

# Service Oriented Computing from a Process Algebraic Perspective<sup>★</sup>

Mario Bravetti      Gianluigi Zavattaro

*Dipartimento di Scienze dell'Informazione, Università degli Studi di Bologna,  
Mura Anteo Zamboni 7, I-40127 Bologna, Italy.  
e-mail:{bravetti,zavattar}@cs.unibo.it*

---

## Abstract

Service Oriented Computing is emerging as a reference model for a new class of distributed computing technologies such as Web Services and the Grid. We discuss three main aspects of Service Oriented Computing (loose coupling, communication latency, and open endedness), and we relate them with traditional process algebra operators. We also indicate some new issues, raising from the combination of these three aspects, that require the investigation of suitable new process algebra operators.

*Key words:* service oriented computing, service composition, process calculi.

---

## 1 Introduction

Service Oriented Computing is nowadays emerging as the new reference computing paradigm for the networked world. Even if this new paradigm has evolved from object-oriented and component computing, it comprises new original basic aspects that enable the agile design and development of interoperable collaborating applications. The basic computational elements in these applications are the services, that are autonomous platform-independent elements that can be described, published and discovered using interoperable standards. Services are more than just software components because of their platform neutrality and their self-describing nature. Platform neutrality is usually achieved thanks to new XML-based technology. The self-describing

---

<sup>★</sup> Research partially funded by EU Integrated Project Sensoria, contract n. 016004.

nature, on the other hand, allows for dynamic and automatic retrieving and composition of services.

Another relevant novelty of service oriented computing, with respect to traditional distributed computing models, is that the global information about the state of the computation is not recorded inside the individual services but it is usually passed within the exchanged messages. This technique is called contextualization because the messages contain additional context information such as cookies or session identifiers. Due to contextualization, the service oriented paradigm is particularly suited to program systems based on a minimal shared knowledge and understanding among the interacting parts. These systems are usually referred to as loosely coupled systems.

The most prominent service oriented technologies are Web Services and the Grid that support the deployment of services that can interact via basic interaction patterns. Complex service interactions, which cannot be trivially encoded in the basic patterns, require a so-called service orchestration. Service orchestration is usually achieved adding new components (called the orchestrators) that do not actually perform computation, but simply manage the flow of invocation of the services involved in the collaboration. A plethora of languages (comprising e.g. XLANG [17], BizTalk [16], WS-BPEL [18], and WS-CDL [21]) has been recently defined to specify and program orchestrators. All these languages combine workflow constructs and communication primitives. The workflow constructs are used to describe the flow of execution of the orchestration activities, while the communication primitives correspond to the basic service interaction patterns. Some of these languages (such as XLANG and WS-CDL) have explicitly taken inspiration from process algebras such as CSP or the  $\pi$ -calculus. Nevertheless, due to peculiarities of service oriented computing, some constructs and primitives differ from the traditional operators of process algebras. Three of these peculiarities are:

**Loose coupling.** Orchestrators have a minimal control on the orchestrated services, for instance, a service can autonomously exit the orchestration without any previous notice.

**Communication latency.** The transport layer responsible for the exchange of messages, among the orchestrator and the services, does not give guarantees about the reliability and timing of remote message delivery.

**Open-endedness.** An orchestrator can dynamically, i.e. at run time, retrieve new services to be involved in the orchestration; for instance, this could be useful to replace services that autonomously leave the orchestration.

The literature already includes process calculi comprising new operators that take inspiration from at least one of these three aspects. For instance, the process calculi in [11,4,10,9,8] comprise operators inspired by long running transactions, the most relevant programming construct used to deal with loose

coupling in the context of service oriented computing. As far as communication latency is concerned, timeout is the typical programming mechanism exploited in order avoid that programs wait too long for delayed service responses. Timeouts are well studied in the context of process calculi, see e.g. [1] for a comprehensive presentation of these calculi, and some preliminary work about their exploitation in the context of service oriented computing can be found for instance in [15]. Finally, open endedness has been already considered in recent work such as [7] where the problem of balancing the distribution of the service workload has been analysed in a process algebraic setting.

Nevertheless, all these process calculi deal with only one of the three above peculiarities taken in isolation. In this paper, for the first time, we present a process calculus that combines all of these three peculiar aspects of service oriented computing. The remainder of the paper is organized as follows: in Sections 2, 3, and 4 we focus separately on the three above aspects, while in Section 5 we conclude discussing interesting issues raised by their combination.

## 2 Loose coupling

In order to cope with loose coupling, orchestration languages usually provide linguistic constructs to program the so-called loosely coupled transactions. Traditional database transactions guarantees the ACID properties: Atomicity, Consistency, Isolation, and Durability. When the activities involved in a transaction are loosely coupled the ACID properties adapt badly. In particular, Isolation usually requires to lock resources. In Web Services applications, for instance, resources may belong to different companies and there is no chance for an orchestrator to lock resources of other companies. Additionally, transactions may last long periods of time, and it is not feasible to block resources so long.

The loosely coupled transactions weaken the notion of rollback: a service might decide that rollback will not cancel all the activities carried out. The cancellation of an airplane booking, for instance, may lead to the payment of a fee. Services that do not support an “absolute” mechanism of rollback, make failures extremely complicated, to be dealt with ad-hoc rollbacks. These ad-hoc rollback processes are called compensations.

The notion of compensation is the key aspect of several recent process algebras defined on purpose to formalize the semantics of compensation execution, and to reason about properties of compensation policies. The first proposal in this direction is StAC, a calculus with an explicit compensation operator whose operational semantics has been formalized in [10]. StAC has recently inspired also a new CSP dialect, called cCSP [11], whose semantics is defined

denotationally in terms of traces. An alternative proposal is represented by the SAGAS calculi [9] that defines a concurrent big-step semantics for sequential, parallel, and nested compensatable transactions. Recently, in [8] cCSP and the SAGAS calculi have been thoroughly compared discussing how to encode (fragments of) the former in (some of) the latter calculi, and vice versa. Compensations have been formalized and investigated also in the context of  $\pi$ -calculus in [4], where a calculus inspired by the compensation policy of BizTalk is presented.

In this Section we focus on the basic mechanisms required to run compensations. Compensations are usually activated in case an unexpected event occurs such as the reception of a negative response or the unavailability of a service. Usually, when one of these events occurs, some activity must be interrupted (because it has failed) and some other should be activated instead (responsible for executing the ad-hoc rollback procedure). In order to investigate these aspects formally, we define a process calculus comprising a new operator that combines the possibility to interrupt a process with the possibility to activate an alternative compensating process. A similar operator has been already investigated in [15] in the context of the asynchronous  $\pi$ -calculus. However, there are significant differences with that paper. First of all, in this paper interrupt signals can be produced only internally from the activity to be interrupted. Moreover, in case an interruptible activity contains other nested interruptible activities, the inner ones are blocked in case of interruption of the outer one; this was not the case in [15] where inner transactions and messages are not involved in transaction abort operations.

In this paper we include the process interruption operator in a process calculus based on asynchronous shared dataspace communication: processes interact by producing tuples (i.e. ordered sequences of data) that are stored in a shared repository called the data space where they can be subsequently retrieved (either read or consumed) by means of read or input coordination primitives. Shared dataspace communication has revealed a natural choice for modeling asynchronously interacting services. In particular, as will be made more clear in the next Section, we will be able to model in a rather simple manner networks of remotely interacting services exploiting the notion of distributed tuple spaces.

### *2.1 A Basic Calculus with Interruptible Processes*

The coordination primitives that we consider to access the shared data space are:  $out(e)$ ,  $in(t)$  and  $rd(t)$ . The output operation  $out(e)$  inserts a tuple  $e$  in the data space. Primitive  $in(t)$  is the blocking input operation: when an occurrence of a tuple  $e$  matching with  $t$  (denoting a template) is found in the

data space, it is removed from the data space and the primitive returns the tuple. The read primitive  $rd(t)$  is the blocking read operation that, differently from  $in(t)$ , returns the matching tuple  $e$  without removing it from the data space.

In languages based on shared data space communication, tuples are ordered and finite sequences of typed fields, while template are ordered and finite sequences of fields that can be either *actual* or *formal* (see [13]): a field is actual if it specifies a type and a value, whilst it is formal if the type only is given. For the sake of simplicity, in the formalization we are going to present, fields are not typed.

Formally, let  $Mess$ , ranged over by  $m, m', \dots$ , be a denumerable set of messages and  $Var$ , ranged over by  $x, y, \dots$ , be the set of data variables. In the following, we use  $\vec{x}, \vec{y}, \dots$ , to denote finite sequences  $x_1; x_2; \dots; x_n$  of variables. We consider also expressions taken from a generic set  $Exp$  that are terms over a syntax where variables and messages may occur. In particular,  $m$  and  $x$  are special kinds of expressions (formally,  $Mess \subset Exp, Var \subset Exp$ ). Moreover, we assume a partial evaluation function  $Eval : Exp \dashrightarrow Mess$  that computes, only for those expressions that do not contain variables, the associated value.

Tuples, denoted by  $e, e', \dots$ , are finite and ordered sequences of data fields (we use  $arity(e)$  to denote the number of fields of  $e$ ), whilst templates, denoted by  $t, t', \dots$ , are finite and ordered sequences of fields that can be either data or wildcards (used to match with any message).

Formally, tuples are defined as follows:

$$e = \langle \vec{d} \rangle$$

where  $\vec{d}$  is a term of the following grammar:

$$\begin{aligned} \vec{d} &::= d \mid d; \vec{d} \\ d &::= exp. \end{aligned}$$

We overload the evaluation function and apply it also to tuples with the expected meaning;  $Eval(e)$  returns the tuple obtained by evaluation of the fields of  $e$  that contain an expression.

The definition of template follows:

$$t = \langle \vec{dt} \rangle$$

where  $\vec{dt}$  is a term of the following grammar:

$$\begin{aligned} \vec{dt} &::= dt \mid dt; \vec{dt} \\ dt &::= d \mid null. \end{aligned}$$

The additional value *null* for data fields denotes the wildcard, whose meaning

is the same of formal fields, i.e. it matches with any field value. With abuse of notation we apply the *Eval* function also to templates. In the following, the set *Tuple* (resp. *Template*) denotes the set of tuples (resp. templates) containing only messages (resp. messages and *null*).

The matching rule between tuples and templates we consider is as follows.

**Definition 2.1 Matching rule** - Let  $e = \langle d_1; d_2; \dots; d_n \rangle \in \textit{Tuple}$  be a tuple,  $t = \langle dt_1; dt_2; \dots; dt_m \rangle \in \textit{Template}$  be a template; we say that  $e$  matches  $t$  (denoted by  $e \triangleright t$ ) if the following conditions hold:

- (1)  $m = n$ .
- (2)  $dt_i = d_i$  or  $dt_i = \textit{null}$ ,  $1 \leq i \leq n$ .

Condition 1. checks if  $e$  and  $t$  have the same arity, whilst 2. tests if each non-wildcard field of  $t$  is equal to the corresponding field of  $e$ .

Processes, denoted by  $P, Q, \dots$ , are defined as follows:

$P, Q, \dots ::=$	<i>commit</i>	commit command
	<i>abort</i>	abort command
	<i>out</i> ( $e$ ). $P$	output
	<i>rd</i> $t(\vec{x})$ . $P$	read
	<i>in</i> $t(\vec{x})$ . $P$	input
	$P \mid Q$	parallel composition
	$!P$	replication
	$P \leftrightarrow Q$	interruption

where we assume that each prefix of kind *rd/in*  $\langle \vec{dt} \rangle(\vec{x})$  is such that the variables  $\vec{x}$  and the data  $\vec{dt}$  have the same arity.

A process can complete its computation entering either in a *commit* or *abort* state; these two states are denoted by the *commit* and *abort* processes, respectively. The other processes are prefix forms  $\mu.P$ , the parallel composition of two programs, the replication of a program, or an interruptible process  $P \leftrightarrow Q$ . The prefix  $\mu$  can be one of the following coordination primitives: i) *out* ( $e$ ), that writes the tuple  $e$  in the data space; ii) *rd*  $t(\vec{x})$ , that given a template  $t$  reads a matching tuple  $e$  in the data space and stores the return value in  $\vec{x}$ ; iii) *in*  $t(\vec{x})$ , that given a template  $t$  consumes a matching tuple  $e$  in the data space and stores the return value in  $\vec{x}$ . In both the *rd*  $t(\vec{x})$  and *in*  $t(\vec{x})$  operations ( $\vec{x}$ ) is a binder for the variables in  $\vec{x}$  that will be bound to the actual fields of the matching tuple. The parallel composition  $P \mid Q$  of two processes  $P$  and  $Q$  behaves as two processes running in parallel. Infinite behaviours can be expressed using the replication operator  $!P$ . Replication is

a typical operator used in process calculi to denote the parallel composition of an unbounded amount of instances of the same process. The last operator is used to program interruptible activities. In the term  $P \leftrightarrow Q$  the process  $P$  executes its operation until a *abort* command is executed. After execution of the *abort* command the process  $Q$  is activated as interrupt handler.

To shorten the notation we usually omit trailing *commit*, e.g., we write the process  $out(e).commit$  simply as  $out(e)$ .

In the following,  $P[d/x]$  denotes the process that behaves as  $P$  in which all free occurrences of  $x$  (also inside expressions) are replaced with  $d$ . We also use  $P[\vec{d}/\vec{x}]$  to denote the process obtained by replacing in  $P$  all occurrences of variables in  $\vec{x}$  with the corresponding value in  $\vec{d}$ , i.e.  $P[d_1; d_2; \dots; d_n/x_1; x_2; \dots; x_n] = P[d_1/x_1][d_2/x_2] \dots [d_n/x_n]$ .

In the following we will consider only processes that are *closed*, i.e., all the variables (also those occurring inside expressions) are included in the scope of a binder. In the following, we denote with *Process* the set of such processes.

Let *DSpace*, ranged over by  $DS, DS', \dots$ , be the set of possible configurations of the data space, that is  $DSpace = \mathcal{M}_{fin}(Tuple)$ , where  $\mathcal{M}_{fin}(S)$  denotes the set of all the possible finite multisets on  $S$ . In the following, we use  $DS(e)$  to denote the number of occurrences of  $e$  within  $DS \in DSpace$ . The set  $System = \{\langle P, DS \rangle \mid P \in Process, DS \in DSpace\}$ , ranged over by  $s, s', \dots$ , denotes the possible configurations of systems.

The semantics we use to describe processes interacting via coordination primitives is defined by means of a structural congruence relation  $\equiv$  that equates processes that we do not want to distinguish.  $\equiv$  is defined to be the minimal congruence relation over processes such that

$$\begin{aligned} P|Q &\equiv Q|P & P|(Q|R) &\equiv (P|Q)|R & !P &\equiv !P|P \\ P|commit &\equiv P & commit \leftrightarrow P &\equiv commit \end{aligned}$$

The equivalence relation  $\equiv$  is extended to system configurations by means of the rule

$$\text{if } P \equiv Q \text{ then } \langle P, DS \rangle \equiv \langle Q, DS \rangle$$

Semantics of processes is defined in terms of a transition system over equivalence classes of system configurations, i.e.  $[s]_{\equiv}$  for some  $s \in System$ . In the rest of this section we will use  $[s]$  as a shorthand for  $[s]_{\equiv}$  and  $[P, DS]$  as a shorthand for  $[\langle P, DS \rangle]_{\equiv}$ . More precisely the operational semantics is defined to be  $(System/\equiv, \longrightarrow)$ , where:  $System/\equiv$ , i.e. the equivalence classes of  $\equiv$  over  $System$ , is the set of states and  $\longrightarrow \subseteq System/\equiv \times System/\equiv$  is the minimal relation satisfying the axioms and rules of Table 1.  $([s], [s']) \in \longrightarrow$

(also denoted by  $[s] \longrightarrow [s']$ ) means that a system (configuration)  $s$  can evolve (performing a single action) in the system (configuration)  $s'$ . When evaluating the semantics of a process  $P$ , we consider  $[P, \emptyset] \in System/\equiv$  to be the initial state.

$(1) \frac{Eval(e) = e'}{[out(e).P, DS] \longrightarrow [P, DS \oplus e']}$	$(2) \frac{Eval(t) = t' \quad e \in DS : e \triangleright t'}{[in\ t(\vec{x}).P, DS] \longrightarrow [P[e/\vec{x}], DS - e]}$
$(3) \frac{Eval(t) = t' \quad e \in DS : e \triangleright t'}{[rd\ t(\vec{x}).P, DS] \longrightarrow [P[e/\vec{x}], DS]}$	$(4) \frac{[P, DS] \longrightarrow [P', DS']}{[P \mid Q, DS] \longrightarrow [P' \mid Q, DS']}$
$(5) [abort \mid P \leftrightarrow Q, DS] \longrightarrow [Q, DS]$	$(6) \frac{[P, DS] \longrightarrow [P', DS']}{[P \leftrightarrow Q, DS] \longrightarrow [P' \leftrightarrow Q, DS']}$

Table 1

Semantics of the basic calculus with interruptible processes.

Rule (1) describes the output operation that produces a new occurrence of the tuple  $e'$  in the shared space  $DS$  ( $DS \oplus e'$  denotes the multiset obtained by  $DS$  increasing by 1 the number of occurrences of  $e'$ ). Rules (2) and (3) describe the *in* and the *rd* operations, respectively: if a matching  $e$  tuple is currently available in the space, it is returned at the process invoking the operation and, in the case of *in*, it is also removed from the space ( $DS - e$  denotes the removal of an occurrence of  $e$  from the multiset  $DS$ ). Rule (4) represents a local computation of processes. Rule (5) defines how the *abort* command can be used to interrupt a process and activate the interrupt handler instead. Finally, rule (6) is used to “close” the transition system to the interruptable process composition operator.

### 3 Communication Latency

Orchestration languages support a time aware programming style. For instance, in the visual orchestration language BizTalk [16], timed activities can be programmed which are interrupted in case they do not complete within a predefined period of time. Similarly, in WS-BPEL [18], it is possible to program signals that are raised at specific time instant, and to install handlers that are triggered by these signals.

Timed process algebras are an extremely powerful tool for modeling and

analysing timed systems. There exists numerous models of time inspired by different intuitions and abstractions, see e.g. [1] for a comprehensive overview. According to the traditional taxonomy of timed process algebras, the model of time that we adopt can be classified as *discrete*, *synchronous* and without *maximal progress* assumption [6]. More precisely, timed actions of different machines are assumed to take the same amount of time (one time unit) and are executed synchronously; moreover, in the choice between a timed action and a reduction no priority among them is considered. The unique time aware operator is  $P \leftarrow_n Q$ , a timed version of the interrupt operator that permits to interrupt the process  $P$  in the case it does not complete within a predefined number of time intervals. This number of intervals is quantified by the timeout  $n$  which is a strictly positive natural number (or infinity). Standard operations (output, read and input) take no time. An additional timed relation is defined in order to model the effect of the elapse of one time interval that require to decrease by one unit the timeouts. If one timeout cannot be decreased because it is already equal to 1, the corresponding activity  $P$  is interrupted and the interrupt handler  $Q$  is activated instead.

Besides timed interruptions, we consider also process and dataspace distribution. Time and distribution are strongly related concepts due to communication latency. In fact, messages exchanged among remote services are delivered after an unpredictable time delay, and services usually do not wait indefinitely for such these messages. For instance, an orchestrator that sends requests to two services, e.g. an hotel and an airplane reservation service, cannot indefinitely wait for the two answers; in case one of the two reservations does not complete in due time, the whole orchestration is aborted.

Distribution is achieved simply adding the notion of location  $\langle P, DS \rangle_l$  which is a triple composed of a process  $P$ , a dataspace  $DS$  and a location identifier  $l$ . A tuple  $e$  can be sent towards a remote location  $l$  simply by performing the *out* ( $e@l$ ) operation. Locations can be composed in parallel to model network of located machines.

### 3.1 A Calculus with Distribution and Time

We consider the new set of location  $Loc$ , ranged over by  $l, l', \dots$ , to denote location names. A location name is a special kind of message that can be communicated, formally  $Loc \subset Mess$ .

Processes are the same except that the parameter of the output operation may be also of the kind  $e@l$ ; such a message is inserted in the local data space but will be subsequently sent towards its destination injecting it in the network. Moreover, the transaction operator has an additional timeout

$P \leftarrow_n Q$  where  $n$  is a strictly positive natural number or  $\infty$  (for which  $\infty+1 = \infty$ ) representing the absence of timeout, i.e. retrieving the behaviour of the previous  $P \leftarrow Q$  operator.

Machines, denoted by  $M, N, \dots$ , are defined as follows:

$$\begin{array}{lcl}
M, N, \dots ::= & & \\
& \langle P, DS \rangle_l & \text{location} \\
& | \quad e@l & \text{message} \\
& | \quad M\|N & \text{network}
\end{array}$$

We define a predicate  $P \downarrow$  that verifies whether a process has committed

$$\begin{array}{l}
\text{commit} \downarrow \\
\text{if } P \downarrow \text{ then } P \leftarrow_n Q \downarrow \\
\text{if } P \downarrow \text{ and } Q \downarrow \text{ then } P|Q \downarrow
\end{array}$$

We need to extend the notion of structural congruence to machines also.  $\equiv$  is now defined as the minimal congruence relation over machines that includes equivalence over configurations of locations defined in the previous section (we consider  $\langle P, DS \rangle_l \equiv \langle Q, DS \rangle_l$  when  $\langle P, DS \rangle \equiv \langle Q, DS \rangle$ ) and that is such that

$$M\|N \equiv N\|M \quad M\|(N\|L) \equiv (M\|N)\|L$$

We are now ready to define the function on processes  $\phi(\cdot)$  that models the effect of the passing of one time unit

$$\begin{array}{l}
\text{if } P \not\downarrow \text{ then } \phi(P \leftarrow_{n+1} Q) = \phi(P) \leftarrow_n Q \\
\text{if } P \not\downarrow \text{ then } \phi(P \leftarrow_1 Q) = Q \\
\phi(P|Q) = \phi(P)|\phi(Q) \\
\phi(P) = P \text{ in all other cases}
\end{array}$$

Operational semantics is now defined in terms of two transition systems, one describing the execution of operations, another one describing the effect of the passing of one time unit. More precisely semantics of machines is defined in terms of a timed transition system  $(Machine/\equiv, \longrightarrow, \overset{\checkmark}{\longrightarrow})$ , where:  $Machine/\equiv$  is the set of states;  $\longrightarrow \subseteq Machine/\equiv \times Machine/\equiv$  is the minimal relation satisfying the set of axioms and rules in Table 1 (where  $l$  is added as subscript of systems, i.e.,  $[P, DS]$  becomes  $[\langle P, DS \rangle_l]$ , and timeout  $n$  is added as index of the interruptible processes, i.e.,  $P \leftarrow Q$  becomes  $P \leftarrow_n Q$ ) and in Table 2; and  $\overset{\checkmark}{\longrightarrow} \subseteq Machine/\equiv \times Machine/\equiv$  is the minimal relation

satisfying the set of axioms and rules in Table 3.  $([M], [M']) \in \overset{\checkmark}{\longrightarrow}$  (also denoted by  $[M] \overset{\checkmark}{\longrightarrow} [M']$ ) means that a machine  $M$  after one time tick evolves into a machine  $M' \in \text{Machine}$ . When evaluating the semantics of a machine  $M$ , we consider  $[M] \in \text{Machine}/\equiv$  to be the initial state of the transition system.

$(7) \quad \frac{[M] \longrightarrow [M']}{[M\ N] \longrightarrow [M'\ N]}$	$(8) \quad [\langle P, DS \oplus e@l \rangle_\nu] \longrightarrow [\langle P, DS \rangle_\nu \mid e@l]$
$(9) \quad [\langle P, DS \rangle_l \mid e@l] \longrightarrow [\langle P, DS \oplus e \rangle_l]$	

Table 2

Semantics of the calculus with distribution and time: reduction relation.

$(10) \quad [\langle P, DS \rangle_l] \overset{\checkmark}{\longrightarrow} [\langle \phi(P), DS \rangle_l]$	$(11) \quad [e@l] \overset{\checkmark}{\longrightarrow} [e@l]$
$(12) \quad \frac{[M] \overset{\checkmark}{\longrightarrow} [M'] \quad [N] \overset{\checkmark}{\longrightarrow} [N']}{[M\ N] \overset{\checkmark}{\longrightarrow} [M'\ N']}$	

Table 3

Semantics of the calculus with distribution and time: timed relation.

The most interesting new rules of the reduction relation are (8) and (9). The former models the injection on the network of a message to be sent to a remote location, while the latter models the delivery of the message. Rule (7) simply lifts the reductions for machines to an entire network of machines. As far as the rules of Table 3 are concerned, rule (10) indicates that the elapsing of one time unit requires the application on processes of the function  $\phi(\cdot)$  defined above, rule (11) shows that messages injected in the network are not affected by time passing, and rule (12) states that the effect of the elapsing of one time unit on an entire network is given by the effect of time passing on each of its nodes.

## 4 Open-endedness

Open-endedness is an inherent characteristics in orchestration of services retrieved from the internet: new services may appear and disappear at run-time,

available services (or their efficiency) may depend on their current location or on the current location of the orchestrator (if we deal with mobile entities), requests towards services offering the same service (where “same” is established in terms of some semantical definition of its behavior) may be distributed so to have a balanced workload. Assuming that we know available services and we bind them when the orchestrator is created (i.e. at “compile-time”) is not realistic in this context.

Expressing open-endedness in process algebra requires evolved mechanism for channel retrieval to access services. In particular the retrieval should be based on requirements on the desired service, e.g. on some abstraction of its behavior. This can be done in several ways: by using matching rules on tuples of data (formed e.g. by one element representing the channel and others describing the service and its behavior) as in Linda [13] or by using direct subtyping on channels themselves [12]. Note that in this context process algebra may be involved even in the description of the desired behavior of services itself. For example [14] uses abstract process algebraic descriptions as types of systems (services in our case) which are expressed in a more complex process algebra.

In this Section we extend the process calculus previously defined in order to model dynamic retrieval of services. The basic idea is to exploit tuples to describe available services, that can be retrieved using the read and/or input operations. Inspired by the standard UDDI protocol [20], we model a service registry as a node in the network that can be used to publish and discovery new services. One of the main limitation of shared data space used as service description repository is that it has no structure: all tuples in the data space have the same relevance, thus it is rather complex to cope with contexts in which there are services that are more important than other because, for instance, provides more powerful resources or connection with a larger bandwidth. To address this limitation, we follow the approach initiated in [5] that consists of adding quantitative information to tuples in order to quantify the relevance of the tuple via weights and priority. Among the matching tuples, only those with the higher priority level can be retrieved. Among them, the higher is the weight of a tuple, the higher is the probability for that tuple to be retrieved. In this paper, for the sake of simplicity we just consider weights; in particular we first extend the calculus with weights and probabilities, then we formally define how to model a service registry in the new calculus.

#### 4.1 A Calculus with Weighted Tuples

Let *Weight*, ranged over by  $w, w', \dots$ , be the set of the possible weights. We assume to use positive (non-zero) real numbers as weights, thus *Weight* coincides with  $\mathfrak{R}_+ \setminus \{0\}$ . We also assume, that a weight is a special kind of

message, formally  $Weight \subset Mess$ .

Tuples are now defined as follows:

$$e = \langle \vec{d} \rangle [exp]$$

where  $exp$  is an expression that, when evaluated, returns a weight. To cope with this extended form of weighted tuples, we assume that  $Eval(e)$  evaluates also the expression  $exp$  in  $\langle \vec{d} \rangle [exp]$ , thus returning a complete tuple of the form  $\langle \vec{d} \rangle [w]$ . Also the set  $Tuple$  containing the possible entries in the data space is updated accordingly: it now contains all the terms  $\langle \vec{d} \rangle [w]$  composed of only messages.

We also define  $\tilde{\cdot}$  as the function that, given a tuple  $e$ , returns its sequence of data fields (e.g. if  $e = \langle \vec{d} \rangle [w]$  then  $\tilde{e} = \vec{d}$ ). In the following, we denote with  $W$  the function that, given a tuple, returns its weight (e.g., if  $e = \langle \vec{d} \rangle [w]$  then  $W(e) = w$ ). Weights are not considered in the matching rule whose definition is unchanged.

The semantics replaces the standard non-deterministic choice of a tuple among the matching ones in the data space, with a probabilistic choice exploiting weights.

We consider probability distributions taken from the set  $Prob = \{\rho \mid \rho : Machine/\equiv \rightarrow [0, 1] \wedge supp(\rho) \text{ is finite} \wedge \sum_{[M] \in Machine/\equiv} \rho([M]) = 1\}$ , where  $supp(\rho) = \{[M] \mid \rho([M]) > 0\}$ .

The operational semantics is defined in terms of a probabilistic and timed transition system  $(Machine/\equiv, \rightarrow, \overset{\vee}{\rightarrow})$ , where:  $Machine/\equiv$  is the set of states;  $\rightarrow \subseteq Machine/\equiv \times Prob$  is the minimal relation satisfying the set of axioms and rules that are obtained from those included in Tables 1 and 2 by updating some rules as described in Table 4; and  $\overset{\vee}{\rightarrow}$  is the minimal relation satisfying the set of axioms and rules in Table 3.  $([M], \rho) \in \rightarrow$  (also denoted by  $[M] \rightarrow \rho$ ) means that a machine  $M$  can evolve (performing a single action) into a probability distribution  $\rho$  over machines, such that the machine  $M' \in Machine$  is reached with a probability equal to  $\rho([M'])$ . We use  $[M] \rightarrow [M']$  to denote  $[M] \rightarrow \rho$ , with  $\rho$  the trivial distribution which gives probability 1 to  $[M']$  and probability 0 to all other states. When evaluating the semantics of a machine  $M$ , we consider  $[M] \in Machine/\equiv$  to be the initial state.

Note that, a machine  $M$  can evolve into several probability distributions, i.e. it may be that  $[M] \rightarrow \rho$  for several different  $\rho$ . This means that (like in the simple model of [19]) whenever the system is in state  $[M]$ , first a non-deterministic choice is performed which decides which of the several probability distributions  $\rho$  must be considered, then the next state is probabilistically determined

(2')	$\frac{Eval(t) = t' \quad \exists e \in DS : e \triangleright t'}{[\langle in\ t(\vec{x}).P, DS \rangle_l] \longrightarrow \rho_{\langle in\ t(\vec{x}).P, DS \rangle_l}}$
(3')	$\frac{Eval(t) = t' \quad \exists e \in DS : e \triangleright t'}{[\langle rd\ t(\vec{x}).P, DS \rangle_l] \longrightarrow \rho_{\langle rd\ t(\vec{x}).P, DS \rangle_l}}$
(4')	$\frac{[\langle P, DS \rangle_l] \longrightarrow \rho}{[\langle P \mid Q, DS \rangle_l] \longrightarrow  Q(\rho)}$
(6')	$\frac{[\langle P, DS \rangle_l] \longrightarrow \rho}{[\langle P \leftarrow_n Q, DS \rangle_l] \longrightarrow \leftarrow_n Q(\rho)}$
(7')	$\frac{[M] \longrightarrow \rho}{[M \parallel N] \longrightarrow \ N(\rho)}$

Table 4  
Semantics of the calculus with weighted tuples.

by the chosen distribution  $\rho$ . Note that the non-deterministic choice may, e.g., arise from several concurrent *rd* operations which probabilistically retrieve data from the data space.

Table 5 defines: (i) The probability distributions  $\rho_{\langle in\ t(\vec{x}).P, DS \rangle_l}$  and  $\rho_{\langle rd\ t(\vec{x}).P, DS \rangle_l}$  used for *in* and *rd* operations, respectively. It is worth noting that  $\rho_{\langle in\ t(\vec{x}).P, DS \rangle_l}$  and  $\rho_{\langle rd\ t(\vec{x}).P, DS \rangle_l}$  are defined only if  $t \in Template$  and  $DS \in DSspace$  are such that there exists  $e \in DS : e \triangleright t$ , that is the condition reported in rules (2') and (3'). (ii) The operators  $|Q(\rho)$ ,  $\leftarrow_n Q(\rho)$  and  $\|N(\rho)$ , used in the rules (4'), (6') and (7'), that, given  $\rho$ , compute new probability distributions that account for parallel composition with process  $Q$ , composition with interrupting process  $Q$  and parallel composition with machine  $N$ , respectively. Such operators are expressed by defining three partial functions  $f$  of the form  $f : Machine/ \equiv \multimap Machine/ \equiv$  (where  $f$  stands for “ $|Q$ ”, “ $\leftarrow_n Q$ ” and “ $\|N$ ”, respectively) and by lifting them, with a general formula, to functions  $f : Prob \multimap Prob$  that turn probability distributions into probability distributions. The meaning of all the probability distributions defined in Table 5, that are used in rules of Table 4, is commented in the description of the operational semantics that follows.

$\rho_{\langle in\ t(\vec{x}).P, DS \rangle_l} ([M]) =$ $\begin{cases} \frac{W(e) \cdot DS(e)}{\sum_{e' \in DS: e' \triangleright t} W(e') \cdot DS(e')} & \text{if } M \equiv \langle P[\tilde{e}/\vec{x}], DS - e \rangle_l \\ & \text{with } e \in DS \wedge e \triangleright t \\ 0 & \text{o.w.} \end{cases}$ $\rho_{\langle rd\ t(\vec{x}).P, DS \rangle_l} ([M]) =$ $\begin{cases} \frac{\sum_{e' \in DS: e' \triangleright t \wedge P[\tilde{e}'/\vec{x}] \equiv P[\tilde{e}/\vec{x}]} W(e') \cdot DS(e')}{\sum_{e' \in DS: e' \triangleright t} W(e') \cdot DS(e')} & \text{if } M \equiv \langle P[\tilde{e}/\vec{x}], DS \rangle_l \\ & \text{with } e \in DS \wedge e \triangleright t \\ 0 & \text{o.w.} \end{cases}$ $f(\rho) ([M]) = \sum_{[M'] \in dom(f): f([M']) = [M]} \rho([M'])$ <p>where <math>f : Machine/\equiv \rightarrow Machine/\equiv</math>  can be “<math> Q</math>”, “<math>\leftarrow_n Q</math>” or “<math>\ N</math>”, defined by:</p> $ Q ( \quad \langle P, DS \rangle_l ) = \langle P Q, DS \rangle_l$ $\leftarrow_n Q ( \langle P, DS \rangle_l ) = \langle P \leftarrow_n Q, DS \rangle_l$ $\ N ( \quad [M] \quad ) = [M\ N]$
---

Table 5  
Probability distributions.

Rule (2') describes the behaviour of *in* operations; if a tuple  $e$  matching with the evaluation of template  $t$  is available in the  $DS$ , the *in* execution produces the removal from the space of  $e$  and then the process behaves as  $P[\tilde{e}/\vec{x}]$ . The probability of reaching a configuration where a matching tuple  $e$  contained in the  $DS$  is removed is the ratio of the total weight of the several instances of  $e$  in the  $DS$ , to the sum of the total weights of the several instances of the matching tuples currently available in the  $DS$ . In this way, the probability to reach a system configuration takes into account the multiple ways of removing  $e$  due to the several occurrences of  $e$  in the  $DS$ .

The rule (3') describes *rd* operations; if a tuple  $e$  matching with the evaluation of template  $t$  is available in the  $DS$ , then the process behaves as  $P[\tilde{e}/\vec{x}]$ . Differently from the previous axiom, *rd* operations do not modify the data space, i.e. in the reached states the configuration of  $DS$  is not changed, therefore they

are simply differentiated by the continuation  $P[\tilde{e}/\vec{x}]$  of the reading process. For example, let us consider a tuple space composed of two different tuples  $e = \langle d; d_c \rangle[w]$  and  $e' = \langle d'; d_c \rangle[w']$ . Let  $P = rd \langle null; null \rangle(x_1; x_2).out(\langle x_2 \rangle[w])$  be the process performing the read: it is not possible to discriminate between the selection of the two different tuples (not even from the continuations). Therefore, in general, the probability of reaching a configuration  $s$  that is obtained by reading a tuple  $e$  matching with (the evaluation of)  $t$  in the  $DS$  (yielding value  $e$ ) is the ratio of the sum of the total weights associated to the several instances of tuples  $e'$  matching with (the evaluation of)  $t$  in the  $DS$  such that the continuation of the reading process obtained by reading tuple  $e'$  is the same as the one obtained by reading  $e$ , to the sum of the total weights of the several instances of the matching tuples currently available in the  $DS$ .

Rule (4') defines the behaviour of the parallel composition of processes: if states reachable from  $[\langle P, DS \rangle_l]$  are described by the probability distribution  $\rho$ , and  $P$  performs an action in the system  $[\langle P \mid Q, DS \rangle_l]$  (the process that proceeds between  $P$  and  $Q$  is non-deterministically selected), then the reachable states are of the form  $[\langle P' \mid Q, DS' \rangle_l]$ , for some  $P' \in Process$  and  $DS' \in DSpace$ . The probability value of one such state  $[\langle P' \mid Q, DS' \rangle_l]$  does not depend on  $Q$  (that is “inactive”) and is equal to the summation of the probability values  $\rho([\langle P'', DS' \rangle_l])$  for all equivalence classes  $[\langle P'', DS' \rangle_l]$  (among which there is  $[\langle P', DS' \rangle_l]$ ) whose  $P''$  is such that  $[\langle P'' \mid Q, DS' \rangle_l] = [\langle P' \mid Q, DS' \rangle_l]$ <sup>1</sup>. Note that the summation above is necessary, i.e. there exists  $Q$  such that two different equivalence classes  $[\langle P', DS' \rangle_l]$  and  $[\langle P'', DS' \rangle_l]$  merge into a single one when the “ $|Q$ ” context is considered, i.e.  $[\langle P' \mid Q, DS' \rangle_l] = [\langle P'' \mid Q, DS' \rangle_l]$ . For instance, this happens if we take  $Q$  to be  $!P' \mid !P''$ .

The rules (6') and (7') define the behaviour of process interruption and parallel composition of machines in a similar way.

**Example 4.1 Service registry** - *A service description is a pair composed of two information: task and bind. The former describes the kind of task provided and executed by the service. The latter provides the information needed to connect and exploit the service functionality. A service registry can be seen as a machine whose data space is a repository of service descriptions. Two processes run on that machine that are responsible to provide to the client of the registry the publish and discovery operations. The former is used to add a service description in the repository, the latter is used to retrieve and available service description.*

*Following this approach, new services can be published sending a request to the registry. We assume that messages containing a publish request have the format  $\langle \text{PUBREQ}; \text{sndr}; \text{task}; \text{bind} \rangle$  where PUBREQ is a value that indicates that*

<sup>1</sup> This definition is correct because  $\equiv$  is a structural congruence relation, i.e. in the case of parallel of processes,  $P'' \equiv P'$  implies  $P'' \mid Q \equiv P' \mid Q$ .

the tuple is a publish request message, *sndr* is a location where an acknowledgement of the registration will be sent, and *task* and *bind* describes the service to be published. On the other hand, a message containing a discovery request must have the format  $\langle \text{DISREQ}; \text{sndr}; \text{task} \rangle$  where *DISREQ* is the value characterizing the tuples representing discovery requests, *sndr* is the location where the description of the retrieved service will be sent, and *task* describes the functionality expected from the retrieved service.

Namely, the initial state of the service registry is modeled by the machine

$$\langle !\text{Publish} \mid !\text{Discovery}, \emptyset \rangle_{\text{uddi}}$$

where the processes *Publish* and *Discovery* are defined as follows:

$$\begin{aligned} \text{Publish} &= \text{in } \langle \text{PUBREQ}; \text{null}; \text{null}; \text{null} \rangle(x; \text{sndr}; \text{task}; \text{bind}). \\ &\quad \text{out}(\langle \text{task}; \text{bind} \rangle[w(\text{sndr}, \text{task}, \text{bind})]). \text{out}(\langle \text{PUBACK}; \text{uddi} \rangle @ \text{sndr}) \\ \text{Discovery} &= \text{in } \langle \text{DISREQ}; \text{null}; \text{null} \rangle(x; \text{sndr}; \text{task}). \\ &\quad \text{rd } \langle \text{task}; \text{null} \rangle(x; \text{bind}). \text{out}(\langle \text{DISRES}; \text{uddi}; \text{task}; \text{bind} \rangle @ \text{sndr}) \end{aligned}$$

Note that the weight of the tuple that describes the service is computed according to an application dependent expression  $w(\text{sndr}, \text{task}, \text{bind})$ . Moreover, two extra special values *PUBACK* and *DISRES* are used to qualify messages representing publish acknowledgements and discovery responses, respectively.

**Example 4.2 Service publication and discovery** - A process willing to publish a service on a registry defined as in the Example 4.1 can be modeled by the following machine

$$\langle \text{DoPub}_{\text{task}, \text{bind}, \text{uddi}} \mid \text{SuccessPub}, \emptyset \rangle_{\text{user}}$$

where the processes  $\text{DoPub}_{\text{task}, \text{bind}, \text{uddi}}$  and *SuccessPub* are as follows: <sup>2</sup>

$$\begin{aligned} \text{DoPub}_{\text{task}, \text{bind}, \text{uddi}} &= \text{out}(\langle \text{PUBREQ}; \text{user}; \text{task}; \text{bind} \rangle @ \text{uddi}). \\ &\quad (\text{in } \langle \text{PUBACK}; \text{uddi} \rangle(x_1; x_2). \text{out}(\langle \text{SUCCESS} \rangle)) \leftarrow_n \text{Timeout} \\ \text{SuccessPub} &= \text{in}(\langle \text{SUCCESS} \rangle)(x). \text{Normal} \end{aligned}$$

Note that the process is not willing to wait indefinitely for the acknowledgement. The timed interruption operator is used to activate an alternative process *Timeout* responsible to manage situations in which the acknowledgement is not received in due time. On the contrary, if the acknowledgement is successfully received, the execution continues with the process *Normal*.

<sup>2</sup> We assume that “ $\leftarrow_n$ ” has the lowest precedence when writing process terms.

On the contrary, a process willing to discover a new service can be modeled by the following machine

$$\langle DoDis_{task,uddi} \mid SuccessDis, \emptyset \rangle_{user}$$

where the processes  $DoDis_{task,uddi}$  and  $SuccessDis$  are as follows:

$$\begin{aligned} DoDis_{task,uddi} &= out(\langle DISREQ; user; task \rangle @ uddi). \\ & \quad (in \langle DISRES; uddi; task; null \rangle (x_1; x_2; x_3; bind). \\ & \quad \quad out(\langle SUCCESS; bind \rangle) \leftarrow_n Timeout) \\ SuccessDis &= in(\langle SUCCESS; null \rangle (x; bind)).Normal \end{aligned}$$

Also in this case the process is not willing to wait indefinitely for the service binding.

**Example 4.3 Orchestrated discovery** - As a final example we present a process which needs to contact two different services in order to compose them. Consider, e.g. the organization of a travel that requires to book a flight as well as the hotel. Two different discoveries are needed to retrieve two specialized services, one for flight reservation, and another one for hotel reservation. In case one of the two discoveries fail, it is necessary to interrupt the whole orchestration as it is no more possible to complete it successfully.

In order to program such an orchestrator, we consider two processes, similar to  $DoDis_{task,uddi}$  defined in the previous example, that retrieve the flight reservation and hotel reservation services, respectively. The machine performing the orchestrated discovery is modeled by

$$\langle (FlightDis \mid HotelDis \mid SuccessCombined) \leftarrow_\infty Failure, \emptyset \rangle_{user}$$

and the processes  $FlightDis$ ,  $HotelDis$  and  $SuccessCombined$  by

$$\begin{aligned} FlightDis &= out(\langle DISREQ; user; FLIGHT \rangle @ uddi). \\ & \quad (in \langle DISRES; uddi; FLIGHT; null \rangle (x_1; x_2; x_3; bind). \\ & \quad \quad out(\langle FLIGHT; bind \rangle) \leftarrow_n abort) \\ HotelDis &= out(\langle DISREQ; user; HOTEL \rangle @ uddi). \\ & \quad (in \langle DISRES; uddi; HOTEL; null \rangle (x_1; x_2; x_3; bind). \\ & \quad \quad out(\langle HOTEL; bind \rangle) \leftarrow_n abort) \\ SuccessCombined &= in(\langle FLIGHT; null \rangle (x; bind1). \\ & \quad in(\langle HOTEL; null \rangle (y; bind2)).Normal \end{aligned}$$

where *Normal* denotes the behaviour in case of successful discovery, while *Failure* manages the case of failed discovery. The overall discovery fails when one of the two discoveries does not complete in due time. Indeed, in this case the abort process is activated as timed failure handler. The effect of the abort process is to interrupt the overall orchestrated discovery and to activate the *Failure* process.

## 5 Conclusion

In this paper we have discussed some of the interesting issues raised by the combination of three relevant characteristics of service oriented computing: loose coupling, communication latency and open-endedness. In particular, we have modeled basic mechanisms for programming timed loosely coupled transactions and for managing dynamic service retrieval according to quantitative information (the weights) that permits to model run time features of services that are not known at design time.

Other interesting aspects raise, for instance, from the combination of loosely coupled transactions and open endedness. In service oriented computing there is a great interest in transaction protocols. These are used, on the one hand, to select the partners involved in a transaction according to some minimal service requirements and, on the other hand, to ensure that such requirements are guaranteed during the overall service collaboration. Process algebras have been already used to model specific protocols used in the context of distributed commitment. The two-phase commit protocol guaranteeing atomicity is analysed in [2], while the BTP protocol guaranteeing a relaxed notion of partial atomicity –called cohesion– was investigated in [3]. These are only specific cases of protocols; a formal investigation of other protocols such as those based on quality of services, or supporting the dynamic redefinition of the involved partners during the execution of the transaction, is still lacking.

## References

- [1] J.C.M. Baeten and C.A. Middelburg. *Process Algebra with Timing*. EATCS Monograph, Springer Verlag 2002.
- [2] M. Berger and K. Honda. The Two-Phase Commit Protocol in an Extended  $\pi$ -Calculus. In *EXPRESS'00, Proc. of the 7th International Workshop on Expressiveness in Concurrency*, volume 39 of *ENTCS*, 2000.
- [3] L. Bocchi. Compositional Nested Long Running Transactions. In *FASE'04, Proc. of the Fundamental Approaches to Software Engineering*, volume 2984 of *LNCS*,

pages 194–208. 2004.

- [4] L. Bocchi, C. Laneve, and G. Zavattaro. A calculus for long running transactions. In *FMOODS'03, Proc. of the 6th IFIP International Conference on Formal Methods for Open Object-based Distributed Systems*, volume 2884 of *LNCS*, pages 124–138. Springer-Verlag, 2003.
- [5] M. Bravetti, R. Gorrieri, R. Lucchi, G. Zavattaro. “Quantitative Information in the Tuple Space Coordination Model”, *Theoretical Computer Science*, 346:1, pages 28-57, Elsevier, 2005.
- [6] M. Bravetti, R. Gorrieri, “A Complete Axiomatization for Observational Congruence of Prioritized Finite-State Behaviors”, in Proc. of the *27th Int. Colloquium on Automata, Languages and Programming (ICALP 2000)*, U. Montanari, J.D.P. Rolim and E. Welzl editors, *LNCS* 1853:744-755, Geneva (Switzerland), July 2000
- [7] M. Bravetti, R. Lucchi, G. Zavattaro, R. Gorrieri. “Web Services for E-commerce: guaranteeing security access and quality of service”, in Proc. of the *19th ACM Symposium on Applied Computing (SAC'04), special track on E-Commerce Technologies*, ACM Press, pp. 800-806, Nicosia (Cyprus), March 2004
- [8] R. Bruni, M. Butler, C. Ferreira, T. Hoare, H. Melgratti, and U. Montanari. Reconciling two approaches to compensable flow composition. In *CONCUR'05, Proc. of the 16th International Conference on Concurrency Theory*, volume to appear of *LNCS*, 2005.
- [9] R. Bruni, H. Melgratti, and U. Montanari. Theoretical foundations for compensations in flow composition languages. In *POPL'05, Proc. of the 32nd Annual ACM SIGPLAN - SIGACT Symposium on Principles of Programming Languages*, 2005
- [10] M. Butler and C. Ferreira. An operational semantics for StAC, a language for modelling long-running business transactions. In *COORDINATION'04, Proc. of the 6th International Conference on Coordination Models and Languages*, volume 2949 of *LNCS*, pages 87–104. Springer-Verlag, 2004.
- [11] M. Butler, T. Hoare, and C. Ferreira. A trace semantics for long-running transactions. In *Proceedings of 25 Years of CSP*, London, 2004.
- [12] G. Castagna, R. De Nicola, and D. Varacca. Semantic subtyping for the pi-calculus. In *LICS'05, Proc. of Logic in Computer Science*, june 2005.
- [13] D. Gelernter. Generative Communication in Linda. *ACM Transactions on Programming Languages and Systems*, 7(1):80–112, 1985.
- [14] A. Igarashi, N. Kobayashi. A Generic Type System for the Pi-Calculus. In *Theoretical Computer Science*, 311(1-3):121-163, 2004.
- [15] C. Laneve, and G. Zavattaro. Foundations of Web Transactions. In *FOSSACS'05, Proc. of the 8th Foundations of Software Science and Computational Structures*, volume 3441 of *LNCS*, pages 282–298. Springer-Verlag, 2005.

- [16] Microsoft BizTalk Server. [<http://www.microsoft.com/biztalk/default.asp>], Microsoft Corporation.
- [17] Microsoft Corporation. *XLANG: Web Services for Business Process Design*. [http://www.gotdotnet.com/team/xml\\_wsspecs/xlang-c/default.htm](http://www.gotdotnet.com/team/xml_wsspecs/xlang-c/default.htm).
- [18] OASIS. *Web Services Business Process Execution Language Version 2.0, Working Draft*. [<http://www.oasis-open.org/committees/download.php/10347/wsbpel-specification-draft-120204.htm>].
- [19] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1995.
- [20] Universal Description, Discovery and Integration for Web Services (UDDI) V3 Specification.
- [21] W3C. *Web Services Choreography Description Language Version 1.0*. <http://www.w3.org/TR/2004/WD-ws-cdl-10-20041217>.