

A Formal Account of WS-BPEL[★]

Alessandro Lapadula, Rosario Pugliese, and Francesco Tiezzi

Dipartimento di Sistemi e Informatica Università degli Studi di Firenze

Abstract. We introduce *Blite*, a lightweight language for web services orchestration designed around some of WS-BPEL peculiar features like partner links, process termination, message correlation, long-running business transactions and compensation handlers. *Blite* formal presentation helps clarifying some ambiguous aspects of the WS-BPEL specification, which have led to engines implementing different semantics and, thus, have undermined portability of WS-BPEL programs over different platforms. We illustrate the main features of *Blite* by means of many examples, some of which are also exploited to test and compare the behaviour of three of the most known free WS-BPEL engines.

1 Introduction

There is an ever increasing acceptance of WS-BPEL (Web Services Business Process Execution Language, [23]) as a standard language for service composition within and across multiple enterprises. The fact that it has become an OASIS standard, however, has not solved all the difficulties of using the language. Indeed, WS-BPEL comes without a formal semantics and its specification document [23], written in ‘natural’ language, contains a fair number of acknowledged ambiguous aspects that may lead to different interpretations. For example, the relationship between WS-BPEL (multiple) *start activities* and the mechanisms handling race conditions has not been fully explored; moreover, if suitable measures for ‘protecting’ such critical activities as fault and compensation handlers are not taken into account, then subtle behaviours can arise when implementing activities that cause immediate termination of other activities.

The design of WS-BPEL applications is difficult and error-prone also due to the presence of such intricate features as concurrency and race conditions, forced termination, multiple instances and message correlation, long-running business transactions and compensation handlers. It would thus benefit from the use of *formal methods* because these can provide a framework to precisely describe some aspects of an application, to state and prove its properties, and to direct attention towards issues that might otherwise be overlooked.

As a step in this direction, in this paper we introduce *Blite*, a ‘lightweight’ variant of WS-BPEL designed around the above mentioned features. *Blite*, being obtained by dropping redundant features from the full-fledged language, permits to focus on those fragments of the design that are more challenging and need more attention. For example, *Blite* clarifies the relationship between compensation activities and the control flow of the originating process, and illustrates the mechanisms for service instance creation and

[★] This work has been supported by the EU project SENSORIA, IST-2005-016004.

identification, and their interplay. Our study can also contribute to the many discussions on compensation and correlation which have been reported by the WS-BPEL technical committee [22] (see, e.g., discussions related to issues 66, 207 and 271).

Moreover, *Blite*'s formal presentation can help clarifying many ambiguous aspects of the WS-BPEL specification and, thus, can be used prescriptively to drive implementations of future WS-BPEL engines. In fact, by means of several examples, we test and compare three of the most known free BPEL engines, namely ActiveBPEL [1], Apache ODE [2] and Oracle BPEL Process Manager [3]. As a matter of fact, the considered engines exhibit quite different behaviours and diverge from the WS-BPEL specification in many important aspects. This is complicating the task of developing WS-BPEL applications and undermining their portability across different platforms.

We also believe that the formalization of WS-BPEL's operational semantics, through the introduction of *Blite*, can also enable tailoring proof techniques and analytical tools typical of process calculi to the needs of WS-BPEL applications. Indeed, on the one hand, alike other standards enabling the web services technology, WS-BPEL does not provide support for guided forms of application development and analysis. On the other hand, it has been shown that type systems, model checking and (bi)simulation analysis provide adequate tools to address topics relevant to the web services technology (see e.g. [21,26]). In the end, this 'proof technology' can pave the way for the development of (semi-)automatic property validation tools.

The rest of the paper is organized as follows. Section 2 presents *Blite*'s syntax and operational semantics. Section 3 illustrates most of the language features at work on modelling a shipping service scenario borrowed from the official WS-BPEL specification. Section 4 presents many peculiar examples and the results of our experimentation with the three WS-BPEL engines mentioned above. Section 5 touches upon more closely related work and directions for future work.

2 *Blite*: A 'Lightweight' Variant of WS-BPEL

The language *Blite*¹ is a simplification of WS-BPEL designed around some of its peculiar features like partner links, process termination, message correlation, long-running business transactions and compensation handlers. *Blite* is the result of the usual tension between handiness and expressiveness. Therefore, to keep the design of the language manageable, we intentionally left out other important aspects, including timeouts, event and termination handlers, flow graphs, and sophisticated forms of data handling.

Blite provides a formal description of service deployments by only retaining relevant implementation details such as partner links, service definitions and correlation sets. For example, the roles played by service partners in a service interaction are explicitly indicated through *partner links* and *partners*, while such aspects as physical *service binding* described in associated WSDL documents are abstracted away. In request-response interactions, for example, partner links indicate two partners because the requesting partner must provide a callback operation used by the receiving partner to send notifications.

¹ We refer the interested reader to [18] for a deeper presentation of which aspects of WS-BPEL are supported by *Blite* and their mapping.

Table 1. Syntax of *Blite*

<i>Basic activities</i>	$b ::= \text{inv } \ell^i \circ \bar{x} \mid \text{rcv } \ell^r \circ \bar{x} \mid x := e$ empty throw exit	invoke, receive, assign empty, throw, exit
<i>Structured activities</i>	$a ::= b \mid \text{if}(x)\{a_1\}\{a_2\} \mid \text{while}(x)\{a\}$ $a_1 ; a_2$ $\sum_{j \in J} \text{rcv } \ell_j^r \circ_j \bar{x}_j ; a_j$ $a_1 \mid a_2$ $[a \bullet a_f \star a_c]$	basic, conditional, iteration sequence, pick parallel, scope
<i>Start activities</i>	$r ::= \text{rcv } \ell^r \circ \bar{x} \mid \sum_{j \in J} \text{rcv } \ell_j^r \circ_j \bar{x}_j ; a_j$ $r ; a$ $r_1 \mid r_2$ $[r \bullet a_f \star a_c]$	receive, pick sequence, parallel, scope
<i>Services</i>	$s ::= [r \bullet a_f]$ $\mu \vdash a$ $\mu \vdash a, s$	definition, instance, multiset
<i>Deployments</i>	$d ::= \{s\}_c$ $d_1 \parallel d_2$	deployment, composition

Instead, in one-way interactions a partner link indicates a single partner because one of the parties provides all the invoked operations. Besides asynchronous invocation, WS-BPEL also provides a construct for synchronous invocation of remote services. This construct forces the invoker to wait for an answer by the invoked service, that indeed performs a pair of operations *receive-reply*. In *Blite*, this behaviour is rendered in terms of a pair of activities *invoke-receive* executed by the invoker and a pair of activities *receive-invoke* executed by the invoked service.

An important aspect is that, in general, the information provided by partner links is not enough to deliver messages to a service. Indeed, since services are instantiated to serve the received requests, messages need to be delivered not only to the correct partner, but also to the correct instance of the service that the partner provides. To achieve this, WS-BPEL relies on the business data exchanged rather than on specific mechanisms, such as *WS-Addressing* [9] or low-level correlation methods based on SOAP headers. Specifically, *Blite* exploits *correlation variables* that permit to declare the parts of a message that are instance dependent, e.g. *order number* or *client id*, so that a message can be routed to the correct service instance on the basis of the values of the correlation variables it provides, independently of any routing mechanism.

Syntax. The syntax of *Blite* is given in Table 1. Services are *structured activities* built from *basic activities* by exploiting operators for conditional choice $\text{if}(\cdot)\{\cdot\}\{\cdot\}$, iteration $\text{while}(\cdot)\{\cdot\}$, sequential composition $\cdot ; \cdot$, pick $\sum_{j \in J} \text{rcv } \cdot \cdot ; \cdot$ (i.e., external choice with the constraint that $|J| > 1$), parallel composition $\cdot \mid \cdot$ and scope $[\cdot \bullet \star \cdot]$. A scope activity $[a \bullet a_f \star a_c]$ groups a primary activity a together with a fault handling activity a_f and a compensation activity a_c . *Start activities* r are structured activities that initially can only execute receive activities.

In the sequel, we shall use $\cdot + \cdot$ to abbreviate binary external choice. We let sequence have higher priority (i.e. bind more tightly) than parallel composition and external choice, i.e. $a_1 ; a_2 \mid a_3 ; a_4$ stands for $(a_1 ; a_2) \mid (a_3 ; a_4)$ and $a_1 ; a_2 + a_3$ stands for $(a_1 ; a_2) + a_3$. Moreover, we adopt the convention that fault and compensation activities may be omitted from a scope construct, in which case they are intended to be throw and empty, respectively.

Data can be shared among different activities through *shared variables* (ranged over by x, x', \dots). The set of manipulable values (ranged over by v, v', \dots) is left unspecified; however, we assume that it includes the set of *partner names* (ranged over by p, q, \dots) and the set of *operation names* (ranged over by o, o', \dots). We use u to range over partners and variables and w to range over values and variables. *Expressions* (ranged over by e, e', \dots) are left unspecified but contain, at least, values and variables.

Notation $\bar{\tau}$ stands for tuples of objects, e.g. \bar{x} is a compact notation for denoting the tuple of variables $\langle x_1, \dots, x_h \rangle$ (with $h \geq 0$). We assume that variables in the same tuple are pairwise distinct. The special notation $\bar{\tau}$ stands for tuples of one or two objects, e.g. \bar{p} denotes either $\langle p_1, p_2 \rangle$ or $\langle p_1 \rangle$. Tuples can be constructed using a concatenation operator $\cdot : \cdot$, i.e. $\langle p, u \rangle : \langle x_1, \dots, x_h \rangle$ returns $\langle p, u, x_1, \dots, x_h \rangle$. We will write $Z \triangleq W$ to assign a symbolic name Z to the term W .

Partner links ℓ^r of receive activities can be either $\langle p, u \rangle$ or $\langle p \rangle$, where p is the partner providing the operation and u is a partner or variable used to send messages in reply. Indeed, service partners used for receiving messages must be known at design-time, while the partners used to send messages in reply may be dynamically determined. Partner links ℓ^i within invoke activities can be either $\langle u, p \rangle$ or $\langle u \rangle$, where u is the partner providing the operation and, possibly, p is a partner used to receive messages in reply. Like before, this latter partner must be statically known, thus it cannot be a variable.

Deployments are finite compositions of multisets of service *instances* $\mu \vdash a$, containing at most one service *definition* $[r \bullet a_f]$ and associated to a *correlation set* c , namely a (possibly empty) set of *correlation variables*. A service definition provides a ‘top-level’ scope (i.e. a scope that cannot be compensated) and offers a choice of alternative receives among multiple start activities. Each service instance $\mu \vdash a$ has its own (private) state μ . States are (partial) functions mapping variables to values and are written as collections of pairs of the form $\{x \mapsto v\}$. The state obtained by updating μ with μ' , written as $\mu \circ \mu'$, is inductively defined by: $\mu \circ \mu'(x) = \mu'(x)$ if $x \in \text{dom}(\mu')$ (where $\text{dom}(\mu)$ denotes the domain of μ) and $\mu(x)$ otherwise. The empty state is denoted by \emptyset . In the sequel, we will only consider *well-formed* deployments, i.e. compositions where the sets of partners used for handling requests within different deployments are pairwise disjoint. The rationale is that each service definition has its own partner names and all its instances run within the same deployment where the definition resides.

Operational Semantics. The semantics is defined over an enriched syntax that also includes *protected activities* $\langle a \rangle$, *unsuccessful termination* stop , *messages* $\langle \bar{p} : o : \bar{v} \rangle$ and *scopes* of the form $[a \bullet a_f \star a_c \Delta a_d]$. The first three ‘auxiliary’ activities are used to replace, respectively, unsuccessfully completed scopes (with their protected default compensation), compulsorily or faultily terminated services (with stop), and invoke activities (with the message they produced). Instead, such scopes as $[a \bullet a_f \star a_c \Delta a_d]$ are dynamically generated to store in a_d the compensation activities of the immediately enclosed scopes that have successfully completed, together with the order in which they must be executed. In the sequel, *empty*, *exit*, *throw*, *stop* and *messages* will be called *short-lived* activities and will be generically indicated by sh .

The operational semantics of *Blite* deployments is defined in terms of a structural congruence and a reduction relation. The *structural congruence*, written \equiv , identifies syntactically different terms which intuitively represent the same term. It is defined as

Table 2. Structural congruence for *Blite* activities and deployments

$a \mid \text{empty} \equiv a$	$\text{empty}; a \equiv a; \text{empty} \equiv a$	$\text{stop} \mid \text{stop} \equiv \text{stop}$	$\text{stop}; a \equiv \text{stop}$
$\langle\langle a \rangle\rangle \equiv \langle a \rangle$	$\langle\langle \text{sh} \rangle\rangle \equiv \text{sh}$	$\langle\langle \tilde{p}:o:\tilde{v} \rangle\rangle \mid a \equiv \langle\langle \tilde{p}:o:\tilde{v} \rangle\rangle \mid \langle a \rangle$	
$[a \bullet a_f \star a_c] \equiv [a \bullet a_f \star a_c \triangle \text{empty}]$		$\langle\langle \tilde{p}:o:\tilde{v} \rangle\rangle \mid a_1; a_2 \equiv \langle\langle \tilde{p}:o:\tilde{v} \rangle\rangle \mid (a_1; a_2)$	
$\langle\langle \tilde{p}:o:\tilde{v} \rangle\rangle \mid a \bullet a_f \star a_c \triangle a_d \equiv \langle\langle \tilde{p}:o:\tilde{v} \rangle\rangle \mid [a \bullet a_f \star a_c \triangle a_d]$ if $\neg a \Downarrow_{\text{throw}}$			
$\frac{a \equiv a' \quad a_f \equiv a'_f \quad a_c \equiv a'_c \quad a_d \equiv a'_d}{[a \bullet a_f \star a_c \triangle a_d] \equiv [a' \bullet a'_f \star a'_c \triangle a'_d]}$			
$\frac{r \equiv r' \quad a_f \equiv a'_f}{\{[r \bullet a_f], s\}_c \equiv \{s, [r' \bullet a'_f]\}_c}$		$\frac{a \equiv a'}{\{\mu \vdash a, s\}_c \equiv \{s, \mu \vdash a'\}_c}$	
$d_1 \parallel d_2 \equiv d_2 \parallel d_1$	$(d_1 \parallel d_2) \parallel d_3 \equiv d_1 \parallel (d_2 \parallel d_3)$	$\{\mu \vdash \text{empty}, s\}_c \equiv \{s\}_c$	
$\{\mu \vdash \text{stop}, s\}_c \equiv \{s\}_c$		$\{\mu \vdash \text{empty}\}_c \parallel d \equiv d$	$\{\mu \vdash \text{stop}\}_c \parallel d \equiv d$

the least congruence relation induced by a given set of equational laws. In Table 2, we explicitly show, in the upper part, the laws for **empty**, **stop**, protected activities, messages and scopes, and, in the lower part, the laws for services and deployments. Standard laws stating, e.g., that sequence is associative, parallel composition is commutative and associative, are omitted. A few observations on the structural laws are in order. Activity **empty** acts as the identity element both for sequence and parallel composition. Multiple **stop** in parallel are equivalent to just one **stop**, moreover **stop** disables subsequent activities. The protection operator is idempotent, and short-lived activities are implicitly protected, thus messages can go in/out of the scope of a protection operator. Default compensation is initially **empty**. Messages do not block subsequent activities and scope completion, except when **throw** is active in the scope (this is checked by predicate $\cdot \Downarrow_{\text{throw}}$ that will be explained later on). Structural congruence is extended to scopes, instances and deployments in the obvious way. Moreover, the order in which definition and instances occur within a deployment does not matter, and deployment composition is commutative and associative. Instances like $\mu \vdash \text{empty}$ and $\mu \vdash \text{stop}$ are terminated and, thus, can be removed. Similarly, deployments only containing terminated instances are terminated too and can be removed.

The *reduction relation* over deployments, written $\succ\!\!\rightarrow$, exploits a labelled transition relation over structured activities, written $\xrightarrow{\alpha}$, where α is generated by the grammar:

$$\alpha ::= \tau \mid \mathbf{x} \leftarrow \mathbf{v} \mid !\tilde{p}:o:\tilde{v} \mid ?\ell^r:o:\tilde{x} \mid \boxtimes \mid \uparrow \mid (a)$$

The meaning of labels is as follows: τ indicates message productions, guard evaluations for conditional and iteration or installation/activation of compensations; $\mathbf{x} \leftarrow \mathbf{v}$ indicates assignment of value \mathbf{v} to variable \mathbf{x} ; $!\tilde{p}:o:\tilde{v}$ and $?\ell^r:o:\tilde{x}$ indicate execution of invoke and receive activities for operation o , where \tilde{p} and \tilde{v} match with ℓ^r and \tilde{x} , respectively; \boxtimes indicates forced termination of a service instance; \uparrow indicates production of a fault

Table 3. Basic, auxiliary and structured activities

$\mu \vdash \text{inv } \ell^i \circ \bar{x} \xrightarrow{\tau} \ll \mu(\ell^i) : \circ : \mu(\bar{x}) \gg$ (inv)	$\text{rcv } \ell^r \circ \bar{x} \xrightarrow{? \ell^r : \circ : \bar{x}} \text{empty}$ (rec)
$\mu \vdash x := e \xrightarrow{x \leftarrow \mu(e)} \text{empty}$ (asg)	$\text{throw} \xrightarrow{\uparrow} \text{stop}$ (thr)
$\text{exit} \xrightarrow{\boxplus} \text{stop}$ (term)	$\ll \bar{p} : \circ : \bar{v} \gg \xrightarrow{! \bar{p} : \circ : \bar{v}} \text{empty}$ (msg)
$\frac{\mu \vdash a \xrightarrow{\alpha} a'}{\mu \vdash (a) \xrightarrow{\alpha} (a')}$ (prot)	$\frac{\mu \vdash a \xrightarrow{\alpha} a'}{\mu \vdash (a) \xrightarrow{\alpha} (a')}$ (prot)
$\frac{\mu \vdash a_1 \xrightarrow{\alpha} a'_1}{\mu \vdash a_1 ; a_2 \xrightarrow{\alpha} a'_1 ; a_2}$ (seq)	$\sum_{j \in J} \text{rcv } \ell_j^r \circ_j \bar{x}_j ; a_j \xrightarrow{? \ell_h^r : \circ_h : \bar{x}_h} a_h, h \in J$ (pick)
$a = \begin{cases} a_1 & \text{if } \mu(x) = \text{tt} \\ a_2 & \text{if } \mu(x) = \text{ff} \end{cases}$ (if)	$a' = \begin{cases} a ; \text{while}(x) \{a\} & \text{if } \mu(x) = \text{tt} \\ \text{empty} & \text{if } \mu(x) = \text{ff} \end{cases}$ (while)
$\frac{\mu \vdash \text{if}(x)\{a_1\}\{a_2\} \xrightarrow{\tau} a}{\mu \vdash a_1 \xrightarrow{\alpha} a'_1 \quad \alpha \notin \{\boxplus, \uparrow\} \quad \neg(a_2 \Downarrow_{\text{throw}} \vee a_2 \Downarrow_{\text{exit}})}$ (par ₁)	$\frac{a_1 \xrightarrow{\alpha} a'_1 \quad \alpha \in \{\boxplus, \uparrow\}}{a_1 a_2 \xrightarrow{\alpha} a'_1 \text{end}(a_2)}$ (par ₂)
$\frac{\mu \vdash a_1 \xrightarrow{\alpha} a'_1 \quad \alpha \notin \{\boxplus, \uparrow\} \quad \neg(a_2 \Downarrow_{\text{throw}} \vee a_2 \Downarrow_{\text{exit}})}{\mu \vdash a_1 a_2 \xrightarrow{\alpha} a'_1 a_2}$ (par ₁)	$\frac{a_1 \xrightarrow{\alpha} a'_1 \quad \alpha \in \{\boxplus, \uparrow\}}{a_1 a_2 \xrightarrow{\alpha} a'_1 \text{end}(a_2)}$ (par ₂)
$[\text{empty} \bullet a_f \star a_c \Delta a_d] \xrightarrow{(a_c)} \text{empty}$ (done ₁)	$[\text{stop} \bullet a_f \star a_c \Delta a_d] \xrightarrow{\tau} (a_d ; a_f)$ (done ₂)
$\frac{\mu \vdash a \xrightarrow{\alpha} a' \quad \alpha \notin \{\uparrow, (a'')\}}{\mu \vdash [a \bullet a_f \star a_c \Delta a_d] \xrightarrow{\alpha} [a' \bullet a_f \star a_c \Delta a_d]}$ (exec)	$\frac{a \xrightarrow{(a'')}}{[a \bullet a_f \star a_c \Delta a_d] \xrightarrow{\tau} [a' \bullet a_f \star a_c \Delta a'' ; a_d]}$ (done ₃)
$\frac{a \xrightarrow{\uparrow} a'}{[a \bullet a_f \star a_c \Delta a_d] \xrightarrow{\tau} [a' \bullet a_f \star a_c \Delta a_d]}$ (fault)	$\frac{a \xrightarrow{\uparrow} a'}{[a \bullet a_f \star a_c \Delta a_d] \xrightarrow{\tau} [a' \bullet a_f \star a_c \Delta a_d]}$ (fault)

signal from inside a scope; (a) indicates successful completion of a scope that can be compensated by the structured activity a .

The relation $\xrightarrow{\alpha}$ is defined by the rules in Table 3 with respect to a state μ , that is omitted when unnecessary (writing $a \xrightarrow{\alpha} a'$ instead of $\mu \vdash a \xrightarrow{\alpha} a'$). Before commenting the rules, we introduce the auxiliary functions and predicates they exploit. Specifically, the predicates $a \Downarrow_{\text{exit}}$ and $a \Downarrow_{\text{throw}}$ check the ability of a of performing `exit` or `throw`, respectively. They are defined inductively on the syntax of activities and act as an homomorphism in all cases, but for conditional choice and iteration for which they hold false, and for the following cases

$$\text{exit} \Downarrow_{\text{exit}} \quad \text{throw} \Downarrow_{\text{throw}} \quad \frac{a_1 \Downarrow_{\text{exit}}}{a_1 ; a_2 \Downarrow_{\text{exit}}} \quad \frac{a_1 \Downarrow_{\text{throw}}}{a_1 ; a_2 \Downarrow_{\text{throw}}} \quad \frac{a \Downarrow_{\text{exit}}}{[a \bullet a_f \star a_c \Delta a_d] \Downarrow_{\text{exit}}}$$

The function $\text{end}(\cdot)$, given an activity a , returns the activity obtained by only retaining short-lived and protected activities inside a . It is defined inductively on the syntax of activities, the most significant cases being

Table 4. Matching rules / Is there an active receive along \tilde{p} and o matching \tilde{v} ?

$\text{match}(\mathbf{c}, \mu, \mathbf{x}, \mathbf{v}) = \begin{cases} \{\mathbf{x} \mapsto \mathbf{v}\} & \text{if } \mathbf{x} \notin \mathbf{c} \vee (\mathbf{x} \in \mathbf{c} \wedge \mathbf{x} \notin \text{dom}(\mu)) \\ \emptyset & \text{if } \mathbf{x} \in \mathbf{c} \wedge \{\mathbf{x} \mapsto \mathbf{v}\} \in \mu \end{cases}$		
$\text{match}(\mathbf{c}, \mu, \mathbf{v}, \mathbf{v}) = \emptyset \quad \frac{\text{match}(\mathbf{c}, \mu, \mathbf{w}_1, \mathbf{v}_1) = \mu' \quad \text{match}(\mathbf{c}, \mu, \bar{\mathbf{w}}_2, \bar{\mathbf{v}}_2) = \mu''}{\text{match}(\mathbf{c}, \mu, (\mathbf{w}_1, \bar{\mathbf{w}}_2), (\mathbf{v}_1, \bar{\mathbf{v}}_2)) = \mu' \circ \mu''}$		
$\frac{ \text{match}(\mathbf{c}, \mu, \ell^r : \mathbf{o} : \bar{\mathbf{x}}, \tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}) < n}{\mu \vdash \text{rcv } \ell^r \circ \bar{\mathbf{x}} ; \mathbf{a} \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}$	$\frac{\exists h \in J. \text{match}(\mathbf{c}, \mu, \ell_h^r : \mathbf{o}_h : \bar{\mathbf{x}}_h, \tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}) < n}{\mu \vdash \sum_{j \in J} \text{rcv } \ell_j^r \circ_j \bar{\mathbf{x}}_j ; \mathbf{a}_j \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}$	
$\frac{\mu \vdash \mathbf{a}_1 \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}{\mu \vdash \mathbf{a}_1 ; \mathbf{a}_2 \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}$	$\frac{\mu \vdash \mathbf{a}_1 \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n} \quad \vee \quad \mu \vdash \mathbf{a}_2 \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}{\mu \vdash \mathbf{a}_1 \mid \mathbf{a}_2 \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}$	
$\frac{\mu \vdash \mathbf{a} \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}{\mu \vdash (\mathbf{a}) \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}$	$\frac{\mu \vdash \mathbf{a} \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}{\mu \vdash [\mathbf{a} \bullet \mathbf{a}_f \star \mathbf{a}_c \Delta \mathbf{a}_d] \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}$	$\frac{\mu \vdash \mathbf{a} \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n} \quad \vee \quad \mathbf{s} \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}{\mu \vdash \mathbf{a}, \mathbf{s} \Downarrow_{\tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}}}^{c,n}}$

$$\begin{aligned} \text{end}(\text{sh}) &= \text{sh} & \text{end}([\mathbf{a}]) &= ([\mathbf{a}]) & \text{end}(\mathbf{a}_1 ; \mathbf{a}_2) &= \text{end}(\mathbf{a}_1) \\ \text{end}([\mathbf{a} \bullet \mathbf{a}_f \star \mathbf{a}_c \Delta \mathbf{a}_d]) &= [\text{end}(\mathbf{a}) \bullet \mathbf{a}_f \star \mathbf{a}_c \Delta \mathbf{a}_d] \end{aligned}$$

where \mathbf{a}_1 may not be congruent to `empty` or to $\ll \tilde{\mathbf{p}} : \mathbf{o} : \bar{\mathbf{v}} \gg$, or to parallel compositions of them. In the remaining cases, `end`(\cdot) returns `stop`, except for parallel composition for which it acts as an homomorphism.

We now briefly comment on the rules in Table 3. Rules (`inv`) and (`asg`) state that invoke and assign activities can proceed only if their arguments are *closed* expressions (i.e. expressions without uninitialized variables) and can be evaluated (i.e. $\mu(\cdot)$ returns a value). By rule (`rec`), a receive activity offers an invocable operation along a given partner link. Rules (`thr`) and (`term`) report production of fault and forced termination signals, respectively. Auxiliary activities behave as expected: a message can always be delivered (rule (`msgg`)) and the protected activity (`(a)`) behaves like `a` (rule (`prot`)). Rule (`seq`) takes care of activities executed sequentially, while rule (`pick`) permits to choose among alternative receive activities. Rules for conditional choice and iteration (`(if)` and (`while`), resp.) are standard. Execution of parallel activities is interleaved (rules (`par1`) and (`par2`)), except when a terminate/fault activity can be executed (rule (`par3`)), in which case all parallel activities must immediately terminate except for short-lived activities and protected fault/compensation handlers. In other words, termination activities `throw` and `exit` are executed eagerly.

By rules (`done1`) and (`done3`), scope completions can be compensated according to the WS-BPEL *default* compensation behaviour (i.e. in the reverse order of completion) by the immediately enclosing scope. Notably, scopes like `[empty \bullet $\mathbf{a}_f \star \mathbf{a}_c \Delta \mathbf{a}_d$]` have not completed yet and when a scope completes, the default compensation \mathbf{a}_d of inner scopes is not passed to the enclosing scope (rule (`done1`)). Rule (`exec`) permits to

Table 5. Reduction rules for *Blite* deployments (where $t_1 = \ell^r : o : \bar{x}$ and $t_2 = \tilde{p} : o : \bar{v}$)

$\frac{a_1 \xrightarrow{?t_1} a'_1 \quad a_2 \xrightarrow{!t_2} a'_2 \quad \text{match}(c_1, \mu_1, t_1, t_2) = \mu'_1 \quad \neg(\mu_1 \vdash a_1, s_1 \Downarrow_{t_2}^{c_1, \mu'_1})}{\{\mu_1 \vdash a_1, s_1\}_{c_1} \parallel \{\mu_2 \vdash a_2, s_2\}_{c_2} \succ \{\mu_1 \circ \mu'_1 \vdash a'_1, s_1\}_{c_1} \parallel \{\mu_2 \vdash a'_2, s_2\}_{c_2}} \quad (\text{com})$
$\frac{[r \bullet a_f \star \text{empty}] \xrightarrow{?t_1} a_1 \quad a_2 \xrightarrow{!t_2} a'_2 \quad \text{match}(c_1, \emptyset, t_1, t_2) = \mu_1 \quad \neg(s_1 \Downarrow_{t_2}^{c_1, \mu_1})}{\{[r \bullet a_f], s_1\}_{c_1} \parallel \{\mu_2 \vdash a_2, s_2\}_{c_2} \succ \{\mu_1 \vdash a_1, [r \bullet a_f], s_1\}_{c_1} \parallel \{\mu_2 \vdash a'_2, s_2\}_{c_2}} \quad (\text{new})$
$\frac{\mu \vdash a \xrightarrow{x \leftarrow v} a' \quad \text{match}(c, \mu, x, v) = \mu'}{\{\mu \vdash a, s\}_c \succ \{\mu \circ \mu' \vdash a', s\}_c} \quad (\text{var}) \quad \frac{d_1 \succ d'_1}{d_1 \parallel d_2 \succ d'_1 \parallel d_2} \quad (\text{part})$
$\frac{\mu \vdash a \xrightarrow{\alpha} a' \quad \alpha \notin \{?t_1, !t_2, x \leftarrow v\}}{\{\mu \vdash a, s\}_c \succ \{\mu \vdash a', s\}_c} \quad (\text{pass}) \quad \frac{d \equiv d_1 \quad d_1 \succ d_2 \quad d_2 \equiv d'}{d \succ d'} \quad (\text{cong})$

perform any action of the primary activity a except for fault emission and scope completion. In particular, inner forced terminations are propagated externally outside the scope. Differently from forced termination, faults arising within a scope are managed internally (rule (fault)), and the corresponding handler is installed when the main activity completes (rule (done₂)). By rule (done₂), default compensation is performed *after* termination of the primary activity and before fault handling. Note that compensation activities do not store any state with them: hence, if the state changes between the compensation being stored and executed, the current state is used.

A few auxiliary functions are also used in the semantics of deployments defined in Table 5. The rules for communication and updating of variables ((com), (new) and (var)) need a mechanism for checking if an assignment of some values \bar{v} to \bar{w} complies with the constraints imposed by the given correlation set c and state μ and, in case of success, returns a state μ' for the variables in \bar{w} that records the effect of the assignment. This mechanism is implemented by the function $\text{match}(\cdot, \cdot, \cdot, \cdot)$ defined through the rules in the upper part of Table 4. Notice that $\text{match}(\cdot, \cdot, \cdot, \cdot)$ is undefined when \bar{w} and \bar{v} have different length or when $x \in c$ and $\{x \mapsto v'\} \in \mu$ for some $v' \neq v$ (since the state $\{x \mapsto v\}$ does not comply with c and μ). Rules (com) and (new) also use the auxiliary predicate $s \Downarrow_{\tilde{p}:o:\bar{v}}^{c,n}$, defined inductively on the syntax of s in the lower part of Table 4, that checks the ability of s of performing a receive on the operation o exploiting the partner link \tilde{p} , matching the tuple of values \bar{v} and generating a state with fewer pairs than n that complies with c and the current state of the activity performing the receive.

Finally, we linger on the rules in Table 5. By rule (com), communication can take place when two service instances perform matching receive and invoke activities complying with the correlation set of the receiving instance. Notice that matching covers both partner link \tilde{p} and business data \bar{v} . Communication generates a state that updates the state of the receiving instance. If more than one matching receive activity is able to process a given invoke, then only the more defined one (i.e. the receive that generates the ‘smaller’ state) progresses (predicate $\cdot \Downarrow \cdot$ serves this purpose). This mechanism permits

to correlate messages to different service instances and to model the precedence of an existing service instance over a new service instantiation (rule (new), see also the *Multiple start and conflicting receive activities* example in Section 4). In rules (com) and (new), the assumption about *well-formedness* of deployments finds full employment, because it avoids to check every single deployment for possible conflicting receive activities. By rule (new), service instantiation can take place when a service definition and a service instance perform matching receive and invoke activities, respectively. By rule (var), correlation variables cannot be reassigned if the new value does not match with the old one. Moreover, if an assignment takes place, its effect is global to the instance, i.e. the state is updated. By rule (pass), execution of activities different from communications or assignments can always proceed. If part of a larger deployment evolves, the whole composition evolves accordingly (rule (part)) and, as usual, structural congruent deployments have the same reductions (rule (cong)).

3 A Shipping Service Scenario

We consider an extended version of the shipping service described in the official specification of WS-BPEL [23] (Section 15.1). This example will allow us to illustrate most of the language features, including correlation sets, shared variables, flow control structures, fault and compensation handling. We will see that, in particular, scope activities are especially useful for modelling fault handling and compensation behaviours, while exit activities are useful to exit from while loops and terminate the customer instance.

The shipping service handles the shipment of orders. From the service point of view, orders are composed of a number of items. The service offers two types of shipment: shipments where the items are held and shipped together and shipments where the items are shipped piecemeal until the order is fulfilled. The service specification in *Blite* is

$$\begin{aligned} s_{ship} &\triangleq [\text{rcv} \langle p_{ship}, x_{cust} \rangle \text{o}_{req} \langle x_{id}, x_c, x_{items} \rangle ; \\ &\quad \text{if} (x_c) \{ \text{inv} \langle x_{cust} \rangle \text{o}_{notice} \langle x_{id}, x_{items} \rangle \} \{ [a_{ship} \bullet \text{inv} \langle x_{cust} \rangle \text{o}_{err} \langle x_{id}, \text{"sorry"} \rangle] \}] \\ a_{ship} &\triangleq [a_{priceCalc} \star a_{comp}] ; x_{shipped} := 0 ; \\ &\quad \text{while} (x_{shipped} < x_{items}) \{ \\ &\quad \quad x_{count} := \text{rand}() ; \\ &\quad \quad \text{if} (x_{count} \leq 0) \{ x_{ratio} := x_{shipped} / x_{items} ; \text{throw} \} \\ &\quad \quad \{ \text{inv} \langle x_{cust} \rangle \text{o}_{notice} \langle x_{id}, x_{count} \rangle ; x_{shipped} := x_{shipped} + x_{count} \} \} \end{aligned}$$

p_{ship} is the partner associated to the shipping service, o_{req} is the operation used to receive the shipping request, and $\langle x_{id}, x_c, x_{items} \rangle$ is the tuple of variables used for the request shipping message: x_{id} stores the order identifier, that is used to correlate the ship notice(s) with the ship order, x_c stores a boolean indicating whether the order is to be shipped complete or not, and x_{items} stores the number of items in the order. Shipping notices and error messages to customers are sent using the partner stored in x_{cust} and the operations o_{notice} and o_{err} , respectively. A notice message is a tuple composed of the order identifier and the number of items in the shipping notice. When partial shipment is acceptable, $x_{shipped}$ is used to record the number of items already shipped.

Our example extends that in [23] by allowing the service to generate a fault in case the shipping company has ended the stock of items (this is modelled by function $\text{rand}()$)

returning an integer less than or equal to 0). The fault is handled by sending an error message to the customer and by compensating the inner scope, that has already completed successfully. Function *rand()* returns a random integer number and represents an internal interaction with a back-end system. For the sake of simplicity, we do not further describe this interaction. Moreover, we do not show services $\mathbf{a}_{priceCalc}$ and \mathbf{a}_{comp} . Basically, the former calculates the shipping price according to the value assigned to \mathbf{x}_{items} and sends the result to the accounts department. The latter is the corresponding compensation activity, that sends information about the non-shipped items to the accounts department and sends a refund to the customer according to the ratio (stored in \mathbf{x}_{ratio}) between the shipped items (stored in $\mathbf{x}_{shipped}$) and the required ones (stored in \mathbf{x}_{items}). Now, consider the following composition of a deployment containing the shipping service definition and a deployment containing a customer's invocation of the service

$$\{\mathbf{s}_{ship}\}_{\{\mathbf{x}_{id}\}} \parallel \{\mu_{cust} \vdash \text{inv} \langle \mathbf{p}_{ship}, \mathbf{p}_{cust} \rangle \mathbf{o}_{req} \langle \mathbf{y}_{id}, \mathbf{y}_c, \mathbf{y}_{items} \rangle ; \mathbf{a}_{cust}\}_{\{\mathbf{y}_{id}\}}$$

where $\mu_{cust} = \{\mathbf{y}_{id} \mapsto 123, \mathbf{y}_c \mapsto \text{ff}, \mathbf{y}_{items} \mapsto 50\}$ and \mathbf{a}_{cust} is the following term

$$\begin{aligned} \mathbf{y}_{shipped} := 0 ; \text{while} (\mathbf{y}_{shipped} < \mathbf{y}_{items}) \{ \\ \quad \text{rcv} \langle \mathbf{p}_{cust} \rangle \mathbf{o}_{notice} \langle \mathbf{y}_{id}, \mathbf{y}_{count} \rangle ; \mathbf{y}_{shipped} := \mathbf{y}_{shipped} + \mathbf{y}_{count} \\ \quad + \text{rcv} \langle \mathbf{p}_{cust} \rangle \mathbf{o}_{err} \langle \mathbf{y}_{id}, \mathbf{y}_{err} \rangle ; \text{exit} \} \end{aligned}$$

In the first computational step, the customer's invocation is consumed and an instance of the shipping service is created. Thus the overall computation becomes

$$\{\mathbf{s}_{ship}, \mu_{ship} \vdash [\mathbf{a}_{ship} \bullet \text{inv} \langle \mathbf{x}_{cust} \rangle \mathbf{o}_{err} \langle \mathbf{x}_{id}, \text{"sorry"} \rangle]\}_{\{\mathbf{x}_{id}\}} \parallel \{\mu_{cust} \vdash \mathbf{a}_{cust}\}_{\{\mathbf{y}_{id}\}}$$

where μ_{ship} is $\{\mathbf{x}_{id} \mapsto 123, \mathbf{x}_c \mapsto \text{ff}, \mathbf{x}_{items} \mapsto 50, \mathbf{x}_{cust} \mapsto \mathbf{p}_{cust}\}$. The computation can now go on, e.g., with the inner scope $[\mathbf{a}_{priceCalc} \star \mathbf{a}_{comp}]$ that successfully completes while its continuation fails, e.g., *rand()* returns an integer less than or equal to 0.

4 Evaluations of BPEL Engines

We now present some illustrative examples and use them to test and compare the behaviour of three well-known free WS-BPEL engines, namely Oracle BPEL Process Manager [3], ActiveBPEL Engine [1], and Apache ODE [2] (the latter two are open source projects, whereas Oracle BPEL is distributed under the Oracle Technology Network Developer License). For our evaluation, we have taken into account fundamental features of WS-BPEL that remained unchanged since its initial version. Due to lack of space, we refer the interested reader to [18] for further details and examples.

Example 1 (Message correlation). For our simplification purposes, tuples can be used to represent XML messages by adopting the convention that the first field of each tuple acts as a 'tag' (like originally proposed in the coordination language Linda [12]). Tuples

plus correlation variables can be exploited to correlate, by means of their same contents, different service interactions logically forming a same ‘session’. For example, consider the two uncorrelated receive activities of the following service definition:

$$\{ [\text{rcv } \langle p \rangle \text{ o } \langle x \rangle ; \text{rcv } \langle p \rangle \text{ o}' \langle y \rangle ; \mathbf{a}] \}_{\{x,y\}}$$

The fact that the messages for operations o and o' are uncorrelated implies that, e.g., if there are concurrent instances then successive invocations for the same instance can be mixed up and be delivered to different instances. If one thinks it right, this behavior can be prevented simply by correlating consecutive messages by means of some correlation data, e.g. the first received value as in the following modified service definition:

$$\{ [\text{rcv } \langle p \rangle \text{ o } \langle x \rangle ; \text{rcv } \langle p \rangle \text{ o}' \langle x, y \rangle ; \mathbf{a}] \}_{\{x,y\}}$$

A particular case is when the two previous receives are identical, i.e. when we have:

$$\{ [\text{rcv } \langle p \rangle \text{ o } \langle x \rangle ; \text{rcv } \langle p \rangle \text{ o } \langle x \rangle ; \mathbf{a}] \}_{\{x\}}$$

Note that the WS-BPEL specification permits to consecutively receive a same request on a specific partner and operation ([23], Section 10.4), and does not mention that possible conflicting receives could arise. To illustrate, include a client process as follows:

$$\{ [\text{rcv } \langle p \rangle \text{ o } \langle x \rangle ; \text{rcv } \langle p \rangle \text{ o } \langle x \rangle ; \mathbf{a}] \}_{\{x\}} \parallel \{ \{y \mapsto v\} \vdash \text{inv } \langle p \rangle \text{ o } \langle y \rangle ; \text{inv } \langle p \rangle \text{ o } \langle y \rangle \}_0$$

The client process performs two requests that, according to the semantics of *Blite*, trigger only one instantiation of the service. Thus, the only possible evolution leads to

$$\{ [\text{rcv } \langle p \rangle \text{ o } \langle x \rangle ; \text{rcv } \langle p \rangle \text{ o } \langle x \rangle ; \mathbf{a}] , \{x \mapsto v\} \vdash [\mathbf{a}] \}_{\{x\}}$$

Differently from *Blite*, when executing this example, Oracle BPEL creates two instances, one for each received request. An important consequence, and an unexpected side effect, is that the created instances are in conflict and, then, will never be executed. Instead, ActiveBPEL and Apache ODE, just like *Blite*, exploit the received data to correlate two consecutive receives and, thus, prevent creation of a new instance. However, if the client performs a third invocation $\text{inv } \langle p \rangle \text{ o } \langle y \rangle$, Apache ODE is not able to serve this last request, while ActiveBPEL behaves properly.

Example 2 (Persistent messages). In service-oriented systems communication paradigms are usually asynchronous (mainly for scalability reasons [5]), in the sense that there may be an arbitrary delay between the sending and the receiving of a message, the ordering in which messages are received may differ from that in which they were produced, and a sender cannot determine if and when a sent message will be received. We can guess from [23], Section 10.4, that this is also the case of WS-BPEL. To illustrate, consider the following *Blite* term:

$$\begin{aligned} & \{ [\text{rcv } \langle p \rangle \text{ o}_1 \langle x \rangle ; \text{rcv } \langle p \rangle \text{ o}_2 \langle x, z \rangle ; \mathbf{a}] \}_{\{x\}} \\ & \parallel \{ \{y_1 \mapsto v, y_2 \mapsto v'\} \vdash \text{inv } \langle p \rangle \text{ o}_2 \langle y_1, y_2 \rangle ; \text{inv } \langle p \rangle \text{ o}_1 \langle y_1 \rangle \}_0 \end{aligned}$$

After the message $\ll \langle p \rangle : \text{o}_2 : \langle v, v' \rangle \gg$ is produced by the first invoke activity, a service instance is created as a result of consumption of the message produced by the second invoke activity.

$$\begin{aligned}
& \{ [\text{rcv} \langle p \rangle o_1 \langle x \rangle ; \text{rcv} \langle p \rangle o_2 \langle x, z \rangle ; a] \}_{|x|} \\
& \| \{ \{ y_1 \mapsto v, y_2 \mapsto v' \} \vdash \ll \langle p \rangle : o_2 : \langle v, v' \rangle \gg \mid \text{inv} \langle p \rangle o_1 \langle y_1 \rangle \}_0 \xrightarrow{\quad} \\
& \{ [\text{rcv} \langle p \rangle o_1 \langle x \rangle ; \text{rcv} \langle p \rangle o_2 \langle x, z \rangle ; a], \{ x \mapsto v \} \vdash [\text{rcv} \langle p \rangle o_2 \langle x, z \rangle ; a] \}_{|x|} \\
& \| \{ \{ y_1 \mapsto v, y_2 \mapsto v' \} \vdash \ll \langle p \rangle : o_2 : \langle v, v' \rangle \gg \}_0
\end{aligned}$$

Now, the first produced message is not considered expired and, thus, can be consumed by the newly created service instance.

$$\begin{aligned}
& \{ [\text{rcv} \langle p \rangle o_1 \langle x \rangle ; \text{rcv} \langle p \rangle o_2 \langle x, z \rangle ; a], \{ x \mapsto v, z \mapsto v' \} \vdash [a] \}_{|x|} \\
& \| \{ \{ y_1 \mapsto v, y_2 \mapsto v' \} \vdash \text{empty} \}_0
\end{aligned}$$

All the examined BPEL engines ‘tacitly’ agree with this communication paradigm, although no explicit requirement is reported in the WS-BPEL specification.

Example 3 (Multiple start and conflicting receive activities). The WS-BPEL specification permits to use multiple start activities ([23], Section 10.4), however it is not clear how conflicting receive activities must be handled. The following example shows that conflicting receive activities can be enabled when a service definition with multiple start activities is instantiated. Consider the three composed deployments

$$\begin{aligned}
& \{ [(\text{rcv} \langle p_1 \rangle o \langle x \rangle \mid \text{rcv} \langle p_2 \rangle o \langle x, z \rangle) ; a] \}_{|x|} \| \{ \{ y \mapsto v \} \vdash \text{inv} \langle p_1 \rangle o \langle y \rangle \}_0 \\
& \| \{ \{ y_1 \mapsto v, y_2 \mapsto v' \} \vdash \text{inv} \langle p_2 \rangle o \langle y_1, y_2 \rangle \}_0
\end{aligned}$$

After message $\ll \langle p_1 \rangle : o : \langle v \rangle \gg$, produced by invocation $\text{inv} \langle p_1 \rangle o \langle y \rangle$, has been processed by $\text{rcv} \langle p_1 \rangle o \langle x \rangle$, the overall composition becomes

$$\begin{aligned}
& \{ [(\text{rcv} \langle p_1 \rangle o \langle x \rangle \mid \text{rcv} \langle p_2 \rangle o \langle x, z \rangle) ; a], \{ x \mapsto v \} \vdash [\text{rcv} \langle p_2 \rangle o \langle x, z \rangle ; a] \}_{|x|} \\
& \| \{ \{ y_1 \mapsto v, y_2 \mapsto v' \} \vdash \text{inv} \langle p_2 \rangle o \langle y_1, y_2 \rangle \}_0
\end{aligned}$$

Now, the definition and the instance of the service compete for receiving the same message sent by the invoke activity $\text{inv} \langle p_2 \rangle o \langle y_1, y_2 \rangle$. In cases like this, the WS-BPEL specification requires that the invocation is only delivered to the existing instance, which prevents creation of a new instance. In fact, in *Blite* the above term can only reduce to

$$\{ [(\text{rcv} \langle p_1 \rangle o \langle x \rangle \mid \text{rcv} \langle p_2 \rangle o \langle x, z \rangle) ; a], \{ x \mapsto v, z \mapsto v' \} \vdash [a] \}_{|x|}$$

In case of conflicting receives, the WS-BPEL specification document prescribes to raise the standard fault `bpel:conflictingReceive`. For example, this situation readily occurs when a service exploits multiple start activities, because of race conditions on incoming messages among the service definition and the created instances. However, in such cases, it does not seem fair to raise a fault because the correlation data contained within each incoming message should be sufficient to decide if the message has to be routed to a specific instance or to the service definition. This is indeed a tricky question. For example, Oracle BPEL raises the fault `bpel:conflictingReceive` also in these situations. ActiveBPEL behaves differently and, just like *Blite*, exploits correlation to restrict instantiation to one service instance, whereas multiple start activities are not currently supported by Apache ODE.

Example 4 (Scheduling for parallel execution). While using the BPEL engines, we have also experimented that they implement the parallel operator in a different manner. For example, in WS-BPEL, the expected behaviour of the following term:

$$x_1 := v_1 \mid x_2 := v_2 \mid x_3 := v_3$$

is that the three assignments are executed in an unpredictable order that may change in different executions. In fact, only Apache ODE implements this semantics, while the other two engines execute the assignments in an order fixed in advance (that is from left to right in case of ActiveBPEL and from right to left in case of Oracle BPEL).

Example 5 (Forced termination). The WS-BPEL specification ([23], Section 12.6) says: “The <sequence> and <flow> constructs *must* be terminated by terminating their behavior and applying termination to all nested activities currently active within them”. This definition is ambiguous because it is not clear what “nested activities currently active” means in case of termination due to <exit> or <throw> activities. For example, Oracle BPEL interprets the behaviour $\text{end}(a_1 ; a_2)$ as it were $a_1 ; a_2$ if it is prompted by activity <exit>, and as $\text{end}(a_1)$, if it is prompted by activity <throw>. ActiveBPEL is more faithful to WS-BPEL and *Blite* for which all currently running activities are terminated as soon as possible without any fault handling or compensation ([23], Section 10.10). But, differently from *Blite*, ActiveBPEL does not distinguish short-lived from basic activities and makes them terminate in the same way. Finally, Apache ODE is compliant with *Blite*, because function $\text{end}(\cdot)$ retains short-lived activities.

Example 6 (Eager execution of termination activities). As previously stated, in order to be compliant with the WS-BPEL requirement stating that termination activities must end immediately all currently running activities ([23], Section 10.10), in the semantics of *Blite* activities *throw* and *exit* have higher priority than the remaining ones. E.g., consider the following structured activity:

$$a \triangleq \text{throw} \mid sh_1 ; sh_2 \mid \text{rcv}\langle p \rangle o\langle x \rangle ; a'$$

In *Blite*, by executing the activity *throw*, this term can only reduce to:

$$\text{stop} \mid \text{end}(sh_1 ; sh_2) \mid \text{end}(\text{rcv}\langle p \rangle o\langle x \rangle ; a') \equiv \text{stop} \mid sh_1$$

While ActiveBPEL agrees with this requirement, Oracle BPEL and Apache ODE do not implement any prioritized behavior for termination activities. Thus, for example, the above term *a* can evolve by firstly performing the activity sh_1 and then the activity *throw*; this way, the activity sh_2 is not terminated.

Example 7 (Protected handlers). The following structured activity consists of a top-level scope with two inner parallel scopes, one of which being a sequence of two scopes.

$$a \triangleq [([[a_1 \bullet \text{throw} \star a_c] ; [a_2 \bullet \text{throw} \star \text{empty}] \bullet \text{throw} \star \text{empty}] \mid [a_3 \bullet \text{throw} \star \text{empty}]) \bullet \text{empty}]$$

For the sake of presentation, suppose that a_1 performs an assignment and completes, say $a_1 \triangleq x_1 := v_1$, while both activities a_2 and a_3 perform an assignment and reduce to the *throw* activity, say $a_i \triangleq x_i := v_i ; \text{throw}$ for $i = 2, 3$. Now, consider a deployment containing a service instance $\mu \vdash a$ such that variables x_1 , x_2 and x_3 are not in $\text{dom}(\mu)$. A possible computation is the following one

$$\begin{aligned}
\{\mu \vdash \mathbf{a}\}_c &\xrightarrow{(1)} \{\mu_1 \vdash [([[\text{empty} \bullet \text{throw} \star \mathbf{a}_c]; \\
&\quad [\mathbf{a}_2 \bullet \text{throw} \star \text{empty}] \bullet \text{throw} \star \text{empty} \triangle \text{empty}] \\
&\quad | [\mathbf{a}_3 \bullet \text{throw} \star \text{empty}]) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(2)} \{\mu_1 \vdash [([[\mathbf{a}_2 \bullet \text{throw} \star \text{empty}] \bullet \text{throw} \star \text{empty} \triangle (\mathbf{a}_c ; \text{empty})] \\
&\quad | [\mathbf{a}_3 \bullet \text{throw} \star \text{empty}]) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(3)} \{\mu_2 \vdash [([[\text{throw} \bullet \text{throw} \star \text{empty}] \bullet \text{throw} \star \text{empty} \triangle \mathbf{a}_c] \\
&\quad | [\mathbf{a}_3 \bullet \text{throw} \star \text{empty}]) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(4)} \{\mu_2 \vdash [([[\text{stop} \bullet \text{throw} \star \text{empty}] \bullet \text{throw} \star \text{empty} \triangle \mathbf{a}_c] \\
&\quad | [\mathbf{a}_3 \bullet \text{throw} \star \text{empty}]) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(5)} \{\mu_2 \vdash [([(\text{throw}) \bullet \text{throw} \star \text{empty} \triangle \mathbf{a}_c] \\
&\quad | [\mathbf{a}_3 \bullet \text{throw} \star \text{empty}]) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(6)} \{\mu_2 \vdash [([(\text{stop}) \bullet \text{throw} \star \text{empty} \triangle \mathbf{a}_c] \\
&\quad | [\mathbf{a}_3 \bullet \text{throw} \star \text{empty}]) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(7)} \{\mu_2 \vdash [(([\mathbf{a}_c ; \text{throw}] | [\mathbf{a}_3 \bullet \text{throw} \star \text{empty}]) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(8)} \{\mu_3 \vdash [(([\mathbf{a}_c ; \text{throw}] | [\text{throw} \bullet \text{throw} \star \text{empty}]) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(9)} \{\mu_3 \vdash [(([\mathbf{a}_c ; \text{throw}] | [\text{stop} \bullet \text{throw} \star \text{empty}]) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(10)} \{\mu_3 \vdash [(([\mathbf{a}_c ; \text{throw}] | (\text{throw})) \bullet \text{empty} \star \text{empty}]]_c \\
&\xrightarrow{(11)} \{\mu_3 \vdash [(\text{end}([\mathbf{a}_c ; \text{throw}]) | (\text{stop})) \bullet \text{empty} \star \text{empty}]]_c \\
&\equiv \{\mu_3 \vdash [(([\mathbf{a}_c ; \text{throw}] | \text{stop}) \bullet \text{empty} \star \text{empty}]]_c
\end{aligned}$$

where the reductions are labelled by numbers indicating the corresponding steps. When \mathbf{a}_1 completes, the compensation handler \mathbf{a}_c is inserted into the default compensation activities of its enclosing scope (1-2). When execution of \mathbf{a}_2 rises a fault, then the fault is caught by the corresponding fault handler (3-7) that activates the default compensation $\mathbf{a}_c ; \text{throw}$. This activity is protected, by using the auxiliary operator $(\cdot|)$, from the effect of the forced termination triggered by the parallel scope $[\mathbf{a}_3 \bullet \text{throw} \star \text{empty}]$ (7-11).

We end by remarking two aspects of the compensation mechanism prescribed by the WS-BPEL specification ([23], Sections 12.5 and 10.10). First, compensation handlers of faultily terminated scopes should not be installed. Second, fault and compensation handlers should not be affected by the activities causing the forced termination. However, both aspects are not faithfully implemented in Oracle BPEL, while ActiveBPEL and Apache ODE meet these specific requirements and adhere to *Blite* semantics.

Evaluation Results. The results of our experiments, summarized in Table 6, point out that the engines we have experimented with are not fully compliant with *Blite*, that, in our opinion, faithfully represents the intended semantics of WS-BPEL. This is also a consequence of the lack of a formal semantics for WS-BPEL, that would have disambiguated the intricate and complex features of the language. We believe that *Blite*, and

Table 6. *Blite* compliance

	Oracle BPEL	ActiveBPEL	Apache ODE
Message correlation (Ex. 1)	+	+	+
Consecutive conflicting receives (Ex. 1)	-	+	+/-
Persistent messages (Ex. 2)	+	+	+
Multiple start (Ex. 3)	-	+	-
Parallel execution (Ex. 4)	-	-	+
Short-lived activities (Ex. 5)	+	-	+
Function end(\cdot) (Ex. 5)	-	+	+
Eager execution (Ex. 6)	-	+	-
Protected handlers (Ex. 7)	-	+	+
Compensation handler installation (Ex. 7)	-	+	+

works with similar goals, other than as a guide for the development of faithful implementations since the early stages, can be also used to make future versions of existing implementations more compatible.

5 Concluding Remarks

We have introduced *Blite*, a significative and non-redundant fragment of WS-BPEL, designed around some of its peculiar features like partner links, process termination, message correlation, long-running business transactions and compensation handlers. Our formal presentation of *Blite* helps clarifying some undefined/ambiguous aspects of the WS-BPEL specification. For example, we have formalized the close relationship between multiple start activities and race conditions. By means of several examples, we have also pointed out that the behaviour of three of the most used free BPEL engines (namely, ActiveBPEL, Apache ODE and Oracle BPEL Process Manager) differs from each other and from the WS-BPEL specification in many important aspects.

Several formal semantics of WS-BPEL were proposed in the literature (for an overview see [24]). Many of these efforts aim at formalizing a *complete* semantics for WS-BPEL using Petri nets [24,19], but do not cover such dynamical aspects as service instantiation and message correlation. Other works [11,14] using process calculi focus instead on small and relatively simple subsets of WS-BPEL. Another bunch of related works [15,20] formalize the semantics of WS-BPEL by encoding parts of the language into more foundational orchestration languages. Our work differs for the number of features that are simultaneously modelled and for the fact that dynamical aspects are fully taken into account. Recently, a very general and flexible framework for error recovery has been introduced in [13]; this framework extends [14] with dynamic compensation, modelling in particular the dependency between fault handling and the request-response communication pattern.

Some other relevant related works are [7,6,4]. In the first two, the authors propose a formal approach to model compensation in transactional calculi and present a detailed comparison with [8]. The third is an extension of the asynchronous π -calculus with long-running (scoped) transactions. The language has a scope construct which plays

a role similar to the scope activity presented in our semantics, but it is not aimed at capturing the order in which compensations should be activated. On the contrary, the semantics we propose faithfully captures the intended semantics of WS-BPEL, thus for example compensations are activated in the reverse order w.r.t. the order of completion of the original scopes.

Our programme is to provide a framework for the design and the verification of WS-BPEL applications that supports analysis of service orchestration. As a further step in this direction, in [18] we have also defined an encoding from *Blite* to COWS [16], a calculus for orchestration of web services that we recently proposed, and we have formalized the properties enjoyed by the encoding. By relying on these results, we plan to devise methods to analyze *Blite* specifications (and the WS-BPEL applications they model) by exploiting the analytical tools already developed for COWS, such as the stochastic extension defined in [25] that enables quantitative reasoning on service behaviours, the type system introduced in [17] that permits to check confidentiality properties, and the logic and model checker presented in [10] that permits expressing and checking functional properties of services.

Acknowledgements. We thank the anonymous referees for their useful comments.

References

1. ActiveBPEL 4.1 (September 2007), <http://www.active-endpoints.com>
2. Apache ODE 1.1.1 (August 2007), <http://ode.apache.org>
3. Oracle BPEL Process Manager 10.1.3 (December 2007), <http://www.oracle.com/technology/bpel>
4. Bocchi, L., Laneve, C., Zavattaro, G.: A calculus for long-running transactions. In: Najm, E., Nestmann, U., Stevens, P. (eds.) FMOODS 2003. LNCS, vol. 2884, pp. 124–138. Springer, Heidelberg (2003)
5. Brown, A., Johnston, S., Kelly, K.: Using service-oriented architecture and component-based development to build web service applications, TR, Rational Software Corp. (2002)
6. Bruni, R., Butler, M.J., Ferreira, C., Hoare, C.A.R., Melgratti, H.C., Montanari, U.: Comparing two approaches to compensable flow composition. In: Abadi, M., de Alfaro, L. (eds.) CONCUR 2005. LNCS, vol. 3653, pp. 383–397. Springer, Heidelberg (2005)
7. Bruni, R., Melgratti, H.C., Montanari, U.: Theoretical foundations for compensations in flow composition languages. In: POPL, pp. 209–220. ACM, New York (2005)
8. Butler, M.J., Ferreira, C.: An operational semantics for StAC, a language for modelling long-running business transactions. In: De Nicola, R., Ferrari, G.L., Meredith, G. (eds.) COORDINATION 2004. LNCS, vol. 2949, pp. 87–104. Springer, Heidelberg (2004)
9. Box, D.: et al. Web services addressing. W3C member submission, August 10 (2004)
10. Fantechi, A., Gnesi, S., Lapadula, A., Mazzanti, F., Pugliese, R., Tiezzi, F.: A model checking approach for verifying COWS specifications. In: FASE. LNCS, Springer, Heidelberg (to appear, 2008)
11. Geguang, P., Xiangpeng, Z., Shuling, W., Zongyan, Q.: Semantics of BPEL4WS-like fault and compensation handling. In: Fitzgerald, J.S., Hayes, I.J., Tarlecki, A. (eds.) FM 2005. LNCS, vol. 3582, pp. 350–365. Springer, Heidelberg (2005)
12. Gelernter, D.: Generative communication in Linda. ACM TOPLAS 7(1), 80–112 (1985)

13. Guidi, C., Lanese, I., Montesi, F., Zavattaro, G.: On the interplay between fault handling and request-response service invocations. In: ACSD, IEEE CS Press, Los Alamitos (to appear, 2008)
14. Guidi, C., Lucchi, R., Gorrieri, R., Busi, N., Zavattaro, G.: SOCK: a calculus for service oriented computing. In: Dan, A., Lamersdorf, W. (eds.) ICSOC 2006. LNCS, vol. 4294, pp. 327–338. Springer, Heidelberg (2006)
15. Laneve, C., Zavattaro, G.: Foundations of web transactions. In: Sassone, V. (ed.) FOSSACS 2005. LNCS, vol. 3441, pp. 282–298. Springer, Heidelberg (2005)
16. Lapadula, A., Pugliese, R., Tiezzi, F.: A Calculus for Orchestration of Web Services. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 33–47. Springer, Heidelberg (2007)
17. Lapadula, A., Pugliese, R., Tiezzi, F.: Regulating data exchange in service oriented applications. In: Arbab, F., Sirjani, M. (eds.) FSEN 2007. LNCS, vol. 4767, pp. 223–239. Springer, Heidelberg (2007)
18. Lapadula, A., Pugliese, R., Tiezzi, F.: A formal account of WS-BPEL (full version), Technical report, Univ. Firenze (2008), <http://rap.dsi.unifi.it/cows>
19. Lohmann, N.: A feature-complete Petri net semantics for WS-BPEL 2.0. In: Web Services and Formal Methods. LNCS, vol. 4937, pp. 77–91. Springer, Heidelberg (2008)
20. Mazzara, M., Lucchi, R.: A pi-calculus based semantics for WS-BPEL. *Journal of Logic and Algebraic Programming* 70(1), 96–118 (2006)
21. Meredith, L.G., Bjorg, S.: Contracts and types. *Commun. ACM* 46(10), 41–47 (2003)
22. OASIS WSBPEL TC. WS-BPEL issues list, http://www.oasis-open.org/committees/download.php/20228/WS_BPEL_issues_list.html
23. OASIS WSBPEL TC. Web Services Business Process Execution Language Version 2.0 (April 2007), <http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.html>
24. Ouyang, C., van der Aalst, W.M.P., Breutel, S., Dumas, M., ter Hofstede, A.H.M., Verbeek, H.M.W.: Formal semantics and analysis of control flow in WS-BPEL (revised version). Technical report, BPM Center Report (2005), <http://www.bpmcenter.org>
25. Prandi, D., Quaglia, P.: Stochastic COWS. In: Krämer, B.J., Lin, K.-J., Narasimhan, P. (eds.) ICSOC 2007. LNCS, vol. 4749, pp. 245–256. Springer, Heidelberg (2007)
26. van Breugel, F., Koshkina, M.: Models and verification of BPEL. Technical report (2006), <http://www.cse.yorku.ca/~franck/research/drafts/tutorial.pdf>