Durham Research Online

Deposited in DRO:

11 December 2009

Version of attached file:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Luo, C. and Qin, S. and Qiu, Z. (2008) 'Verifying BPEL-like programs with Hoare logic.', *Frontiers of computer science in China*, 2 (4). pp. 344-356.

Further information on publisher's website:

<http://dx.doi.org/10.1007/s11704-008-0039-2>

Publisher's copyright statement:

The original publication is available at www.springerlink.com

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

Frontiers of Computer Science in China

Verifying BPEL-like Programs with Hoare Logic

Chenguang LUO¹, Shengchao QIN¹, Zongyan QIU²

¹ Department of Computer Science, Durham University, Durham DH1 3LE, United Kingdom

² LMAM and Department of Informatics, School of Math. Sciences, Peking University, Beijing 100871, China

© Higher Education Press and Springer-Verlag 2008

Abstract The WS-BPEL language has recently become a *de facto* standard for modeling Web-based business processes. One of its essential features is the fully programmable compensation mechanism. To understand it better, many recent works have mainly focused on formal semantic models for WS-BPEL. In this paper, we make one step forward by investigating the verification problem for business processes written in BPEL-like languages. We propose a set of proof rules in Hoare-logic style as an axiomatic verification system for a BPEL-like core language containing key features such as data states, fault and compensation handling. We also propose a big-step operational semantics which incorporates all these key features. Our verification rules are proven sound with respect to this underlying semantics. The application of the verification rules is illustrated via the proof search process for a nontrivial example.

Keywords WS-BPEL, compensation mechanism, operational semantics, axiomatic verification system, soundness

1 Introduction

The Internet is now developing at a high speed supported by the web technology. As a result, many web-based applications, such as Web services, begin to flourish and play a more and more significant role in various application areas. Web services boost a new approach to the construction of business processes where many basic functions are encapsulated and provided as individual services on the web, which later may be composed to form complex services according

to diverse clients' demands. To cater for the description of Web service composition, researchers and industrial practitioners have proposed several Web service orchestration languages such as XLANG [1], WSFL [2], StAC [3], and WS-BPEL [4, 5].

Among these orchestration languages, WS-BPEL has now become a *de facto* standard. One important feature of WS-BPEL, as well as some other similar languages, is its mechanism for supporting long run transactions (LRTs). In any single step of an LRT, a fault may occur and appropriate compensation actions may be required. To address such demand, WS-BPEL provides a set of scope-based fault handling and compensation mechanisms to deal with faults and potential undoing of some already completed business activities. The compensation mechanisms are fully programmable, and thus allow users to define any application-specific compensation rules. Nevertheless, these mechanisms, despite very flexible and powerful, also bring intricacies into the WS-BPEL language specification. As a result, it becomes a challenging issue to formalize and reason about WS-BPEL processes.

Many recent works focused mostly on the formal semantics for WS-BPEL, e.g. [6–10]. These pioneering works are very important for reducing possible ambiguity in the language specification and also for better understanding of the language. In this paper we will target at an orthogonal but equally important problem, the partial correctness of WS-BPEL processes. To make the presentation simple, we shall focus on a subset of WS-BPEL. However, our core language will take into account most of the important language features of WS-BPEL, including data state, fault handling and compensation mechanism. We will design a concise yet novel operational semantics for our language, and propose a Hoare logic style verification system on top of it, which will be proven sound with respect to the underlying semantics. Due

to the complexity of web-based business processes, the correctness of such programs remains as a challenge. Our verification system for BPEL-like language makes one step forward towards tackling this challenging problem. To the best of our knowledge, this is the first axiomatic verification system for a language with data states, scope-based fault and compensation handling mechanisms. The main contributions of this paper can be summarized as follows:

- We propose a concise yet novel operational semantics for a BPEL-like core language. Although there are some semantic works with similar topics, our semantics is interesting in that it integrates features like scopes, data states, fault handling and compensation in a very simple way.
- We design an assertion language for specifying certain safety properties for BPEL-like processes, and also propose a set of axioms and inference rules in Hoare logic style to form an axiomatic verification system for the language. The pre- and postconditions are formulas expressed in our assertion language.
- We state and prove the soundness of our axiomatic verification system with respect to the semantics. That is, provable specifications are all semantically valid. A non-trivial example is presented to illustrate the application of the verification rules.

The remainder of this paper is organized as follows. Sec introduces our language $BPEL^*$ which is a core subset of WS-BPEL. A new operational semantics for $BPEL^*$ is then presented in Sec . Sec 2 is devoted to the Hoare logic style verification system for $BPEL^*$. Sec 3 deals with the soundness of our verification system, while Sec 3 gives a nontrivial example proof using our verification system. Related work and concluding remarks follow afterwards.

2 The $BPEL^*$ Language

To concentrate on the main aim of this study, we take into account a core subset of the WS-BPEL language, called $BPEL^*$, which comprises not only the important fault and compensation handling mechanisms but also data states of WS-BPEL processes.

The abstract syntax of $BPEL^*$ is given in Figure 1. Note that a program written in $BPEL^*$ is called a *business process* (denoted as BP) which may contain an activity A and a fault handler F . We may sometimes use the general term *process* to refer to an activity A , a compensation handler C , or a fault handler F . The set of all processes is denoted as $\mathbb{P}$.

In Figure 1, x and y stand for variable names, e represents arithmetic expressions, b is for boolean expressions, and n for scope names. A denotes a general activity, while C and F are for compensation and fault handlers, respectively. It is worth noting that the compensation activity $\uparrow n$ can only appear in these two constructs.

Generally a business process has an activity to perform its

BP	$::= \{ A : F \}$	(business process)
A	$::= \text{skip}$	(do nothing)
	$ x := e$	(assignment)
	$ \text{inv } a \ x \ y$	(invoke)
	$ \text{rec } a \ y$	(receive)
	$ \text{rep } a \ x$	(reply)
	$ \text{throw}$	(throw a fault)
	$ A; A$	(sequence)
	$ \text{if } b \text{ then } A \text{ else } A$	(conditional)
	$ A \parallel A$	(flow)
	$ n : \{ A ? C : F \}$	(scope)
C, F	$::= \uparrow n$	(compensation)
	$ \dots$	(similar as A)

Fig. 1 The Syntax of $BPEL^*$

normal work, and once an error occurs it uses the fault handler to deal with it. As for a general activity, the skip does nothing to a process, and an assignment simply overwrites a variable's value to user's intention.

The inv , rec and rep constitute our abstract model of communication with external Web services. The a here is an abstraction of call port of some Web service beyond the current process we are interested in. In our work we have taken a most general web model, where “ a ” is assumed to have some arbitrary behavior as far as the current business process is concerned, as it either returns an arbitrary value or fails. It is also possible to take more specialized models, which is out the consideration of this paper.

The throw throws a fault to prevent any other activity in the current scope from being processed, until the end of the business process — or it is captured somewhere by a fault handler.

The $A \parallel A$ is a simplification of activities' parallel composition (flow). To focus more on the novel aspects of WS-BPEL, including the fault and compensation handling, we put some restrictions over this construct so that links between its components (i.e. additional control-flow restrictions) are disallowed in $BPEL^*$. We can do so because this issue is almost orthogonal to our focus in this paper and it has already been well investigated by researchers, e.g. [11, 12].

Inside a scope $n : \{ A ? C : F \}$, A is the normal activity, C is the compensation handler, and F is the fault handler. In $BPEL^*$, we assume all names for variables defined in a business process are distinct, so are the scope names. This is just for simplicity and does not lose generality as we can easily achieve this by a pre-processing step. Under such assumptions, we can refer to a variable or a scope simply by its name, with no need of mentioning its enclosing context. We also assume that the processes under consideration have been statically checked to meet certain basic well-formedness conditions. For instance, the compensation activity $\uparrow n$ will only

occur in the immediate enclosing scope of the scope n .

3 Dynamic Semantics

In this section, we propose a big-step operational semantics for $BPEL^*$. The semantics not only serves as a runtime model for the language, but also acts as a reference semantics in the soundness proof for our axiomatic verification system. In what follows, we will define the runtime states used for the semantics and then depict the semantic rules.

3.1 Runtime States

The nontrivial business processes need often to support long-running transactions (LRTs), where the exceptional faults are unavoidable, and as a result the partially completed tasks may need to be revoked accordingly. This kind of processes are hard to describe without language support. WS-BPEL deals with this necessity with its scope and compensation mechanism, which can be invoked to reverse partially completed transactions. Since a fault may happen from time to time, the WS-BPEL specification advocates to keep records of state snapshots for the successfully completed scopes, as the associated compensation handlers may refer to such completion states when the compensation is invoked. Our semantics will record those successfully completed scope snapshots in the runtime state, similar to the way used in Qiu *et al.* [6] for recording compensation closures. To facilitate the handling of faults, we also instrument the runtime state with a boolean value to indicate whether the current state is a normal state or a faulty state. The formal notations we use are as follows:

$$\begin{aligned} f \in Status &=_{df} \{fail, norm\} \\ s \in Val &=_{df} Var \rightarrow Value \\ \alpha, [\delta, \dots, \delta] \in CPCTX &=_{df} seq\ CPCTX \\ \delta, \langle n, s, \alpha \rangle \in CPCL &=_{df} ScopeN \times Val \times CPCTX \\ \sigma, (f, s, \alpha) \in \Sigma &=_{df} Status \times Val \times CPCTX \end{aligned}$$

In the semantic model, a runtime state $\sigma = (f, s, \alpha)$ is composed of three elements, where f indicates whether the current state is normal ($f = norm$) or of a fault ($f = fail$), and the s records current snapshot for the values of all variables in the process. The third element α is the compensation context used to record the state snapshots and relative compensation information for successfully completed scopes.

When a compensation activity $\hookrightarrow n$ runs, the code to be executed (i.e. the compensation handler defined in scope n) is statically determined. However, the behavior of the compensation will depend on not only the scope snapshot of n , but also the dynamic execution of the normal activity in scope n that yields the state snapshot. This is due to the fact that (1) the current compensation may invoke compensation handlers from the immediate sub-scopes of n , so its behavior will depend on whether or not each of the sub-scopes has completed successfully (thus the associative compensation handler has

been installed) and (2) such information is determined dynamically during the execution of the normal activity of scope n . To record such information along with the scope snapshot, we define the compensation context α as a (possibly empty) sequence of compensation closures $[\delta_1, \delta_2, \dots, \delta_n]$, whereby a compensation closure $\delta_i = \langle n, s, \alpha_1 \rangle$ is a nested structure which records the state snapshot s for scope n (i.e., the data state at the end of the normal execution of scope n). The third element α_1 is the compensation context accumulated during the execution of the normal activity of scope n . It includes all the compensation closures for those normally completed immediately-enclosed sub-scopes. When the compensation handler of n is invoked, both the scope snapshot s and the enclosed context α_1 are restored for the compensation activity.

We do not record the handlers in the context as such information can be statically determined for a given business process. Instead, we assume the availability of a mapping to fetch the corresponding handlers:

$$\mathcal{C} : ScopeN \rightarrow \mathbb{P}$$

where $ScopeN$ is the set of scope names. For a valid scope name $n \in \text{dom}(\mathcal{C})$, $\mathcal{C}(n) \in \mathbb{P}$ is the compensation handler defined in scope n .

We will make use of standard sequence operators given below (where $\alpha_1 = [\delta_1, \dots, \delta_m]$ and $\alpha_2 = [\delta'_1, \dots, \delta'_n]$):

$$\begin{aligned} \delta_0 \cdot \alpha_1 &= [\delta_0, \delta_1, \dots, \delta_m] \\ \text{hd}(\alpha_1) &= \delta_1 \\ \text{tl}(\alpha_1) &= [\delta_2, \dots, \delta_m] \\ \alpha_1 \frown \alpha_2 &= [\delta_1, \dots, \delta_m, \delta'_1, \dots, \delta'_n] \end{aligned}$$

We define a membership relation as follows:

$$\begin{aligned} \delta \in \alpha &=_{df} \begin{cases} \text{false} & \text{if } \alpha = [] \\ \text{true} & \text{if } \text{hd}(\alpha) = \delta \\ \delta \in \text{tl}(\alpha) & \text{else} \end{cases} \\ \delta \notin \alpha &=_{df} \neg(\delta \in \alpha) \end{aligned}$$

Based on it we can define the following analogous relation:

$$\begin{aligned} n \in \alpha &=_{df} \exists s, \alpha_1 \bullet \langle n, s, \alpha_1 \rangle \in \alpha \\ n \notin \alpha &=_{df} \neg(n \in \alpha) \end{aligned}$$

where n is a scope name and α is a compensation context. Informally, $n \in \alpha$ indicates that the compensation handler for the scope n has been installed (and hence n 's scope snapshot appears in α).

3.2 Operational Semantics

In this subsection, we present the semantic rules for the processes in $BPEL^*$. The big-step operational semantics for $BPEL^*$ is defined by a set of rules of the form

$$\langle A, \sigma \rangle \rightsquigarrow \sigma'$$

where A is a process, while σ and σ' denote the initial and final states, respectively.

When a fault has occurred, the process to be executed will do nothing but propagate the fault. The rule below describes this scenario:

$$\frac{\sigma = (\text{fail}, s, \alpha)}{\langle A, \sigma \rangle \rightsquigarrow \sigma}$$

The following rules define the behavior of skip, assignment, and throw activities from normal states:

$$\begin{aligned} \langle \text{skip}, (\text{norm}, s, \alpha) \rangle &\rightsquigarrow (\text{norm}, s, \alpha) \\ \langle x := e, (\text{norm}, s, \alpha) \rangle &\rightsquigarrow (\text{norm}, s \oplus \{x \mapsto s(e)\}, \alpha) \\ \langle \text{throw}, (\text{norm}, s, \alpha) \rangle &\rightsquigarrow (\text{fail}, s, \alpha) \end{aligned}$$

where $s \oplus s'$ is a state formed by s and s' :

$$(s \oplus s')(x) =_{df} \begin{cases} s'(x) & \text{when } x \in \text{dom } s' \\ s(x) & \text{otherwise} \end{cases}$$

With $s(e)$ to denote the value of expression e under state s , the skip and assignment are analogous to normal imperative language. The throw here changes the process faulty state to fail immediately, resulting in its propagation to all following activities until the end of the enclosing scope or the whole business process, where it will be dealt with by the fault handler.

When synchronized communication activity $\text{inv } a \ x \ y$ succeeds, the value received from a , the other end of communication, is assigned to y ; while failed communication also makes the process fail.

$$\begin{aligned} \langle \text{inv } a \ x \ y, (\text{norm}, s, \alpha) \rangle &\rightsquigarrow (\text{norm}, s \oplus \{y \mapsto \nu\}, \alpha) \\ \langle \text{inv } a \ x \ y, (\text{norm}, s, \alpha) \rangle &\rightsquigarrow (\text{fail}, s, \alpha) \end{aligned}$$

where ν is the value achieved through the communication.

The rules for the one-way communications $\text{rec } a \ y$ and $\text{rep } a \ x$ are as follows:

$$\begin{aligned} \langle \text{rec } a \ y, (\text{norm}, s, \alpha) \rangle &\rightsquigarrow (\text{norm}, s \oplus \{y \mapsto \nu\}, \alpha) \\ \langle \text{rec } a \ y, (\text{norm}, s, \alpha) \rangle &\rightsquigarrow (\text{fail}, s, \alpha) \\ \langle \text{rep } a \ x, (f, s, \alpha) \rangle &\rightsquigarrow (f, s, \alpha) \end{aligned}$$

Note that the one-way communications provide an invocation mechanism for external Web services. The $\text{rec } a \ y$ is used to retrieve parameters from other Web services (a). Its effect is to update variable y using the value received from the external Web service. On the contrary, the $\text{rep } a \ x$ replies to other external Web services (a) with the value of x . Thus its effect is just like a skip to the current process.

Rules for sequence and conditional activities are routine:

$$\begin{aligned} &\frac{\langle A_1, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (f_1, s_1, \alpha_1) \quad \langle A_2, (f_1, s_1, \alpha_1) \rangle \rightsquigarrow (f_2, s_2, \alpha_2)}{\langle A_1; A_2, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (f_2, s_2, \alpha_2)} \\ &\frac{s(b) = \text{true} \quad \langle A_1, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (f_1, s_1, \alpha_1)}{\langle \text{if } b \text{ then } A_1 \text{ else } A_2, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (f_1, s_1, \alpha_1)} \\ &\frac{s(b) = \text{false} \quad \langle A_2, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (f_1, s_1, \alpha_1)}{\langle \text{if } b \text{ then } A_1 \text{ else } A_2, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (f_1, s_1, \alpha_1)} \end{aligned}$$

The rule for the parallel composition is as follows:

$$\frac{\begin{aligned} (s_1, s_2) &= \text{split}(s, \text{Var}(A_1), \text{Var}(A_2)) \\ \langle A_1, (\text{norm}, s_1, []) \rangle &\rightsquigarrow (f_1, s'_1, \alpha_1) \\ \langle A_2, (\text{norm}, s_2, []) \rangle &\rightsquigarrow (f_2, s'_2, \alpha_2) \\ f' &= f_1 \wedge f_2 \quad s' = s'_1 \cup s'_2 \quad \alpha' = \text{interleave}(\alpha_1, \alpha_2) \wedge \alpha \end{aligned}}{\langle A_1 \parallel A_2, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (f', s', \alpha')}$$

where the splitting of variable mappings is based on the separation of variable names:

$$\begin{aligned} \text{split}(s, \text{Var}(A_1), \text{Var}(A_2)) &=_{df} \\ &(\{x_1 \mapsto e_1 \mid x_1 \in \text{Var}(A_1)\} \wedge x_1 \mapsto e_1 \in s, \\ &\{x_2 \mapsto e_2 \mid x_2 \in \text{Var}(A_2) \wedge x_2 \mapsto e_2 \in s\}). \end{aligned}$$

in which $\text{Var}(A_1) \cap \text{Var}(A_2) = \emptyset$. And for f_1 and f_2 , $f_1 \wedge f_2$ is defined as

$$f_1 \wedge f_2 =_{df} \begin{cases} \text{norm}, & \text{if } f_1 = \text{norm} \text{ and } f_2 = \text{norm}; \\ \text{fail}, & \text{otherwise.} \end{cases}$$

The initial sub-states s_1 and s_2 are obtained from the overall state s via a splitting operation whose definition is straightforward given that A_1 and A_2 do not share variables, i.e., $\text{Var}(A_1) \cap \text{Var}(A_2) = \emptyset$. The function $\text{interleave}(\alpha_1, \alpha_2)$ returns a merged sequence of α_1 and α_2 by arbitrarily interleaving elements of α_1 and α_2 :

$$\text{interleave}(\alpha_1, \alpha_2) =_{df} [\delta_1, \delta_2, \dots, \delta_{m+n}]$$

where we denote $\alpha_1 = [\delta_{1,1}, \delta_{1,2}, \dots, \delta_{1,m}]$ and $\alpha_2 = [\delta_{2,1}, \delta_{2,2}, \dots, \delta_{2,n}]$, and then the following holds: $\delta_i \in \alpha_1 \wedge \alpha_2, i = 1, 2, \dots, m+n$; $\forall 1 \leq i < j \leq m+n$, if $\delta_i, \delta_j \in \alpha_1, \delta_i = \delta_{1,s}$ and $\delta_j = \delta_{1,t}$, then $s < t$; and the same condition for α_2 .

The execution of a scope $n : \{A ? C : F\}$ may result in two different situations: the execution of A may complete successfully or raise a fault. For the former, the compensation handler will be installed by adding the compensation closure into the compensation context. For the latter, the fault handler is invoked instead.

$$\begin{aligned} &\frac{\langle A, (\text{norm}, s, []) \rangle \rightsquigarrow (\text{norm}, s_1, \alpha_1) \quad s' = s_1 \upharpoonright_{V(n)}}{\langle n : \{A ? C : F\}, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (\text{norm}, s_1, \langle n, s', \alpha_1 \rangle \cdot \alpha)} \\ &\frac{\langle A, (\text{norm}, s, []) \rangle \rightsquigarrow (\text{fail}, s_1, \alpha_1) \quad \langle F, (\text{norm}, s_1, \alpha_1) \rangle \rightsquigarrow (f_2, s_2, \alpha_2)}{\langle n : \{A ? C : F\}, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (f_2, s_2, \alpha_2)} \end{aligned}$$

Here $V(n)$ denotes the set of local variables of scope n , and $s_1 \upharpoonright_{V(n)}$ takes the part of state local to n , which is the snapshot of scope n when it completes execution.

Note that the scope is the only part in the model to deal with faults. Once a fault is propagated from an activity A to its enclosing scope, it will be caught by the relevant fault handler F . If the fault handler of the immediately enclosing scope of A throws the fault again rather than completes the handling, the fault continues its propagation to the next fault

handler, or meets the end of the process. This is elaborated in the rules defined above.

Next comes the definition of compensation. According to the WS-BPEL Specification [4], our compensation looks for the installed compensation closure of corresponding scope, removes it from the compensation context and runs its handler. If the closure is not installed, the invocation behaves like a skip. Since we have actually accumulated the compensation contexts, it turns out simple to execute the handler as below:

$$\frac{n \notin \alpha}{\langle \uparrow n, (\text{norm}, s, \alpha) \rangle \rightsquigarrow (\text{norm}, s, \alpha)} \\ \frac{\sigma = (\text{norm}, s, \alpha_1 \frown [\langle n, s', \beta \rangle] \frown \alpha_2) \quad \langle \mathcal{C}(n), (\text{norm}, s \oplus s', \beta) \rangle \rightsquigarrow (f_1, s_1, \gamma)}{\langle \uparrow n, \sigma \rangle \rightsquigarrow (f_1, s_1, \alpha_1 \frown \alpha_2)}$$

Note that $n \notin \alpha$, defined in last section, means that the compensation handler for n is not installed (hence the closure for n does not appear in α).

The rules for the whole business process are as follows:

$$\frac{\langle A, \sigma \rangle \rightsquigarrow (\text{norm}, s_1, \alpha_1)}{\langle \{ A : F \}, \sigma \rangle \rightsquigarrow (\text{norm}, s_1, \alpha_1)} \\ \frac{\langle A, \sigma \rangle \rightsquigarrow (\text{fail}, s_1, \alpha_1) \quad \langle F, (\text{norm}, s_1, \alpha_1) \rangle \rightsquigarrow (f_2, s_2, \alpha_2)}{\langle \{ A : F \}, \sigma \rangle \rightsquigarrow (f_2, s_2, \alpha_2)}$$

There is no top-level compensation handler in the business process because no one could invoke it if there were any.

4 An Axiomatic System for *BPEL**

As a first step to support mechanized verification for *BPEL** processes, we propose in this section a set of inference rules in the style of a Floyd-Hoare logic.

4.1 Assertion Language

To specify properties for *BPEL** processes, apart from the usual logical operations, we shall make use of some logical constructs that are specific for compensation related reasoning. The syntax for the assertion language *Assn* is:

$$P \in \text{Assn} \\ P ::= \text{true} \mid \text{false} \mid \text{normal} \mid x \odot e \mid \sim P \mid P_\epsilon \mid P \downarrow_V \mid \\ P_{+n} \mid P_{-n} \mid P \uparrow_n \mid P_{*n} \mid P \parallel P \mid P \star P \mid P * P \mid \\ \neg P \mid P \wedge P \mid P \vee P \mid P \Rightarrow P$$

Note that x , e and n denote a variable name, an expression and a scope name, respectively. The $\odot$ denotes a relational operator in $\{=, <, >, \leq, \geq\}$.

In the axiomatic system, each assertion is viewed as a set of states that satisfy the assertion. The semantics for all assertions is given in Figure 2.

Among all assertion constructs, true and false are modeled as the whole and empty sets of states, respectively. Semantics

$$\begin{aligned} \llbracket \text{true} \rrbracket &= \Sigma \\ \llbracket \text{false} \rrbracket &= \emptyset \\ \llbracket x \rrbracket \sigma &= \sigma.2(x) \\ \llbracket \text{normal} \rrbracket &= \{ \sigma \mid \sigma.1 = \text{norm} \} \\ \llbracket e \rrbracket \sigma &= \sigma.2(e) \text{ which is the evaluation result of } e \text{ under state } \sigma \\ \llbracket x \odot e \rrbracket &= \{ \sigma \mid \llbracket x \rrbracket \sigma \odot \llbracket e \rrbracket \sigma \}, \text{ where } \odot \text{ has the semantics of the relational operator} \\ \llbracket \sim P \rrbracket &= \{ (\neg \sigma.1, \sigma.2, \sigma.3) \mid \sigma \in \llbracket P \rrbracket \} \\ \llbracket P_\epsilon \rrbracket &= \{ (\sigma.1, \sigma.2, []) \mid \sigma \in P \} \\ \llbracket P \downarrow_V \rrbracket &= \{ (\sigma.1, \sigma.2 \downarrow_V, \sigma.3) \mid \sigma \in \llbracket P \rrbracket \} \\ \llbracket P_{+n} \rrbracket &= \{ (\sigma.1, \sigma.2, \langle n, \sigma.2 \downarrow_V(n), \sigma.3 \rangle) \mid \sigma \in \llbracket P \rrbracket \} \\ \llbracket P_{-n} \rrbracket &= \{ (\sigma.1, \sigma.2, \alpha) \mid \sigma \in \llbracket P \rrbracket \wedge \alpha = \text{before}(n, \sigma.3) \frown \text{after}(n, \sigma.3) \} \\ \llbracket P \uparrow_n \rrbracket &= \{ \sigma \mid \sigma \in \llbracket P \rrbracket \wedge n \in \sigma.3 \} \\ \llbracket P_{*n} \rrbracket &= \{ \text{firstof}(n, \sigma) \mid \sigma \in \llbracket P \rrbracket \wedge n \in \sigma.3 \} \\ \llbracket P \parallel Q \rrbracket &= \{ (\sigma.1 \wedge \sigma'.1, \sigma.2 \cup \sigma'.2, \alpha) \mid \sigma \in \llbracket P \rrbracket \wedge \sigma' \in \llbracket Q \rrbracket \wedge \alpha = \text{interleave}(\sigma.3, \sigma'.3) \} \\ \llbracket P \star Q \rrbracket &= \{ (\sigma.1.1, \sigma.1.2, \sigma.1.3 \frown \sigma.2.3) \mid \sigma_1 \in \llbracket P \rrbracket \wedge \sigma_2 \in \llbracket Q \rrbracket \} \\ \llbracket P * Q \rrbracket &= \{ (\sigma.1.1, \sigma.1.2, \sigma.2.3) \mid \sigma_1 \in \llbracket P \rrbracket \wedge \sigma_2 \in \llbracket Q \rrbracket \} \\ \llbracket \neg P \rrbracket &= \Sigma \setminus \llbracket P \rrbracket \\ \llbracket P \wedge Q \rrbracket &= \llbracket P \rrbracket \cap \llbracket Q \rrbracket \\ \llbracket P \vee Q \rrbracket &= \llbracket P \rrbracket \cup \llbracket Q \rrbracket \\ \llbracket P \Rightarrow Q \rrbracket &= \llbracket \neg P \vee Q \rrbracket \end{aligned}$$

Fig. 2 Semantics for Assertions

for normal contains all states without fault. Assertion $x \odot e$ can be in forms $x < e$, $x = e$, $x > e$ and so forth, to model the relationship between variable x and expression e .

To facilitate the description, we use here (and below) $\sigma.i$ to denote the i -th element of tuple σ . For instance, given $\sigma = (f, s, \alpha)$, we will have $\sigma.1 = f$, $\sigma.2 = s$ and $\sigma.3 = \alpha$. In the definition, $n \in \sigma.3$, defined in last section, denotes that the compensation handler for scope n is installed in σ . We also use three operations to extract information w.r.t. scope n from compensation context α :

$$\begin{aligned} \text{firstof}(n, \sigma) &=_{df} (\text{norm}, \sigma.2 \oplus s, \beta) \\ &\quad \text{if } \sigma.3 = \alpha_1 \frown [\langle n, s, \beta \rangle] \frown \alpha_2 \wedge n \notin \alpha_1 \\ \text{before}(n, \alpha) &=_{df} \begin{cases} \alpha, & \text{if } n \notin \alpha \\ \alpha_1, & \text{if } \alpha = \alpha_1 \frown [\langle n, s, \beta \rangle] \frown \alpha_2 \wedge n \notin \alpha_1 \end{cases} \\ \text{after}(n, \alpha) &=_{df} \begin{cases} [], & \text{if } n \notin \alpha \\ \alpha_2, & \text{if } \alpha = \alpha_1 \frown [\langle n, s, \beta \rangle] \frown \alpha_2 \wedge n \notin \alpha_1 \end{cases} \end{aligned}$$

Operation $\text{firstof}(n, \sigma)$ extracts from $\alpha = \sigma.3$ the first state snapshot for n , and merges it with $\sigma.2$. In the case $n \notin \sigma.3$, $\text{firstof}(n, \sigma)$ is undefined. $\text{before}(n, \alpha)$ returns the largest prefix of α which contains no closure for scope n , and $\text{after}(n, \alpha)$ returns the sub-sequence of α after the first closure for scope n , or the empty sequence when there is no such closure in α .

Among the semantics for the assertions, some relating to flow, scope, and compensation are worth illustration.

The assertions $P \downarrow_V$ and $P \parallel Q$ are used in verification of flow constructs. In the first one, V is a set of variables and $P \downarrow_V$ restricts the domain of variable mapping $\sigma.2$ (where $\sigma \in \llbracket P \rrbracket$) to V . For example, $(x > 0 \wedge y \leq 0) \downarrow_{\{x\}} = x > 0$. The second one, $P \parallel Q$, enumerates all possible interleaving cases of compensation contexts of states in $\llbracket P \rrbracket$ and $\llbracket Q \rrbracket$, respectively.

The following assertions mainly concern scope and compensation. $\sim P$ reverses all the faulty states in each $\sigma \in \llbracket P \rrbracket$ (from norm to fail and vice versa). This corresponds to the verification of throw activity and fault handler which change the process faulty state. P_ϵ reserves the first and second components of states but empties their compensation contexts. This is useful for verifying scopes whose inner activity A begins with empty compensation context.

Assertion P_{+n} extracts each state σ from set $\llbracket P \rrbracket$, sets its compensation context to the closure $\langle n, \sigma.2 \downarrow_{V(n)}, \sigma.3 \rangle$, and forms a new set with all of these states.

As its form suggests, P_{-n} performs an “elimination” of scope name n “from” the elements in $\llbracket P \rrbracket$. It extracts first the compensation context α from each state of $\llbracket P \rrbracket$, then finds the first compensation closure with name n , and removes it to form a new context α . If there is no such closure found, then α will be the original context. The semantics of P_{-n} is the set of states with these newly formed α .

What $P_{\uparrow n}$ does is, informally, to “restrict” $\llbracket P \rrbracket$ to the set of states in which the compensation context contains a closure with name n , P_{*n} “locates” the first occurrence of the closure with name n in each state in $\llbracket P \rrbracket$, and forms a set of states from these closures.

$P \star Q$ and $P \ast Q$ are for compensation contexts concatenation and replacement between assertions, respectively. The first appends the compensation contexts within Q ’s model to those of P ’s, to accumulate new compensation closure based on old ones, according to scope’s behavior. The second discards directly the compensation contexts of the states in P ’s semantics, because of the manner of compensation handlers.

An assertion is modeled as a set containing all the states which satisfy it. Thus we define

$$\sigma \models P =_{df} \sigma \in \llbracket P \rrbracket.$$

A specification in our system takes the ordinary form $\{P\} A \{Q\}$, where $P, Q \in Assn$ and $A \in \mathbb{P}$ is an activity.

One thing notable is that a business process may communicate via activities *inv*, *rec* and *rep* with external processes, which are essentially other Web services within the same application or from third party. As a result, whether a business process behaves in a desired way might depend on the external processes being interacted with. Hence, a business process is more like an open system which makes the verification problem rather challenging. Our proposal is to verify each business process separately according to certain dependency order in the first step. We assume that specifications for communication activities are available in the verification of one business process. When all relevant business processes have been verified separately, we can then check the consistency of

all the assumptions made on communication activities. In this paper, we focus only on the verification of individual business processes.

Remembering that in the operational semantics for communications with external Web services, we have addressed that their behaviors can be arbitrary, either to deliver a value or to fail. However, to verify a business process involving communications more precisely, we need to put more restrictions over semantics of the communications. These restrictions take the form of a set of specifications $\{P\} c \{Q\}$, where each c is any one of *inv* $a \ x \ y$, *rec* $a \ y$, or *rep* $a \ x$, representing a communication that might be executed by the process with the environment. We use T to denote a set of such specifications and use it as a context of the verification rules. For example, for a specification $\{P\} \text{inv } a \ x \ y \{Q\} \in T$, the precondition P acts as an assertion imposed on the current process to ensure that information sent out (the value of x) satisfies the requirement of the environment, while Q acts as an assumption made on the environment: the result sent back by the environment (final value of y) satisfies the constraint described by Q , with a possible substitution of the communication channel and variable names.

The proof rules in our verification system are of the form $\mathcal{C}, T \vdash \{P\} A \{Q\}$, where $\mathcal{C}$, defined earlier, is the mapping from scope names to associated compensation handlers, and T is the set of specifications defined above. We shall now present the syntax-directed proof rules in our logic.

4.2 Consequence Rule

The only structural rule in our axiomatic system is the consequence rule for precondition weakening and postcondition strengthening:

$$\frac{P \Rightarrow P' \quad \mathcal{C}, T \vdash \{P'\} A \{Q'\} \quad Q' \Rightarrow Q}{\mathcal{C}, T \vdash \{P\} A \{Q\}} \quad (conseq)$$

4.3 BPEL*-specific Rules

The rules for skip and assignment are simple:

$$\begin{aligned} \mathcal{C}, T \vdash \{P\} \text{skip} \{P\} & \quad (skip) \\ \mathcal{C}, T \vdash \{\text{normal} \wedge P[e/x]\} x := e \{P\} & \quad (assign) \end{aligned}$$

The rule for throw is clear too:

$$\mathcal{C}, T \vdash \{P\} \text{throw} \{\neg \text{normal} \wedge (P \vee \sim P)\} \quad (throw)$$

Here we do not need to care whether the pre-condition is normal, because the type of fault is not in the range of our current consideration.

For the basic communication activities, the rules need to use their assumed specifications in T . For the convenience of description, we assume the variable names in the pre- and postconditions are correspondent with those used in the invocations. Meanwhile, as is stated in former section, in the verification of the process, a triple $\{P\} A \{Q\}$ in T can

also be used to verify a triple whose pre- and postcondition have the same denotation of compensation contexts, such as $\{P \star R\} A \{Q \star R\}$. And in this situation it must be guaranteed that the denotations of compensation contexts in both pre- and postcondition are the same.

If the environment can be modeled as a subset of normal, then *rec* sets the variable's value to what the specification denotes. Or it just propagates the fault.

$$\frac{\begin{array}{c} \{\text{normal}\} \text{rec } a \ v \ \{Q\} \in T \\ \neg \text{normal} \Rightarrow Q[v/y] \end{array}}{\mathcal{C}, T \vdash \{\text{true}\} \text{rec } a \ y \ \{Q[v/y]\}} \quad (\text{rec})$$

where $Q[v/y]$ is an assertion formed by substituting each occurrence of v in Q by y , for filling the gap between the specification in T and the current process. Because of *rep*'s analogous behavior to *skip*, its rule is also the same.

$$\mathcal{C}, T \vdash \{P\} \text{rep } a \ x \ \{P\} \quad (\text{rep})$$

The semantics of two-way invocation is simple:

$$\frac{\{P\} \text{inv } a \ u \ v \ \{Q\} \in T}{\mathcal{C}, T \vdash \{P\} \text{inv } a \ x \ y \ \{Q[u, v/x, y]\}} \quad (\text{inv})$$

Note that these rules depend on T – the set of specifications assumed on communication activities.

The rules for control structures are as follows.

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow Q \\ \mathcal{C}, T \vdash \{\text{normal} \wedge P\} A \{R\} \\ \mathcal{C}, T \vdash \{R\} B \{Q\} \end{array}}{\mathcal{C}, T \vdash \{P\} A; B \{Q\}} \quad (\text{seq})$$

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow Q \\ \mathcal{C}, T \vdash \{\text{normal} \wedge P \wedge b\} A \{Q\} \\ \mathcal{C}, T \vdash \{\text{normal} \wedge P \wedge \neg b\} B \{Q\} \end{array}}{\mathcal{C}, T \vdash \{P\} \text{if } b \text{ then } A \text{ else } B \{Q\}} \quad (\text{if})$$

where b is a boolean expression of the form $x \odot e$.

Since we assume that the different parallel flows share no variables, the rule for the parallel structures is given as

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow (Q_1 \parallel Q_2) \star P \\ \mathcal{C}, T \vdash \{P_\epsilon\}_{V_1} A \{Q_1\} \\ \mathcal{C}, T \vdash \{P_\epsilon\}_{V_2} B \{Q_2\} \end{array}}{\mathcal{C}, T \vdash \{P\} A \parallel B \{(Q_1 \parallel Q_2) \star P\}} \quad (\text{flow})$$

where V_1 and V_2 are disjoint variable sets and A and B only modify variables in V_1 and V_2 , respectively.

Now we present the two most significant rules, which reveal the essential features of our language. The rule for scopes is as follows:

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow Q \\ \mathcal{C}, T \vdash \{\text{normal} \wedge P_\epsilon\} A \{R\} \\ (\text{normal} \wedge R)_{+n} \star P \Rightarrow Q \\ \mathcal{C}, T \vdash \{\sim(\neg \text{normal} \wedge R)\} F \{Q\} \end{array}}{\mathcal{C}, T \vdash \{P\} n : \{A ? C : F\} \{Q\}} \quad (\text{scope})$$

Note that the rule (*scope*) captures two cases. One stands for the scenario where a fault occurs in A . In that case the control transfers to the fault handler, and the compensation handler for scope n is not installed. The other is for the normal completion of A and the concatenation of this scope's compensation context to the process state.

Then the most intricate rule in our system, the named compensation, comes as follows:

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow Q \\ \neg P_{\downarrow n} \wedge P \Rightarrow Q \\ \mathcal{C}, T \vdash \{(P_{\downarrow n})_{*n}\} C(n) \{R\} \\ R \star P_{-n} \Rightarrow Q \end{array}}{\mathcal{C}, T \vdash \{P\} \uparrow_n \{Q\}} \quad (\text{compensate})$$

In this rule, the behavior of a named compensation is depicted with the relevant compensation handler. If the pre-condition does not entail a scope name n , the post-condition must be automatically satisfied. Otherwise, the snapshots' set (as the pre-condition for the compensation handler) is extracted and the post-condition is a combination of the fault and variable mapping states after the handler's execution, and the compensation context with the elimination of the first compensation closure named n .

At last is the rule for the whole business process:

$$\frac{\begin{array}{c} \mathcal{C}, T \vdash \{P\} A \{R\} \\ (\text{normal} \wedge R) \Rightarrow Q \\ \mathcal{C}, T \vdash \{\sim(\neg \text{normal} \wedge R)\} F \{Q\} \end{array}}{\mathcal{C}, T \vdash \{P\} \llbracket A : F \rrbracket \{Q\}} \quad (\text{bp})$$

5 Soundness

This section is devoted to the soundness of our verification system. We will first give two definitions and then formalize the soundness theorem and its proof.

Definition 1 (Validity). We denote that a triple $\{P\} A \{Q\}$ is *valid* under $\mathcal{C}, T$, i.e. $\mathcal{C}, T \models \{P\} A \{Q\}$, if for all $\sigma \in \Sigma$, if $\sigma \models P$ and $\langle A, \sigma \rangle \rightsquigarrow \sigma'$ for some σ' , then $\sigma' \models Q$.

Definition 2 (Soundness). Our verification system for BPEL^* is *sound* if all provable specifications are indeed valid, that is, if $\mathcal{C}, T \vdash \{P\} A \{Q\}$, then $\mathcal{C}, T \models \{P\} A \{Q\}$.

The theorem for soundness can be stated as below:

Theorem 1. *The Hoare logic for BPEL^* presented in this paper is sound.*

As is indicated by Definition 2 above, we need to show that, for any P, A, Q , if $\mathcal{C}, T \vdash \{P\} A \{Q\}$, then $\mathcal{C}, T \models \{P\} A \{Q\}$. The proof can be accomplished by structural induction over A .

Proof. The verification of $\mathcal{C}, T \vdash \{P\} A \{Q\}$ (denoted as t) can be a process such as

$$\frac{\begin{array}{c} \text{(some premises)} \\ \vdots \\ \hline \end{array} \quad \dots \quad \begin{array}{c} \text{(some other premises)} \\ \vdots \\ \hline \end{array}}{\mathcal{C}, T \vdash \{P\} A \{Q\}}$$

From the perspective of backwards reasoning, a rule r should be utilized on t according to A 's structure, and from this rule some other premises need to be verified with similar backward verifications until all the premises are axioms or known facts. As an illustration, if A is $\mathcal{C}, T \vdash \{P\} \llbracket A_1 : F_1 \rrbracket \{Q\}$, then we must verify $\mathcal{C}, T \vdash \{P\} A \{R\}$, ($\text{normal} \wedge R \Rightarrow Q$) and $\mathcal{C}, T \vdash \{\sim(\neg \text{normal} \wedge R)\} F \{Q\}$, according to the (bp) rule. Hence the last rule r used to verify t depends on the structure of the activity A . Therefore, the following cases are organized according to the structure of A , which is equivalent to r to some extent.

- Case (skip). The last rule r for this is (skip):

$$\mathcal{C}, T \vdash \{P\} \text{skip} \{Q\}$$

Since $\langle \text{skip}, \sigma \rangle \rightsquigarrow \sigma$, it is easy to see that rule (skip) is sound in our system.

- Case ($x := e$). The corresponding rule is (assign):

$$\mathcal{C}, T \vdash \{\text{normal} \wedge P[e/x]\} x := e \{P\}$$

The proof for rule (assign) simply follows the canonical Hoare logic's proof using the Substitution Theorem and thus is omitted here.

- Case (throw). The last rule to apply is (throw):

$$\mathcal{C}, T \vdash \{P\} \text{throw} \{\neg \text{normal} \wedge (P \vee \sim P)\}$$

Take any σ such that $\sigma \models P$. If $\sigma.1 = \text{fail}$, then we have $\langle \text{throw}, \sigma \rangle \rightsquigarrow \sigma$ and $\sigma \models \neg \text{normal} \wedge P$. Otherwise, if $\sigma.1 = \text{norm}$, then we have $\langle \text{throw}, \sigma \rangle \rightsquigarrow \sigma'$ where $\sigma' = (\text{fail}, \sigma.2, \sigma.3)$, and $\sigma' \models \neg \text{normal} \wedge \sim P$. So we get $\mathcal{C}, T \models \{P\} \text{throw} \{\neg \text{normal} \wedge (P \vee \sim P)\}$.

- Case ($\text{rec } a \ y$).

$$\frac{\begin{array}{c} \{\text{normal}\} \text{rec } a \ v \ \{Q\} \in T \\ \neg \text{normal} \Rightarrow Q[v/y] \end{array}}{\mathcal{C}, T \vdash \{\text{true}\} \text{rec } a \ y \ \{Q[v/y]\}}$$

For the proof of the rule (rec), if $\sigma \models \text{normal}$, then since $\{\text{normal}\} \text{rec } a \ v \ \{Q\}$ is already known for the communication, the model of postcondition $Q[v/y]$ should contain the final state transited from σ (either $(\text{norm}, \sigma.2 \oplus \{y \mapsto v\}, \sigma.3)$ or $(\text{fail}, \sigma.2, \sigma.3)$, according to the communication's behavior). Otherwise if $\sigma \models \neg \text{normal}$, then from the semantics for $\neg \text{normal} \Rightarrow Q[v/y]$ we know $\sigma \models Q[v/y]$. Therefore we conclude in this case.

- Case ($\text{rep } a \ x$).

$$\mathcal{C}, T \vdash \{P\} \text{rep } a \ x \ \{P\}$$

Since the communication of reply does not change the process status, rule (rep) shares the same proof of skip 's.

- Case ($\text{inv } a \ x \ y$).

$$\frac{\{P\} \text{inv } a \ u \ v \ \{Q\} \in T}{\mathcal{C}, T \vdash \{P\} \text{inv } a \ x \ y \ \{Q[u, v/x, y]\}}$$

The proof can be completed in the similar way as that of $\text{rec } a \ y$.

- Case ($A; B$). The rule applied in this case is (seq):

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow Q \\ \mathcal{C}, T \vdash \{\text{normal} \wedge P\} A \{R\} \\ \mathcal{C}, T \vdash \{R\} B \{Q\} \end{array}}{\mathcal{C}, T \vdash \{P\} A; B \{Q\}}$$

The proof for rule (seq) is classical, except that the faulty state is taken into consideration first. That is, for any state $\sigma \models P$, if $\sigma.1 = \text{fail}$, then $\sigma \models \neg \text{normal} \wedge P$ and thus $\sigma \models Q$. If not, then take σ^* as $\langle A, \sigma \rangle \rightsquigarrow \sigma^*$, we have $\sigma^* \models R$. And from $\langle B, \sigma^* \rangle \rightsquigarrow \sigma'$ and the inductive assumption, it holds that $\sigma' \models Q$.

- Case ($\text{if } b \text{ then } A \text{ else } B$). In this case the condition rule (if) is applied:

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow Q \\ \mathcal{C}, T \vdash \{\text{normal} \wedge P \wedge b\} A \{Q\} \\ \mathcal{C}, T \vdash \{\text{normal} \wedge P \wedge \neg b\} B \{Q\} \end{array}}{\mathcal{C}, T \vdash \{P\} \text{if } b \text{ then } A \text{ else } B \{Q\}}$$

The proof of the condition rule is also similar as the classical one. Except for the abnormal state, consider any σ where $\sigma.1 = \text{norm}$. Then no matter whether $\sigma \models \text{normal} \wedge P \wedge b$ or $\sigma \models \text{normal} \wedge P \wedge \neg b$, for some σ' and σ^* that $\langle A, \sigma \rangle \rightsquigarrow \sigma'$ in the first case and $\langle B, \sigma \rangle \rightsquigarrow \sigma^*$ in the second, we always get $\sigma' \models Q$ and $\sigma^* \models Q$ from inductive assumption.

- Case ($A \parallel B$). The last rule used is (flow):

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow (Q_1 \parallel Q_2) * P \\ \mathcal{C}, T \vdash \{P_{\epsilon_1} \}_{v_1} A \{Q_1\} \\ \mathcal{C}, T \vdash \{P_{\epsilon_2} \}_{v_2} B \{Q_2\} \end{array}}{\mathcal{C}, T \vdash \{P\} A \parallel B \{(Q_1 \parallel Q_2) * P\}}$$

Take any $\sigma \models \text{normal} \wedge P$ (the case for $\sigma.1 = \text{fail}$ is like other rules), from the premises and the inductive assumption we know that $\langle A, (\sigma.1, \sigma.2]_{s_1}, [] \rangle \rightsquigarrow \sigma'_1$ and $\langle B, (\sigma.1, \sigma.2]_{s_2}, [] \rangle \rightsquigarrow \sigma'_2$, for some $\sigma'_1 \models Q_1$, $\sigma'_2 \models Q_2$. Hence $(\sigma'_1.1 \wedge \sigma'_2.1, \sigma'_1.2 \cup \sigma'_2.2, \text{interleave}(\sigma'_1.3, \sigma'_2.3) \wedge \sigma.3) \models (Q_1 \parallel Q_2) * P$, and thus we conclude in this case.

- Case $(n : \{A ? C : F\})$. Rule (*scope*) is the last rule applied in the proof for $\mathcal{C}, T \vdash \{P\} n : \{A ? C : F\} \{Q\}$:

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow Q \\ \mathcal{C}, T \vdash \{\text{normal} \wedge P_\epsilon\} A \{R\} \\ (\text{normal} \wedge R)_{+n} \star P \Rightarrow Q \\ \mathcal{C}, T \vdash \{\sim(\neg \text{normal} \wedge R)\} F \{Q\} \end{array}}{\mathcal{C}, T \vdash \{P\} n : \{A ? C : F\} \{Q\}}$$

The following cases are discussed for all $\sigma \models P$.

- If $\sigma.1 = \text{fail}$, from inductive assumption and the premise $\neg \text{normal} \wedge P \Rightarrow Q$, we have $\sigma \models \neg \text{normal} \wedge P$, and thus $\sigma \models Q$.
- If $\sigma.1 = \text{norm}$, then take $\sigma_\epsilon = (\sigma.1, \sigma.2, [])$, and hence we have $\sigma \models \text{normal} \wedge P_\epsilon$. With inductive assumption and the premise, denoting σ'_ϵ as $\langle A, \sigma_\epsilon \rangle \rightsquigarrow \sigma'_\epsilon$, then $\sigma'_\epsilon \models R$ is achieved.
 - * If $\sigma'_\epsilon.1 = \text{norm}$, then $\sigma'_\epsilon \models \text{normal} \wedge R$, and $\sigma'_{+n} = (\sigma'_\epsilon.1, \sigma'_\epsilon.2, \langle n, \sigma'_\epsilon.2 \upharpoonright_{V(n)}, \sigma'_\epsilon.3 \rangle) \models (\text{normal} \wedge R)_{+n}$, and still $\sigma' = (\sigma'_{+n}.1, \sigma'_{+n}.2, \sigma'_{+n}.3 \cdot \sigma.3) \models (\text{normal} \wedge R)_{+n} \star P$. We get $\sigma' \models Q$ from the last implication.
 - * If $\sigma'_\epsilon.1 = \text{fail}$, then $\sigma'_F \models \sim(\neg \text{normal} \wedge R)$ where $\sigma'_F = (\text{norm}, \sigma'_\epsilon.2, \sigma'_\epsilon.3)$. From the inductive assumption and the semantics $\langle F, \sigma'_F \rangle \rightsquigarrow \sigma'$ for some σ' , we have $\sigma' \models Q$. This completes our proof for *scope*.

- Case $(\uparrow n)$:

$$\frac{\begin{array}{c} \neg \text{normal} \wedge P \Rightarrow Q \\ \neg P_{\uparrow n} \wedge P \Rightarrow Q \\ \mathcal{C}, T \vdash \{(P_{\uparrow n})_{*n}\} \mathcal{C}(n) \{R\} \\ R \star P_{-n} \Rightarrow Q \end{array}}{\mathcal{C}, T \vdash \{P\} \uparrow n \{Q\}}$$

For the rule of compensation, consider any $\sigma \models P$ in the following cases.

- If $\sigma.1 = \text{fail}$, then directly we have $\sigma \models Q$.
 - If $\sigma.1 \neq \text{fail}$ and there are no compensation closures named n in σ 's compensation context, then $\sigma \models \neg P_{\uparrow n} \wedge P$ by definition, and thus $\sigma \models Q$ which conforms to the operational semantics.
 - Otherwise, we need to run the compensation handler named n . Denote $\sigma_n = \text{firstof}(n, \sigma)$, and we have $\sigma_n \models (P_{\uparrow n})_{*n}$ and hence $\langle \mathcal{C}(n), \sigma_n \rangle \rightsquigarrow \sigma'_n$ for some σ'_n , while $\sigma'_n \models R$. Then take $\sigma = (f, \sigma, \alpha_1 \wedge [\langle n, \sigma^*, \beta \rangle] \wedge \alpha_2)$, and we have $\sigma' = (\sigma'_n.1, \sigma'_n.2, \alpha_1 \wedge \alpha_2) \models R \star P_{-n}$, and thus $\sigma' \models Q$. From all discussion above, we conclude this case.
- Case $(\llbracket A : F \rrbracket)$. The last rule applied in the proof for the whole business process will be the rule (*bp*):

$$\frac{\begin{array}{c} \mathcal{C}, T \vdash \{P\} A \{R\} \\ (\text{normal} \wedge R) \Rightarrow Q \\ \mathcal{C}, T \vdash \{\sim(\neg \text{normal} \wedge R)\} F \{Q\} \end{array}}{\mathcal{C}, T \vdash \{P\} \llbracket A : F \rrbracket \{Q\}}$$

It is similar as the scope rule with compensation handler eliminated. For any $\sigma \models P$ and $\langle A, \sigma \rangle \rightsquigarrow \sigma'$ for some σ' there are the following two cases:

- $\sigma'.1 = \text{norm}$. From $(\text{normal} \wedge R) \Rightarrow Q$, we know that $\sigma' \models Q$.
- $\sigma'.1 = \text{fail}$. If $\langle F, (\text{norm}, \sigma'.2, \sigma'.3) \rangle \rightsquigarrow \sigma^*$ for some σ^* , then we have $\sigma^* \models Q$ from the premise $\mathcal{C}, T \vdash \{\sim(\neg \text{normal} \wedge R)\} F \{Q\}$.

- Besides the aforesaid A 's possible structures directly related to rules, sometimes we may be not able to verify $\mathcal{C}, T \vdash \{P\} A \{Q\}$ with an existing rule but can verify $\mathcal{C}, T \vdash \{P'\} A \{Q'\}$ where P' is weaker than P and/or Q' is stronger than Q . Thus the structural rule (*conseq*) is employed in such cases:

$$\frac{P \Rightarrow P' \quad \mathcal{C}, T \vdash \{P'\} A \{Q'\} \quad Q' \Rightarrow Q}{\mathcal{C}, T \vdash \{P\} A \{Q\}}$$

For all $\sigma \models P$ and $\langle A, \sigma \rangle \rightsquigarrow \sigma'$ for some σ' , we have $\sigma \models P'$ from $P \Rightarrow P'$ and also $\langle A, \sigma \rangle \rightsquigarrow \sigma^*$ for some $\sigma^* \models Q'$. Then from $Q' \Rightarrow Q$ we get $\sigma^* \models Q$. Hence σ^* is the σ' we need and the proof for this rule is completed.

Above are all the cases of our structural induction, and each of them is proven to be sound. Hence this completes our proof for the soundness. $\square$

6 Example

In this section a purchase example is exhibited to illustrate the verification of a real business process, which is a modified version of that in [13].

The general flow of the example is as follows. First the process receives the price for each single item (stored in variable p) and the class of the customer from other service with communication (into variable y). Then it decides the discount ratio according to the customer class, and receives the amount of items to store in t . After having all the items purchased, it computes the shipping fare according to the value of t . At last the real average price (including shipping cost) for each item is calculated and replied, which may incur fault and hence call for compensation.

This business process, denoted as BP , is written in $BPEL^*$ below.

```

{
  n1 : {rec a p; q := p ? p := -p : skip};
  rec b y;
  if y = 1 then
    n2 : {p := p × 0.5 ? p := p × 2 : skip}
  else
    n3 : {p := p × 0.8 ? p := p × 1.25 : skip};
  n4 : {rec c t; p := p × t ? p := p/t : skip};
  if t > 500 then
    n5 : {p := p + 500 ? p := p - 500 : skip}
}

```

```

else
   $n_6 : \{p := p + t ? p := p - t : \text{skip}\};$ 
  if  $t > 0$  then  $p := p/t$ ; rep  $d$   $p$  else throw
  :  $\neg n_6; \neg n_5; \neg n_4; \neg n_3; \neg n_2; \neg n_1$ 
}

```

The specification for us to verify is $\{\text{normal}\} BP \{Q\}$ where Q is $p=q/2+500/t \vee p=0.8q+500/t \vee p=q/2+1 \vee p=0.8q+1 \vee p=-q$. The first four parts of the disjunctions in Q present the different situations of discount ratio and shipping fare, while the last $p=-q$ is the case where a fault is compensated. This specification states that, if BP starts in a normal state and terminates at last, it should establish the postcondition Q , provided that the specifications of the communication activities are as follows:

$$\begin{aligned} & \{\text{normal}\} \text{rec } a \ y \ \{\text{normal} \wedge y > 0\} \\ & \{\text{normal}\} \text{rec } b \ y \ \{\text{normal} \wedge (y=1 \vee y=2)\} \\ & \{\text{normal}\} \text{rec } c \ y \ \{\text{normal} \wedge y \neq 0\} \end{aligned}$$

Here we give an outline of the verification for BP with the backwards searching strategy. First, for the whole business process, we use the rule of bp to get three subgoals G_1, G_2, G_3 for further verification:

$$\begin{aligned} G_1 &: \mathcal{C}, T \vdash \{\text{normal}\} A \{R\} \\ G_2 &: (\text{normal} \wedge R) \Rightarrow Q \\ G_3 &: \mathcal{C}, T \vdash \{\sim(\neg \text{normal} \wedge R)\} F \{Q\} \end{aligned}$$

where A stands for the codes before the last : in the process, Q is the postcondition we want to verify, F represents the six compensations ($\neg n_6; \neg n_5; \neg n_4; \neg n_3; \neg n_2; \neg n_1$), and R should be both strong enough to derive Q and still sufficiently weak as F 's precondition to get Q . A possible R can be the assertion below:

$$(\text{normal} \wedge (p=q/2+500/t \vee p=0.8q+500/t \vee p=q/2+1 \vee p=0.8q+1)) \vee (\neg \text{normal} \wedge \sim P_6)$$

where P_6 is

$$\begin{aligned} & \text{normal} \wedge \\ & ((t \leq 500 \Rightarrow ((y=1 \Rightarrow p=0.5qt + t) \wedge \\ & \quad (y \neq 1 \Rightarrow p=0.8qt + t))_{+n_5} \wedge \\ & \quad t > 500 \Rightarrow ((y=1 \Rightarrow p=0.5qt + 500) \wedge \\ & \quad (y \neq 1 \Rightarrow p=0.8qt + 500))_{+n_6}) \star P_5 \end{aligned}$$

where P_5 stands for the set of states for compensation accumulated in the previous execution of the process, from a semantics perspective (and so are the other assertions to be depicted below); its definition is

$$\text{normal} \wedge (((y=1 \Rightarrow p=0.5qt) \wedge (y \neq 1 \Rightarrow p=0.8qt))_{+n_4} \star P_4)$$

where P_4 is

$$\text{normal} \wedge (((y=1 \Rightarrow (p=0.5q)_{+n_2}) \wedge (y \neq 1 \Rightarrow (p=0.8q)_{+n_3})) \star P_2)$$

where P_2 is

$$(\text{normal} \wedge p=q)_{+n_1} \star \text{normal}$$

and the semantics and the derivations of these assertions will be introduced in the following descriptions. With them as bridges we will try to verify the three subgoals separately.

For the first subgoal G_1 , it can still be divided into six subgoals, since A is a sequence made up of six other activities, including two scopes, one basic communication activity and three conditional judgments. We will denote these activities as $A_1, A_2, \dots, A_6$ according to their original orders in BP , and call these six subgoals $G_{1,i}, i = 1, 2, \dots, 6$, defined as below:

$$G_{1,i} : \mathcal{C}, T \vdash \{P_i\} A_i \{P_{i+1}\}$$

where $i = 1, 2, \dots, 6, P_1 = \text{normal}$ and $P_7 = R$. (Note that the $P_2, \dots, P_5$ are what have been described above.) We will demonstrate the verification of each subgoal.

For $G_{1,1}$, since A_1 is the scope n_1 and the precondition is normal, we will use the *scope* rule to divide it further into three subgoals:

$$\begin{aligned} G_{1,1,1} &: \{\text{normal}_\epsilon\} \text{rec } a \ p; \ q := p \{P_{1,1}\} \\ G_{1,1,2} &: (\text{normal} \wedge P_{1,1})_{+n_1} \star \text{normal} \Rightarrow P_2 \\ G_{1,1,3} &: \{\sim(\neg \text{normal} \wedge P_{1,1})\} \text{skip} \{P_2\} \end{aligned}$$

For the first subgoal, the rules *seq*, *rec* and *assign* are used, to get that $P_{1,1}$ is $\text{normal} \wedge p=q$. With this result and the second subgoal the strongest P_2 is derived as $(\text{normal} \wedge p=q)_{+n_1} \star \text{normal}$. For the third subgoal, since $\neg \text{normal} \wedge P_{1,1} = \text{false}$, it holds automatically. Then $G_{1,1}$ is verified with the postcondition P_2 , that is, $(\text{normal} \wedge p=q)_{+n_1} \star \text{normal}$.

Next we will examine $G_{1,2}$. Here the rule of *rec* is applied to P_2 and *rec* $b \ y$, with the result of postcondition P_3 which is $(y=1 \vee y=2) \wedge P_2$.

Subgoal $G_{1,3}$ concerns the first if construct of the process, and its verification is an application of rule *if* with the result of two other subgoals

$$\begin{aligned} G_{1,3,1} &: \mathcal{C}, T \vdash \{\text{normal} \wedge P_3 \wedge y=1\} A_{1,3} \{P_4\} \\ G_{1,3,2} &: \mathcal{C}, T \vdash \{\text{normal} \wedge P_3 \wedge \neg y=1\} B_{1,3} \{P_4\} \end{aligned}$$

where $A_{1,3}$ and $B_{1,3}$ are scopes n_2 and n_3 , respectively. They can be verified similarly as n_1 (using rules *scope* and *assign*), and we get the postcondition P_4 :

$$\begin{aligned} & \text{normal} \wedge q=X \wedge \\ & ((y=1 \Rightarrow (p=0.5q)_{+n_2}) \wedge \\ & \quad y \neq 1 \Rightarrow (p=0.8q)_{+n_3}) \star P_2 \end{aligned}$$

$G_{1,4}$ is to verify the Hoare triple for scope n_4 . Following similar way of $G_{1,1}$ it can be verified with P_5 as $\text{normal} \wedge (((y=1 \Rightarrow p=0.5qt) \wedge (y \neq 1 \Rightarrow p=0.8qt))_{+n_4} \star P_4)$.

$G_{1,5}$ again seeks the verification of the second if construct. With the approach like that of $G_{1,3}$ (splitting it into two subgoals) we can achieve P_6 :

$$\begin{aligned} & \text{normal} \wedge \\ & ((t \leq 500 \Rightarrow ((y=1 \Rightarrow p=0.5qt + t) \wedge \\ & \quad (y \neq 1 \Rightarrow p=0.8qt + t))_{+n_5} \wedge \\ & \quad t > 500 \Rightarrow ((y=1 \Rightarrow p=0.5qt + 500) \wedge \\ & \quad (y \neq 1 \Rightarrow p=0.8qt + 500))_{+n_6}) \star P_5 \end{aligned}$$

The last subgoal, $G_{1,6}$, is slightly different from the former two if's. It is also first divided into two subgoals:

$$\begin{aligned} G_{1,6,1} : \mathcal{C}, T \vdash \{ \text{normal} \wedge P_6 \wedge t > 0 \} p := p/t; \text{rep } d \ p \ \{ R \} \\ G_{1,6,2} : \mathcal{C}, T \vdash \{ \text{normal} \wedge P_6 \wedge \neg t > 0 \} \text{throw} \{ R \} \end{aligned}$$

in which the first subgoal's verification is as the former ones, with the postcondition $p = q/2 + 500/t \vee p = 0.8q + 500/t \vee p = q/2 + 1 \vee p = 0.8q + 1$. (Note that we omit the part for compensation and present a weaker assertion here.) However, the second one uses the *throw* rule to force a conjunction of $\neg \text{normal}$ with the precondition. Therefore the whole postcondition for A, R , is as follows:

$$(\text{normal} \wedge (p = q/2 + 500/t \vee p = 0.8q + 500/t \vee p = q/2 + 1 \vee p = 0.8q + 1)) \vee (\neg \text{normal} \wedge \sim P_6)$$

It is clear that a conjunction of *normal* and this R automatically implies Q , which is demanded in the subgoal G_2 . So the remaining work is to verify the subgoal G_3 .

G_3 equals to $\mathcal{C}, T \vdash \{ t < 0 \wedge P_6 \} F \{ Q \}$, where F is the sequence of six compensations. Similarly, this can be divided into six subgoals using the *seq* rule, and each subgoal is solved equally with rule *compensate*. We will illustrate its usage with the first two compensations for scopes n_6 and n_5 , and the others are the same as these two.

Since $t < 0$ implies that $t \leq 500$, it can be deduced that the compensation context for n_6 must be installed, and thus we have only two possible cases to consider ($y = 1; y \neq 1$). We now take the first case as an example, in which the precondition of these compensations can be reduced as

$$\begin{aligned} & \text{normal} \wedge y = 1 \wedge t \leq 500 \wedge \\ & (p = 0.5qt + t)_{+n_6} \star (p = 0.5qt)_{+n_4} \star \\ & (p = 0.5q)_{+n_2} \star (p = q)_{+n_1} \star \text{normal} \end{aligned}$$

where we denote the $\star$ as right-associative to prevent excess parentheses. Using once the rule *compensate* we get

$$\begin{aligned} & \text{normal} \wedge y = 1 \wedge t \leq 500 \wedge \\ & (p = 0.5qt)_{+n_4} \star (p = 0.5q)_{+n_2} \star \\ & (p = q)_{+n_1} \star \text{normal} \end{aligned}$$

to remove the compensation context of n_6 from the states (on the level of semantics).

Then for the subgoal concerning n_5 , since in this case it is not installed in the compensation context (which can be seen from the structure of the assertion), its effect, due to rule *compensate*, is like a skip.

Therefore we use the rule four times more on γn_i , $i = 4 \dots 1$ respectively, to verify each remaining subgoal and gain the final assertion

$$\text{normal} \wedge y = 1 \wedge t \leq 500 \wedge p = -q$$

which implies that $p = -q$, and hence Q . This completes the whole process' verification.

nested transactions [15]. There are a few attempts to formalize workflow languages [3, 16–18], and still many of them are about compensation.

On the semantics of such languages there are many investigations. Qiu and *et al.* [6] provided a formal operational semantics for a simplified version of BPEL4WS to specify the execution path of a process with possible compensation behavior. In that work, they adopted an abridgment of BPEL4WS which was followed by a series of works with similar objective. Then they provided a formal semantics for fault handling, compensations, and parallel processes with respect to the original informal description by the BPEL4WS Specification [19], with a clarification to some of its elusive parts. Meanwhile they proved the completeness of the deduction of their semantics. Besides their work, there are still many others to formalize the semantics of these languages from different perspectives. Bruni and *et al.* [20] presented an operational semantics for a series of languages, which included the compensation concept. Pu and *et al.* [7] also presented an abridged edition of WS-BPEL, discussed its operational semantics, and defined the equivalence between two processes with its proposed n -bi-simulation. He and *et al.* [9] focused on the process equivalence from the perspective of an observation-oriented model and its algebraic laws. Zhu and *et al.* [10] made a link among different semantics (operational, denotational and algebraic) of the WS-BPEL language with the approach of the unifying theories of programming. Some of the semantics proposed in these works may also be enhanced as suitable underlying semantics for our verification system, while our semantics is mainly distinguished from theirs in that it incorporates variable mappings in both program runtime states and compensation contexts, enabling concrete variable valuations to be stored in the semantics.

Apart from the work on semantic models, researchers have also tried to model and verify the WS-BPEL processes. Duan and *et al.* [21] introduced a logic model to formally specify the semantics of workflow and its composite tasks described as WS-BPEL abstract processes, and made a deduction of the weakest precondition for workflow. Another work by Duan and *et al.* [22] put some restrictions to such model and found an algorithm to proof the abstract processes' correctness. Fu and *et al.* [23] proposed some techniques and related tools to analyze interactions of composite Web services written in BPEL4WS. The BPEL4WS specifications [19] are first translated into an intermediate representation, and then verified use SPIN [24]. Hamadi and Benatallah [17] transformed the formal semantics of the WS-BPEL composition operators to an expression of Petri nets, and hence allowed the verification of properties and the detection of inconsistencies both within and between services. Pu and *et al.* [25] adopted a similar method by using model checker UPPAAL [26] to verify the correctness of BPEL4WS program including temporal properties. However, none of these works have attempted in verifying WS-BPEL-like fault handling and compensation as we

7 Related Work

The concept of compensation dates back to Sagas [14] and

have done in this paper.

8 Conclusions and Future Work

In this paper we proposed an axiomatic system to verify the correctness of *BPEL** processes. Here we have concentrated on an important core subset of WS-BPEL, namely, *BPEL**. This subset reflects the key features of the language appealing to us, say, fault states, variable mappings and compensation contexts; meanwhile it keeps our mechanism both precise and concise, rather than builds up a huge and complicated system trying to cover all aspects of WS-BPEL. We have formalized these features into program runtime states, and created *BPEL**'s operational semantics with state transition rules, according to the OASIS Standard [4]. Based on this, we have set up the assertions to abstract and express the novel language features, leaving the Hoare triples for verification and their semantics as a natural result. The verification rules for *BPEL** are also formalized after the underlying operational semantics. With respect to such semantics, we have proven the soundness of this system by structural induction on *BPEL**'s constructs, and provided an example as an illustration to the verification process of our system.

Our possible future works following this mainly include two aspects. The first is to extend the logic to cover more language features of WS-BPEL. As our original intention is to propose a concise yet novel system to verify WS-BPEL's fault handling and compensation mechanism, we omitted some language constructs which may cause the verification rule to be uncontrollable under current model, say, partner links, all compensation activities and while loops. These language features may require further invention of verification techniques such as invariants on compensation contexts, and are worth being a natural subsequence of our further work. The second aspect aims at mechanizing the verification system for practical use, which involves some kind of verification condition generators to create the verification conditions, some reasoners to discharge the produced subgoals, and some verification algorithm to integrate these together.

Acknowledgements We appreciate the precious comments from the anonymous reviewers. This work is supported by UK EPSRC project EP/E021948/1 and China NNSF project 60773161.

References

- Thatte S. XLANG: web service for business process design. http://www.gotdotnet.com/team/xml_wsspecs/xlang-c/default.htm, Microsoft, 2001
- Leymann F. WSFL: web services flow language. <http://www-4.ibm.com/software/solutions/webservices/pdf/WSFL.pdf>, IBM, 2001
- Butler M, Ferreira C. An operational semantics for stac, a language for modelling long-running business transactions. In: Proceedings of Sixth International Conference on Coordination Models and Languages, LNCS 2949, 2004, 87–104
- Alves A, Arkin A, Askary S, et al. Web services business process execution language version 2.0. <http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.html>, OASIS Standard, 2007
- Barreto C, Bullard V, Erl T, et al. Web services business process execution language version 2.0 primer. <http://docs.oasis-open.org/wsbpel/2.0/Primer/wsbpel-v2.0-Primer.html>, OASIS Standard, 2007
- Qiu Z, Wang S, Pu G, et al. Semantics of bpel4ws-like fault and compensation handling. In: Proceedings of International Symposium of Formal Methods Europe, LNCS 3582, 2005, 350–365
- Pu G, Zhu H, Qiu Z, et al. Theoretical foundation of scope-based compensable flow language for web service. In: Proceedings of International Conference on Formal Methods for Open Object-Based Distributed Systems (FMOODS'06), LNCS 4037, 2006, 251–266
- Qiu Z, Zhao X, Cai C, et al. Towards the theoretical foundation of choreography. In: Proceedings of Sixteenth International World Wide Web Conference (WWW 2007), 2007, 973–982
- He J, Zhu H, Pu G. A model for bpel-like languages. Frontiers of Computer Science in China, 2007, 1(1):9–19
- Zhu H, He J, Li J, et al. Algebraic approach to linking the semantics of web services. In: Proceedings of Fifth IEEE International Conference on Software Engineering and Formal Methods, 2007, 315–328
- Xu Q, de Roeper W, He J. The rely-guarantee method for verifying shared variable concurrent programs. Formal Aspects of Computing, 1997, 9(2):149–174
- Zhu H. Linking the semantics of a multithreaded discrete event simulation language. PhD thesis, London South Bank University, 2005
- Fowler M, Scott K. UML distilled : a brief guide to the standard object modeling language. Addison-Wesley, 2000
- Garcia-Molina H, Salem K. Sagas. In: Proceedings of the Association for Computing Machinery Special Interest Group on Management of Data 1987 Annual Conference, 1987, 249–259
- Moss J. Nested Transactions: An Approach to Reliable Distributed Computing. PhD thesis, Massachusetts Institute of Technology, 1981
- Aalst W, Dumas M, Hofstede A, et al. Analysis of web services composition languages: The case of bpel4ws. In: Proceedings of the 22nd International Conference on Conceptual Modeling, 2003, 200–215
- Hamadi R, Benatallah B. A petri net-based model for web service composition. In: Proceedings of the 14th Australasian Database Conference, 2003, 191–200
- Brogi A, Canal C, Pimentel E, et al. Formalizing web service choreographies. Electronic Notes of Theoretical Computer Science, 2004, 105:73–94
- Andrews T, Curbera F, Dholakia H, et al. Business process execution language for web services 1.1. <http://download.boulder.ibm.com/ibmdl/pub/software/dw/specs/ws-bpel/ws-bpel.pdf>, 2003
- Bruni R, Melgratti H, Montanari U. Theoretical foundations for compensations in flow composition languages. In: Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, 2005, 209–220

21. Duan Z, Bernstein A, Lewis P, et al. Semantics based verification and synthesis of bpel4ws abstract processes. In: Proceedings of IEEE International Conference on Web Services, 2004, 734–737
22. Duan Z, Bernstein A, Lewis P, et al. A model for abstract process specification, verification and composition. In: Proceedings of the 2nd International Conference on Service Oriented Computing, 2004, 232–241
23. Fu X, Bultan T, Su J. Analysis of interacting bpel web services. In: Proceedings of Thirteenth International World Wide Web Conference (WWW 2004), 2004, 621–630
24. Holzmann G. The Spin Model Checker — Primer and Reference Manual. Addison-Wesley, 2003
25. Pu G, Zhao X, Wang S, et al. Towards the semantics and verification of bpel4ws. In: Proceedings of the International Workshop on Web Languages and Formal Methods (WLFM 2005), 2005
26. Bengtsson J, Larsen K, Larsson F, et al. Uppaal — a tool suite for automatic verification of real-time systems. In: Proceedings of the DIMACS/SYCON Workshop on Hybrid Systems III : Verification and Control, 1996, 232–243