

Adding Sessions to BPEL

Jonathan Michaux

Télécom ParisTech
Paris, France

michaux@telecom-paristech.fr

Elie Najm

Télécom ParisTech
Paris, France

najm@telecom-paristech.fr

Alessandro Fantechi

Università degli Studi di Firenze
Florence, Italy

fantechi@dsi.unifi.it

By considering an essential subset of the BPEL orchestration language, we define SeB, a session based style of this subset. We discuss the formal semantics of SeB and we present its main properties. We use a new approach to address the formal semantics, based on a translation into so-called control graphs. Our semantics handles control links and addresses the static semantics that prescribes the valid usage of variables. We also provide the semantics of collections of networked services.

Relying on these semantics, we define precisely what is meant by interaction safety, paving the way to the formal analysis of safe interactions between BPEL services.

1 Introduction

In service-oriented computing, services are exposed over a network via well defined interfaces and specific communication protocols. The design of software as an orchestration of services is an active topic today. A service orchestration is a local view of a structured set of interactions with remote services.

In this context, our endeavour is to guarantee that services interact safely. The elementary construct in a Web service orchestration is a message exchange between two partner services. The message specifies the name of the operation to be invoked and bears arguments as its payload. An interaction can be long-lasting because multiple messages of different types can be exchanged in both directions before a service is delivered.

The set of interactions supported by a service defines its behavior. We argue that the high levels of concurrency and complex behavior found in orchestrations make them susceptible to programming errors. Widely adopted standards such as the Web Service Description Language (WSDL) [5] provide support for syntactical compatibility analysis by defining message types in a standard way. However, WSDL defines one-way or request-response exchange patterns and does not support the definition of more complex behavior. Relevant behavioral information is exchanged between participants in human-readable forms, if at all. Automated verification of behavioral compatibility is impossible in such cases.

The present paper is a first step towards addressing the problem of behavioral compatibility of Web services. To that end, we follow the promising session based approach. Indeed, the session paradigm is now an active area of research with potential to improve the quality and correctness of software. We chose to adapt and sessionize a significant subset of the industry standard orchestration language BPEL [18]. We call the resulting language SeB (Sessionized BPEL). It supports the same basic constructs as BPEL, but being a proof of concept, it does not include the non basic BPEL constructs such as exception handling. These differences are explained in more detail in section 2.

On the other hand, SeB extends BPEL by featuring sessions as first class citizens. A client wishing to interact with a service begins by opening a session with this service. The set of possible interactions with the service forms the behaviour of the session. In the present paper, we concentrate on the definition of untyped SeB. In the future we plan to define *session types* in order to allow us to verify interaction safety by means of type verification.

In order to define the SeB language, we give a formal semantics that is a novel contribution in itself as it takes into account both the graph nature of the language and the static semantics that define how variables are declared and used, and this is applicable to other BPEL like languages. Indeed, previous approaches either resort to process algebraic simplifications, thus neglecting control links which, in fact, are an essential part of BPEL; or are based on Petri nets and thus do not properly cover the static semantics that regulate the use of variables.

In our approach, the operational semantics is obtained in two steps. The first consists in the creation of what we have called a control graph. This graph takes into account the effect of the evaluation of the control flow and of join-conditions. Control graphs contain symbolic actions and no variables are evaluated in the translation into control graphs. The second step in the operational semantics describes the execution of services when they are part of an assembly made of a client and other Web services. Based on this semantics we formalize the concepts of interaction error and of interaction safety. These concepts will be further developed in future work on verification of session typing.

The rest of this paper is organized as follows. Section 2 provides an informal introduction to the SeB language and contrasts its features with those of BPEL. Sections 3 and 4 give the syntax and semantics of untyped SeB. These sections are self-contained and do not require any previous knowledge of BPEL. Section 5 presents the semantics of networked service configurations described in SeB, and the concepts of interaction error and of interaction safety. Relevant related work is surveyed in section 6 and the paper is concluded in section 7 along with a brief discussion on how session types might be used in SeB in order to prove interaction safety.

2 Informal introduction to SeB

Session initiation. The main novelty in SeB, compared to BPEL, resides in the addition of the session initiation, a new kind of atomic activity, and in the way sessions impact the invoke and receive activities. The following is a typical sequence of three atomic SeB activities that can be performed by a client (we use a simplified syntax): $s@p; s!op_1(x); s?op_2(y)$. This sequence starts by a session initiation activity, $s@p$, where s is a session variable and p a service location variable (this corresponds to a BPEL partnerlink). The execution of $s@p$ by the client and by the target service (the one whose address is stored in p) has the following effects: (i) a fresh session id is stored in s , (ii) a new service instance is created on the service side and is dedicated to interact with the client, (iii) another fresh session id is created on the service instance side and is bound to the one stored in s on the client side. The second activity, $s!op_1(x)$, is the sending of an invocation operation, op_1 , with argument x . The invocation is sent precisely to this newly created service instance. The third activity of the sequence, $s?op_2(y)$, is the reception of an invocation operation op_2 with argument y that comes from this same service instance. Note that invocation messages are all one way and asynchronous: SeB does not provide for synchronous invocation. Furthermore SeB does not provide for explicit correlation sets as does BPEL. But, on the other hand, sessions are to be considered as implicit correlation sets and, indeed, they can be emulated by them. But, in this paper, we preferred to have sessions as an explicit primitive of the language so as to better discuss and illustrate their contribution. Hence, in SeB, session ids are the only means to identify source and target service instances. This is illustrated in the above example where the session variable s is systematically indicated in the invoke and receive activities. Moreover, sessions involve two and only two partners and any message sent by one partner over a session is targeted at the other partner. Biparty sessions are less expressive than correlation sets. Nevertheless, at the end of the paper, we will give some ideas as to how this limitation can be lifted.

Structured activities. SeB has the principal structured activities of BPEL, i.e., flow, for running activ-

ities in parallel; sequence, for sequential composition of activities, and pick, which waits for multiple messages, the continuation behavior being dependent on the received message. SeB also inherits the possibility of having control links between different subactivities contained in a flow, as well as adding a join condition to any activity. As in BPEL, a join condition requires that all its arguments have a defined value (true or false) and must evaluate to true in order for the activity to be executable. SeB also implements so-called dead path elimination whereby all links outgoing from a canceled activity, or from its subactivities, have their values set to false.

Imperative Computations. Given that SeB is a language designed as a proof of concept, we wished to limit its main features to interaction behavior. Hence, imperative computation and branching are not part of the language. Instead, they are assumed to be performed by external services that can be called upon as part of the orchestration. This approach is similar to languages like Orc [10] where the focus is on providing the minimal constructs that allow one to perform service orchestration functions and where imperative computation and boolean tests are provided by external *sites*. In particular, the original *do until* iteration of BPEL is replaced in SeB by a structured activity called “repeat”, given by the syntax: (*do pic₁ until pic₂*). The informal meaning of repeat is: perform *pic₁* repeatedly until the arrival of an invocation message awaited for in *pic₂*.

Example Service. The QuoteComparer is an example of a service written in SeB that will be used throughout the paper. Given an item description from a client, the purpose of the service is to offer a comparison of quotes from different providers for the item while at the same time reserving the item with the best offer. Figure 1 contains a graphical representation of the QuoteComparer implementation. The representation is partial due to space constraints. The service waits for a client to invoke its *s₀?searchQuote(desc)* operation with string parameter *desc* containing an item description. Note the use of the special session variable *s₀*, called the root session variable. By accepting the initial request from the client, the service implicitly begins an interaction with the client over session *s₀*. The service then compares quotes for the item from two different providers (*EZshop* and *QuickBuy*) by opening sessions with each of these providers. Here, the sessions are explicitly opened: *s@EZshop* is the opening of a session with the service having its address stored in variable *EZshop*. The execution of *s@EZshop* will result in a root session being initiated at the *EZshop* service and a fresh session id being stored in *s*.

The service then behaves in the following way: depending on the returns made by the two providers, the QuoteComparer service either returns the best quote, the only available quote, or indicates to the client that no quote is available. The control links and join conditions illustrated in Figure 1 implement this behaviour. For example, control link *l₁* is set to true if *EZshop* returns message *quote(quote1)*, meaning *EZshop* has an offer for the item. The value of this link and others will be later used to determine which provider’s quote should be chosen. Indeed, the join condition (given by $JCD = (l_1 \text{ and } l_2) \text{ or } l_3$) of the bottom left box that deals with reserving an item with EZshop depends on the value of control link *l₁*. It also depends on control link *l₃* which comes from the bottom central box that decides, in the case where both providers offer a quote, which quote is most advantageous. Link *l₂* will be set to true only if provider *QuickBuy* does not return a quote. Hence, the join condition of the bottom left box indicates that it should begin executing either if EZShop’s quote is favourable (the control link *l₃* is set to true), or if EZShop returns a quote and QuickBuy returns *noQuote* (*l₁* and *l₂* are set to true). The conditions for reserving with the *QuickBuy* provider are symmetrical.

If neither *QuickBuy* nor *EZShop* return a quote for the requested item, then control links *l₆* and *l₇* are set to true. This results in the execution of the central invocation operation *s₀!notFound*, i.e the client is informed that no quotes were found for the item description *desc*.

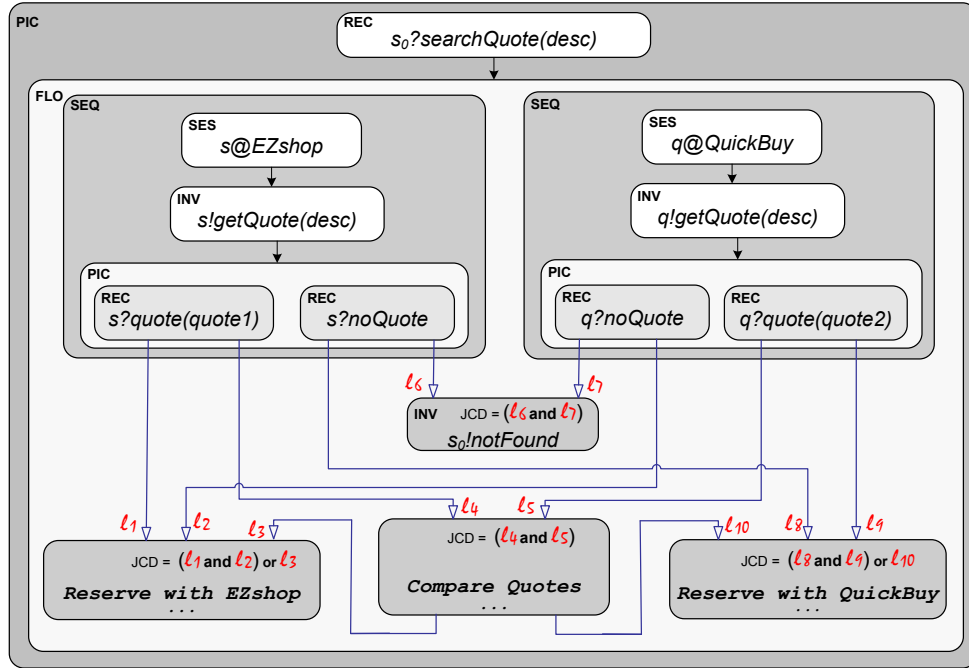


Figure 1: The QuoteComparer Service

3 Syntax of SeB

3.1 Basic Sets

SeB assumes three categories of basic sets: values, variables and others. They are introduced hereafter where, for each set, a short description is provided as are the names of typical elements. All the sets are pairwise disjoint unless stated otherwise.

Set	Description	Ranged over by
<i>DatVal</i>	Data Values	u, u', u_i, \dots
<i>SrvVal</i>	Service Locations	$\pi, \pi', \pi_i \dots$
<i>SesVal</i>	Session Ids	$\alpha, \alpha', \alpha_i \dots \beta \dots$
<i>ExVal</i>	Exchangeable Values	w, w', w_i, \dots
<i>LocVal</i>	All Locations	$\delta, \delta', \delta_i, \dots$
<i>Val</i>	All Values	v, v', v_i, \dots

Table 1: Values

Table 1 presents the various sets of values used in SeB. Data Values (*DatVal*), Service Locations (*SrvVal*) and Sessions Ids (*SesVal*) are ground sets. The set of Exchangeable Values *ExVal* is given by $ExVal = DatVal \uplus SrvVal$, which means that both data values and service locations can be passed between services. Hence, SeB services may dynamically discover other services and may interact with them. The set of all locations *LocVal* is given by $LocVal = SrvVal \uplus SesVal$. Hence, session ids are used to locate service instances. The set of all values *Val* is given by $Val = ExVal \uplus SesVal$.

Table 2 presents the various sets of variables. We can note two distinguished variables, p_0 and s_0 : p_0 , is the service location variable that is dedicated to holding a service's own location; s_0 is a session

<i>Set</i>	<i>Description</i>	<i>Ranged over by</i>
<i>DatVar</i>	Data Variables	y
<i>SrvVar</i>	Service Location Variables	$p_0, p, p', p_i \dots q \dots$
<i>ExVar</i>	Variables of Exchangeable Values	$x, x', x_i \dots$
<i>SesVar</i>	Session Variables	$s_0, s, s', s_i \dots r \dots$
<i>Var</i>	All Variables	$z, z', z_i \dots$

Table 2: **Variables**

variable for accepting sessions initiated by clients. The use of p_0 and s_0 will be described in detail later on in the paper. The set of variables of exchangeable values *ExVar* and the set of all variables *Var* are defined by: $ExVar = DatVar \uplus SrvVar$ and $Var = ExVar \uplus SesVar$.

<i>Set</i>	<i>Description</i>	<i>Ranged over by</i>
<i>Op</i>	Operation Names	$op, op', op_i \dots$
<i>Lnk</i>	Control Links	$l, l', l_i \dots$
2^{Lnk}	subsets of <i>Lnk</i>	$L, L_s, l_T \dots$
<i>Jcd</i>	Join Conditions	$e, e', e_i \dots f \dots$

Table 3: **Miscellaneous Sets**

Table 3 presents the other basic sets. *Lnk* is the set of all control links and *Jcd* is the set of join conditions, i.e., boolean expressions over control links. An example of a join condition: $e = (l_1 \text{ and } l_2) \text{ or } l_3$.

Reconsidering our previous example, *EZShop* and *QuickBuy* are service location variables (elements of set *SrvVar*, with values taken in the set of service locations *SrvVal*); *desc* is a data variable (with values in the set *DatVal*); s_0, s and q are session variables (elements of *SesVar*, taking their values in the set of session ids *SesVal*); l_1, l_2 and l_3 are control links; and finally, expression $(l_3 \text{ and } l_2) \text{ or } l_3$ is a join condition.

3.2 Syntax of SeB Activities

SeB being a dialect of BPEL, XML would have been the most appropriate metalanguage for encoding its syntax. However, for the purpose of this paper, we have adopted a syntax based on records (à la Cardelli and Mitchell [4]) as it is better suited for discussing the formal semantics and properties of the language. By virtue of this syntax, all SeB activities, except **nil**, are records having the following predefined fields: **KND** which identifies the kind of the activity, **BEH**, which gives its behavior, **SRC** (respectively **TGT**), which contains a declaration of the set of control links for which the activity is the source (respectively target), **JCD** which contains the activity's join condition, i.e., a boolean expression over control link names (those given in field **TGT**). Moreover, the *flow* activity has an extra field, **LNK**, which contains the set of links that can be used by the subactivities contained in this activity. Field names are also used to extract the content of a field from an activity, e.g., if *act* is an activity, then *act*.**BEH** yields its behavior. For example: a *flow* activity is given by the record $\langle \mathbf{KND} = \mathbf{FLO}, \mathbf{BEH} = my_behavior, \mathbf{TGT} = L_T, \mathbf{SRC} = L_S, \mathbf{JCD} = e, \mathbf{LNK} = L \rangle$ where L_T, L_S and L are sets of control link names, and e is a boolean expression over control link names. Finally, for the sake of conciseness, we will often drop field names in records and instead we will associate a fixed position to each field. Hence, the *flow* activity given above becomes: $\langle \mathbf{FLO}, my_behavior, L_T, L_S, e, L \rangle$.

We let *ACT* be the set of all activities and *act* a running element of *ACT*, the syntax of activities is given in the following table:

act	::=	nil	(* nil activity *)
		ses inv rec	(* atomic activities *)
		seq flo pic rep	(* structured activities *)
ses	::=	$\langle \text{SES}, s@p, L_r, L_s, e \rangle$	(* session init *)
inv	::=	$\langle \text{INV}, s!op(x_1, \dots, x_n), L_r, L_s, e \rangle$	(* invocation *)
rec	::=	$\langle \text{REC}, s?op(x_1, \dots, x_n), L_r, L_s, e \rangle$	(* reception *)
seq	::=	$\langle \text{SEQ}, act_1; \dots; act_n, L_r, L_s, e \rangle$	(* sequence *)
flo	::=	$\langle \text{FLO}, act_1 \dots act_n, L_r, L_s, e, L \rangle$	(* flow *)
pic	::=	$\langle \text{PIC}, rec_1; act_1 + \dots + rec_n; act_n, L_r, L_s, e \rangle$	(* pick *)
rep	::=	$\langle \text{REP}, do\ pic_1\ until\ pic_2, L_r, L_s, e \rangle$	(* repeat *)

Note that in the production rule for flo, “|” is to be considered merely as a token separator. It is preferred over comma because it is more visual and better conveys the intended intuitive meaning of the flo activity, which is to be the container of a set of sub activities that run in parallel. The same remark applies to symbols “;”, “+”, “do” and “until” which are used as token separators in the production rules for seq, pic and rep to convey their appropriate intuitive meanings. Note that “|” and “+” are commutative. As a final note, the number n appearing in the rules for seq, flo and pic is such that $n \geq 1$.

Returning to Figure 1, the syntax of the central **INV** activity is given by the following expression: $\langle \text{INV}, s_0!notFound, \{l_6, l_8\}, \emptyset, l_6 \text{ and } l_8 \rangle$, and the syntax of the seq activity at the top left of the example is given by the following expression:

$$\langle \text{SEQ}, \langle \text{SES}, s@EZshop, \emptyset, \emptyset, true \rangle; \langle \text{INV}, s!getQuote(desc), \emptyset, \emptyset, true \rangle; \langle \text{PIC}, \langle \text{REC}, s?quote(quote1), \emptyset, \{l_1, l_4\}, true \rangle + \langle \text{REC}, s?noQuote, \emptyset, \{l_6, l_8\}, true \rangle, \emptyset, \emptyset, true \rangle \rangle$$

Note: Syntax simplification - when an activity has no incoming or outgoing control links, we may omit its encapsulating record and represent it just by its behavior. Hence, for example, we may write $act; act'$ instead of $\langle \text{SEQ}, act; act', \emptyset, \emptyset, true \rangle$, and we will write $s!op(x)$ instead of $\langle \text{INV}, s!op(x), \emptyset, \emptyset, true \rangle$.

Definition 3.1. Subactivities - For an activity act , we define the sets of activities, \widehat{act} and $\widehat{\widehat{act}}$:

- $\widehat{act} \triangleq \{act\}$ if act is an atomic activity, else $\widehat{act} \triangleq \{act\} \cup \widehat{act.BEH}$
- $\widehat{\widehat{act}} \triangleq \emptyset$ if act is an atomic activity, else $\widehat{\widehat{act}} \triangleq \widehat{act.BEH}$

Informally, \widehat{act} is the set of activities transitively contained in act and $\widehat{\widehat{act}}$ is the set of activities strictly and transitively contained in act .

Definition 3.2. Precedence relation between activities - For an activity act , we define relation $pred$ on the set \widehat{act} as follows:

$$act_1\ pred\ act_2 \quad \text{iff} \quad (act_1.SRC \cap act_2.TGT \neq \emptyset) \quad \text{or} \quad (\exists seq \in \widehat{act} \text{ with } seq.BEH = \dots act_1; act_2 \dots)$$

Informally, relation $pred$ implies that act_1 precedes act_2 either in some seq activity or because act_1 has at least one outgoing control link targeting act_2 .

Activities need to comply to certain constraints concerning the declaration and usage of their control links. Firstly, control links should be well scoped, unique and should form no cycles. Secondly, all rep subactivities should be well-formed in terms of incoming and outgoing control links. An activity that complies with these constraints is said to be well-formed.

Definition 3.3. rep-well-formed activity - A SeB activity act_0 is rep-well-formed iff any activity $rep = \langle \mathbf{REP}, do\ pic_1\ until\ pic_2, L_r, L_s, e \rangle$ of $\widehat{act_0}$ satisfies the following 3 conditions:

(1) $pic_1.TGT = pic_2.TGT = \emptyset$, (2) $pic_1.SRC = \emptyset$, and (3) $\widehat{\widehat{pic_1}}.SRC = \widehat{\widehat{pic_1}}.TGT$.

Informally, (1) implies that pic_1 and pic_2 have no incoming links, (2) states that pic_1 has no outgoing links and (3) states that all control links of activities (strictly) contained in pic_1 are (strictly) internal to pic_1 .

Definition 3.4. Well-structured activity - A SeB activity act_0 is well structured iff the control links occurring in any activity of $\widehat{act_0}$ satisfy the unicity, scoping and non cyclicity conditions given below, where act , act' , act'' and seq are subactivities of act_0 .

1. Control links unicity - Given any control link l , and any pair of activities act and act' :
if $(l \in act.LNK \cap act'.LNK)$ or $(l \in act.SRC \cap act'.SRC)$ or $(l \in act.TGT \cap act'.TGT)$ then $act = act'$.
2. Control links scoping - If $l \in act.SRC$ (respectively if $l \in act.TGT$) then $\exists act', act''$ with $act \in \widehat{act''}$ and $act' \in \widehat{act''}$ and with $l \in act''.LNK$ and $l \in act'.TGT$. (respectively $l \in act'.SRC$).
3. Control links non cyclicity:
 - (i) Relation $pred$ is acyclic, and
 - (ii) $\forall act, act' \in \widehat{act_0} \quad act' \in \widehat{act} \Rightarrow (act.SRC \cap act'.TGT = \emptyset)$ and $(act'.SRC \cap act.TGT = \emptyset)$

Informally, a well-structured activity is such that all of its control links (including those in subactivities) are well scoped (i.e., are within the scope of one flow subactivity), unique (each target declaration corresponds to one and only one source declaration and vice versa), and do not form any causality cycles, either directly (from source to target of activities) or through activities that are chained within some sequence activity or through the containment relation between activities.

Definition 3.5. Well-formed activity - an activity is well formed iff it is both well structured and rep-well-formed.

4 Semantics of SeB

Here we define the notion of control graphs and then provide the SOS rules that define a translation from a well-formed activity into a control graph. We then study the properties of control graphs. The semantics of SeB activities is obtained by applying a series of transformations to this first control graph that lead to a final control graph, noted $\mathbf{cg}(act)$.

4.1 Definitions

4.1.1 Control Graphs

Definition 4.1. Observable Actions - The set, $ACTIONS$, of observable actions is defined by:

$ACTIONS =_{def} \{ a \mid a \text{ is any action of the form: } s@p, s!op(x_1, \dots, x_n) \text{ or } s?op(x_1, \dots, x_n) \}$

Definition 4.2. All actions - We define the set $ACTIONS_\tau$ of all actions (ranged over by σ): $ACTIONS_\tau =_{def} ACTIONS \cup \{\tau\}$ where τ denotes the unobservable (or silent) action.

Definition 4.3. Control Graphs - A control graph, $\Gamma = \langle G, g_0, \mathcal{A}, \rightarrow \rangle$, is a labelled transition system

- where:
- G is a set of states, called control states
 - g_0 is the initial control state
 - \mathcal{A} is a set of actions ($\mathcal{A} \subset ACTIONS_\tau$)
 - $\rightarrow \subset G \times \mathcal{A} \times G$

4.1.2 Control Link Maps

We now define the control part of activities where we consider only the values of control links. Activities are given a map that stores the running values of control links, which are initially undefined values.

Definition 4.4. - Control Link Maps - A Control link map c is a partial function from the set of control links, Lnk , to the set of boolean values extended with the undefined value. $c : Lnk \rightarrow \{\mathbf{true}, \mathbf{false}, \perp\}$

Definition 4.5. Initial Control Links Map - For an activity act we define c_{act} , the initial control links map: $dom(c_{act}) = \{l \in Lnk \mid l \text{ occurs in } \widehat{act}\}$ and $\forall l \in dom(c_{act}) : c_{act}(l) = \perp$

Definition 4.6. Evaluation of a Join Condition - If L is a set of control links, e a boolean expression over L and c a control links map, then the evaluation of e in the context of c is written: $c \triangleright e(L)$. Furthermore, we consider that this evaluation is defined only when $\forall l \in L, c(l) \neq \perp$.

4.2 From Activities to Control Graphs: Structured Operational Semantics Rules

$\frac{\boxed{\text{SES}} \quad c \triangleright e(L_T) = \mathbf{true}}{(c, \langle \text{SES}, s@p, L_T, L_S, e \rangle)} \quad \downarrow s@p \quad (c[\mathbf{true}/L_S], \mathbf{nil})$	$\frac{\boxed{\text{INV}} \quad c \triangleright e(L_T) = \mathbf{true}}{(c, \langle \text{INV}, s!op(\bar{x}), L_T, L_S, e \rangle)} \quad \downarrow s!op(\bar{x}) \quad (c[\mathbf{true}/L_S], \mathbf{nil})$	$\frac{\boxed{\text{REC}} \quad c \triangleright e(L_T) = \mathbf{true}}{(c, \langle \text{REC}, s?op(\bar{x}), L_T, L_S, e \rangle)} \quad \downarrow s?op(\bar{x}) \quad (c[\mathbf{true}/L_S], \mathbf{nil})$	$\frac{\boxed{\text{DPE}} \quad c \triangleright e(L_T) = \mathbf{false}}{(c, \langle *, \mathbf{act}, L_T, L_S, e \rangle)} \quad \downarrow \tau \quad (c[\mathbf{false}/\widehat{\mathbf{act.SRC}}], \mathbf{nil})$
$\frac{\boxed{\text{PIC}} \quad c \triangleright e(L_T) = \mathbf{true} \quad (c, \text{rec}) \xrightarrow{\sigma} (c', \mathbf{nil})}{(c, \langle \text{PIC}, \text{rec}; \mathbf{act} + \sum \text{rec}_i; \mathbf{act}_i, L_T, L_S, e \rangle)} \quad \downarrow \sigma \quad (c'[\mathbf{false}/(\sum \widehat{\text{rec}}_i; \widehat{\mathbf{act}}_i).\text{SRC}], \langle \mathbf{FLO}, \mathbf{act}, L_T, L_S, e \rangle)$	$\frac{\boxed{\text{REP1}} \quad c \triangleright e(L_T) = \mathbf{true} \quad (c, \text{pic}_1) \xrightarrow{\sigma} (c', \mathbf{act})}{(c, \langle \text{REP}, \text{do pic}_1 \text{ until pic}_2, L_T, L_S, e \rangle)} \quad \downarrow \sigma \quad (c', \langle \text{UNF}, \text{do act then pic}_1 \text{ until pic}_2, L_T, L_S, e \rangle)$		
$\frac{\boxed{\text{FLO1}} \quad c \triangleright e(L_T) = \mathbf{true} \quad (c, \mathbf{act}_i) \xrightarrow{\sigma} (c', \mathbf{act}')}{(c, \langle \mathbf{FLO}, \mathbf{act}_1 \dots \mathbf{act}_i \dots \mathbf{act}_n, L_T, L_S, e \rangle)} \quad \downarrow \sigma \quad (c', \langle \mathbf{FLO}, \mathbf{act}_1 \dots \mathbf{act}' \dots \mathbf{act}_n, L_T, L_S, e \rangle)$	$\frac{\boxed{\text{REP2}} \quad c \triangleright e(L_T) = \mathbf{true} \quad (c, \text{pic}_2) \xrightarrow{\sigma} (c', \mathbf{act})}{(c, \langle \text{REP}, \text{do pic}_1 \text{ until pic}_2, L_T, L_S, e \rangle)} \quad \downarrow \sigma \quad (c', \langle \mathbf{FLO}, \mathbf{act}, L_T, L_S, e \rangle)$		
$\frac{\boxed{\text{FLO2}} \quad c \triangleright e(L_T) = \mathbf{true}}{(c, \langle \mathbf{FLO}, \mathbf{act}_1 \dots \mathbf{nil} \dots \mathbf{act}_n, L_T, L_S, e \rangle)} \quad \downarrow \tau \quad (c, \langle \mathbf{FLO}, \mathbf{act}_1 \dots \mathbf{act}' \dots \mathbf{act}_n, L_T, L_S, e \rangle)$	$\frac{\boxed{\text{UNF1}} \quad c \triangleright e(L_T) = \mathbf{true} \quad (c, \mathbf{act}) \xrightarrow{\sigma} (c', \mathbf{act}')}{(c, \langle \text{UNF}, \text{do act then pic}_1 \text{ until pic}_2, L_T, L_S, e \rangle)} \quad \downarrow \sigma \quad (c', \langle \text{UNF}, \text{do act}' \text{ then pic}_1 \text{ until pic}_2, L_T, L_S, e \rangle)$		
$\frac{\boxed{\text{FLO3}} \quad c \triangleright e(L_T) = \mathbf{true}}{(c, \langle \mathbf{FLO}, \mathbf{nil}, L_T, L_S, e \rangle)} \quad \xrightarrow{\tau} (c[\mathbf{true}/L_S], \mathbf{nil})$	$\frac{\boxed{\text{UNF2}} \quad c \triangleright e(L_T) = \mathbf{true}}{(c, \langle \text{UNF}, \text{do nil then pic}_1 \text{ until pic}_2, L_T, L_S, e \rangle)} \quad \downarrow \tau \quad (c', \langle \text{REP}, \text{do pic}_1 \text{ until pic}_2, L_T, L_S, e \rangle)$ <p style="margin-left: 20px;">where</p> $c' = [\overset{\perp}{\widehat{\text{pic}}_1.\text{SRC}}, \overset{\perp}{\widehat{\text{pic}}_1.\text{TGT}}, \overset{\text{true}}{\widehat{\text{pic}}_1.\text{SRC}}]$		

Table 4: Structured Operational Semantics

In Table 4, we provide the sos rules defining a translation from activities to control graphs. Some comments are in order concerning this table:

- the notation for value substitution in control link maps requires some explanation: $c[\mathbf{true}/l] =_{def} c'$

where $c'(l') = c(l)$ for $l \neq l'$ and $c'(l) = \mathbf{true}$. By abuse of notation, we also apply value substitution to sets of control links. Hence, if Π is a set of activities, then, e.g., $c[\mathbf{false}/\widehat{\Pi}.\text{SRC}]$ is the substitution whereby any control link occurring as source of an activity in $\widehat{\Pi}$ has its value set to \mathbf{false} .

- the rules for the seq activity have been skipped as for any seq one can construct a behaviorally equivalent flo activity having the same set of subactivities. This is merely achieved by introducing the appropriate control links between any two consecutive subactivities.
- in the rules for activity flo we have dropped the field LNK since its value is constant (LNK is used to define a scope for control link variables).
- in the rules for the repeat activity rep, we have introduced a new activity, unf (unfold). Its syntax is: $\text{unf} ::= \langle \text{UNF}, \text{do act then pic}_1 \text{ until pic}_2, L_r, L_s, e \rangle$. Activity unf is introduced as a result of the execution of rule REP1. This mirrors the unfolding of an iteration by activity pic₁. Rule UNF1 represents the execution of an action by the unfolded activity and rule UNF2 represents the termination of the iteration whereby unf is transformed back into a rep activity identical to the original one. Note how, in this transformation, all links (strictly) contained in pic₁ are reset to the undefined value. The rep activity as a whole terminates by means of rule REP2 representing the activation of activity pic₂.
- as a notational convention, we have used “*” to denote a wildcard activity, as seen in the rule for dead-path elimination (DPE).

When applied to the initial control state $(c_{\text{act}}, \text{act})$ of a well formed activity act, the sos rules defined in Table 4 yield a first control graph, the raw control graph, that we note $\mathbf{r-cg}(\text{act})$. A transition in this control graph will be denoted by $(c', \text{act}') \xrightarrow[\text{act}]{\sigma} (c'', \text{act}'')$.

4.3 Properties of Raw Control Graphs

Property 1. *The set of states of $\mathbf{r-cg}(\text{act})$ is finite.*

Proof. We can structurally define function $\#_{ub}(\text{act})$ that gives an upper bound for the number of states generated from an activity act. To obtain this upper bound we consider an empty set of control links so as to allow the maximum possible interleaving and hence generating the maximum number of states. This upper bound is structurally defined as follows:

$$\begin{aligned} \#_{ub}(\text{nil}) &= 1 & \#_{ub}(\text{inv}) &= \#_{ub}(\text{rec}) = \#_{ub}(\text{ses}) = 2 & \#_{ub}(\langle \mathbf{REP}, \text{do pic}_1 \text{ until pic}_2, , , \rangle) &= \#_{ub}(\text{pic}_1) + \#_{ub}(\text{pic}_2) + 1 \\ \#_{ub}(\langle \mathbf{FLO}, \text{act}_1 | \dots | \text{act}_n, , , \rangle) &= (\#_{ub}(\text{act}_1) + 1) \times \dots \times (\#_{ub}(\text{act}_n) + 1) + 1 \\ \#_{ub}(\langle \mathbf{PIC}, \text{rec}_1; \text{act}_1 + \dots + \text{rec}_n; \text{act}_n, , , \rangle) &= \#_{ub}(\text{act}_1) + 2 + \dots + \#_{ub}(\text{act}_n) + 2 + 1 \end{aligned}$$

□

Property 2. *$\mathbf{r-cg}(\text{act})$ is free of τ loops.*

Proof. The only case in which a τ loop could appear is in the pic₁ of a repeat activity. But since pic necessarily contains an initial receive activity, any potential τ loop is broken by this receive activity. □

Property 3. *$\mathbf{r-cg}(\text{act})$ is τ -confluent, i.e.:*

$$g \xrightarrow[\text{act}]{\tau} g_1 \text{ and } g \xrightarrow[\text{act}]{\sigma} g_2 \Rightarrow \exists g' \text{ where } g_1 \xrightarrow[\text{act}]{\sigma} g' \text{ and } g_2 \xrightarrow[\text{act}]{\tau} g'$$

Proof. (sketch) Let us consider the situation where from some state (c, act_0) we can fire two transitions: (1) $(c, \text{act}_0) \xrightarrow[\text{act}]{\tau} (c_1, \text{act}_1)$ and (2) $(c, \text{act}_0) \xrightarrow[\text{act}]{\sigma} (c_2, \text{act}_2)$. Since one of the actions is τ then there is necessarily a subactivity $\text{flo} \in \widehat{\text{act}}_0$ with $\text{flo} = \langle \mathbf{FLO}, \dots \text{act}'_1 | \dots | \text{act}'_2 \dots, L_r, L_s, e \rangle$ where transition (1) is produced by act'_1 and transition (2) by act'_2 . Since transitions (1) and (2) are not conflicting, then both sequences (1) then (2) and (2) then (1) are possible and reach the same state. □

Property 4.

(i) All sink states of $\mathbf{r}\text{-cg}(\text{act})$ (i.e. states with no outgoing transitions) are of the form (c, \mathbf{nil}) . (Hence sink states may differ only by their control link maps),

(ii) for any state of $\mathbf{r}\text{-cg}(\text{act})$, there exists a path that leads to a sink state.

Proof. (sketch)

(i) Let us consider the non trivial case where act is not an atomic activity, and a state (c, act_0) reachable from the initial state $(c_{\text{act}}, \text{act})$. If $\text{act}_0.\text{BEH} = \mathbf{nil}$ then (c, act_0) cannot be a sink state since it can make a transition to (c, \mathbf{nil}) . If $\text{act}_0.\text{BEH} \neq \mathbf{nil}$ then surely there is one activity $\text{act}' \in \text{act}_0.\text{BEH}$ which is *head* and ready for execution in the context of control link map c . This stems from the non cyclicity property which, it can be proven, is preserved along the execution path from the initial state. This means that all the activities in $\widehat{\text{act}}$ that precede act' have been either executed or cancelled. Hence, (c, act_0) can make a transition (the one derived from act'), and thus cannot be a sink state.

(ii) the inspection of all SOS rules shows that for all states (c, act_0) , where act_0 is not a repeat activity, there is a transition to some other state (c', act') where act' is strictly (syntactically) simpler (i.e., smaller in size) than act_0 . In the case of repeat, there is also always a path leading to a syntactically simpler activity, which is through the pre-empting activity pic_2 . \mathbf{nil} being the simplest activity, thus all activities must reduce to \mathbf{nil} . □

4.4 Control Graph Transformations

The first transformation applied to $\mathbf{r}\text{-cg}(\text{act})$ is τ -prioritization and results in control graph $\tau\mathbf{p}\text{-cg}(\text{act})$. Then we apply τ -compression, resulting in control graph $\tau\mathbf{c}\text{-cg}(\text{act})$ that is free of τ transitions. Next comes a transformation into run-to-completion semantics resulting in $\mathbf{rtc}\text{-cg}(\text{act})$. Finally, strong equivalence minimization is applied resulting in $\mathbf{cg}(\text{act})$, a graph with a single sink state.

4.4.1 τ -Prioritisation and τ -Compression

In [8], the authors give an efficient algorithm that, given a graph in which the τ transitions are confluent and in which there are no τ loops, prioritises τ transitions in the graph while preserving branching equivalence. Given Property 3, i.e that the τ transitions of any $\mathbf{r}\text{-cg}(\text{act})$ are confluent, and Property 2, i.e there are no τ loops in any $\mathbf{r}\text{-cg}(\text{act})$, we can apply the algorithm given in [8] and obtain a τ -prioritised control graph that we call $\tau\mathbf{p}\text{-cg}(\text{act})$. In such a graph, the outgoing transitions from a state are either all τ transitions, or they are all observable actions. In [8], an algorithm for τ -compression is also given, essentially leaving the graph τ transition free. We can therefore apply τ -compression to $\tau\mathbf{p}\text{-cg}(\text{act})$ and obtain a τ -free branching equivalent control graph that we call $\tau\mathbf{c}\text{-cg}(\text{act})$.

4.4.2 Applying run-to-completion semantics

Here we process control graphs so as to verify the *run-to-completion* property. This property implies the following behavior of control graphs: after a receive, all possible invocations, session initiations, or τ actions are executed before another message can be received.

In order to give SeB a run-to-completion semantics, we give a priority order to the transitions in a τ -prioritised control graph $\tau\mathbf{p}\text{-cg}(\text{act})$ (this also applies for τ -compressed control graphs). We consider the outgoing transitions from a state and we define the following priority order between the actions that label these transitions: $s!\text{op}(x) > s@\text{p} > s?\text{op}(x)$. This means that when two transitions labelled with actions with different priorities are possible from a given state, then the one having a lower priority is pruned. All states that have become unreachable from the initial state are also pruned.

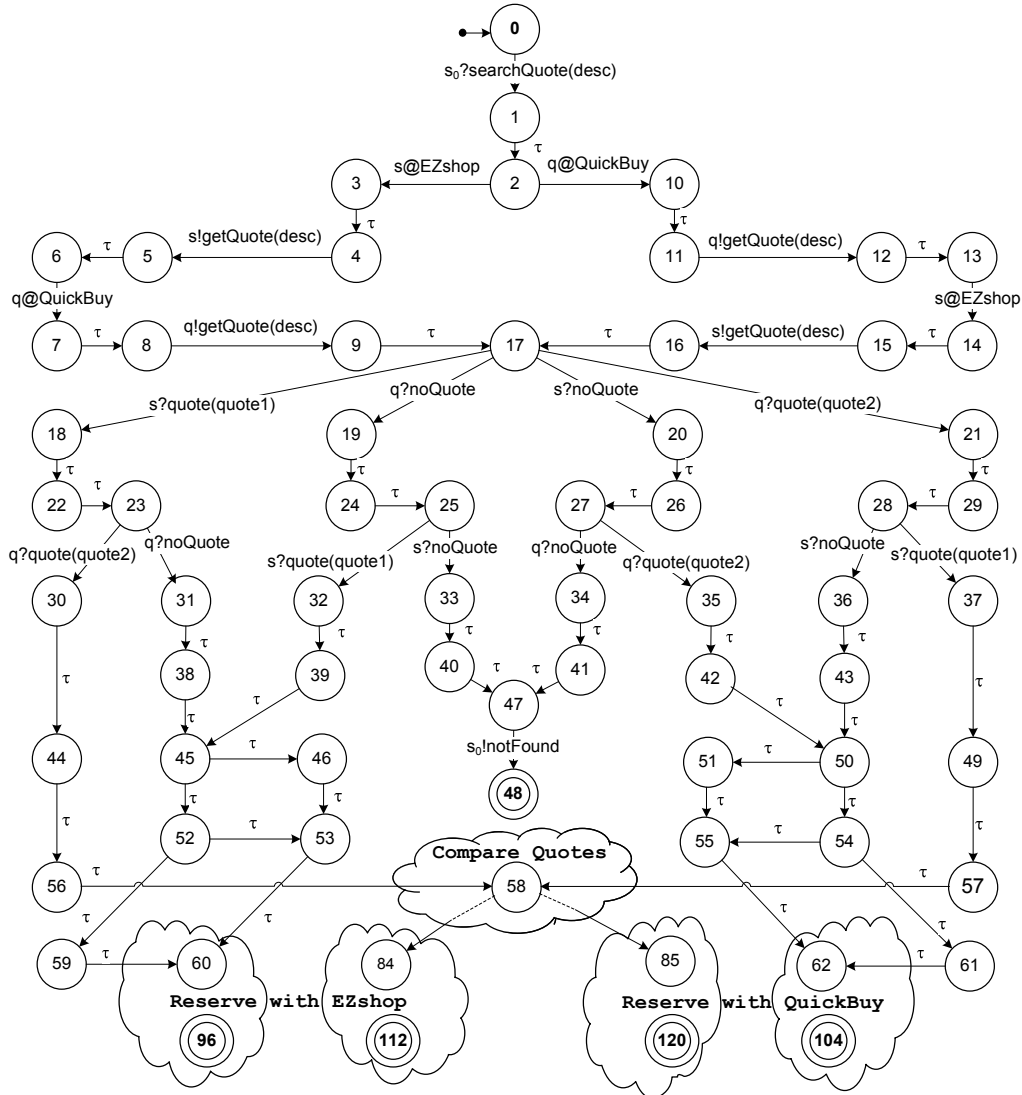


Figure 2: The QuoteComparer Control Graph

Figure 2 shows the control graph generated for the QuoteComparer SeB service shown in Figure 1 in which the τ -priorisation and run-to-completion semantics transformations have been applied. τ -compression was not applied here for the sake of illustration. In this figure, the computation between states 0 and 17 represents an example of run-to-completion.

Note that the control graph of Figure 2 contains 5 terminal states. They all correspond to couples of the form (c_i, nil) , in accordance with Property 4. They differ only by their control link maps. For example, state 48 corresponds to the couple (c_{48}, nil) where c_{48} is given by $c_{48}(l_6) = c_{48}(l_7) = c_{48}(l_4) = c_{48}(l_5) = \text{true}$ and where c_{48} is false for the remaining links.

4.4.3 Minimisation

When strong (or branching) equivalence minimisation is applied to control graph $\text{rtc-cg}(\text{act})$, a minimal control graph $\text{cg}(\text{act})$ is produced that has one and only sink state (because all sink states are equivalent).

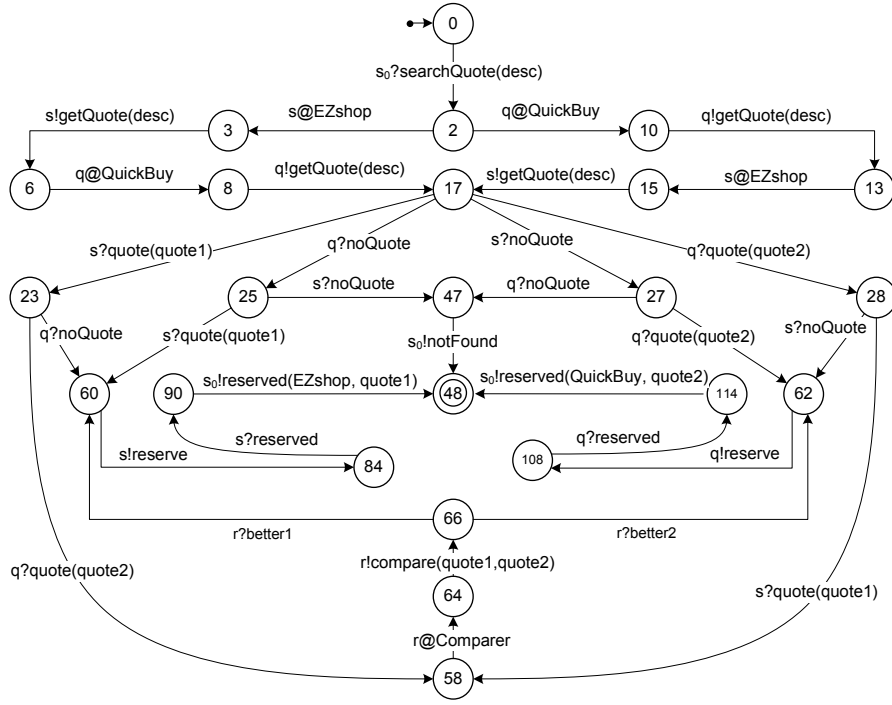


Figure 3: The QuoteComparer Reduced Control Graph

Figure 3 shows the minimised control graph of the QuoteComparer example. We used the CADP [7] toolset to perform some of the transformations defined in this section on the QuoteComparer example.

4.5 Semantics of activities

Henceforth, when we write $\mathbf{cg}(\text{act})$ for any well-formed activity act , we will consider that we are dealing with the τ -compressed, run-to-completion and minimised control graph, and we name its unique sink state $\mathbf{term}(\text{act})$, to be referred to as the terminal state. We also adopt the following notations: $\mathbf{init}(\text{act})$ denotes the initial state of $\mathbf{cg}(\text{act})$; $\mathbf{states}(\text{act})$ is the set of states of $\mathbf{cg}(\text{act})$; $\mathbf{trans}(\text{act})$ is the set of transitions of $\mathbf{cg}(\text{act})$. A transition in this final control graph $\mathbf{cg}(\text{act})$ will be denoted by $g \xrightarrow[\text{act}]{\sigma} g'$.

Definition 4.7. Open for reception - A state g of $\mathbf{cg}(\text{act})$ is said to be open for reception on session s and we note $\text{open}(\text{act}, g, s)$, iff state g has at least one outgoing transition labeled with a receive action on session s . More formally: $\text{open}(\text{act}, g, s) =_{\text{def}} \exists op, x_1 \dots x_n, g'$ such that $g \xrightarrow[\text{act}]{s?op(x_1, \dots, x_n)} g'$.

4.6 Free, bound, usage and forbidden occurrences of variables

Definition 4.8. Variables of an activity - For an activity act we define the set of variables occurring in act : $V(\text{act}) =_{\text{def}} \{z \mid z \text{ occurs in } \text{act}\}$.

Definition 4.9. Binding occurrences - For variables $y \in V(\text{act})$, $s \in V(\text{act})$ and $p \in V(\text{act})$, the following underlined occurrences are said to be binding occurrences in act : $\underline{s}@p$, $s?op(\dots, \underline{y}, \dots)$ and $s?op(\dots, \underline{p}, \dots)$. We denote $BV(\text{act})$ the set of variables having a binding occurrence in act .

Definition 4.10. Usage occurrences - For variables $y \in V(\text{act})$, $s \in V(\text{act})$ and $p \in V(\text{act})$, the following underlined occurrences are said to be usage occurrences in act : $s@p$, $\underline{s}^?op(\dots)$, $\underline{s}!op(\dots)$, $\underline{s}!op(\dots, \underline{p}, \dots)$ and $\underline{s}!op(\dots, \underline{y}, \dots)$. We denote $UV(\text{act})$ the set of variables having at least one usage occurrence in act .

Definition 4.11. Free occurrences - A variable $z \in V(\text{act})$ is said to occur free in act , iff there is a path in $\mathbf{cg}(\text{act})$: $\text{init}(\text{act}) \xrightarrow{\sigma_1}_{\text{act}} g_1, \dots, g_{n-1} \xrightarrow{\sigma_n}_{\text{act}} g_n$ where z has a usage occurrence in σ_n and has no binding occurrence in any of $\sigma_1, \dots, \sigma_{n-1}$. We denote $FV(\text{act})$ the set of variables having at least one free occurrence in act .

Definition 4.12. Forbidden occurrences - $op^?(\dots p_0 \dots)$ and $s_0@p$ are forbidden occurrences. As we shall explain later, p_0 is reserved for the own location of the service, while s_0 is a reserved session variable that receives a session id implicitly at service instantiation time.

5 Syntax and Semantics of Service Configurations

SeB activities can become deployable services that can be part of configurations of services. These configurations have a dynamic semantics, based on which we can define the property of safe interaction.

Let M be a partial map from variables Var to $Val \cup \{\perp\}$, the set of values augmented with the undefined value. Henceforth, we consider couples (M, act) where $\text{dom}(M) = V(\text{act})$.

5.1 Deployable Services

Definition 5.1. Deployable services - The couple $((M, \text{pic}))$ is a deployable service iff:

- $p_0 \in FV(\text{pic})$, $s_0 \in FV(\text{pic})$
- $FV(\text{pic}) \cap \text{SesVar} = \{s_0\}$
- $\text{pic.BEH} = \sum s_0^?op_i(\tilde{x}_i); \text{act}_i$
- $\text{dom}(M) = V(\text{pic})$
- $\forall z \in V(\text{pic}) \setminus FV(\text{pic}) : M(z) = \perp$
- $\forall z \in FV(\text{pic}) \setminus \{s_0\} : m(z) \neq \perp$

Informally, (M, act) is a deployable service if act is a pic , s_0 is its only free session variable, the initial receptions of pic are on session s_0 , its location is defined (variable p_0 is set in M), and all its free variables, except s_0 , have defined values in M .

Definition 5.2. Running service instances - The running state of a service instance derived from the deployable service $((M, \text{pic}))$ is the triple $(m, \text{pic} \blacktriangleright g)$ where: $g \in \mathbf{states}(\text{pic})$ and $\text{dom}(m) = \text{dom}(M)$. The initial state of this running service instance is the triple $(M[\beta/s_0], \text{pic} \blacktriangleright \mathbf{init}(\text{pic}))$ with β fresh. A deployed service behaves like a factory creating a new running service instance each time it receives a session initiation request.

5.2 Service Configurations

Configurations are built from deployable services. We provide an abstraction of the communication bus that is necessary to formalize and prove the desirable properties of service configurations. Service instances exchange messages through FIFO queues, and session bindings are set up in order to establish the corresponding queues.

Definition 5.3. Service configurations - When deployed, a set of deployable services yields a configuration noted: $((M_1, \text{pic}_1) \diamond \dots \diamond (M_k, \text{pic}_k))$. The symbol \diamond denotes the associative and commutative deployment operator, meaning that services are deployed together and share the same address space.

Definition 5.4. Well-partnered service configuration - A service configuration $((M_1, pic_1)) \diamond \dots \diamond ((M_k, pic_k))$ is said to be well partnered iff:

- $\forall i, j : i \neq j \Rightarrow M_i(p_0) \neq M_j(p_0)$ • $\forall i, p : M_i(p) \neq \perp \Rightarrow \exists j \text{ with } M_i(p) = M_j(p_0)$

That is, any two services have different location addresses, and any partner required by one service is present in the set of services.

5.3 Running Configurations

Definition 5.5. Message queues - \mathcal{Q} is a set made of message queues with $\mathcal{Q} ::= q \mid q \diamond \mathcal{Q}$, where q is an individual FIFO message queue of the form $q ::= \delta \leftarrow \widetilde{Mes}$ with \widetilde{Mes} a possibly empty list of ordered messages and δ the destination of the messages in the queue. The contents of \widetilde{Mes} depend on the kind of the destination δ . If δ is a service location, \widetilde{Mes} contains only session initiation requests of the form $new(\alpha)$. However if δ is a session id, then \widetilde{Mes} contains only operation messages of the form $op(\tilde{w})$.

Definition 5.6. Session bindings - A session binding is an unordered pair of session ids (α, β) . A running set of session bindings is noted \mathcal{B} and has the syntax $\mathcal{B} ::= (\alpha, \beta) \mid (\alpha, \beta) \diamond \mathcal{B}$. If $(\alpha, \beta) \in \mathcal{B}$ then α and β are said to be bound and messages sent on local session id α are routed to a partner holding local session id β , and vice-versa.

Definition 5.7. Running configurations - A running configuration, C , is a configuration made of services, service instances, queues and bindings all running concurrently and sharing the same address space: $C = \mathcal{C}_{serv} \diamond (m_1, act_1, g_1) \dots (m_k, act_k, g_k) \diamond \mathcal{Q} \diamond \mathcal{B}$ where $\mathcal{C}_{serv} = ((M_1, pic_1)) \dots ((M_n, pic_n))$ is a well-partnered service configuration, and $(m_1, act_1, g_1) \dots (m_k, act_k, g_k)$ are service instances.

Again, operator \diamond is associative and commutative hence the order of services, instances, bindings and queues is irrelevant. Also if the sets of bindings or queues are empty, they are omitted.

Definition 5.8. Initial Running configuration - A service configuration cannot bootstrap itself as this requires at least one client instance that begins by opening a session. (m, act, g) is one such client where $act.BEH = s@p; act'$ and $\exists i$ such that $m(p) = M_i(p_0)$. Hence the minimal initial running configuration is: $C = ((M_1, pic_1)) \dots ((M_n, pic_n)) \diamond (m, act, g)$.

5.4 Semantics of Service Configurations

Here we provide a full semantics for SeB that allows us to formalize the property that we want to assess in SeB programs. The service configuration semantics is defined using the four sos rules in Table 5.

SES1 applies when a service instance has a session initiation transition, $s@p$, where $m(p)$ is the address of the remote service. The result is that the two fresh session ids α and β are created for the local and distant session ends, and message $new(\beta)$ is added to the tail of the message queue targeting service π . SES2 applies when message $new(\beta)$ is at the head of the input queue of the service located at π (for which $m(p_0) = \pi$). The result is that the message is consumed and a new service instance is created with root session s_0 set to β . INV states that when a service instance is ready to send an invocation message over session s , then the message is appended to the queue whose target is β which is bound to $m(s)$. REC is symmetrical to INV.

5.5 Interaction-safe Service Configurations

The property of *interaction safety* is verified when no service instance ever reaches a state in which it is awaiting a message, but the message at the head of the corresponding input queue is not expected. The definition that follows is given in two parts, the second depends on the first.

SES1	$g \xrightarrow[\text{act}]{s@p} g' \quad \alpha, \beta \text{ fresh}$
$\dots (m, \text{act} \blacktriangleright g) \dots \diamond \dots (m(p) \leftrightarrow \widetilde{Mes}) \dots \diamond \mathcal{B} \longrightarrow$ $\dots (m[\alpha/s], \text{act} \blacktriangleright g') \dots \diamond \dots (m(p) \leftrightarrow \widetilde{Mes} \cdot \text{new}(\beta)) \dots \diamond \mathcal{B} \diamond (\alpha, \beta)$	
SES2	$M(p_0) = \pi$
$\dots ((M, \text{pic})) \dots \diamond \dots (\pi \leftrightarrow \text{new}(\beta) \cdot \widetilde{Mes}) \dots \diamond \mathcal{B} \longrightarrow$ $\dots ((M, \text{pic}) \diamond (M[\beta/s_0], \text{pic} \blacktriangleright \text{init}(\text{pic}))) \dots \diamond \dots (\pi \leftrightarrow \widetilde{Mes}) \dots \diamond \mathcal{B}$	
INV	$g \xrightarrow[\text{act}]{s!op(x_1, \dots, x_n)} g' \quad (m(s), \beta) \in \mathcal{B}$
$\dots (m, \text{act} \blacktriangleright g) \dots \diamond \dots (\beta \leftrightarrow \widetilde{Mes}) \dots \diamond \mathcal{B} \longrightarrow$ $\dots (m, \text{act} \blacktriangleright g') \dots \diamond \dots (\beta \leftrightarrow \widetilde{Mes} \cdot op(m(x_1), \dots, m(x_n))) \dots \diamond \mathcal{B}$	
REC	$g \xrightarrow[\text{act}]{s?op(x_1, \dots, x_n)} g' \quad m(s) = \beta$
$\dots (m, \text{act} \blacktriangleright g) \dots \diamond \dots (\beta \leftrightarrow op(w_1, \dots, w_n) \cdot \widetilde{Mes}) \dots \diamond \mathcal{B} \longrightarrow$ $\dots (m[w_1/x_1, \dots, w_n/x_n], \text{act} \blacktriangleright g') \dots \diamond \dots (\beta \leftrightarrow \widetilde{Mes}) \dots \diamond \mathcal{B}$	

Table 5: sos Rules for Service Configurations

Definition 5.9. One-step Interaction-safe Running Configurations - A running configuration

$\mathcal{C} = ((M_1, \text{pic}_1)) \dots ((M_n, \text{pic}_n)) \diamond (m_1, \text{act}_1 \blacktriangleright g_1) \dots (m_k, \text{act}_k \blacktriangleright g_k) \diamond \mathcal{Q} \diamond \mathcal{B}$
is one-step interaction-safe iff for any (service or client) instance j , any session variable s , and any operation op the following implication holds:

$$m_j(s) \leftrightarrow op(w_1, \dots, w_l) \cdot \widetilde{Mes} \in \mathcal{Q} \quad (\text{for some values } w_1, \dots, w_l), \quad \text{and} \quad \text{open}(\text{act}_j, g_j, s)$$

$$\Rightarrow g_j \xrightarrow[\text{act}_j]{s?op(x'_1, \dots, x'_l)}$$

Definition 5.10. Interaction-safe Service Configuration - A service configuration C_{serv} is interaction-safe iff for any client instance (m, act, g) and any running configuration C reachable from $C_{serv} \diamond (m, \text{act}, g)$, C is one-step interaction-safe.

6 Related work

The potential of sessions in programming languages has gained recognition recently, and languages such as Java [17] and Erlang [16] have been extended in order to support them. In the service orchestration community, a significant body of work looks at formal models that support sessions for services as a first-class element of the language, such as in the Service-Centered Calculus (SCC) [1], CaSPiS [2], and Orcharts [6], among others. To the best of our knowledge, our work on SeB is the first that has introduced sessions into the widely adopted BPEL [18] orchestration language.

Other formalizations of BPEL have been suggested. For instance, in [11] the authors defined an algorithm to derive data-links from BPEL code by evaluating the control flow of processes as described by their control links. In our work, we have taken a step further and given an overall static semantics of variables in BPEL, which has allowed us to define the notion of a well structured activity. BPEL models

are quite often based on Petri nets [19, 15], which lend themselves well to the task of formalizing control links. By separating our semantics into two steps, we were able to propose a formal specification of BPEL with control links more concisely than with Petri nets, notably when it comes to capturing the behavior of dead-path elimination. Other work suggests a formalization that takes into account typing of BPEL process with WSDL descriptors [13]. The session types defined in the present paper can model WSDL descriptors, and they add the possibility to model a service's behavior.

BPEL [18] uses correlation sets rather than sessions to relate messages belonging to a particular instance of an interaction. Messages that are logically related are identified as sharing the same *correlation data* [12]. We could have defined a “session oriented” style for BPEL implemented with correlation sets, but this approach would not lend itself easily to behavioural typing which is our objective for SeB. While BPEL correlation sets allow for multi-party choreographies to be defined, we argue that similar expressivity is attainable with session types by extending them to support multi-party sessions. This challenge for future work has been addressed outside the context of BPEL in [9, 3]. In [20] a calculus based on process algebra and enhanced with a system for correlating related operations is presented. This calculus was shown to be able to reach a certain degree of BPEL-like expressiveness.

Blite [14] is a formalized subset of BPEL with operational semantics that take into account correlation sets. The authors have also written an associated translator that converts Blite code into executable BPEL. Blite uses a WSDL-typing system, but it does not feature control links.

7 Conclusion

In order to provide a basis for formal reasoning and verification of service orchestrations, we have adapted and formalized a subset of the widely adopted orchestration language BPEL. The resulting formalism, that we call SeB for Sessionized BPEL, supports sessions as first class citizens of the language. The separation of the proposed operational semantics into two steps has allowed a relatively concise semantics to be provided when compared to previous approaches. Furthermore our semantics take into account the effect of BPEL control links, which are an essential and often neglected part of the language.

In the sequel, we plan to use *session types* as a means of prescribing the correct structure of an interaction between two partner services during the fulfillment of a service. A typed SeB service will declare the session types that it can provide to prospective partners, while also declaring its required session types. A verification step will be needed to see whether or not a service is well-typed, hence answering the question of whether or not the service respects its required and provided types.

We will also investigate and show how the interaction safety property, that we defined on the basis of the semantics of SeB, can be guaranteed in configurations of compatible and well typed services. The formal approach taken with SeB as presented in this paper also opens up the possibility of defining and proving other properties of Web service interactions, such as controllability and progress properties. The study of all these verification aspects is left to future work.

References

- [1] M. Boreale, R. Bruni, L. Caires, R. De Nicola, I. Lanese, M. Loreti, F. Martins, U. Montanari, A. Ravara, D. Sangiorgi, V. Vasconcelos & G. Zavattaro (2006): *A Service Centered Calculus*. In M. Bravetti, M. Nez & G. Zavattaro, editors: *Web Services and Formal Methods, Lecture Notes in Computer Science* 4184, Springer Berlin / Heidelberg, pp. 38–57, doi:10.1007/11841197_3.
- [2] M. Boreale, R. Bruni, R. De Nicola & M. Loreti (2008): *Sessions and Pipelines for Structured Service Programming*. In G. Barthe & F. de Boer, editors: *Formal Methods for Open Object-Based Distributed Systems, LNCS* 5051, Springer Berlin / Heidelberg, pp. 19–38, doi:10.1007/978-3-540-68863-1_3.

- [3] R. Bruni, I. Lanese, H. Melgratti & E. Tuosto (2008): *Multiparty Sessions in SOC*. In D. Lea & G. Zavattaro, editors: *Coordination Models and Languages, Lecture Notes in Computer Science 5052*, Springer Berlin / Heidelberg, pp. 67–82, doi:10.1007/978-3-540-68265-3_5.
- [4] L. Cardelli & J. Mitchell (1990): *Operations on records*. In M. Main, A. Melton, M. Mislove & D. Schmidt, editors: *Mathematical Foundations of Programming Semantics, Lecture Notes in Computer Science 442*, Springer Berlin / Heidelberg, pp. 22–52, doi:10.1007/BFb0040253.
- [5] R. Chinnici, J.-J. Moreau, A. Ryman & S. Weerawarana (2007): *Web Service Definition Language (WSDL) Version 2.0*. Technical Report. Available at <http://www.w3.org/TR/wsdl20>.
- [6] A. Fantechi & E. Najm (2008): *Session Types for Orchestration Charts*. In D. Lea & G. Zavattaro, editors: *Coordination Models and Languages, Lecture Notes in Computer Science 5052*, Springer Berlin / Heidelberg, pp. 117–134, doi:10.1007/978-3-540-68265-3_8.
- [7] J. Fernandez, H. Garavel, A. Kerbrat, L. Mounier, R. Mateescu & M. Sighireanu (1996): *CADP a protocol validation and verification toolbox*. In R. Alur & T. Henzinger, editors: *Computer Aided Verification, Lecture Notes in Computer Science 1102*, Springer Berlin / Heidelberg, pp. 437–440, doi:10.1007/3-540-61474-5_97.
- [8] J. Groote & J. van de Pol (2000): *State Space Reduction Using Partial τ -Confluence*. In M. Nielsen & B. Rovan, editors: *Mathematical Foundations of Computer Science 2000, Lecture Notes in Computer Science 1893*, Springer Berlin / Heidelberg, pp. 383–393, doi:10.1007/3-540-44612-5_34.
- [9] K. Honda, N. Yoshida & M. Carbone (2008): *Multiparty asynchronous session types*. *SIGPLAN Not.* 43, pp. 273–284, doi:10.1145/1328897.1328472.
- [10] D. Kitchin, A. Quark, W. Cook & J. Misra (2009): *The Orc Programming Language*. In D. Lee, A. Lopes & A. Poetzsch-Heffter, editors: *Formal Techniques for Distributed Systems, Lecture Notes in Computer Science 5522*, Springer Berlin / Heidelberg, pp. 1–25, doi:10.1007/978-3-642-02138-1_1.
- [11] O. Kopp, R. Khalaf & F. Leymann (2008): *Deriving Explicit Data Links in WS-BPEL Processes*. In: *IEEE International Conference on Services Computing, 2008. SCC '08.*, 2, pp. 367–376, doi:10.1109/SCC.2008.122.
- [12] A. Lapadula, R. Pugliese & F. Tiezzi (2007): *A Calculus for Orchestration of Web Services*. In R. De Nicola, editor: *Programming Languages and Systems, Lecture Notes in Computer Science 4421*, Springer Berlin / Heidelberg, pp. 33–47, doi:10.1007/978-3-540-71316-6_4.
- [13] A. Lapadula, R. Pugliese & F. Tiezzi (2011): *A WSDL-based type system for asynchronous WS-BPEL processes*. *Formal Methods in System Design* 38(2), pp. 119–157, doi:10.1007/s10703-010-0110-0.
- [14] A. Lapadula, R. Pugliese & F. Tiezzi (2012): *Using formal methods to develop WS-BPEL applications*. *Sci. Comput. Program.* 77(3), pp. 189–213, doi:10.1016/j.scico.2011.03.002.
- [15] A. Martens (2005): *Analyzing Web Service Based Business Processes*. In M. Cerioli, editor: *Fundamental Approaches to Software Engineering, Lecture Notes in Computer Science 3442*, Springer Berlin / Heidelberg, pp. 19–33, doi:10.1007/978-3-540-31984-9_3.
- [16] D. Mostrous & V. Vasconcelos (2011): *Session Typing for a Featherweight Erlang*. In W. De Meuter & G.-C. Roman, editors: *Coordination Models and Languages, Lecture Notes in Computer Science 6721*, Springer Berlin / Heidelberg, pp. 95–109, doi:10.1007/978-3-642-21464-6_7.
- [17] N. Ng, N. Yoshida, O. Pernet, R. Hu & Y. Kryftis (2011): *Safe Parallel Programming with Session Java*. In: *Coordination Models and Languages, Lecture Notes in Computer Science 6721*, Springer Berlin / Heidelberg, pp. 110–126, doi:10.1007/978-3-642-21464-6_8.
- [18] OASIS (2007): *Web Services Business Process Execution Language (WS-BPEL) Version 2.0*. Available at <http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.html>.
- [19] C. Ouyang, E. Verbeek, W. M. P. van der Aalst, S. Breutel, M. Dumas & A. H. M. ter Hofstede (2007): *Formal semantics and analysis of control flow in WS-BPEL*. *Sci. Comput. Program.* 67(2-3), pp. 162–198, doi:10.1016/j.scico.2007.03.002.
- [20] M. Viroli (2007): *A core calculus for correlation in orchestration languages*. *Journal of Logic and Algebraic Programming* 70(1), pp. 74 – 95, doi:10.1016/j.jlap.2006.05.006.