E} .
2. A *knowledge repository* K , managing both application data and awareness data (following the approach of [4]), together with the specific handling mechanism.
3. A *process* P , together with a set of process definitions. Processes may execute local computations, coordinate local and remote interaction with a knowledge repository, or perform adaptation and reconfiguration.

Processes

PROCESSES are the active computational units. Each process is built up from the *inert* process \mathbf{nil} via *action prefixing* ($a.P$), *nondeterministic choice* ($P_1 + P_2$), *parallel composition* ($P_1 \mid P_2$), *process variable* (X), and *parameterised process invocation* ($A(\bar{p})$). We feel free to omit trailing occurrences of \mathbf{nil} , writing e.g. a instead of $a.\mathbf{nil}$, whenever there is no confusion arising.

Process variables can support *higher-order* communication, namely the capability to exchange (the code of) a process, and possibly execute it, by first adding an item containing the

process to a knowledge repository and then retrieving/withdrawing this item while binding the process to a process variable. We assume that A ranges over a set of parameterised *process identifiers* that are used in recursive process definitions. We also assume that each process identifier A has a *single* definition of the form $A(\bar{f}) \triangleq P$ where all free variables in P are contained in \bar{f} and all occurrences of process identifiers in P are within the scope of an action prefixing. \bar{p} and \bar{f} denote lists of actual and formal parameters, respectively. In the sequel we will use $Proc$ to denote the set of processes, ranged over by variables $P, Q, \dots, P_1, Q_1, \dots, P', Q', \dots$.

Actions and targets

Processes can perform three different kinds of ACTIONS: **get**(T)@ c , **qry**(T)@ c and **put**(t)@ c , used to act over shared knowledge repositories by, respectively, withdrawing, retrieving, and adding information items from/to the knowledge repository identified by c .

These actions exploit templates T as patterns to select knowledge items t in the repositories. The precise syntax of templates and knowledge items depends on the specific instance of *knowledge repository* that is used. Indeed, in Example 1 we provided the syntax for items (\mathbb{I}) and templates (\mathbb{T}) for one possible instance of the repository. In the next section we show how STOCS is in fact parametric with respect to different types of *knowledge repository*.

In the following, we identify the component executing an action as the *source*, and the component (whose knowledge state is) affected by the action as the *destination*. In order for this synchronization to occur, the *output* action (performed on the source side) has a corresponding *input* action (performed on the destination side). In terms of action labels, they are denoted by $\bar{\alpha}$ and α , respectively, while their synchronization is denoted by $\overleftarrow{\alpha}$.

For the sake of simplicity, in this report we restrict targets c to the distinguished variable **self**, that is used by processes to refer to the component hosting it, and to component *predicates* \mathfrak{p} , i.e. formulas on component attributes. A component $I[K, P]$ is *identified* by a predicate \mathfrak{p} if $I(K) \models \mathfrak{p}$, that is, the interpretation defined by the evaluation of I in the knowledge state K is a model of the formula \mathfrak{p} . Note that here we are assuming a fixed interpretation for functions and predicate symbols that are not within the attributes (\mathbb{A}). E.g. $battery < 3$ is a possible predicate, where $<$ and 3 have a fixed interpretation, while the value of *battery* depends on the specific component addressed.

The informal, abstract, semantics of the actions is the following:

- **put**(t)@ c is non-blocking, its execution causes knowledge item t be added to the knowledge repository of *all* the components (the interface of which is) identified by c , *if any*;
- **get**(T)@ c (**qry**(T)@ c , respectively) is blocking, it causes a knowledge item t matching pattern T be withdrawn (retrieved, respectively) from the knowledge repository of *any* of the components (the interface of which is) identified by c , *if any*. If no such component/item is available, the process executing it is *blocked* in a waiting state. The two actions differ for the fact that **get** removes the requested item from the knowledge repository while **qry** leaves the target repository unchanged.

The set of components satisfying a given target c of a communication action can be considered as the *ensemble* with which the process performing the action intends to interact.

Knowledge behavior

Since STOCS is parametric w.r.t. the specific knowledge repository used in a specification, we provide no specific syntax/semantics for knowledge repositories. We only require that a *knowledge repository type* is completely described by a tuple $(\mathbb{K}, \mathbb{I}, \mathbb{T}, \oplus, \ominus, \vdash)$ where \mathbb{K} is the set of possible *knowledge states* (the variables K, K_1, \dots, K', \dots range over \mathbb{K}), \mathbb{I} is the set

of *knowledge items* (the variables t, t_1, \dots, t', \dots range over \mathbb{I}) and \mathbb{T} is the set of *knowledge templates* (the variables T, T_1, \dots, T', \dots range over \mathbb{T}). Knowledge items have no variable, while knowledge templates have. We assume to have a partial function $\text{match} : \mathbb{T} \times \mathbb{I} \rightarrow \text{Subst}(\mathbb{I})$ (where $\text{Subst}(X)$ is the set of substitutions with range in X) and we denote as $\text{match}(T, t) = \vartheta$ the substitution obtained by matching the pattern T against the item t , if any. By a small abuse of notation, we write $\neg\text{match}(T, t)$ to denote that $\text{match}(T, t)$ is undefined.

The operators \oplus, \ominus, \vdash are used to add, withdraw, and infer knowledge items to/from knowledge repositories in \mathbb{K} , respectively. These functions have the following signature, where $\text{Dist}(X)$ denotes the class of probability distributions on set X with finite support:

- $\oplus : \mathbb{K} \times \mathbb{I} \rightarrow \text{Dist}(\mathbb{K})$.
- $\ominus : \mathbb{K} \times \mathbb{T} \leftrightarrow \text{Dist}(\mathbb{K} \times \mathbb{I})$;
- $\vdash : \mathbb{K} \times \mathbb{T} \leftrightarrow \text{Dist}(\mathbb{I})$;

Function \oplus is *total* and defines how a knowledge item can be inserted into a knowledge repository: $K \oplus t = \pi$ is the probability distribution over knowledge states obtained as the effect of adding t . If the item addition operation is modeled in a deterministic way, then the distribution π is a Dirac function. One advantage of allowing a probabilistic item addition operation is, for example, the ability of modeling possible failures in the item addition. We will make use of this feature in the stochastic semantics of STOCS.

Function \ominus is *partial* and computes the result of withdrawing a template from a knowledge state in terms of a probability distribution $K \ominus T$ over the set of pairs $(K, t) \in (\mathbb{K} \times \mathbb{I})$ such that the item t matches the template T . Intuitively, if $K \ominus T = \pi$ and $\pi(K', t) = p$ then, when one tries to remove an item matching template T from K , with probability p item t is obtained and the resulting knowledge state is K' . If a tuple matching template T is not found in K then $K \ominus T$ is undefined, which is indicated by $K \ominus T = \perp$.

Function \vdash is *partial* and computes (similarly to \ominus) a probability distribution over the possible knowledge items matching template T that can be inferred from K . Thus, if $K \vdash T = \pi$ and $\pi(t) = p$ then the probability of inferring t when one tries to infer from K a tuple matching T is p . If no tuple matching T can be inferred from K then $K \vdash T$ is undefined, which is indicated by $K \vdash T = \perp$.

Example 2 (K as a tuple store). Consider the sets \mathbb{I} of items and \mathbb{T} of templates defined in Example 1. A knowledge state K is a multi-set over a set $M \subseteq \mathbb{I}$, that is a pair $\langle M, f_K \rangle$, where $f_K : M \rightarrow \mathbb{N}$ is the so-called characteristic function of K . Given two multisets A and B , the functions $A \uplus B$ (B/A) denote the addition (resp. removal) of the elements of A and (resp. from) B . The restriction $M|_T$ of a multiset M to a pattern T is the sub-multiset of M of those t that match T , that is: $M|_T = \{t \mid t \in M \wedge \exists \vartheta. \text{match}(T, t) = \vartheta\}$. In the following we assume that $[x_1 \mapsto y_1, \dots, x_n \mapsto y_n]$ is the function mapping x_i to y_i for $i = 1, \dots, n$. Furthermore, we use standard conditional notation to specify y_i . Finally, given to functions f and g , we let $(f + g)x = fx + gx$. For any $t \in \mathbb{I}$, $T \in \mathbb{T}$ and knowledge state K :

$$K \oplus t = [K \uplus \{t\} \mapsto 1]$$

$$K \oplus_p t = [K \uplus \{t\} \mapsto p, K \mapsto 1-p]$$

$$K \ominus T = \sum_{t \in \mathbb{I}} [(K/\{t\}, t) \mapsto \begin{cases} 0 & \text{if } \neg\text{match}(T, t) \\ \frac{f_K(t)}{\text{card}(M|_T)} & \text{otherwise} \end{cases}]$$

$$K \vdash T = \sum_{t \in \mathbb{I}} [t \mapsto \begin{cases} 0 & \text{if } \neg\text{match}(T, t) \\ \frac{f_K(t)}{\text{card}(M|_T)} & \text{otherwise} \end{cases}]$$

Concerning \ominus and \vdash , one can choose a different probability distribution over matching items that the one presented in this example (which is a uniform distribution over multi-set elements), depending on the specific modeling domain.

2.2 Informal Timed Semantics - Four Variants

The semantics of SCEL does not consider any time related aspect of computation. More specifically, the execution of an action of the form $\mathbf{act}(T)@c.P$ (for **put/get/query** actions) is described by a *single* transition of the underlying SCEL LTS semantics. In the system state reached by such a transition it is guaranteed that the process which executed the action is in its local state P and that the knowledge repositories of all components involved in the action execution have been modified accordingly. In particular, SCEL abstracts away details concerning:

1. when the execution of the action starts;
2. if c is a predicate \mathbf{p} , when the possible destination components are required to satisfy \mathbf{p} ;
3. when the process executing the action resumes execution (i.e. becomes P);

and their consequent time relationship. If we want to extend SCEL with an explicit notion of (stochastic) time, we need to take into account the time-related issues mentioned above. These issues can be addressed at different levels of abstraction, reflecting a different choice of details that are considered in modeling SCEL actions. In this section, we discuss and motivate several choices we take in the design of STOCS and we also discuss alternative ones. In order to obtain an underlying CTMC semantics, we model state residence times in a Markovian way. Therefore, in the following, whenever we indicate that an action has rate λ , roughly speaking we mean that the duration of the action (or, equivalently, the state residence time before action execution) is modeled by a random variable (RV, in the sequel) with negative exponential distribution having rate λ . Strictly speaking, the actual residence-time depends also on other conflicting actions the process may be engaged in, and the resulting race-condition.

Concerning point (1) above, since actions may involve one or more components (except those executing on self), it is reasonable to distinguish the period in which the source component determines the set of target components (*observation*) from the period in which these target components are involved (*action realization*).

Point (2), in the case c is the predicate \mathbf{p} , requires to define *when* a component satisfies \mathbf{p} with respect to a process executing an action, when time and possibly space are taken into consideration. We assume that source components are not aware of which are the components satisfying predicate \mathbf{p} . Therefore, we define the notion of *observation* of the component by the process, the result of which allows to establish whether the component satisfies the predicate or not. In the context of distributed systems this is very often realized by means of a message sent by the process to the component. According to this view, the check whether a component satisfies predicate \mathbf{p} is performed *when the message reaches it*. This means that a STOCS action may require broadcast communication to be executed, even if its effect involves a few (and possibly no) components. In distributed systems different components may have different response times depending on different network conditions and one can model explicitly the message delivery, taking into account the time required to reach the component.

Finally, point (3) raises the issue on when source component execution is to be resumed. In particular, it is necessary to identify how the source component is made aware that its role in the communication has been completed. Get/query actions are blocking and they terminate when the source receives a knowledge item from any component. A reasonable choice is that further responses received are ignored. We assume appropriate mechanisms that ensure no

confusion arises between distinct actions and corresponding messages. Put actions are non-blocking, so it is sufficient that the source component is aware that all reachable components have been involved in the evaluation of the predicate. A possible choice is to set-up the transmission of one request of predicate evaluation for each component and then terminate the execution on the source side immediately. On the target side, it is necessary to model the reception time as well as subsequent evaluation and corresponding knowledge repository modification.

Depending on the degree of detail in modeling these aspects, we will define four different semantics, that we call: *network-oriented* (NET-OR), *action-oriented* (ACT-OR), *interaction-oriented* (INT-OR), and *activity-oriented* (ACTIV-OR). These semantics have an increasing level of abstraction, which allows us to manage the complexity of the model according to the application of interest. In the remaining part of this section we informally describe these four variants of the stochastic semantics and their motivations. Then, in Sections 6, 4, 5, and 7, we provide the corresponding operational semantics. The sections containing the operational semantics are ordered with the purpose of presenting the simpler semantics first (which is the action-oriented one), and then the others in an incremental way, as refinements or modifications of the previous ones.

2.2.1 Network-oriented Semantics

This semantics takes into account all of the assumptions and observations described previously, which entails that actions are *non-atomic*. Indeed, they are executed through several intermediate steps, each of which requires appropriate time duration modeling. In particular, **put** actions are realized in two steps: (1) an envelope preparation and shipping (one for each component in the system, other than the source), (2) envelope delivery, with its own delivery time, test of the truth value of the communication predicate, and update of the knowledge-state. The actions **get/qry** are realized in two steps: (1) initiation of the item retrieval by a source component by entering in a waiting state, (2) synchronization with a destination component and exchange of the retrieved item. Since actions are not executed atomically, their execution is interleaved with that of other actions executed in parallel. This kind of semantics is appropriate for models with spatial aspects, where distribution is a sensible aspect influencing the duration of communications on the basis of the location of components. A reasonable assumption is that rates of locally executed actions are rather high, if compared to rates of remote actions. If the rates of local and remote actions differ orders of magnitudes, this may lead to the so-called “stiff” Markov models. An efficient analysis of such models may require a multi-scale (in time) analysis approach.

In Section 6 we provide an operational semantics for this variant of the language.

2.2.2 Action-oriented

In this simplified semantics, an action of the form $\mathbf{act}(T)@c$ (for **put/get/qry** actions) is described by a *single* transition and has a state residence time provided by a function \mathcal{R} that takes into account the *source* component, the *target* component, the *cost* of retrieved/transmitted knowledge item, and possibly other parameters. Thus, we can expect this rate to (partially) take into account different locations and thus different response times of components. Upon item retrieval (by a **get/qry** action), eligible components (in terms of predicate satisfaction and availability of requested item) are in a race for response, with rate assigned component-wise. The underlying stochastic semantics drives the outcome of the race, with appropriate weighting depending on the rates. Even if this semantics does not consider the realistic aspects of predicate-based communication, it is simpler and can be used in all of those scenarios where the average execution time of actions does not depend on the

number of components involved in the communication.

In Section 4 we define an operational semantics based on these ideas.

2.2.3 Interaction-oriented

This is based on the action-oriented semantics and distinguishes local and remote actions, by assuming that local actions are executed instantaneously. In some scenarios, local actions happen in a time-scale which is very different (usually much smaller) from that of remote actions. In these situations it is reasonable to consider as instantaneous the execution of local actions, which is the idea we realize in the definition of the INT-OR semantics. As a useful side effect of ignoring the duration of local actions, we obtain more concise models. This approximation in the modeling can be considered as an approach to reducing multi-scale models to single-scale models. In these single-scale models, the macro-scale of inter-component communication drives the execution of macro-actions. A similar idea is explored for Bio-PEPA models in [9] and used to abstract away from fast reactions in biochemical networks. There, under the so-called Quasi-Steady-State Assumption of the system, it is also defined a form of bisimilarity between the abstract and the concrete model.

The assumption we make for defining this semantics is that each remote (timed) action is followed by a (possibly empty) sequence of local (probabilistic) actions. We ensure this assumption is satisfied by imposing syntactic restrictions on processes. Then, by realizing a form of maximal progress [10] we execute a timed action and all of its subsequent probabilistic actions in a single transition of the STOCS LTS. Note that, similarly to the NET-OR semantics, also in this semantics we assume an error probability (called p_{err}) modeling failed delivery of the **put** action.

In Section 5 we define an operational semantics based on these ideas.

2.2.4 Activity-oriented

This semantics is very abstract and allows to explicitly declare as atomic an entire sequence of actions, by assigning to it an execution rate that models the duration of the entire sequence. Since the execution of the sequence of actions is atomic, it allows no interleaving of other actions. As an interesting consequence of this, we have a significant reduction in the state-space of the system.

This variant of the semantics is motivated by the fact that STOCS only provides primitives for asynchronous communication. Synchronization, if needed, has to be encoded through a protocol using asynchronous communication primitives [14]. Whatever is the adopted semantics among the previous ones (for example, the Action-oriented semantics), the protocol execution for the synchronization action is interleaved with the execution of other actions. This leads to have unclear dependencies of protocol execution times from the environment. Therefore, the Activity-oriented semantics allows us to declare as atomic an entire sequence of actions and to assign a rate to it. More in general, the purpose of this semantics is to have a very high-level abstraction of the interaction mechanisms. Of course, this must be handled with care as potentially relevant system behaviors (and interleavings) may be no longer present in the model. Therefore, properties of the model are not necessarily satisfied also by the system. Also in this semantics we assume an error probability (called p_{err}) modeling failed delivery of the **put** action.

3 Preliminary Definitions for Operational Semantics

In this section we provide preliminary notions to support the presentation of the four stochastic semantics of STOCS formalizing the ideas described in the previous section. The semantics

definition is given in the FUTSs style [8] and, in particular, using its Rate Transition Systems (RTS) instantiation [7].

3.1 Preliminaries

In RTSs a transition is a triple of the form (P, α, \mathcal{P}) , the first and second components of which are the source state and the transition label, as usual, and the third component \mathcal{P} is the *continuation function* that associates a real non-negative value with each state P' . A non-zero value represents the rate of the exponential distribution characterizing the time needed for the execution of the action represented by α , necessary to reach P' from P via the transition. Whenever $\mathcal{P}(P') = 0$, this means that P' is not reachable from P via α . RTS continuation functions are equipped with a rich set of operations that help to define these functions over sets of processes, components, and systems. Below we show the definition of those functions that we use in this paper, after having recalled some basic notation, and we define them in an abstract way, with respect to a generic set X .

Let $\mathbf{TF}(X, \mathbb{R}_{\geq 0})$ denote the set of *total* functions from X to $\mathbb{R}_{\geq 0}$, and $\mathcal{F}, \mathcal{P}, \mathcal{Q}, \mathcal{R}, \dots$ range over it. We define $\mathbf{FTF}(X, \mathbb{R}_{\geq 0})$ as the subset of $\mathbf{TF}(X, \mathbb{R}_{\geq 0})$ containing only functions with *finite support*: function \mathcal{F} is an element of $\mathbf{FTF}(X, \mathbb{R}_{\geq 0})$ if and only if there exist $\{d_1, \dots, d_m\} \subseteq X$, the *support* of \mathcal{F} , such that $\mathcal{F} d_i \neq 0$ for $i = 1 \dots m$ and $\mathcal{F} d = 0$ for all $d \in X \setminus \{d_1, \dots, d_m\}$. We equip $\mathbf{FTF}(X, \mathbb{R}_{\geq 0})$ with the operators defined below. The resulting *algebraic structure* of the set of finite support functions will be crucial for the compositional features of our approach.

Def. 3.1. Let X be a set, and d, d_1, \dots, d_m be distinct elements of X , $\gamma_1, \dots, \gamma_m \in \mathbb{R}_{\geq 0}$, let furthermore $\bullet : X \times X \rightarrow X$ be an injective binary operator, \mathcal{F}_1 and \mathcal{F}_2 in $\mathbf{FTF}(X, \mathbb{R}_{\geq 0})$:

1. $[d_1 \mapsto \gamma_1, \dots, d_m \mapsto \gamma_m]$ denotes the following function:

$$[d_1 \mapsto \gamma_1, \dots, d_m \mapsto \gamma_m] d =_{\text{def}} \begin{cases} \gamma_i & \text{if } d = d_i \in \{d_1, \dots, d_m\}, \\ 0 & \text{otherwise.} \end{cases}$$

the 0 constant function in $\mathbf{FTF}(X, \mathbb{R}_{\geq 0})$ is denoted by $[]$;

2. Function $+$ is defined as $(\mathcal{F}_1 + \mathcal{F}_2) d =_{\text{def}} (\mathcal{F}_1 d) + (\mathcal{F}_2 d)$;
3. We lift operator \bullet to $\mathbf{FTF}(X, \mathbb{R}_{\geq 0})$ as follows (with a bit of overloading):

$$(\mathcal{F}_1 \bullet \mathcal{F}_2) s =_{\text{def}} \begin{cases} (\mathcal{F}_1 d_1) \cdot (\mathcal{F}_2 d_2) & \text{if } \exists d_1, d_2 \in X . d = d_1 \bullet d_2, \\ 0 & \text{otherwise.} \end{cases}$$

4. The characteristic function $\mathcal{X} : X \rightarrow \mathbf{FTF}(X, \mathbb{R}_{\geq 0})$ with $\mathcal{X} d =_{\text{def}} [d \mapsto 1]$.

Def. 3.2. An A -labelled Rate Transition System (RTS) is a tuple $(S, A, \mathbb{R}_{\geq 0}, \mapsto)$ where S and A are countable, non-empty, sets of states and transition labels, respectively, and $\mapsto \subseteq S \times A \times \mathbf{FTF}(S, \mathbb{R}_{\geq 0})$ is the A -labelled transition relation.

In order to distinguish and identify the rules of the semantics definition, we label them by unique names. Note that a rule with name r may have one or more associated blocking rules r_B which have the role of allowing the execution of no actions other than those explicitly allowed by existing inference rules. These B-rules will not be further commented in the following sections.

4 Action-oriented Operational Semantics

To facilitate the presentation of the STOCS variants, we start with the semantics called Action-oriented (abbreviated to ACT-OR), associating a single transition to each STOCS action. The more complex semantics, presented in Sections 5, 6, and 7, will be described as modifications of this simpler one.

Before discussing the semantics in detail, we want to discuss the role of interface evaluations. We have explained, in Sec. 2, that interface evaluations are used to establish the truth value of predicates. Now we want to discuss how they are also used to compute appropriate action rates. Indeed, interface evaluations are included within the transition labels to exchange information about source/destination components in a synchronization action, as we now explain.

Interface evaluations (i.e. the values of the component attributes at a given time) are used within a *rate function* $\mathcal{R} : \mathbb{E} \times Act \times \mathbb{E} \rightarrow \mathbb{R}_{\geq 0}$, which takes the interface evaluation of the *source*, an action in the set of labels

$$Act = \{\mathbf{put}(t)@c, \mathbf{get}(T : t)@c, \mathbf{qry}(T : t)@c \mid t \in \mathbb{I} \text{ and } T \in \mathbb{T} \text{ and } c \in \text{TARGET}\}$$

and the interface evaluation of the *destination*, and returns a value in $\mathbb{R}_{\geq 0}$, which is the rate of execution of the given action with counterparts having those interface evaluations. Note that **get**/**qry** actions argument has been set to $T : t$ (rather than T as in Table 1) because the *labels* of **get**/**qry** transition will contain also the matching/retrieved term t . We do not explicitly give the definition of this function, as it is a parameter of the language¹. Considering interface evaluations in the rate functions, together with the executed action, allow us to keep into account, in the computation of the action rate, various aspects depending on the component state such as the position/distance, as well as other time-dependent (i.e. knowledge state dependent) parameters such as the load of the component.

We also assume to have an *error function* $f_{\text{err}} : \mathbb{E} \times Act \times \mathbb{E} \rightarrow [0, 1]$ that, similarly to the function \mathcal{R} , computes the probability of an error in the delivery of messages. Also for this function we do not explicitly give the definition, as it is a parameter of the language, but we expect it to be defined through an appropriate syntax.

Interface evaluations, the rate function \mathcal{R} , and the error function f_{err} have the roles described above also in the other semantics, presented in Sections 5, 6, and 7, so we avoid to mention and explain them again in those sections.

A further clarification we want to make in this premise concerns synchronization of input and output actions. As mentioned in Sec. 2, an input action α and an output action $\bar{\alpha}$ synchronize into a $\overleftrightarrow{\alpha}$ action. These synchronizations can happen within a component, for *local* actions (@self), between a source and a destination component, in *remote* **get**/**qry** actions (@p), and between a source component and all the remaining components, in *remote* **put** actions (@p). In the case of **get**/**qry** actions, the interface transmitted in the synchronization label will be that of the destination, which allows to have a different rate for each possible destination. In the case of **put** action, the interface transmitted in the synchronization label will be that of the source, since there are many possible destinations which are all involved (at once) in the computation. Therefore, for the **put** action, we abstract from the actual destinations and we compute the rate only on the basis of the source.

Finally, to simplify the synchronization of input and output actions, we force input actions to be *probabilistic*, and output actions to be *stochastic*. Therefore their composition can be performed through a simple multiplication.

¹We plan, however, to extend the language with an appropriate syntax for defining the rate function in terms of sender/receiver component attributes and sent messages.

Inactive process:

$$\frac{}{\mathbf{nil} \xrightarrow{\alpha} []} \text{ (NIL)}$$

Actions (where, $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$ and c is a TARGET):

$$\frac{\lambda = \mathcal{R}(\sigma, \mathbf{put}(t)@c, -)}{\mathbf{put}(t)@c . P \xrightarrow{\mathbf{put}(t)@c}_{\sigma} [P \mapsto \lambda]} \text{ (PUT)}} \quad \frac{\alpha \neq \overline{\mathbf{put}(t)@c}}{\mathbf{put}(t)@c . P \xrightarrow{\alpha} []} \text{ (PUT}_B)$$

$$\frac{\text{match}(T, t) = \vartheta \quad \lambda = \mathcal{R}(\sigma, \mathbf{gq}(T : t)@c, \delta)}{\mathbf{gq}(T)@c . P \xrightarrow{\delta : \mathbf{gq}(T : t)@c}_{\sigma} [P \vartheta \mapsto \lambda]} \text{ (GQ)}$$

$$\frac{\neg \text{match}(T, t)}{\mathbf{gq}(T)@c . P \xrightarrow{- : \mathbf{gq}(T : t)@c}_{\sigma} []} \text{ (GQB}_1)} \quad \frac{\alpha \neq - : \mathbf{gq}(T : t)@c}{\mathbf{gq}(T)@c . P \xrightarrow{\alpha} []} \text{ (GQB}_2)$$

Choice, definition, and parallel composition:

$$\frac{P \xrightarrow{\alpha}_{e'} \mathcal{P} \quad Q \xrightarrow{\alpha}_{e'} \mathcal{Q}}{P + Q \xrightarrow{\alpha}_{e'} \mathcal{P} + \mathcal{Q}} \text{ (CHO)} \quad \frac{A(\vec{x}) \stackrel{\text{def}}{=} P \quad P[\vec{v}/\vec{x}] \xrightarrow{\alpha}_{e'} \mathcal{P}}{A(\vec{v}) \xrightarrow{\alpha}_{e'} \mathcal{P}} \text{ (DEF)}$$

$$\frac{P \xrightarrow{\alpha}_{e'} \mathcal{P} \quad Q \xrightarrow{\alpha}_{e'} \mathcal{Q}}{P \mid Q \xrightarrow{\alpha}_{e'} \mathcal{P} \mid (\mathcal{X} Q) + (\mathcal{X} P) \mid \mathcal{Q}} \text{ (PAR)}$$

Table 2: Operational semantics of STOCS processes (ACT-OR).

4.1 Operational semantics of processes

The ACT-OR semantics of STOCS *processes* is the RTS $(Proc, Act_{Proc}, \mathbb{R}_{\geq 0}, \rightarrow_e)$ where $Proc$ is the set of process terms defined according to the syntax of STOCS given in Table 1. The set Act_{Proc} of labels is defined according to the grammar below (where e' is the evaluation of an interface, $t \in \mathbb{I}$, $T \in \mathbb{T}$, and c is a TARGET) and it is ranged over by α, α', \dots :

$$Act_{Proc} ::= \overline{\mathbf{put}(t)@c} \mid e' : \overline{\mathbf{get}(T : t)@c} \mid e' : \overline{\mathbf{qry}(T : t)@c}$$

and $\rightarrow_e \subseteq Proc \times Act_{Proc} \times \mathbf{FTF}(Proc, \mathbb{R}_{\geq 0})$ is the least relation satisfying the rules of Table 2, which we describe in the rest of this section (\rightarrow_e is parameterized by an interface evaluation e : we feel free to omit the parameter if unnecessary).

(NIL) \mathbf{nil} is the terminated process, since no process is reachable from it via any action.

(PUT) allows a process to execute an action \mathbf{put} over the knowledge state of a target c with rate λ , which is computed by the rate function \mathcal{R} . The execution of a $\mathbf{put}(t)@c$ action depends on the source component and *all* the other components in the system, which are involved as potential destinations. Consequently, the execution rate λ can be seen as a function of the action and of the source component only (whose interface evaluation is obtained as a parameter of the \rightarrow relation); in particular, the action rate *does not* depend on (the interface evaluation of) a specific (destination) component (since the destination is not unique because \mathbf{put} is a broadcast action); this is represented by using the symbol $-$ in the destination argument of \mathcal{R} .

(GQ) allows a process to execute a $\mathbf{get}/\mathbf{qry}$ action over the knowledge state of a target c with rate λ , computed by the function \mathcal{R} considering (i) the interface evaluation σ of

the source component (obtained as a parameter of the \rightarrow relation), (ii) the interface evaluation δ of the destination (received in the synchronization label), and (iii) the sent template T and the retrieved item t . The transition is labeled $\delta : \mathbf{gq}(T : t)@c$, which indicates a request for a knowledge item t matching template T from a component identified by c . Note that, in the case of $\mathbf{gq}(T)@c$ the execution rate depends also on the destination interface evaluation because only one destination will be involved in the completion of the execution of the action.

- (CHO) accumulates the relevant rates by the application of the *sum* operator $+$ on the continuation \mathcal{P} of P and \mathcal{Q} of Q , according to the race condition principle of CTMCs;
- (DEF) allows the instantiation of a process definition;
- (PAR) realizes process parallel composition $P \mid Q$ and uses Def. 3.1, item (3) applied to the process parallel composition syntactic constructor \mid (which is obviously injective).

Concerning the rule (PAR), given two functions \mathcal{R}_1 and \mathcal{R}_2 , the function $\mathcal{R}_1 \mid \mathcal{R}_2$ applied to process term R returns the product $(\mathcal{R}_1 R_1) \cdot (\mathcal{R}_2 R_2)$, whenever R is of the form $R_1 \mid R_2$, for some terms R_1 and R_2 , and 0 otherwise. In the rule, also the characteristic function \mathcal{X} is used. Function $\mathcal{P} \mid (\mathcal{X}Q)$ applied to R returns $\mathcal{P} R'$ if $R = R' \mid Q$ for some R' and 0 otherwise; i.e. the function behaves as the continuation \mathcal{P} of P for terms where Q does not progress (for one step). In conclusion, $\mathcal{P} \mid (\mathcal{X}Q) + (\mathcal{X}P) \mid \mathcal{Q}$ correctly represents process interleaving, keeping track of the relevant rates.

4.2 Operational semantics of components and systems

The ACT-OR semantics of STOCS *systems* is the RTS $(Sys, Act_{Sys}, \mathbb{R}_{\geq 0}, \rightarrow)$ where Sys is the set of system terms defined according to the syntax of STOCS given in Table 1. The set Act_{Sys} of labels is defined according to the grammar below (where $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$, e' is the evaluation of an interface, $t \in \mathbb{I}$, $T \in \mathbb{T}$, and \mathbf{p} is a PREDICATE):

$$\begin{array}{lll}
 Act_{Sys} ::= & e' : \mathbf{put}(t)@p & \mid \quad e' : \mathbf{gq}(T : t)@p & \mid \quad \text{(input actions)} \\
 & \overline{e' : \mathbf{put}(t)@p} & \mid \quad \overline{e' : \mathbf{gq}(T : t)@p} & \mid \quad \text{(output actions)} \\
 & \overleftarrow{e' : \mathbf{put}(t)@self} & \mid \quad \overleftarrow{e' : \mathbf{gq}(T : t)@c} & \mid \quad \text{(synchronizations)}
 \end{array}$$

and $\rightarrow \subseteq Sys \times Act_{Sys} \times \mathbf{FTF}(Sys, \mathbb{R}_{\geq 0})$ is the least relation satisfying the rules of Tables 3 and 4, where the process relation \rightarrow_e defined in Table 2 is also used.

In Table 3, rules for components are grouped by action type:

- (C-PUTL) allows a local process P to execute a local **put** action. Let $I[K, P]$ be a component; this rule states that P executes action $\mathbf{put}(t)@self$ with local interface evaluation $\sigma = I(K)$ and evolves to \mathcal{P} , then a local execution of the action can occur and the entire component evolves with label $\sigma : \mathbf{put}(t)@self$ to $I[\pi, \mathcal{P}]$, where $\pi = K \oplus t$ is a probability distribution over the possible knowledge states obtained from K by adding the knowledge item t , while $I[\pi, \mathcal{P}]$ is the function which maps any term of the form $I[K, P]$ to $(\pi K) \cdot (\mathcal{P}P)$ and any other term to 0.
- (C-PUTO) is used when the target of a **put** is not **self** but a predicate \mathbf{p} . This rule simply lifts an output **put** action from the process level to the component level and transmits to its counterpart its current interface evaluation σ by including it in the transition label (which will be used in rule (C-PUTI) to compute the communication error probability).

put actions:

$$\begin{array}{c}
\frac{\sigma = I(K) \quad P \xrightarrow{\overline{\mathbf{put}(t)@\mathbf{self}}}_{\sigma} \mathcal{P} \quad K \oplus t = \pi}{I[K, P] \xrightarrow{\overleftarrow{\sigma:\mathbf{put}(t)@\mathbf{self}}} I[\pi, \mathcal{P}]} \quad (\text{C-PUTL}) \\
\\
\frac{\sigma = I(K) \quad P \xrightarrow{\overline{\mathbf{put}(t)@\mathbf{p}}}_{\sigma} \mathcal{P}}{I[K, P] \xrightarrow{\overleftarrow{\sigma:\mathbf{put}(t)@\mathbf{p}}} I[(\mathcal{X}K), \mathcal{P}]} \quad (\text{C-PUTO}) \\
\\
\frac{\delta = I(K) \quad \delta \models \mathbf{p} \quad K \oplus t = \pi \quad p_{\text{err}} = f_{\text{err}}(\sigma, \mathbf{put}(t)@\mathbf{p}, \delta)}{I[K, P] \xrightarrow{\overleftarrow{\sigma:\mathbf{put}(t)@\mathbf{p}}} [I[K, P] \mapsto p_{\text{err}}] + I[\pi, (\mathcal{X}P)] \cdot (1 - p_{\text{err}})} \quad (\text{C-PUTI}) \\
\\
\frac{I(K) \not\models \mathbf{p}}{I[K, P] \xrightarrow{\overleftarrow{\sigma:\mathbf{put}(t)@\mathbf{p}}} [I[K, P] \mapsto 1]} \quad (\text{C-PUTIR})
\end{array}$$

get/qry actions (where, $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$):

$$\begin{array}{c}
\frac{\sigma = I(K) \quad P \xrightarrow{\overline{\sigma:\mathbf{get}(T:t)@\mathbf{self}}}_{\sigma} \mathcal{P} \quad K \ominus T = \pi}{I[K, P] \xrightarrow{\overleftarrow{\sigma:\mathbf{get}(T:t)@\mathbf{self}}} I[\pi(t), \mathcal{P}]} \quad (\text{C-GETL}) \quad \frac{K \ominus T = \perp}{I[K, P] \xrightarrow{\overleftarrow{\sigma:\mathbf{get}(T:t)@\mathbf{self}}} \square} \quad (\text{C-GETLB}) \\
\\
\frac{\sigma = I(K) \quad P \xrightarrow{\overline{\sigma:\mathbf{qry}(T:t)@\mathbf{self}}}_{\sigma} \mathcal{P} \quad K \vdash T = \pi}{I[K, P] \xrightarrow{\overleftarrow{\sigma:\mathbf{qry}(T:t)@\mathbf{self}}} I[(\mathcal{X}K) \cdot \pi(t), \mathcal{P}]} \quad (\text{C-QRYL}) \quad \frac{K \vdash T = \perp}{I[K, P] \xrightarrow{\overleftarrow{\sigma:\mathbf{qry}(T:t)@\mathbf{self}}} \square} \quad (\text{C-QRYLB}) \\
\\
\frac{\sigma = I(K) \quad P \xrightarrow{\overline{\delta:\mathbf{gq}(T:t)@\mathbf{p}}}_{\sigma} \mathcal{P}}{I[K, P] \xrightarrow{\overleftarrow{\delta:\mathbf{gq}(T:t)@\mathbf{p}}} I[K, \mathcal{P}]} \quad (\text{C-GQO}) \\
\\
\frac{\delta = I(K) \quad \delta \models \mathbf{p} \quad K \ominus T = \pi}{I[K, P] \xrightarrow{\overleftarrow{\delta:\mathbf{get}(T:t)@\mathbf{p}}} I[\pi(t), (\mathcal{X}P)]} \quad (\text{C-GETI}) \quad \frac{\delta \neq I(K) \vee I(K) \not\models \mathbf{p} \vee K \ominus T = \perp}{I[K, P] \xrightarrow{\overleftarrow{\delta:\mathbf{get}(T:t)@\mathbf{p}}} \square} \quad (\text{C-GETIB}) \\
\\
\frac{\delta = I(K) \quad \delta \models \mathbf{p} \quad K \vdash T = \pi}{I[K, P] \xrightarrow{\overleftarrow{\delta:\mathbf{qry}(T:t)@\mathbf{p}}} [I[K, P] \mapsto \pi(t)]} \quad (\text{C-QRYI}) \quad \frac{\delta \neq I(K) \vee I(K) \not\models \mathbf{p} \vee K \vdash T = \perp}{I[K, P] \xrightarrow{\overleftarrow{\delta:\mathbf{qry}(T:t)@\mathbf{p}}} \square} \quad (\text{C-QRYIB})
\end{array}$$

Table 3: Operational semantics of STOCS components (ACT-OR).

(C-PUTI) models a component *accepting* a **put**. The rule is applied when the component satisfies the predicate \mathbf{p} . When predicate \mathbf{p} is satisfied by $I(K)$ and $K \oplus t = \pi$, component $I[K, P]$ evolves to $[I[K, P] \mapsto p_{\text{err}}] + I[\pi, (\mathcal{X}P)] \cdot (1 - p_{\text{err}})$. The first term, that is selected with probability p_{err} , models a failure in the action execution, for instance due to a communication error. Value p_{err} is computed by the function f_{err} taking into account the source, the action performed, and the destination. The second term identifies the different configurations the component can reach when knowledge item t is added to the knowledge repository K .

(C-PUTIR) accepts an input **put** action producing no effect, under the assumption that the predicate \mathbf{p} is not satisfied (the component allows synchronization as required by the broadcast scheme).

(C-GETL)/(C-QRYL) realize the local **get** (**qry**) action retrieving an item $t \in \mathbb{I}$ matching the

pattern T , if possible, in the execution of process P (in the label of the process action we include the item t and we include the interface evaluation σ of the component for computing the action rate) and, since the **get** (**qry**) action may result in several distinct knowledge bases, these need to be summed together considering all possibilities: π is a distribution over pairs (knowledge base and knowledge item) and the possible components in the continuation are weighted by using π ;

(C-GQO) lifts a **get**/**qry** output action to the level of the component, similarly to (C-PUTO), and produces a continuation computed by the \rightarrow relation (in rule GQ) using the source interface evaluation (σ) and the destination interface evaluation (δ , taken from the synchronization label);

(C-GETI)/(C-QRYI) under the condition that a component satisfies the predicate \mathbf{p} , allow the component to accept an input **get** (**qry**) action (respectively), by computing the continuation on the basis of the distribution π over possible pairs of knowledge state and retrieved item (retrieved item, respectively), computed by using the \ominus (\vdash , respectively) operator, but they do not compute any rate as this is computed on the source side by rules (GQ/C-GQO); notice that the component communicates its interface evaluation δ (destination) in the synchronization label.

It is worth noting that rules (C-PUTI) and (C-PUTIR) deal with probability only. In fact, the actual rate of the action is the one which will result from system synchronization (Rules (S-PO) and (S-PI) in Table 4) on the basis of the rates settled by the rule (PUT) of Table 2.

Concerning (C-PUTL), (C-PUTI), (C-GETL), and (C-GETI), **put**/**get** actions may result in several distinct knowledge states and the rate functions need to be summed together considering all possible knowledge states. This is addressed by using Def. 3.1, item (3) applied to the component syntactic constructor $I[\cdot, \cdot]$, which is obviously injective. Thus, for $\pi \in \text{Dist}(\mathbb{K})$ and $\mathcal{P} \in \mathbf{FTF}(Proc, \mathbb{R}_{\geq 0})$, function $I[\pi, \mathcal{P}] \in \mathbf{FTF}(Sys, \mathbb{R}_{\geq 0})$ returns $(\mathcal{P} P) \cdot \pi(K)$ on systems $I[K, P]$, for any P , and 0 on any other system term. On the contrary, when we consider rules (C-QRYL) and (C-QRYI), even though **qry** may return several retrieved knowledge items, the actual transition fixes a precise item t , whose retrieval probability $\pi(t)$ is combined with the component continuation.

Let us now consider systems of components: similarly to $|$ in the case of components, also the definition of the semantics of system parallel composition $S_1 \parallel S_2$ uses Def. 3.1, item (3) applied to the system parallel composition syntactic constructor \parallel , which is obviously injective. As usual, interleaving is modeled as a combination of lifted \parallel , $+$ on functions and the characteristic function. In Table 4 we give the rules for systems, grouped by action type:

(S-PO)/(S-PI) realize the broadcast communication of **put**: (S-PO) ensures that if any subsystem executes an output **put** action (i.e. it executes a $\sigma : \mathbf{put}(t)@p$ labeled transition), then the remaining subsystem must execute the corresponding input **put** action (i.e. it should execute a $\sigma : \mathbf{put}(t)@p$ labeled transition). the composed system does not exhibit a synchronization label, but it rather propagates the output $\sigma : \mathbf{put}(t)@p$ to allow further synchronization with all the other components in parallel; in the computation of the final rate it is necessary to consider output on the left sub-system and input on the right as well as the symmetric case. Rule (S-PI) ensures that all subcomponents of a system synchronize, all together, on a (specific) input **put** action, completing the broadcast communication. Note that each component is constantly enabled on the input label for any **put** action (rules (C-PUTI) and (C-PUTIR)).

(S-GQS) realizes one-to-one synchronization of **get**/**qry** actions (which are not broadcast), denoted by a $\overleftarrow{e : \mathbf{gq}(T : t)@p}$ label, and performs aggregation of: (1) the synchronization rate of the left (right) sub-system, with the right (resp. left) subsystem that must

put synchronization:

$$\frac{S_1 \xrightarrow{\overline{\sigma : \text{put}(t)@p}} \mathcal{S}_1^o \quad S_1 \xrightarrow{\sigma : \text{put}(t)@p} \mathcal{S}_1^i \quad S_2 \xrightarrow{\overline{\sigma : \text{put}(t)@p}} \mathcal{S}_2^o \quad S_2 \xrightarrow{\sigma : \text{put}(t)@p} \mathcal{S}_2^i}{S_1 \parallel S_2 \xrightarrow{\overline{\sigma : \text{put}(t)@p}} \mathcal{S}_1^o \parallel \mathcal{S}_2^i + \mathcal{S}_1^i \parallel \mathcal{S}_2^o} \quad (\text{S-PO})$$

$$\frac{S_1 \xrightarrow{\sigma : \text{put}(t)@p} \mathcal{S}_1 \quad S_2 \xrightarrow{\sigma : \text{put}(t)@p} \mathcal{S}_2}{S_1 \parallel S_2 \xrightarrow{\sigma : \text{put}(t)@p} \mathcal{S}_1 \parallel \mathcal{S}_2} \quad (\text{S-PI})$$

get/qry synchronization ($\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$):

$$\frac{S_1 \xrightarrow{\overleftarrow{\delta : \mathbf{gq}(T:t)@p}} \mathcal{S}_1^s \quad S_1 \xrightarrow{\overline{\delta : \mathbf{gq}(T:t)@p}} \mathcal{S}_1^o \quad S_1 \xrightarrow{\delta : \mathbf{gq}(T:t)@p} \mathcal{S}_1^i}{S_2 \xrightarrow{\overleftarrow{\delta : \mathbf{gq}(T:t)@p}} \mathcal{S}_2^s \quad S_2 \xrightarrow{\overline{\delta : \mathbf{gq}(T:t)@p}} \mathcal{S}_2^o \quad S_2 \xrightarrow{\delta : \mathbf{gq}(T:t)@p} \mathcal{S}_2^i} \quad (\text{S-GQS})$$

$$S_1 \parallel S_2 \xrightarrow{\overleftarrow{\delta : \mathbf{gq}(T:t)@p}} \mathcal{S}_1^s \parallel (\mathcal{X} S_2) + \mathcal{S}_1^o \parallel \mathcal{S}_2^i + \mathcal{S}_1^i \parallel \mathcal{S}_2^o + (\mathcal{X} S_1) \parallel \mathcal{S}_2^s$$

$$\frac{S_1 \xrightarrow{\delta : \mathbf{gq}(T:t)@p} \mathcal{S}_1 \quad S_2 \xrightarrow{\delta : \mathbf{gq}(T:t)@p} \mathcal{S}_2}{S_1 \parallel S_2 \xrightarrow{\delta : \mathbf{gq}(T:t)@p} \mathcal{S}_1 \parallel (\mathcal{X} S_2) + (\mathcal{X} S_1) \parallel \mathcal{S}_2} \quad (\text{S-GQI})$$

Internal actions ($\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$):

$$\frac{S_1 \xrightarrow{\overleftarrow{e : \text{put}(t)@self}} \mathcal{S}_1 \quad S_2 \xrightarrow{\overleftarrow{e : \text{put}(t)@self}} \mathcal{S}_2}{S_1 \parallel S_2 \xrightarrow{\overleftarrow{e : \text{put}(t)@self}} \mathcal{S}_1 \parallel (\mathcal{X} S_2) + (\mathcal{X} S_1) \parallel \mathcal{S}_2} \quad (\text{S-SPL})$$

$$\frac{S_1 \xrightarrow{\overleftarrow{e : \mathbf{gq}(T:t)@self}} \mathcal{S}_1 \quad S_2 \xrightarrow{\overleftarrow{e : \mathbf{gq}(T:t)@self}} \mathcal{S}_2}{S_1 \parallel S_2 \xrightarrow{\overleftarrow{e : \mathbf{gq}(T:t)@self}} \mathcal{S}_1 \parallel (\mathcal{X} S_2) + (\mathcal{X} S_1) \parallel \mathcal{S}_2} \quad (\text{S-SGQL})$$

Table 4: Operational semantics of STOCS systems (ACT-OR).

not progress (this is realized using the \mathcal{X} characteristic function), and (2) output rates of the left sub-system and input rates of the right subsystem (as well as the symmetric case), combined with the \parallel operator.

(S-GQI) realizes the input **get** (resp. **qry**) action for systems in which *one* component, among those satisfying the target predicate and having a matching knowledge item, can answer.

(S-SPL/S-SGQL) allow a system to execute a local **put/get/qry** action and exposes the label denoting the type of action to allow appropriate aggregation of the observed rates.

As a final observation, we want to underline that the interleaving semantics for concurrency is realized, for **put** actions, by rule (S-PO) and, for **get/qry** actions, by rule (S-GQS).

5 Interaction-oriented Operational Semantics

In this section we illustrate the Interaction-oriented semantics (abbreviated to INT-OR), which is based on the idea of having different semantics for local and remote actions (as defined in

Section 4). In particular, remote actions maintain their stochastic semantics defined in ACT-OR, while local actions are instantaneous and their non-determinism is resolved by relying on probability distributions computed by applying the knowledge base operators \oplus , \ominus , and \vdash . Since we are interested in obtaining a CTMC semantics also for this variant of STOCS, we require that the execution of a remote action α_R by a component C includes the execution of *all* the subsequent local actions $\alpha_{L_1} \cdots \alpha_{L_n}$ in the behavior of C , until a new remote action or a **nil** is found, by applying a sort of maximal execution [10]. As a consequence, we consider as *atomic* the execution of $\alpha_R \alpha_{L_1} \cdots \alpha_{L_n}$ by the component C . To this end, we impose restrictions on process syntax that avoid the possibility of realizing *infinite* sequences of local actions. That is, sequences of local actions are always finite and either terminate in a **nil** or in a remote action. The stochastic semantics of the (atomic) sequence $\alpha_R \alpha_{L_1} \cdots \alpha_{L_n}$ of actions is obtained by appropriate composition of the ACT-OR stochastic semantics of α_R with the *probabilistic* semantics of $\alpha_{L_1} \cdots \alpha_{L_n}$, which we shall define in the rest of this section. For this purpose, we will use the notion of *convolution*.

Convolution. Given a relation $\circ \subseteq X \times \mathbf{FTF}(X, \mathbb{R}_{\geq 0})$, we define the convolution $\circ^{cv} \subseteq \mathbf{FTF}(X, \mathbb{R}_{\geq 0}) \times \mathbf{FTF}(X, \mathbb{R}_{\geq 0})$ as follows:

$$\delta_1 \circ^{cv} \delta_2 \quad \text{iff} \quad \forall x \in X \exists \delta_x \in \mathbf{SubDist}(X) \left(x \circ \delta_x \wedge \delta_2 = \sum_{y \in X} \delta_x \cdot \delta_1(y) \right)$$

Example 3. Consider a sub-probability distribution $\delta_1 = [a \mapsto 0.2, b \mapsto 0.3, c \mapsto 0.5]$ and a relation $\circ = \{(a, [a \mapsto 0.4, b \mapsto 0.6]), (b, [b \mapsto 0.5, c \mapsto 0.5]), (b, [a \mapsto 0.3, b \mapsto 0.7]), (c, [a \mapsto 0.8, c \mapsto 0.2])\}$ on a set $X = \{a, b, c\}$. Then, $\delta_1 \circ^{cv} \delta_2$ for $\delta_2 = [a \mapsto 0.48, b \mapsto 0.27, c \mapsto 0.25]$ and also $\delta_1 \circ^{cv} \delta_3$ for $\delta_3 = [a \mapsto 0.57, b \mapsto 0.33, c \mapsto 0.1]$.

Before describing the details of the semantics, let us now introduce informally a bit of notation concerning transition relations. We distinguish between (i) *remote-enabled* process states, that have *only* remote ($@p$) actions activated, *if any*, and (ii) *local-enabled* process states that consist of a non-empty sequence of local actions ($@self$) followed by a remote-enabled process state. At the process level, we change the ACT-OR transition relation \rightarrow_e as follows. First, we limit the applicability of the process relation \rightarrow_e to labels pertaining to *remote* actions, so that it applies only to remote-enabled process states. Then, we define an additional transition relation $\xrightarrow{_}$ for executing local actions ($@self$). At the component level, the \rightarrow relation is modified as follows: (1) \rightarrow is restricted to remote actions, (2) input remote actions remain as in ACT-OR while output remote actions semantics is changed, and (3) local actions are *not* allowed by \rightarrow . In order to execute local actions, we introduce a transition relation $\xrightarrow{_}$, which essentially lifts the relation $\xrightarrow{_}$ to the level of components. For realizing the maximal execution of the $\alpha_{L_1} \cdots \alpha_{L_n}$ action sequence, we introduce the relation $\xrightarrow{\text{SL}}$ which concatenates local transitions, executed by using $\xrightarrow{_}$, until no local action remains to be executed. In this concatenation, the convolution operator \cdot^{cv} is used to lift $\xrightarrow{\text{SL}}$ from a component-to-rate-function relation to a rate-function-to-rate-function relation $\xrightarrow{\text{SL}}^{cv}$. We discuss this aspect in more details in the rest of this section. The final step to define the new semantics of the \rightarrow relation for components is that each sequence $\alpha_R \alpha_{L_1} \cdots \alpha_{L_n}$ is executed by appropriate composition of the stochastic transition relation \rightarrow_e , for the execution of the remote action α_R , and the probabilistic transition relation $\xrightarrow{\text{SL}}$, for the execution of the sequence $\alpha_{L_1} \cdots \alpha_{L_n}$ of local actions. Note that if any of the α_{L_i} cannot be executed, then the entire sequence $\alpha_R \alpha_{L_1} \cdots \alpha_{L_n}$ is not executed.

To complete this informal introduction to the ideas underlying the INT-OR semantics we mention that the new \rightarrow relation for components is designed to be applied only to components with a remote-enabled process state. The effect of its application is to obtain a new

component with a remote-enabled process state (so it is remote-enabled-preserving). A sequence $\alpha_R \alpha_{L_1} \cdots \alpha_{L_n}$ is, therefore, executed entirely in an atomic way, and the label of \rightarrow will be just the remote action α_R , thus hiding the sequence of local actions performed.

In the rest of this section we formalize these ideas. For the sake of conciseness, we show only the modifications to the Action-oriented semantics that are necessary to implement the Interaction-oriented semantics. In particular, in Table 6 we omit the rules that are unchanged w.r.t Table 3 and we just mention their labels. Tables 2 and 4 are unchanged so we omit them.

5.1 Operational semantics of processes

We refine the syntactic category of processes provided in Table 1 by imposing constraints on process syntax. In particular, the syntactic category R below defines the remote-enabled process states, that have only remote ($@p$) actions activated, if any, and include **nil**. The syntactic category L below defines the local-enabled process states, that are made of a non-empty sequence of local ($@self$) actions followed by a remote-enabled process state. We also consider the category of *mixed* processes (M), constructed by parallel composition of local- and remote-enabled processes. This is a syntactic category that will appear in computations as an intermediate product of the semantics. We do not allow to directly specify components behavior using mixed processes. In particular we require that, for the INT-OR semantics, components have their processes specified according to the *remote* syntactic category only.

In this section we put special attention on the use of names for variables ranging on processes. In particular, R, R_1, \dots, R', \dots are variables ranging over remote-enabled processes, L, L_1, \dots, L', \dots are variables ranging over local-enabled processes, and M, M_1, \dots, M', \dots are variables ranging over mixed processes.

$$\begin{aligned} R &::= \mathbf{nil} \mid \mathbf{act}(T)@p.L \mid \mathbf{act}(T)@p.R \mid R + R \mid R \mid R \mid X \mid A(\bar{p}) \\ L &::= \mathbf{act}(T)@self.L \mid \mathbf{act}(T)@self.R \\ M &::= L \mid R \mid M \mid M \end{aligned}$$

Furthermore, process definitions must use remote-enabled processes on their right-hand sides.

The INT-OR semantics of STOCS processes is the RTS $(Proc, Act_{Proc}, \mathbb{R}_{\geq 0}, \rightarrow_e \cup \overline{\top})$, where $Proc$ is the set of process terms defined according to the syntax of STOCS given in Table 1 with the above syntactic restrictions. The set Act_{Proc} of process actions labels is partitioned into the two sets Act_{Proc}^R and Act_{Proc}^L of remote and local process action labels, respectively. Those two sets are defined according to the grammar below (where $t \in \mathbb{I}$, $T \in \mathbb{T}$, c is a TARGET, p is a PREDICATE, and e is the evaluation of an interface):

$$\begin{aligned} Act_{Proc}^R &::= \overline{e : \mathbf{put}(t)@p} \mid \overline{e : \mathbf{get}(T : t)@p} \mid \overline{e : \mathbf{qry}(T : t)@p} \\ Act_{Proc}^L &::= \mathbf{put}(t)@self \mid \mathbf{get}(T)@self \mid \mathbf{qry}(T)@self \mid \checkmark \end{aligned}$$

The transition relation $\rightarrow_e \subseteq Proc \times Act_{Proc}^R \times \mathbf{FTF}(Proc, \mathbb{R}_{\geq 0})$ is the least relation satisfying the rules of Table 2 (note that here *we restrict the action labels* to disallow actions $@self$ on \rightarrow_e). It uses the auxiliary relation $\overline{\top} \subseteq Proc \times Act_{Proc}^L \times \mathbf{FTF}(Proc, \mathbb{R}_{\geq 0})$, which is the the least relation satisfying the rules of Table 5 discussed in the following:

(TPUT/TGQ/TCHO/TPARL) execute the \checkmark action that has the role of allowing local-enabled processes to progress in parallel with remote-enabled and mixed ones, within the (TPAR) rule;

(PUTL/GQL) allow local-enabled processes to execute local actions with rate 1;

Tick (\surd is the tick label, $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$, and \mathbf{p} is a PREDICATE):

$$\frac{}{\mathbf{put}(t)@p . P \xrightarrow[\surd]{L} [\mathbf{put}(t)@p . P \mapsto 1]} \text{ (TPUT)}} \quad \frac{\alpha \neq \surd}{\mathbf{put}(t)@p . P \xrightarrow[\surd]{L} []} \text{ (TPUT}_B\text{)}$$

$$\frac{}{\mathbf{gq}(T)@p . P \xrightarrow[\surd]{L} [\mathbf{gq}(T)@p . P \mapsto 1]} \text{ (TGQ)}} \quad \frac{\alpha \neq \surd}{\mathbf{gq}(T)@p . P \xrightarrow[\surd]{L} []} \text{ (TGQB)}$$

$$\frac{}{R_1 + R_2 \xrightarrow[\surd]{L} [R_1 + R_2 \mapsto 1]} \text{ (TCHO)}} \quad \frac{\alpha \neq \surd}{R_1 + R_2 \xrightarrow[\surd]{L} []} \text{ (TCHO}_B\text{)}$$

$$\frac{R_1 \xrightarrow[\surd]{L} \mathcal{R}_1 \quad R_2 \xrightarrow[\surd]{L} \mathcal{R}_2}{R_1 \mid R_2 \xrightarrow[\surd]{L} \mathcal{R}_1 \mid \mathcal{R}_2} \text{ (TPARL)}$$

Actions (where, $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$, c is a TARGET, and \mathbf{p} is a PREDICATE):

$$\frac{}{\mathbf{put}(t)@self . P \xrightarrow[\surd]{L} [\mathbf{put}(t)@self . P \mapsto 1]} \text{ (PUTL)}} \quad \frac{\alpha \neq \mathbf{put}(t)@self}{\mathbf{put}(t)@self . P \xrightarrow[\surd]{L} []} \text{ (PUTLB)}$$

$$\frac{}{\mathbf{gq}(T)@self . P \xrightarrow[\surd]{L} [\mathbf{gq}(T:t)@self . P \mapsto 1]} \text{ (GQL)}} \quad \frac{\alpha \neq \mathbf{gq}(T:t)@self}{\mathbf{gq}(T)@self . P \xrightarrow[\surd]{L} []} \text{ (GQL}_B\text{)}$$

Choice, definition, and parallel composition:

$$\frac{M_1 \xrightarrow[\surd]{L} \mathcal{M}_1 \quad M_1 \xrightarrow[\surd]{L} \mathcal{M}_1^\surd \quad M_2 \xrightarrow[\surd]{L} \mathcal{M}_2 \quad M_2 \xrightarrow[\surd]{L} \mathcal{M}_2^\surd \quad \alpha \neq \surd}{M_1 \mid M_2 \xrightarrow[\surd]{L} \mathcal{M}_1 \mid \mathcal{M}_2^\surd + \mathcal{M}_1^\surd \mid \mathcal{M}_2} \text{ (TPAR)}$$

Table 5: Operational semantics of STOCS processes (INT-OR), $\xrightarrow[\surd]{L}$ relation only.

(TPAR) applies to mixed processes and is used to postpone the execution of remote-enabled processes by allowing only the execution of local non- \surd actions, indeed \surd actions are used to keep non-local processes unchanged while local ones progress (recall that $\xrightarrow[\surd]{L}$ is defined on Act_{Proc}^L actions).

5.2 Operational semantics of components and systems

The INT-OR semantics of STOCS systems is the RTS $(Sys, Act_{Sys}, \mathbb{R}_{\geq 0}, \rightarrow)$ where Sys is the set of system terms defined according to the syntax of STOCS given in Table 1. The set Act_{Sys} of labels is partitioned into the two sets Act_{Sys}^L and Act_{Sys}^R , defined according to the grammar below (where $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$, e' is the evaluation of an interface, $t \in \mathbb{I}$, $T \in \mathbb{T}$, and \mathbf{p} is a PREDICATE) and it is ranged over by α, α', \dots :

$$\begin{aligned} Act_{Sys}^R ::= & e' : \mathbf{put}(t)@p \quad | \quad e' : \mathbf{gq}(T:t)@p \quad | \quad & \text{(input actions)} \\ & \overline{e' : \mathbf{put}(t)@p} \quad | \quad \overline{e' : \mathbf{gq}(T:t)@p} \quad | \quad & \text{(output actions)} \\ & \overleftarrow{e' : \mathbf{gq}(T:t)@p} & \text{(synchronizations)} \end{aligned}$$

put actions (R is a remote-enabled process):

$$\frac{\sigma = I(K) \quad R \xrightarrow{\overline{\text{put}(t)\text{@p}}}_{\sigma} \mathcal{R} \quad I[K, \mathcal{R}] \xrightarrow{\text{sl}} \mathcal{C}}{I[K, R] \xrightarrow{\overline{\sigma : \text{put}(t)\text{@p}}}_{\sigma} \mathcal{C}} \quad (\text{C-PUTO})$$

(C-PUTI, Table 3) (C-PUTIR, Table 3)

get/qry actions (R is a remote-enabled process):

$$\frac{\sigma = I(K) \quad R \xrightarrow{\overline{\delta : \text{gq}(T:t)\text{@p}}}_{\sigma} \mathcal{R} \quad I[K, \mathcal{R}] \xrightarrow{\text{sl}} \mathcal{C}}{I[K, R] \xrightarrow{\overline{\delta : \text{gq}(T:t)\text{@p}}}_{\sigma} \mathcal{C}} \quad (\text{C-GQO})$$

(C-GETI, Table 3) (C-GETIB, Table 3) (C-QRYI, Table 3) (C-QRYIB, Table 3)

Table 6: Operational semantics of STOCS components (INT-OR).

put local actions (M is a mixed process):

$$\frac{M \xrightarrow{\text{put}(t)\text{@self}}_{\text{L}} \mathcal{M} \quad K \oplus t = \pi}{I[K, M] \xrightarrow{\text{put}(t)\text{@self}}_{\text{L}} I[\pi, \mathcal{M}]} \quad (\text{C-PUTL})$$

get/qry local actions (M is a mixed process):

$$\frac{M \xrightarrow{\text{get}(T)\text{@self}}_{\text{L}} \mathcal{M} \quad K \ominus T = \pi}{I[K, M] \xrightarrow{\text{get}(T)\text{@self}}_{\text{L}} \sum_{t \in \{t \mid \text{match}(T,t)=\vartheta\}} I[\pi(t), \mathcal{M}\vartheta]} \quad (\text{C-GETL})$$

$$\frac{M \xrightarrow{\text{qry}(T)\text{@self}}_{\text{L}} \mathcal{M} \quad K \vdash T = \pi}{I[K, M] \xrightarrow{\text{qry}(T)\text{@self}}_{\text{L}} \sum_{t \in \{t \mid \text{match}(T,t)=\vartheta\}} I[(\mathcal{X}K) \cdot \pi(t), \mathcal{M}\vartheta]} \quad (\text{C-QRYL})$$

$$\frac{K \ominus T = \perp}{I[K, M] \xrightarrow{\text{get}(T)\text{@self}}_{\text{L}} []} \quad (\text{C-GETLB}) \quad \frac{K \vdash T = \perp}{I[K, M] \xrightarrow{\text{qry}(T)\text{@self}}_{\text{L}} []} \quad (\text{C-QRYLB})$$

Maximal execution of local actions:

$$\frac{}{I[K, R] \xrightarrow{\text{sl}} [I[K, R] \vdash 1]} \quad (\text{C-RMT}) \quad R \text{ is a remote-enabled process}$$

$$\frac{I[K, M] \xrightarrow{\alpha}_{\text{L}} \mathcal{C} \quad \mathcal{C} \xrightarrow{\text{sl}}^{cv} \mathcal{D}}{I[K, M] \xrightarrow{\text{sl}} \mathcal{D}} \quad (\text{C-MXT}) \quad \begin{array}{l} M \text{ is a mixed process, } \alpha \in \text{Act}_{Sys}^L, \\ \text{and} \\ \xrightarrow{\text{sl}}^{cv} \text{ is the convolution of } \xrightarrow{\text{sl}} \end{array}$$

Table 7: Operational semantics of STOCS components (INT-OR), $\xrightarrow{\text{L}}$ and $\xrightarrow{\text{sl}}$ transition relations.

$$\text{Act}_{Sys}^L ::= \text{put}(t)\text{@self} \mid \text{get}(T)\text{@self} \mid \text{qry}(T)\text{@self}$$

The transition relation $\rightarrow_{\subseteq} Sys \times Act_{Sys}^R \times \mathbf{FTF}(Sys, \mathbb{R}_{\geq 0})$ for components is the least relation satisfying the rules of Tables 6 and 4. Note that it does not allow to execute local actions. Indeed, these are executed during the execution of a remote action in an atomic way, according to the idea discussed at the beginning of this section. The process transition relation $\rightarrow_{e\subseteq} Proc \times Act_{Proc}^R \times \mathbf{FTF}(Proc, \mathbb{R}_{\geq 0})$ is defined in Table 2 and, again, only allows the execution of remote actions. The transition relation $\xrightarrow{\subseteq}_{SL} Comp \times \text{Dist}(Comp)$ is defined only on the components syntactic category $Comp$ and on the set of local actions Act_{Sys}^L by rules (C-RMT/C-MXT) in Table 7 and it allows the maximal execution of local actions within a component with mixed process state. Let us now discuss the rules in Table 6.

- (C-PUTO) applies to a remote-enabled process (R) and realizes maximal execution of a remote **put** operation by: (i) applying a stochastic **put** transition by using the \rightarrow relation, (ii) applying local transitions as much as possible by using the $\xrightarrow{\subseteq}_{SL}$ relation, and (iii) sending the interface evaluation σ of the source component (itself) in the synchronization label, all this within the atomic component transition $\xrightarrow{\sigma : \mathbf{put}(t) @ \mathbf{p}}$, note that for maximal execution of local actions we consider the entire component, rather than only the behavior process, because we need to take into account also the knowledge state;
- (C-GQO) applies to a remote-enabled process (R) and realizes maximal execution of a remote **get/qry** action by: (i) applying a stochastic **get/qry** transition by using the \rightarrow relation, (ii) applying local transitions as much as possible by using the $\xrightarrow{\subseteq}_{SL}$ relation, and (iii) receiving the interface evaluation δ of the destination component in the synchronization label, all this within an atomic component transition $\xrightarrow{\delta : \mathbf{gq}(T:t) @ \mathbf{p}}$, note that similarly to the previous rule for maximal execution of local actions we consider the entire component, for the same reasons.

Now we discuss the transition relation $\xrightarrow{\subseteq}_{SL} Comp \times \text{Dist}(Comp)$, defined in Table 7 by rules (C-RMT/C-MXT) in a recursive way. This relation realizes maximal executions of the local actions by iterating the application of the component transition relation $\xrightarrow{\subseteq}_{\mathcal{L}} Comp \times Act_{Sys}^L \times \text{Dist}(Comp)$ for local actions. This relation is defined only on the components syntactic category $Comp$ and on the set of local actions Act_{Sys}^L . The relation $\xrightarrow{\subseteq}_{\mathcal{L}}$ is defined in Table 7 by rules (C-PUTL/C-GETL/C-QRYL). Let us now discuss the rules.

- (C-RMT) is the base case of the definition of $\xrightarrow{\subseteq}_{SL}$, it applies to a *remote-enabled* process (R), and, since there is no *local* action to perform, it returns the same component state with rate 1;
- (C-MXT) is the inductive case of the definition, it applies to a mixed process (M) and it executes local actions as much as possible by: (1) applying one local action ($\xrightarrow{\subseteq}_{\mathcal{L}}$), thus obtaining a continuation \mathcal{C} which is a probability distribution over possible components, and (2) applying the convolution $\xrightarrow{\subseteq}_{SL}^{cv} \subseteq \text{Dist}(Comp) \times \text{Dist}(Comp)$ of the relation $\xrightarrow{\subseteq}_{SL} \subseteq Comp \times \text{Dist}(Comp)$ to the continuation \mathcal{C} ;
- (C-PUTL) applies to mixed processes (M) and allows to execute a local **put** action, by using the \oplus operator to add the item t to the local knowledge state and obtaining a probability distribution π on possible new knowledge states which is combined with the continuation \mathcal{M} of the process M , which is simply a Dirac distribution (see Table 5);
- (C-GETL)/(C-QRYL) apply to mixed processes (M) and realize the local **get/qry** actions, by using the \vdash operator to remove/deduce an item t matching the template T from the local knowledge state and combining the resulting probability distribution with the continuation \mathcal{M} of the process M , which is again a Dirac distribution.

The rules for systems are the same as those of the ACT-OR semantics and can be found in Table 4.

6 Network-oriented Operational Semantics

For the sake of conciseness, in this section we illustrate only the modifications to the Action-oriented semantics necessary to model the Network-oriented semantics. In particular, in Tables 8, 9, and 10 we omit the rules that are unchanged w.r.t Tables 2, 3, and 4 and we just leave their label, as usual.

In order to realize this semantics we extend the set of labels of actions performed by processes and systems as described in the following.

6.1 Explanatory example

In the following we illustrate this variant of the semantics through a couple of examples and we point out issues that lead us to necessary simplifications and further design choices.

Let us consider a process P , of the form $\mathbf{put}(v)@p.Q$, and the execution of action $\mathbf{put}(v)@p$, as illustrated in Figure 1. The execution begins with the creation of one special message *for each* component of the system. These messages play the role of *observers*: each of them travels in the system and reaches the component it is associated with. The special message creation and message-component association phase has rate λ : this value is computed as a function of several factors, among which (the size of) t . After message creation, P can proceed—since \mathbf{put} actions are non-blocking—behaving like Q . Each special message has to reach its destination component, which checks if its interface satisfies p , and if so, it delivers t in the knowledge repository of the component. Observer delivery is performed with rate μ , which may depend on t and other parameters like the distance between the component where P resides and the target component. Therefore, a *distinct* rate μ is associated to each target

In practice, one can be interested in modeling also the event of failed delivery of the observers. This is interesting both for producing more realistic models (with unreliable network communication), and for allowing the application of advanced analysis techniques based on fluid approximation [3], such as fluid model-checking [2]. Therefore, we add an error probability to the observers delivery, which we indicate as p_{err} (or simply err , in the figure). This more detailed semantics of the $\mathbf{put}(v)@p$ action is illustrated in Fig. 3 and described below in more detail.

We consider three components: $C_1 = I_1 [K_1, P_1]$, $C_2 = I_2 [K_2, P_2]$, and $C_3 = I_3 [K_3, P_3]$ and we assume process P_1 is defined as $\mathbf{put}(v)@p.Q$ as described above. Note that different components may be in different locations. The interaction we illustrate starts with process P_1 executing the first phase of $\mathbf{put}(v)@p$, i.e. creating two² copies of the special message $\{v@p\}$, one for component C_2 and one for component C_3 , and sending these messages. The time required for this phase (denoted in blue in the figure) is modeled by a RV with rate λ . The time for the two messages to arrive at components C_2 and C_3 is modeled by RVs with rates μ_2 and μ_3 (in the figure this is illustrated by two arrows). The delivery of the two messages fails with probability err_2 and err_3 , respectively, and succeeds with their complement. The execution of component C_1 restarts as soon as the copies of the messages are sent, without waiting for their arrival at the destination components (the red stripe in the figure illustrates the resumed execution of C_1). The evaluation of predicate p is performed when the message arrives at the corresponding component so, for example, it may happen that C_2 satisfies p

²For the sake of notational simplicity, in the present report we assume that predicate p in process actions implicitly refers only to the *other* components, excluding the one where the process is in execution.

at the time the message arrives (so K_2 is updated accordingly), while C_3 does not satisfy \mathbf{p} (thus leaving K_3 unchanged).





Let us now consider the execution of a process P of the form $\mathbf{get}(T)@p.Q$ (or, equivalently, $\mathbf{qry}(T)@p.Q$), as illustrated in Fig. 2. Similarly to the execution of $\mathbf{put}(v)@p$, the first operation performed consists in the creation of a special request message, which is delivered to every component. Since the \mathbf{get} (resp. \mathbf{qry}) action is blocking, the process moves into a waiting state after creation of this message. Each copy of the message is delivered to the corresponding component. Upon message arrival, every component satisfying predicate \mathbf{p} and having in its knowledge an item t matching the template T , is eligible to answer the request with item t . Message preparation, message delivery, as well as response delivery are events that take time (e.g. depending on the size of the item t or the distance of the involved component) and have rates λ , β_j , and μ_j , respectively. Concerning response delivery, there is a race condition and *only one* component succeeds in providing the item t . Once the item has been delivered, process P can restart its execution from $Q\vartheta$, with a suitable variable binding $\mathbf{match}(T, t) = \vartheta$.

In order to simplify the semantics of the $\mathbf{get}/\mathbf{qry}$ actions and to make it more similar to the two-steps semantics of the \mathbf{put} action, we decided to model the two phases of delivery and response as a single one. So, on message creation, the sender is blocked on waiting for some receiver to synchronize with it on the exchange of the retrieved item t matching the template T , as illustrated in Fig. 4. During this synchronization, the predicate \mathbf{p} is also checked, on the side of the receiver, and the knowledge is changed accordingly. In terms of the underlying stochastic model, we are replacing an Erlang RV by an Exponential RV. This choice is also convenient for simplifying the definition of the formal semantics, since it avoids the need of giving a unique id to observer messages, to be used in the subsequent response phase.

In Fig. 4, we illustrate an example of execution of a $\mathbf{get}/\mathbf{qry}$ action. Here we consider four components: $C_1 = I_1[K_1, P_1]$, $C_2 = I_2[K_2, P_2]$, $C_3 = I_3[K_3, P_3]$, and $C_4 = I_4[K_4, P_4]$ and we assume process P_1 is defined as $\mathbf{get}(T)@p.Q$. The interaction we illustrate starts with process P_1 creating the observer messages with rate λ (denoted in blue in the figure). The waiting state of C_1 , after the creation of the messages, is denoted by a dashed line. The messages are delivered to C_2 , C_3 , and C_4 . Upon arrival, components check for satisfaction of predicate \mathbf{p} and availability of an item matching T . Let us assume that C_2 and C_3 satisfy \mathbf{p} and have matching knowledge items t_2 and t_3 , respectively, whereas C_4 either does not satisfy \mathbf{p} or has no matching knowledge item. Therefore, C_2 and C_3 are the only components that can synchronize with C_1 : they perform this synchronization with rates μ_2 and μ_3 , respectively (depending, for example, on the size of t_2 and t_3 , and on their locations). In the figure, this is represented by arrows from C_2 and C_3 to C_1 . In the Markovian semantics, this is realized by a race condition in which just one of the two synchronizations is actually executed. In the figure, it is assumed that component C_2 wins the race, and synchronizes on an item t_2 . Then, the execution of C_1 restarts (the red stripe in the figure illustrates the resumed execution of C_1). The synchronization attempt of C_3 is simply lost. Here we need to be careful in how to handle these message deliveries. If we are considering a \mathbf{get} action, the retrieving of a knowledge item *from* a component requires the item to be *removed* from the knowledge. In the case illustrated in the figure, then, the knowledge K_2 needs to be updated removing t_2 , while the knowledge K_3 must remain unchanged.

6.2 Operational semantics of processes

The NET-OR semantics of STOCS processes is the RTS $(Proc, Act_{Proc}, \mathbb{R}_{\geq 0}, \neg_e)$, where $Proc$ is the set of process terms defined according to the syntax of STOCS given in Table 1 and the set Act_{Proc} of labels is defined according to the grammar below (where $t \in \mathbb{I}$, $T \in \mathbb{T}$, $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$, c is a TARGET, and e is the evaluation of an interface) and it is ranged over

Figure 1: Dynamics of the **put** action.Figure 2: Dynamics of the **get/qry** action.Figure 3: Actual model of **put**.Figure 4: Actual model of **get/qry**.

by α, α', \dots :

$$Act_{Proc} ::= \tau \mid \overline{\{t@p\}} \mid \overline{e : \mathbf{put}(t)@c} \mid \overline{e : \mathbf{gq}(T:t)@c}$$

and $\rightarrow \subseteq Proc \times Act_{Proc} \times \mathbf{FTF}(Proc, \mathbb{R}_{\geq 0})$ is the least relation satisfying the rules of Table 8 (like in the ACT-OR semantics, \rightarrow_e is parametrized by e , which is the interface evaluation of the component in which the process resides: we feel free to omit the parameter, if not used in the rule).

We now briefly illustrate the rules of Table 8. We assume to have additional syntactical terms (not available at the user syntax level) which we call *envelopes*. They are of the form $\{t@p\}_\mu$, can be put in parallel with processes, and denote messages that are currently traveling towards targets. A second syntactical construct we introduce is $\{\mathbf{get}(T)@p\}$ ($\{\mathbf{qry}(T)@p\}$, respectively) which denotes a *waiting* state of the process and it is treated as an action.

(ENV) allows to complete envelope delivery within a duration parametrized by μ ;

(PUT) allows a process to issue a **put** action at a target c with rate 1;

(GQL) allows a process to issue a **get** (**qry**, respectively) action over the local knowledge repository (i.e. with target **self**). The rule models the execution of action $\mathbf{get}(T)@self$ ($\mathbf{qry}(T)@self$, respectively) by process $\mathbf{get}(T)@self.P$ ($\mathbf{qry}(T)@self.P$, respectively). The duration of this action is described by a rate λ computed using the function \mathcal{R}

Inactive process and envelopes:

$$\text{(NIL, Table 2)} \quad \frac{}{\{t@p\}_\mu \xrightarrow{\{t@p\}} [\mathbf{nil} \mapsto \mu]} \text{(ENV)} \quad \frac{\alpha \neq \overline{\{t@p\}}}{\{t@p\}_\mu \xrightarrow{\alpha} []} \text{(ENVB)}$$

Actions (where, $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$, c is a TARGET, and \mathbf{p} is a PREDICATE):

$$\text{(PUT, Table 2)} \quad \text{(PUT}_B\text{, Table 2)}$$

$$\frac{\text{match}(T, t) = \vartheta \quad \lambda = \mathcal{R}(\sigma, \mathbf{gq}(T : t)@self, -)}{\mathbf{gq}(T)@self.P \xrightarrow{\sigma : \mathbf{gq}(T:t)@self} /_\sigma [P\vartheta \mapsto \lambda]} \text{(GQL)}$$

$$\frac{\neg \text{match}(T, t)}{\mathbf{gq}(T)@self.P \xrightarrow{- : \mathbf{gq}(T:t)@self} []} \text{(GQLB}_1\text{)} \quad \frac{\alpha \neq - : \overline{\mathbf{gq}(T : t)@self}}{\mathbf{gq}(T)@self.P \xrightarrow{\alpha} []} \text{(GQLB}_2\text{)}$$

$$\frac{\lambda = \mathcal{R}(\sigma, \mathbf{gq}(T : -)@p, -)}{\mathbf{gq}(T)@p.P \xrightarrow{\tau /_\sigma} [\{\mathbf{gq}(T)@p\}.P \mapsto \lambda]} \text{(GQW)} \quad \frac{\alpha \neq \tau}{\mathbf{gq}(T)@p.P \xrightarrow{\alpha} []} \text{(GQWB)}$$

$$\frac{\text{match}(T, t) = \vartheta \quad \beta = \mathcal{R}(\sigma, \{\mathbf{gq}(T : t)@p\}, \delta)}{\{\mathbf{gq}(T)@p\}.P \xrightarrow{\delta : \{\mathbf{gq}(T:t)@p\}} /_\sigma [P\vartheta \mapsto \beta]} \text{(GQD)}$$

$$\frac{\neg \text{match}(T, t)}{\{\mathbf{gq}(T)@p\}.P \xrightarrow{- : \{\mathbf{gq}(T:t)@p\}} []} \text{(GQDB}_1\text{)} \quad \frac{\alpha \neq - : \{\overline{\mathbf{gq}(T : t)@p\}}}{\{\mathbf{gq}(T)@p\}.P \xrightarrow{\alpha} []} \text{(GQDB}_2\text{)}$$

Choice, definition, and parallel composition:

$$\text{(CHO, Table 2)} \quad \text{(DEF, Table 2)} \quad \text{(PAR, Table 2)}$$

Table 8: Operational semantics of STOCS processes (NET-OR).

depending on the interface evaluation of the source σ (i.e. the container component) and on the action;

(GQW) realizes the first step of a **get** (**qry**, respectively) action over a remote knowledge in a component satisfying a predicate \mathbf{p} , which consists in preparing an envelope $\{\mathbf{get}(T)@p\}$ ($\{\mathbf{qry}(T)@p\}$, respectively), which takes a time interval exponentially distributed with rate λ , and brings process P to a *wait* state $\{\mathbf{get}(T)@p\}.P$ ($\{\mathbf{qry}(T)@p\}.P$, respectively). Recall that **get**/**qry** actions are *blocking* and the execution of P is resumed only when a counterpart satisfying \mathbf{p} has a knowledge item t matching T available and the delivery of t is completed. The duration of this first step is described by a rate λ computed using the function \mathcal{R} depending only on the interface evaluation of the source σ (i.e. the container component) and the sent template T ;

(GQD) realizes the second step of a **get** (**qry**, respectively) action, which consists in the *delivery* of the knowledge item t matching T and has a duration described by a rate β computed by the function \mathcal{R} . Note that in this case the function \mathcal{R} is computed considering interface evaluation of the source σ and the destination δ , as well as the sent item t , which means that this rate can be made dependent (for example) on the distance of the two parties.

| | |
|---|--|
| put actions: | (C-PUTL, Table 3) (C-PUTO, Table 3) |
| $\frac{\delta = I(K) \quad \mu = \mathcal{R}(\sigma, \{t@p\}, \delta) \quad p_{err} = f_{err}(\sigma, \{t@p\}, \delta)}{I[K, P] \xrightarrow{\sigma: \mathbf{put}(t)@p} [I[K, P] \mapsto p_{err}, I[K, P] \{t@p\}_\mu] \mapsto (1 - p_{err})} \quad (\text{C-PUTI})$ | |
| $\frac{P \xrightarrow{\{t@p\}} \mathcal{P} \quad I(K) \models p \quad K \oplus t = \pi}{I[K, P] \xrightarrow{\{t@p\}} I[\pi, \mathcal{P}]} \quad (\text{C-ENVA}) \quad \frac{P \xrightarrow{\{t@p\}} \mathcal{P} \quad I(K) \not\models p}{I[K, P] \xrightarrow{\{t@p\}} I[(\mathcal{X}K), \mathcal{P}]} \quad (\text{C-ENVR})$ | |

| | |
|--|--|
| get/qry actions (where, $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$): | |
| (C-GETL, Table 3) (C-QRYL, Table 3) (C-GETLB, Table 3) (C-GETRB, Table 3) | |
| $\frac{\sigma = I(K) \quad P \xrightarrow{\delta: \{\mathbf{gq}(T:t)@p\}} \sigma \mathcal{P}}{I[K, P] \xrightarrow{\delta: \{\mathbf{gq}(T:t)@p\}} I[(\mathcal{X}K), \mathcal{P}]} \quad (\text{C-GQO})$ | |
| $\frac{\delta = I(K) \quad \delta \models p \quad K \ominus T = \pi}{I[K, P] \xrightarrow{\delta: \{\mathbf{get}(T:t)@p\}} I[\pi(t), (\mathcal{X}P)]} \quad (\text{C-GETI})$ | |
| $\frac{\delta \neq I(K) \quad \vee \quad I(K) \not\models p \quad \vee \quad K \ominus T = \perp}{I[K, P] \xrightarrow{-: \{\mathbf{get}(T:t)@p\}} \square} \quad (\text{C-GETIB})$ | |
| $\frac{\delta = I(K) \quad \delta \models p \quad K \vdash T = \pi}{I[K, P] \xrightarrow{\delta: \{\mathbf{qry}(T:t)@p\}} [I[K, P] \mapsto \pi(t)]} \quad (\text{C-QRYI})$ | |
| $\frac{\delta \neq I(K) \quad \vee \quad I(K) \not\models p \quad \vee \quad K \vdash T = \perp}{I[K, P] \xrightarrow{-: \{\mathbf{qry}(T:t)@p\}} \square} \quad (\text{C-QRYIB})$ | |

| | |
|------------|--|
| τ actions: | $\frac{\rho = I(K) \quad P \xrightarrow{\tau} \rho \mathcal{P}}{I[K, P] \xrightarrow{\tau} I[(\mathcal{X}K), \mathcal{P}]} \quad (\text{C-TAU})$ |
|------------|--|

Table 9: Operational semantics of STOCS components (NET-OR).

6.3 Operational semantics of components and systems

The NET-OR semantics of STOCS systems is the RTS $(Sys, Act_{Sys}, \mathbb{R}_{\geq 0}, \rightarrow)$, where Sys is the set of system terms defined according to the syntax of STOCS given in Table 1 and the set Act_{Sys} of labels is defined according to the grammar below (where $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$, $t \in \mathbb{I}$, $T \in \mathbb{T}$, p is a PREDICATE, and e is the evaluation of an interface):

$$\begin{array}{lcl}
 Act_{Sys} & ::= & e : \mathbf{put}(t)@p \quad | \quad e : \{\mathbf{gq}(T:t)@p\} \quad | \quad (\text{input actions}) \\
 & & \overline{e : \mathbf{put}(t)@p} \quad | \quad \overline{e : \{\mathbf{gq}(T:t)@p\}} \quad | \quad (\text{output actions}) \\
 \tau & | & \overleftarrow{e : \{\mathbf{gq}(T:t)@p\}} \quad | \quad (\text{synchronizations}) \\
 & & \overline{\{t@p\}} \quad (\text{envelopes})
 \end{array}$$

and $\rightarrow \subseteq Sys \times Act_{Sys} \times \mathbf{FTF}(Sys, \mathbb{R}_{\geq 0})$ is the least relation satisfying the rules of Tables 9 and 10, where the process relation defined in Table 8 is also used.

| | | |
|-----------------------------|-----------------|-----------------|
| put synchronization: | (S-PO, Table 4) | (S-PI, Table 4) |
|-----------------------------|-----------------|-----------------|

| | | |
|--|--|--|
| get/qry synchronization ($\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$): | | |
| $S_1 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_1^s$ | $S_1 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_1^o$ | $S_1 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_1^i$ |
| $S_2 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_2^s$ | $S_2 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_2^o$ | $S_2 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_2^i$ |
| <hr/> | | |
| $S_1 \parallel S_2 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_1^s \parallel (\mathcal{X} S_2) + \mathcal{S}_1^o \parallel \mathcal{S}_2^i + \mathcal{S}_1^i \parallel \mathcal{S}_2^o + (\mathcal{X} S_1) \parallel \mathcal{S}_2^s$ | | |
| <hr/> | | |
| $\frac{S_1 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_1 \quad S_2 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_2}{S_1 \parallel S_2 \xrightarrow{\delta:\{\mathbf{gq}(T:t)\@p\}} \mathcal{S}_1 \parallel (\mathcal{X} S_2) + (\mathcal{X} S_1) \parallel \mathcal{S}_2}$ (S-GQI) | | |

| | | |
|--|--|--|
| Internal actions, for $\alpha \in \{\tau, e : \mathbf{put}(t)\@self, e : \mathbf{gq}(T : t)\@self, \{t\@p\}\}$: | | |
| <hr/> | | |
| $\frac{S_1 \xrightarrow{\alpha} \mathcal{S}_1 \quad S_2 \xrightarrow{\alpha} \mathcal{S}_2}{S_1 \parallel S_2 \xrightarrow{\alpha} \mathcal{S}_1 \parallel (\mathcal{X} S_2) + (\mathcal{X} S_1) \parallel \mathcal{S}_2}$ (S-SPL) | | |

Table 10: Operational semantics of STOCS systems (NET-OR).

The definition of the semantics of system parallel composition $S_1 \parallel S_2$ uses Def. 3.1, item (3) applied to the system parallel composition constructor \parallel , which is injective. As usual, interleaving is modelled as a combination of lifted \parallel , $+$ on functions and the characteristic function. In the rules, we also use Def. 3.1, item (3) applied to the component syntactic constructors $I[\cdot, \cdot]$, which is injective.

In Table 9, rules are grouped to illustrate how the various action types are realized.

(C-PUTI) models the *initiation* of the execution of action $\mathbf{put}(t)\@c$, which requires several steps to complete, it allows the reception of a **put** action, and it is responsible for the creation of the envelope (carrying the incoming message) in parallel to the local process of a component, thus modeling its travel towards that component in terms of the time necessary to reach it, parametrized by rate μ (the fact that the envelope is in parallel with the process of the potential receiver component by no means should be interpreted as the representation of the fact that the message reached the component; simply, the association between the message and the component is represented by means of a parallel composition term; in other words, the fact that a specific message is ‘addressed’ to a component is represented syntactically by such a parallel composition); this action is executed with rate λ , computed using the function \mathcal{R} depending on the interface evaluation of the source σ (i.e. the container component) and the sent item t ; this is postulated by the rule (PUT) and realized at system level by the broadcast rules of Table 4.

(C-ENVA)/(C-ENVR) realize envelope delivery by specifying the conditions under which a component *accepts* or *refuses*, respectively, an arriving envelope;

(C-GQO) realizes an output **get/qry** action as in the ACT-OR semantics, but with a different label ($\{\dots\}$) which denotes the synchronization on a *waiting state*;

(C-GETI)/(C-QRYI) realize an input **get/qry** action, again as in the ACT-OR semantics, but with a different label ($\{\dots\}$);

(C-TAU) allows a component to make a τ whenever its process makes such an action.

7 Activity-oriented Operational Semantics

In this section we illustrate the Activity-oriented semantics (abbreviated to ACTIV-OR), which is essentially based on the ACT-OR syntax and semantics. The main novelty w.r.t. that semantics is the *replacement* of STOCS actions by a new syntactic constructor $\langle seq \rangle_\lambda$ that allows a user to specify a sequence seq of STOCS actions to be executed *atomically* and with overall exponential rate λ . The execution rate λ is not influenced by the actions performed, but specified by the user (possibly through a function of the state of the system). In the case of the empty sequence of actions ($\langle \rangle_\lambda$) this constructor can be also used to specify delays with rate λ having no effect on local/remote knowledge states.

The motivation for this semantics, as already discussed in Section 2.2, is to have the possibility of specifying, at a high-level of abstraction, interaction mechanisms which we call *activities*, that are modeled atomically. Such an abstraction provides a significant reduction of the state-space of the system, since the execution of actions within sequences is not interleaved with that of other processes in parallel.

Recall that in ACT-OR local and output actions are stochastic, while input actions are probabilistic. They are composed together into a timed synchronization. In ACTIV-OR, in order to execute action sequences atomically and with a given rate λ , we give a *time-abstract* semantics for local and output actions of *components* (defining the transition relation \xrightarrow{TA}), a *stochastic* semantics for sequences of actions of *components* (defining the transition relation \xrightarrow{T}), and we maintain the *probabilistic* semantics for input actions of *components* (implemented by the transition relation \xrightarrow{PR}). To obtain a time-abstract semantics it is enough to redefine output and local actions in a time-abstract way (i.e. probabilistically). As a simple consequence, the synchronization of output actions and input actions will be itself time-abstract. Then, when the execution of the sequence of actions is terminated, we perform a stochastic delay λ and compose it with the probability distribution obtained by time-abstract execution. In particular, an action sequence $\langle \alpha_1 \cdots \alpha_n \rangle_\lambda$ is rewritten, by a sequence initiation action $\xrightarrow{TA} \frac{\text{init}_s}{TA}$ into a new sequence $\alpha_1 \cdots \alpha_n \text{end}_s \text{delay}(\lambda)$. Then, using the transition relation \xrightarrow{TA} , the sequence $\alpha_1 \cdots \alpha_n$ is executed in a time-abstract fashion (taking only the probabilistic aspect of α_i and combining them by using the convolution operator \cdot^{cv}), and terminated by a sequence termination action $\xrightarrow{TA} \frac{\text{end}_s}{TA}$. The remaining $\text{delay}(\lambda)$ is executed using the stochastic transition $\xrightarrow{\text{delay}}$. At the *system* level, then, the stochastic and probabilistic semantics (\xrightarrow{T} and \xrightarrow{PR} , respectively) are composed together by the transition relation \rightarrow .

7.1 Operational semantics of processes

The *stochastic* ACTIV-OR semantics of STOCS *processes* is the RTS $(ExtProc, ActProc, \mathbb{R}_{\geq 0}, \rightarrow)$ where $ExtProc$ is the set of process terms defined by the following syntax, starting from the symbol EP :

$$\begin{aligned} EP & ::= Seq.\text{end}_s.\text{delay}(\lambda).P \mid P \\ P & ::= \mathbf{nil} \mid \langle Seq \rangle_\lambda.P \mid P + P \mid P \mid P \mid X \mid A(\bar{p}) \\ Seq & ::= \varepsilon \mid a.Seq \end{aligned}$$

The set $Proc$ modifies the syntax of STOCS processes given in Table 1 by *replacing* actions with sequences $\langle \dots \rangle_\lambda$ of actions (where ε denotes the empty sequence of actions and $\langle \varepsilon \rangle \equiv \langle \rangle$). The set $ExtProc$ is introduced only for internal purposes of the time-abstract semantics and it extends the set $Proc$, defined by the above grammar starting from the symbol P , by allowing the prefix $Seq.\text{end}_s.\text{delay}(\lambda)$ of Processes, made of an (un-parenthesized) time-abstract action

Time-Abstract actions (where, $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$ and c is a TARGET):

$$\frac{}{\langle \rangle_\lambda. P \xrightarrow{\langle \rangle} [P \mapsto \lambda]} \text{ (ESEQ)} \quad \frac{}{\text{delay}(\lambda). P \xrightarrow{\text{delay}} [P \mapsto \lambda]} \text{ (DELAY)}$$

$$\frac{\alpha \in \text{OneAct}_{Proc} \cup \{\text{init_s}, \text{end_s}\}}{P \xrightarrow{\alpha} []} \text{ (TIME-ABSTRACT}_B\text{)}$$

Choice, definition, and parallel composition:

$$\frac{P \xrightarrow{\alpha} \mathcal{P} \quad Q \xrightarrow{\alpha} \mathcal{Q}}{P + Q \xrightarrow{\alpha} \mathcal{P} + \mathcal{Q}} \text{ (CHO)} \quad \frac{A(\vec{x}) \stackrel{def}{=} P \quad P[\vec{v}/\vec{x}] \xrightarrow{\alpha} \mathcal{P}}{A(\vec{v}) \xrightarrow{\alpha} \mathcal{P}} \text{ (DEF)}$$

$$\frac{P \xrightarrow{\alpha} \mathcal{P} \quad Q \xrightarrow{\alpha} \mathcal{Q}}{P \mid Q \xrightarrow{\alpha} \mathcal{P} \mid (\mathcal{X} Q) + (\mathcal{X} P) \mid \mathcal{Q}} \text{ (PAR)}$$

Table 11: Operational sem. of STOCS processes (ACTIV-OR), \rightarrow transition relation.

Time-Abstract actions (where, $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$ and c is a TARGET):

$$\frac{}{\mathbf{put}(t)@c. P \xrightarrow{\mathbf{put}(t)@c}_{TA} [P \mapsto 1]} \text{ (PUT-TA)} \quad \frac{\alpha \neq \overline{\mathbf{put}(t)@c}}{\mathbf{put}(t)@c. P \xrightarrow{\alpha}_{TA} []} \text{ (PUT-TAB)}$$

$$\frac{}{\mathbf{gq}(T)@c. P \xrightarrow{\mathbf{gq}(T)@c}_{TA} [P \mapsto 1]} \text{ (GQ-TA)} \quad \frac{\alpha \neq \mathbf{gq}(T)@c}{\mathbf{gq}(T)@c. P \xrightarrow{\alpha}_{TA} []} \text{ (GQ-TAB)}$$

$$\frac{\text{seq} = a_1 \cdots a_n \quad \text{for } n \geq 1}{\langle \text{seq} \rangle_\lambda. P \xrightarrow{\text{init_s}}_{TA} [\text{seq.end_s.delay}(\lambda). P \mapsto 1]} \text{ (INIT-TA)} \quad \frac{}{\text{end_s}. P \xrightarrow{\text{end_s}}_{TA} [P \mapsto 1]} \text{ (END-TA)}$$

$$\frac{\alpha \in \{\langle \rangle, \text{delay}\}}{\alpha. P \xrightarrow{\alpha}_{TA} []} \text{ (TIMED-TAB)}$$

Choice, definition, and parallel composition:

$$\frac{P \xrightarrow{\alpha}_{TA} \mathcal{P} \quad Q \xrightarrow{\alpha}_{TA} \mathcal{Q}}{P + Q \xrightarrow{\alpha}_{TA} \mathcal{P} + \mathcal{Q}} \text{ (CHO-TA)} \quad \frac{A(\vec{x}) \stackrel{def}{=} P \quad P[\vec{v}/\vec{x}] \xrightarrow{\alpha}_{TA} \mathcal{P}}{A(\vec{v}) \xrightarrow{\alpha}_{TA} \mathcal{P}} \text{ (DEF-TA)}$$

$$\frac{P \xrightarrow{\alpha}_{TA} \mathcal{P} \quad Q \xrightarrow{\alpha}_{TA} \mathcal{Q}}{P \mid Q \xrightarrow{\alpha}_{TA} \mathcal{P} \mid (\mathcal{X} Q) + (\mathcal{X} P) \mid \mathcal{Q}} \text{ (PAR-TA)}$$

Table 12: Operational sem. of STOCS processes (ACTIV-OR), \xrightarrow{TA} transition relation.

sequence, a time-abstract sequence terminator end_s , and a stochastic action $\text{delay}(\lambda)$. The set Act_{Proc} of labels is defined according to the grammar below (where $t \in \mathbb{I}$, $T \in \mathbb{T}$, and c is a TARGET) and it is ranged over by α, α', \dots :

$$\text{Act}_{Proc} ::= \text{init_s} \mid \text{end_s} \mid \text{delay}(\lambda) \mid \langle \rangle \mid \text{OneAct}_{Proc}$$

$$\text{OneAct}_{Proc} ::= \mathbf{put}(t)@c \mid \mathbf{get}(T:t)@c \mid \mathbf{qry}(T:t)@c$$

The transition relation $\rightarrow \subseteq \text{ExtProc} \times \text{Act}_{Proc} \times \mathbf{FTF}(\text{Proc}, \mathbb{R}_{\geq 0})$ is the least relation satis-

| | |
|---|-----------|
| Sequence synchronization (where C_1, \dots, C_n are components and seq a non-empty sequence of actions): | |
| $C_i \xrightarrow{\overline{\langle \rangle}} \mathcal{C}_i \quad C_1 \parallel \dots \parallel C_{i-1} \parallel C_{i+1} \parallel \dots \parallel C_n \xrightarrow{\langle \rangle}_{PR} \mathcal{S}_i \quad \text{for } i = 1, \dots, n$ | (S-DELAY) |
| $C_1 \parallel \dots \parallel C_n \xrightarrow{\langle \rangle} \sum_{i=1}^n \mathcal{C}_i \parallel \mathcal{S}_i$ | |
| $C_i \xrightarrow{\overline{\langle seq \rangle}} \mathcal{C}_i \quad C_1 \parallel \dots \parallel C_{i-1} \parallel C_{i+1} \parallel \dots \parallel C_n \xrightarrow{\langle seq \rangle}_{PR} \mathcal{S}_i \quad \text{for } i = 1, \dots, n$ | (S-SSYNC) |
| $C_1 \parallel \dots \parallel C_n \xrightarrow{\langle seq \rangle} \sum_{i=1}^n \mathcal{C}_i \parallel \mathcal{S}_i$ | |

Table 13: Operational semantics of STOCS systems (ACTIV-OR), \rightarrow relation.

fying the rules of Table 11. Note that \rightarrow always produces processes in the set $Proc$.

The *time-abstract* ACTIV-OR semantics of STOCS *processes* is the RTS $(ExtProc, Act_{Proc}, \mathbb{R}_{\geq 0}, \overline{\text{TA}})$ where the transition relation $\overline{\text{TA}} \subseteq ExtProc \times Act_{Proc} \times \text{Dist}(ExtProc)$, is the least relation satisfying the rules of Table 12

7.2 Operational semantics of components and systems

The ACTIV-OR semantics of STOCS *systems* is the RTS $(Sys, Act_{Sys}, \mathbb{R}_{\geq 0}, \rightarrow)$ where Sys is the set of system terms defined according to the syntax of STOCS given in Table 1, with the exception of processes, that are taken from the set $Proc$ defined in the previous sub-section.

The set Act_{Sys} of labels is defined according to the grammar below (where $\mathbf{gq} \in \{\mathbf{get}, \mathbf{qry}\}$, e is the evaluation of an interface, $t \in \mathbb{I}$, $T \in \mathbb{T}$, and \mathbf{p} is a PREDICATE):

$$Act_{Sys} ::= \overline{\langle Seq_{Sys} \rangle}$$

$$Seq_{Sys} ::= \varepsilon \quad | \quad OneAct_{Proc} . Seq_{Sys}$$

and $\rightarrow \subseteq Sys \times Act_{Sys} \times \mathbf{FTF}(Sys, \mathbb{R}_{\geq 0})$ is the least relation satisfying the rules of Table 13. The \rightarrow transition relation combines the stochastic execution of $\overline{\langle Seq_{Sys} \rangle}$ output/local action sequences, performed by using the $\overline{\text{TA}}$, and the probabilistic execution of $\langle Seq_{Sys} \rangle$ input action sequences.

In particular, the $\overline{\text{TA}} \subseteq Comp \times Act_{Sys}^T \times \mathbf{FTF}(Comp, \mathbb{R}_{\geq 0})$ transition relation is defined by the rules of Table 14, where the set Act_{Sys}^T is defined as follows:

$$Act_{Sys}^T ::= \overline{\langle Seq_{Sys} \rangle}$$

The $\overline{\text{TA}}$ transition relation concatenates the probabilistic sequence initiation action $\frac{\text{init}_s}{\text{TA}} \rightarrow$, the probabilistic maximal execution of the action sequence performed by transition relation $\overline{\text{TA}}$ (terminated by the probabilistic sequence termination action $\frac{\text{end}_s}{\text{TA}} \rightarrow$), and the final delay action $\xrightarrow{\text{delay}}^{cv}$ (performed using the lifting $\rightarrow^{cv} \subseteq \text{Dist}(ExtProc) \times Act_{Proc} \times \mathbf{FTF}(Proc, \mathbb{R}_{\geq 0})$ of \rightarrow to probability distributions over extended processes). The transition relation $\overline{\text{TA}} \subseteq ExtComp \times Act_{Sys}^{\text{TA}} \times \text{Dist}(ExtComp)$ as well as its variant $\overline{\text{TA}}^{cv} \subseteq \text{Dist}(ExtComp) \times Act_{Sys}^{\text{TA}} \times \text{Dist}(ExtComp)$, lifted to probability distributions over components, are defined in Table 14. The set Act_{Sys}^{TA} of actions is defined as follows:

$$Act_{Sys}^{\text{TA}} ::= OneAct_{Proc} \quad | \quad \overline{\langle Seq_{Sys} \rangle}$$

The counterpart of $\overline{\text{TA}}$ is the probabilistic transition relation $\overline{\text{PR}} \subseteq Comp \times Act_{Sys}^{\text{PR}} \times$

put actions:

$$\frac{P \xrightarrow[\text{TA}]{\text{put}(t)\text{@self}} \mathcal{P} \quad K \oplus t = \pi}{I[K, P] \xrightarrow[\text{TA}]{\text{put}(t)\text{@self}} I[\pi, \mathcal{P}]} \quad (\text{C-PUTL-TA}) \qquad \frac{P \xrightarrow[\text{TA}]{\text{put}(t)\text{@p}} \mathcal{P}}{I[K, P] \xrightarrow[\text{TA}]{\text{put}(t)\text{@p}} I[(\mathcal{X}K), \mathcal{P}]} \quad (\text{C-PUTO-TA})$$

get/qry actions (where, $\mathbf{gq} \in \{\text{get}, \text{qry}\}$):

$$\frac{P \xrightarrow[\text{TA}]{\text{get}(T)\text{@self}} \mathcal{P} \quad K \ominus T = \pi}{I[K, P] \xrightarrow[\text{TA}]{\text{get}(T)\text{@self}} \sum_{t \in \{t \in \mathbb{I} \mid \text{match}(T, t) = \vartheta\}} I[\pi(t), \mathcal{P}\vartheta]} \quad (\text{C-GETL-TA})$$

$$\frac{P \xrightarrow[\text{TA}]{\text{qry}(T)\text{@self}} \mathcal{P} \quad K \vdash T = \pi}{I[K, P] \xrightarrow[\text{TA}]{\text{qry}(T)\text{@self}} \sum_{t \in \{t \in \mathbb{I} \mid \text{match}(T, t) = \vartheta\}} I[(\mathcal{X}K) \cdot \pi(t), \mathcal{P}\vartheta]} \quad (\text{C-QRYL-TA})$$

$$\frac{K \ominus T = \perp}{I[K, P] \xrightarrow[\text{TA}]{\text{get}(T)\text{@self}} \square} \quad (\text{C-GETL-TA}_B) \qquad \frac{K \vdash T = \perp}{I[K, P] \xrightarrow[\text{TA}]{\text{qry}(T)\text{@self}} \square} \quad (\text{C-QRYL-TA}_B)$$

$$\frac{P \xrightarrow[\text{TA}]{\mathbf{gq}(T)\text{@p}} \mathcal{P} \quad \text{match}(T, t) = \vartheta}{I[K, P] \xrightarrow[\text{TA}]{\mathbf{gq}(T:t)\text{@p}} I[(\mathcal{X}K), \mathcal{P}\vartheta]} \quad (\text{C-GQO-TA}) \qquad \frac{P \xrightarrow[\text{TA}]{\mathbf{gq}(T)\text{@p}} \mathcal{P} \quad \neg \text{match}(T, t)}{I[K, P] \xrightarrow[\text{TA}]{\mathbf{gq}(T:t)\text{@p}} \square} \quad (\text{C-GQO-TA}_B)$$

Time-abstract sequences of actions:

$$\frac{P \xrightarrow[\text{TA}]{\text{end.s}} \mathcal{P}}{I[K, P] \xrightarrow[\text{TA}]{\langle \rangle} I[K, \mathcal{P}]} \quad (\text{C-BASE-TA}) \qquad \frac{\alpha \in \text{OneAct}_{Proc} \quad I[K, P] \xrightarrow[\text{TA}]{\alpha} \mathcal{C} \quad \mathcal{C} \xrightarrow[\text{TA}]{\langle seq \rangle}^{cv} \mathcal{D}}{I[K, P] \xrightarrow[\text{TA}]{\langle \alpha.seq \rangle} \mathcal{D}} \quad (\text{C-REC-TA})$$

Timed sequences of actions:

$$\frac{P \xrightarrow[\text{T}]{\langle \rangle} \mathcal{Q}}{I[K, P] \xrightarrow[\text{T}]{\langle \rangle} I[(\mathcal{X}K), \mathcal{Q}]} \quad (\text{C-ESEQ-T})$$

$$\frac{P \xrightarrow[\text{TA}]{\text{init.s}} [Q \mapsto 1] \quad I[K, Q] \xrightarrow[\text{TA}]{\langle seq \rangle} I[\mathcal{H}, \mathcal{Q}] \quad \mathcal{Q} \xrightarrow[\text{TA}]{\text{delay}}^{cv} \mathcal{R}}{I[K, P] \xrightarrow[\text{T}]{\langle seq \rangle} I[\mathcal{H}, \mathcal{R}]} \quad (\text{C-SEQ-T})$$

Table 14: Operational sem. of STOCS components (ACTIV-OR), $\xrightarrow[\text{TA}]{} \mathcal{P}$ and $\xrightarrow[\text{T}]{} \mathcal{P}$ transition relations.

$\text{Dist}(\text{Comp})$ which is defined by the rules in Table 15 and uses its lifting $\xrightarrow[\text{PR}]{}^{cv} \subseteq \text{Dist}(\text{Comp}) \times \text{Act}_{Sys}^{\text{PR}} \times \text{Dist}(\text{Comp})$ to probability distributions over components. The set $\text{Act}_{Sys}^{\text{PR}}$ of actions is defined as follows:

$$\text{Act}_{Sys}^{\text{PR}} ::= \text{OneAct}_{Proc} \mid \langle \text{Seq}_{Sys} \rangle$$

| | |
|---|---|
| $\frac{I(K) \models \mathbf{p} \quad K \oplus t = \pi}{I[K, P] \xrightarrow[\text{PR}]{\text{put}(t)@p} I[\pi, (\mathcal{X}P)]} \quad (\text{C-PUTI-PR})$ | $\frac{I(K) \not\models \mathbf{p}}{I[K, P] \xrightarrow[\text{PR}]{\text{put}(t)@p} [I[K, P] \mapsto 1]} \quad (\text{C-PUTIR-PR})$ |
| $\frac{I(K) \models \mathbf{p} \quad K \ominus T = \pi}{I[K, P] \xrightarrow[\text{PR}]{\text{get}(T:t)@p} I[\pi(t), (\mathcal{X}P)]} \quad (\text{C-GETI-PR})$ | $\frac{I(K) \not\models \mathbf{p} \vee K \ominus T = \perp}{I[K, P] \xrightarrow[\text{PR}]{\text{get}(T:t)@p} []} \quad (\text{C-GETI-PR}_B)$ |
| $\frac{I(K) \models \mathbf{p} \quad K \vdash T = \pi}{I[K, P] \xrightarrow[\text{PR}]{\text{qry}(T:t)@p} [I[K, P] \mapsto \pi(t)]} \quad (\text{C-QRYI-PR})$ | $\frac{I(K) \not\models \mathbf{p} \vee K \vdash T = \perp}{I[K, P] \xrightarrow[\text{PR}]{\text{qry}(T:t)@p} []} \quad (\text{C-QRYI-PR}_B)$ |
| $\frac{S_1 \xrightarrow[\text{PR}]{\text{put}(t)@p} \mathcal{S}_1 \quad S_2 \xrightarrow[\text{PR}]{\text{put}(t)@p} \mathcal{S}_2}{S_1 \parallel S_2 \xrightarrow[\text{PR}]{\text{put}(t)@p} \mathcal{S}_1 \parallel \mathcal{S}_2} \quad (\text{S-PI-PR})$ | |
| $\frac{S_1 \xrightarrow[\text{PR}]{\text{gq}(T:t)@p} \mathcal{S}_1 \quad S_2 \xrightarrow[\text{PR}]{\text{gq}(T:t)@p} \mathcal{S}_2}{S_1 \parallel S_2 \xrightarrow[\text{PR}]{\text{gq}(T:t)@p} \mathcal{S}_1 \parallel (\mathcal{X} S_2) + (\mathcal{X} S_1) \parallel \mathcal{S}_2} \quad (\text{S-GQI-PR})$ | |

Sequences of actions:

| | |
|--|---|
| $\frac{}{S \xrightarrow[\text{PR}]{\langle \rangle} [S \mapsto 1]} \quad (\text{S-BASE-PR})$ | $\frac{\alpha \in \text{OneAct}_{Proc} \quad S \xrightarrow[\text{PR}]{\alpha} \mathcal{S} \quad \mathcal{S} \xrightarrow[\text{PR}]{\langle seq \rangle}^{cv} \mathcal{T}}{S \xrightarrow[\text{PR}]{\langle \alpha . seq \rangle} \mathcal{T}} \quad (\text{S-REP-PR})$ |
|--|---|

Table 15: Operational sem. of STOCS systems and components (ACTIV-OR), $\xrightarrow[\text{PR}]{} \rightarrow$ transition relation.

8 Case Study

We develop a model of a *bike sharing service*, where we assume a city with m *parking stations*, each one with his *location* $\ell_i \in \text{Loc} = \{\ell_1, \dots, \ell_m\}$, a number of *available bikes* b_i , and a number of *available parking slots* s_i (for $i = 1, \dots, m$). Parking stations are in one-to-one correspondence with the set of possible locations, which should be considered as (disjoint) areas of influence in the city. We also assume to have n *users* of the bike sharing service: at any time, each user is positioned in *one* location and can be in one of the two states *Pedestrian* and *Biker*. In each of the two states, the user moves around the city (with speed depending on the state) according to its preferences, modeled by two *probability* transition matrices Q_b and Q_p of size $m \times m$ for the biker and the pedestrian state, respectively. Then, the user becomes a *Biker* or a *Pedestrian*, using transitions named *Borrow* and *Return*.

In our model we identify two components: parking stations (containing slots and bikes) and users (that may or may not have a bike), whose population can vary. By relying on the distinguishing features of the STOCS language, we model bike/slot retrieval from stations in a distributed and adaptive way (avoiding a central server being aware of bikes/slots availability for each parking). A user in a location ℓ can borrow (or return) a bike by issuing a request (e.g. by means of a mobile phone application) to the bike sharing system for a parking with an available bike (or slot) within a neighborhood of ℓ . The bike sharing system answers with the location of a parking station having an available bike or an available parking slot, within a neighborhood of ℓ specified by a neighborhood condition $\varphi_n(\ell, \ell')$ (modeling, for example, that a parking station ℓ' is easily reachable from ℓ). This flexibility on the choice allows some control on which parking station is selected among those that are in the neighborhood of the

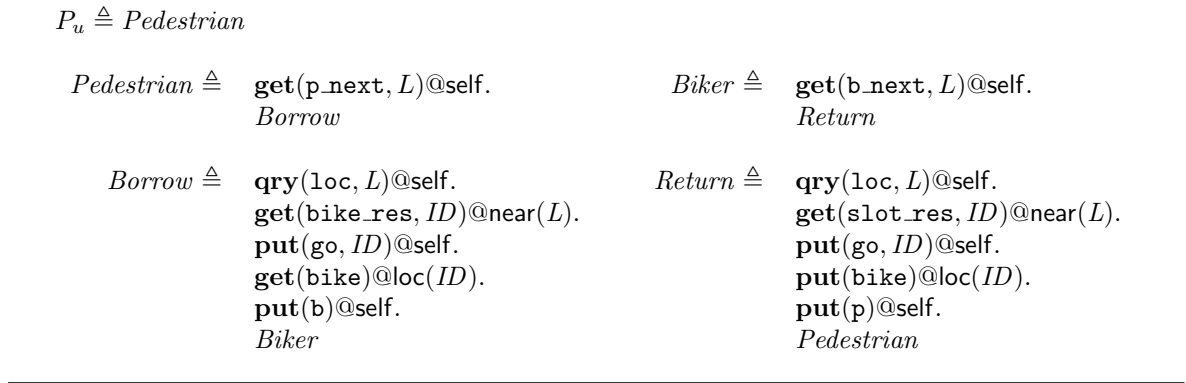


Figure 5: User behavior as a STOCS process.

current user location (including itself), which can be used to re-balance slot/bike availability by redirecting users to parking station that have many available bikes (or slots).

In our model, parking stations are themselves in charge of answering to bike/slot requests. This distributed solution is more efficient, robust, and compositional (i.e. the service can scale without changing the infrastructure) than the centralized one. A further advantage is that we can easily model a control over bike redistribution to minimize the cost of bike reallocation by means of trucks. In particular, the choice between parking stations is realized by using different response rates of parking stations: those that have more bikes available will have higher response rates to bikes requests. Similarly, parking stations that have more slots available will have higher response rates to slots requests. By using a **get** action, we put these responses into a race and, on average, users will be redirected to those services that need bike or slot re-balancing. We will discuss this and other features of the model in the rest of the document. In this example we consider the ACT-OR semantics and we assume that message transmission is free of errors, that is $f_{\text{err}}(\sigma, \alpha, \delta) = 0$ for any interface evaluations σ, δ and action α .

We first describe the behavior of the user, given in Figure 5. Then we describe the knowledge operators, defined by rules in Figure 6, and the actions rates, defined by cases in Figure 7, both for the user and the parking station components. The parking stations have no behavior, i.e. they are passive components.

A single user is represented as a component $I_u [K_u, P_u]$, whose knowledge state K_u is an element $\langle s, \ell \rangle$ in $\{\mathbf{b}, \mathbf{p}\} \times \text{Loc}$ denoting the user state (i.e. either being a pedestrian or a biker) and the user location, and whose interface I_u , which defines the predicates **biker**, **pedestrian**, and **loc**(ℓ) as follows: $I_u(\langle \mathbf{b}, \ell \rangle) \models \text{biker}$, $I_u(\langle \mathbf{p}, \ell \rangle) \models \text{pedestrian}$, and $I_u(\langle s, \ell \rangle) \models \text{loc}(\ell)$, for every $\ell \in \text{Loc}$ and $s \in \{\mathbf{b}, \mathbf{p}\}$. Let us summarize the role of the user knowledge operators, defined in Figure 6. The \oplus operator allows: to change state by $\oplus(\mathbf{b})$ (change to biker state) and by $\oplus(\mathbf{p})$ (change to pedestrian state), and to move to a specified location ℓ' by $\oplus(\text{go}, \ell')$. The \ominus operator allows to move to a location according to the average user behavior in the pedestrian state, by $\ominus(\text{p_next}, L)$, and in the biker state, by $\ominus(\text{b_next}, L)$. Finally, the \vdash operator allows to retrieve the current user location.

The users behaviour is given in Figure 5. Each user starts in the state *Pedestrian*, where movement is possible through a local **get** of the item **p_next**. The choice of the location is resolved internally using the probability matrix Q_p , as shown in the knowledge rule KR_4^u . The effect of the action is the change of location of the user into ℓ_j , which is also returned as a binding for the variable L . Note that the template is $T = (\text{p_next}, L)$ and the retrieved item is $t = (\text{p_next}, \ell_j)$. The location ℓ_j is used to compute the rate of the action (i.e. of the movement), as shown in the definition RR_3^u , specifying the rate function for this specific action.

User:

$$\begin{array}{ll}
\text{KR}_1^u : \langle s, \ell_i \rangle \xrightarrow{\oplus(\mathbf{b})} [\langle \mathbf{b}, \ell_i \rangle \mapsto 1] & \text{KR}_4^u : \langle s, \ell_i \rangle \xrightarrow{\ominus(\mathbf{p_next}, L)} \sum_{j=1}^m [(\langle s, \ell_j \rangle, (\mathbf{p_next}, \ell_j)) \mapsto Q_{\mathbf{p}}(i, j)] \\
\text{KR}_2^u : \langle s, \ell_i \rangle \xrightarrow{\oplus(\mathbf{p})} [\langle \mathbf{p}, \ell_i \rangle \mapsto 1] & \text{KR}_5^u : \langle s, \ell_i \rangle \xrightarrow{\ominus(\mathbf{b_next}, L)} \sum_{j=1}^m [(\langle s, \ell_j \rangle, (\mathbf{b_next}, \ell_j)) \mapsto Q_{\mathbf{b}}(i, j)] \\
\text{KR}_3^u : \langle s, \ell_i \rangle \xrightarrow{\oplus(\mathbf{go}, \ell')} [\langle s, \ell' \rangle \mapsto 1] & \text{KR}_6^u : \langle s, \ell_i \rangle \xrightarrow{\vdash(\mathbf{loc}, L)} [(\mathbf{loc}, \ell_i) \mapsto 1]
\end{array}$$

Parking Station:

$$\begin{array}{ll}
\text{KR}_1^p : \langle b_a, b_r, s_a, s_r, \ell_i \rangle \xrightarrow{\oplus(\mathbf{bike})} [\langle b_a + 1, b_r, s_a, s_r - 1, \ell_i \rangle \mapsto 1] & \text{if } s_r > 0 \\
\text{KR}_2^p : \langle b_a, b_r, s_a, s_r, \ell_i \rangle \xrightarrow{\ominus(\mathbf{bike_res}, L)} [(\langle b_a - 1, b_r + 1, s_a, s_r, \ell_i \rangle, (\mathbf{bike_res}, \ell_i)) \mapsto 1] & \text{if } b_a > 0 \\
\text{KR}_3^p : \langle b_a, b_r, s_a, s_r, \ell_i \rangle \xrightarrow{\ominus(\mathbf{slot_res}, L)} [(\langle b_a, b_r, s_a - 1, s_r + 1, \ell_i \rangle, (\mathbf{slot_res}, \ell_i)) \mapsto 1] & \text{if } s_a > 0 \\
\text{KR}_4^p : \langle b_a, b_r, s_a, s_r, \ell_i \rangle \xrightarrow{\ominus(\mathbf{bike})} [(\langle b_a, b_r - 1, s_a + 1, s_r, \ell_i \rangle, \mathbf{bike}) \mapsto 1] & \text{if } b_r > 0
\end{array}$$

Figure 6: Knowledge behavior (User/Parking Station). For $s \in \{\mathbf{b}, \mathbf{p}\}$, $i = 1, \dots, m$, $\ell' \in \text{Loc}$, and $b_a, b_r, s_a, s_r \in \mathbb{N}$.

Local actions:

$$\begin{array}{l}
\text{RR}_1^u : \mathcal{R}(\sigma, \mathbf{put}(\mathbf{go}, \ell')@self, \sigma) = \begin{cases} \lambda_{\mathbf{b}} \cdot f_{dist}(\ell, \ell') & \text{if } \sigma \models \mathbf{biker} \wedge \mathbf{loc}(\ell) \\ \lambda_{\mathbf{p}} \cdot f_{dist}(\ell, \ell') & \text{if } \sigma \models \mathbf{pedestrian} \wedge \mathbf{loc}(\ell) \end{cases} \\
\text{RR}_2^u : \mathcal{R}(\sigma, \mathbf{put}(\mathbf{p})@self, \sigma) = \mathcal{R}(\sigma, \mathbf{put}(\mathbf{b})@, \sigma) = \lambda_{fast} \\
\text{RR}_3^u : \mathcal{R}(\sigma, \mathbf{get}((\mathbf{p_next}, L) : (\mathbf{p_next}, \ell'))@self, \sigma) = \lambda_{\mathbf{p}} \cdot f_{distance}(\ell, \ell') \quad \text{if } \sigma \models \mathbf{loc}(\ell) \\
\text{RR}_4^u : \mathcal{R}(\sigma, \mathbf{get}((\mathbf{b_next}, L) : (\mathbf{p_next}, \ell'))@self, \sigma) = \lambda_{\mathbf{b}} \cdot f_{distance}(\ell, \ell') \quad \text{if } \sigma \models \mathbf{loc}(\ell) \\
\text{RR}_5^u : \mathcal{R}(\sigma, \mathbf{qry}((\mathbf{loc}, L) : (\mathbf{loc}, \ell'))@self, \sigma) = \lambda_{fast}
\end{array}$$

Remote actions:

$$\begin{array}{l}
\text{RR}_6^u : \mathcal{R}(\sigma, \mathbf{put}(\mathbf{bike})@loc(\ell), \delta) = \mathcal{R}(\sigma, \mathbf{get}(\mathbf{bike})@loc(\ell), \delta) = \lambda_{park} \\
\text{RR}_7^u : \mathcal{R}(\sigma, \mathbf{get}((\mathbf{bike_res}, L) : (\mathbf{bike_res}, \ell'))@near(\ell), \delta) = \frac{b_a}{b_a + s_a} \quad \text{if } \delta \models (\mathbf{bikes} = b_a) \wedge (\mathbf{slots} = s_a) \\
\text{RR}_8^u : \mathcal{R}(\sigma, \mathbf{get}((\mathbf{slot_res}, L) : (\mathbf{slot_res}, \ell'))@near(\ell), \delta) = \frac{s_a}{b_a + s_a} \quad \text{if } \delta \models (\mathbf{bikes} = b_a) \wedge (\mathbf{slots} = s_a)
\end{array}$$

Figure 7: Rate Function, for any $\ell, \ell' \in \text{Loc}$ and $b_a, s_a \in \mathbb{N}$.

We distinguish the movement of the pedestrians from that of the bikers ($\mathbf{p_next}/\mathbf{b_next}$) to allow the more realistic assumption that bikers moves more quickly than pedestrians (in definitions RR_3^u and RR_4^u , as well as the two cases of definition RR_1^u , we use different rates $\lambda_{\mathbf{p}}$ and $\lambda_{\mathbf{b}}$). The role of the information made available in the interfaces is important not only for the evaluation of the predicates, but also for the computation of the rates, as shown in Figure 7.

The process *Borrow* first retrieves the current location L (using rule KR_6^u and a rate given by RR_5^u). Then, it performs a *bike* reservation ($\mathbf{bike_res}$) from a parking station ID satisfying predicate $\mathbf{near}(L)$: this triggers on the side of the parking station a knowledge action that decrements the number of available bikes, increments the number of reserved bikes, and returns the id of the selected parking station (by using rule KR_2^p and with rate RR_7^u). Here the




Figure 8: Simulation of bike sharing service.

template is $T = (\mathbf{bike_res}, L)$ and the retrieved item is $t = (\mathbf{bike_res}, \ell_i)$, which produces the binding $L \mapsto \ell_i$. Note that the actual rate of this action depends on available bikes: the higher is the number of available bikes, the higher is the execution rate. As an effect of this race condition, the *ID* of the near station containing *more bikes* is received by the user with a *higher probability* than a near station with *fewer bikes*, causing a more balanced distribution of bikes in the system. When the parking lot is reserved, the user moves towards the parking station (which modifies the knowledge state according to KR_3^u). The rate of this action depends on the distance between the user and the parking station and on the pedestrian/biker state of the user (according to RR_1^u). After moving, the user takes a **bike**: this operation is performed via a **get** action whose effect on the parking knowledge state is given by rule KR_4^p , that decrements the bikes available and increments the slots available (the action has rate given by RR_6^u). Finally, the user status is updated to biker (using KR_1^u and with rate RR_2^u).

Bikers move around the city and, then, leave their bikes in a parking station by executing the process *Return*. Their behavior is similar to that of pedestrians, except for the fact that they perform complementary operations: they use a local **get** of the item **b_next** for moving (see rule KR_5^u and definition RR_4^u) instead of **p_next**, they perform a *slot* reservation **slot_res** (using rule KR_3^p and definition RR_8^u) instead of a bike reservation **bike_next**, and they return a **bike** to a parking station using a **put** action, rather than a **get** action (using rule KR_1^p and definition RR_6^p). After these actions, they return in state *Pedestrian* (using KR_2^u and with rate RR_2^u).

A parking station is represented as a component $I_p[K_p, \mathbf{nil}]$ that has no behavior (it is passive). Its knowledge state is a vector $\langle b_a, b_r, s_a, s_r, \ell \rangle \in \mathbb{N}^4 \times \text{Loc}$ denoting the number of available bikes (b_a), of reserved bikes (b_r), of available parking slots (s_a), and of reserved parking slots (s_r), as well as the parking location ℓ . The parking station interface I_p defines the predicates $\text{loc}(\ell)$ and $\text{near}(\ell)$ as follows: $I_p(\langle b_a, b_r, s_a, s_r, \ell \rangle) \models \text{loc}(\ell)$ and $I_p(\langle b_a, b_r, s_a, s_r, \ell \rangle) \models \text{near}(\ell')$ if $\varphi_n(\ell, \ell')$ holds, for every $\ell, \ell' \in \text{Loc}$ and $b_a, b_r, s_a, s_r \in \mathbb{N}$.

An initial state of this model is a term

$$\|_{i=1}^m ((I_u[\langle \ell_i, \mathbf{p} \rangle, P_u])[k_i] \parallel I_p[\langle b_i, 0, s_i, 0, \ell_i \rangle, \mathbf{nil}])$$

which denotes, for $i = 1, \dots, m$: (i) k_i pedestrians in locations ℓ_i , and (ii) b_i available bikes and s_i available parking slots in parking station at location ℓ_i . Note that the number of reserved bikes as well as the number of reserved slots is set to zero at the initial state of the system in every parking station. The overall number of bikes in the system is preserved by the knowledge-update rules. Furthermore, each parking station has a slot+bike capacity set initially by the value $b_i + s_i$ which is never exceeded, again thanks to the rules above.

A fully specified model is obtained by fixing the parameters (all of them in $\mathbb{R}_{\geq 0}$): λ_{fast} denoting a fast rate for internal update actions, λ_{park} denoting a (uniform) interaction rate

with the parking (e.g. locking/unlocking the bike), λ_p denoting the rate of movement of a pedestrian, and λ_b denoting the rate of movement of a biker, and also a *distance* function $f_{dist} : Loc \times Loc \rightarrow \mathbb{R}_{\geq 0}$ that encodes the connection topology of the locations. By giving concrete values to those parameters, to populations, and to bikes/slots in the initial state we obtain a finite state CTMC model for the described system.

Preliminary simulation analyses of the considered system have been performed with jRESP³. These simulations compare the case where rates of bikes and slots reservations depend on the number of available resources and the case where these rates are constant, and show that the average number of available bikes/slots per parking station is the same, while variance is lower. That is, when the bikes/slots reservation rate depends on the available resources, the bikes are more evenly distributed over the different parking stations.

9 Conclusions and Future Work

We have introduced STOCS, a stochastic extension of SCEL, for the modeling and analysis of performance aspects of ensemble based autonomous systems. One of the original features of the language is the use of stochastic predicate based multi-cast communication which poses particular challenges concerning stochastically timed semantics. Four variants of the semantics, considering different levels of abstraction, have been presented and their main aspects of the formal semantics have been provided. A case study concerning shared bikes systems was presented to illustrate the use of the various language primitives of STOCS. The development of both numeric and statistical model-checking tools for STOCS is in progress. In particular, (ACT-OR) and (INT-OR) semantics are well suited for formal analysis techniques (e.g. probabilistic model-checking), while the more detailed and complex (NET-OR) can be used for simulation-based techniques (e.g. statistical model-checking). The formal relationship between the different semantics is a non trivial issue and it is left for future work, as well as the development of fluid semantics and verification techniques to address large scale collective systems along the lines of work in [2, 3].

References

- [1] Luca Bortolussi, Vashti Galpin, and Jane Hillston. Hype with stochastic events. In Mieke Massink and Gethin Norman, editors, *QAPL*, volume 57 of *EPTCS*, pages 120–133, 2011.
- [2] Luca Bortolussi and Jane Hillston. Fluid model checking. In Maciej Koutny and Irek Ulidowski, editors, *CONCUR*, volume 7454 of *Lecture Notes in Computer Science*, pages 333–347. Springer, 2012.
- [3] Luca Bortolussi, Jane Hillston, Diego Latella, and Mieke Massink. Continuous approximation of collective system behaviour: A tutorial. *Perform. Eval.*, 70(5):317–349, 2013.
- [4] Roberto Bruni, Andrea Corradini, Fabio Gadducci, Alberto Lluch-Lafuente, and Andrea Vandin. A conceptual framework for adaptation. In Juan de Lara and Andrea Zisman, editors, *FASE*, volume 7212 of *Lecture Notes in Computer Science*, pages 240–254. Springer, 2012.
- [5] Federica Ciocchetta and Jane Hillston. Bio-pepa: A framework for the modelling and analysis of biological systems. *Theor. Comput. Sci.*, 410(33-34):3065–3084, 2009.

³<http://jresp.sourceforge.org>

- [6] Rocco De Nicola, Gian Luigi Ferrari, Michele Loreti, and Rosario Pugliese. A language-based approach to autonomic computing. In Bernhard Beckert, Ferruccio Damiani, Frank S. de Boer, and Marcello M. Bonsangue, editors, *FMCO*, volume 7542 of *Lecture Notes in Computer Science*, pages 25–48. Springer, 2011.
- [7] Rocco De Nicola, Diego Latella, Michele Loreti, and Mieke Massink. Rate-based transition systems for stochastic process calculi. In Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris E. Nikolettseas, and Wolfgang Thomas, editors, *ICALP (2)*, volume 5556 of *Lecture Notes in Computer Science*, pages 435–446. Springer, 2009.
- [8] Rocco De Nicola, Diego Latella, Michele Loreti, and Mieke Massink. A uniform definition of stochastic process calculi. *ACM Comput. Surv.*, 46(1):5:1–5:35, July 2013.
- [9] Vashti Galpin, Jane Hillston, and Federica Ciocchetta. A semi-quantitative equivalence for abstracting from fast reactions. In Ion Petre and Erik P. de Vink, editors, *CompMod*, volume 67 of *EPTCS*, pages 34–49, 2011.
- [10] Holger Hermanns and Joost-Pieter Katoen. The how and why of interactive markov chains. In Frank S. de Boer, Marcello M. Bonsangue, Stefan Hallerstede, and Michael Leuschel, editors, *FMCO*, volume 6286 of *Lecture Notes in Computer Science*, pages 311–337. Springer, 2009.
- [11] Nora Koch, Matthias Hözl, Annabelle Klarl, Philip Mayer, Tomas Bures, Jaques Combaz, Alberto Lluch Lafuente, Rocco De Nicola, Stefano Sebastio, Francesco Tiezzi, Andrea Vandin, Fabio Gadducci, Valentina Monreale, Ugo Montanari, Michele Loreti, Carlo Pinciroli, Mariachiara Puviani, Franco Zambonelli, Nikola Šerbedžija, and Emil Vashev. JD3.2: Software engineering for self-aware SCEs, 2013. ASCENS Deliverable JD3.2.
- [12] Diego Latella, Michele Loreti, Mieke Massink, and Valerio Senni. On StocS: a Stochastic extension of SCEL. Technical Report 11, ASCENS Project, 2013. <http://www.ascens-ist.eu/>.
- [13] Rocco De Nicola, Matthias Hözl, Michele Loreti, Alberto Lluch Lafuente, Ugo Montanari, Emil Vashev, and Franco Zambonelli. JD2.1: Languages and knowledge models for self-awareness and self-expression, 2012. ASCENS Deliverable JD2.1.
- [14] Catuscia Palamidessi. Comparing the expressive power of the synchronous and asynchronous pi-calculi. *Mathematical Structures in Computer Science*, 13(5):685–719, 2003.
- [15] Nikola Šerbedžija, Nicklas Hoch, Carlo Pinciroli, Michal Kit, Tomas Bures, Valentina Monreale, Ugo Montanari, Philip Mayer, and José Velasco. D7.3: Third report on wp7 - integration and simulation report for the ascens case studies, 2013. ASCENS Deliverable D7.3.
- [16] Nikola Šerbedžija, Mieke Massink, Carlo Pinciroli, Manuele Brambilla, Diego Latella, Marco Dorigo, Mauro Birattari, Philip Mayer, José Angel Velasco, Nicklas Hoch, Henry P. Bensler, Dhaminda Abeywickrama, Jaroslav Keznikl, Ilias Gerostathopoulos, Tomas Bures, Rocco De Nicola, and Michele Loreti. D7.2: Second report on wp7 - integration and simulation report for the ascens case studies, 2012. ASCENS Deliverable D7.2.