

A formalisation of Adaptable Pervasive Flows[★]

Antonio Bucchiarone¹, Alberto Lluch Lafuente²,
Annapaola Marconi¹ and Marco Pistore¹

¹FBK-IRST, via Sommarive 18, 38050, Trento, Italy
[bucchiarone,marconi,pistore]@fbk.eu

²Department of Computer Science, University of Pisa
lafuente@di.unipi.it

Abstract. *Adaptable Pervasive Flows* is a novel workflow-based paradigm for the design and execution of pervasive applications, where dynamic workflows situated in the real world are able to modify their execution in order to adapt to changes in their environment. In this paper, we study a formalisation of such flows by means of a formal flow language. More precisely, we define APFoL (Adaptable Pervasive Flow Language) and formalise its textual notation by encoding it in Blite, a formalisation of WS-BPEL. The encoding in Blite equips the language with a formal semantics and enables the use of automated verification techniques. We illustrate the approach with an example of a Warehouse Case Study.

1 Introduction

Flows are models defining a set of activities to be done, and their relations with each other. Flows are deeply seated in many fields, including business processes and service oriented computing. The flow modeling paradigm is often used either implicitly or explicitly in many real life situations. In this paper, we concentrate on a novel usage of flows, which is being investigated by the ALLOW project [1]: the usage of flows as a new programming paradigm for human-oriented pervasive applications. More precisely, *Adaptable Pervasive Flows* (APFs) [10] are proposed as an extension of traditional workflow concepts [19] in order to make them more flexible with respect to their pervasive execution environment. APFs are dynamic workflows situated in the real world that modify their execution in order to adapt to changes in their environment. This requires on the one hand that a flow must be context-aware: during execution it must be possible to obtain information on the underlying environment (e.g. relevant information on world entities, status of other flows, human activities). On the other hand flow models must be flexible enough to allow an easy and continuous adaptation. APFs are based on WS-BPEL [5], a well-known language for specifying flows in a Web Service setting, and extend it in order to implement all the aspects related to pervasive applications.

In [16] the authors define one of these extensions to WS-BPEL. More precisely, they define a set of constructs that allow a convenient way of embedding the adaptation logic within the specification of an APF and show how WS-BPEL can be extended to

[★] Research supported by the EU, STREP project Allow IST-324449 and Sensoria, IST-2005-016004.




Fig. 1. Warehouse structure

support the proposed constructs. These constructs allow for capturing interesting cases of adaptation in pervasive applications that are difficult to address with classical workflows and with the standard WS-BPEL language. In this paper, we extend the work in [16], by providing a formal model for APFs and for the extensions of WS-BPEL related to the adaptation logics. This is a significant extension of the previous work since, due to the high dynamicity of pervasive applications, formal methods become crucial to drive design disciplines, equip existing languages with well-defined semantics and increase the reliability by means of automated verification.

Web services are a good example where many efforts are being invested on the development and application of formal methods. For instance, there have been various proposals to define formal semantics for WS-BPEL, typically by means of process calculi (e.g. [12], [15]), or Petri nets (e.g. [14], [18]), or graphs (e.g. [7]).

In this paper we propose APFoL, a formal language for adaptable pervasive flows. Amongst the various formal approaches to flow languages we have chosen Blite [13] as a starting point. Blite is a process calculus that captures a significant part of the WS-BPEL language. We build our language as an almost straightforward extension of Blite. More precisely, we equip the language with abbreviations to deal with adaptation mechanisms, flow constructs and activity types typical of adaptable pervasive flows. The formal language proposed permit us to formally specify the built-in adaptation constructs informally proposed in [16].

This document is structured as follows. Section 2 introduces a scenario from the Warehouse Case Study of the Allow project [1] while Section 3 presents the background that is used in section 4 where our flow language is introduced. Section 5 illustrates our language with an example drawn from the Warehouse Case Study. Finally, Section 6 draws conclusions and outlines current and future work.




Fig. 2. Good receipt procedure

2 The Running Example: Warehouse Management

We present a scenario from the Warehouse Case Study of the Allow project [1], namely the management of a warehouse that we will use as a reference for all the examples within this document. The main aim of warehouse management is to organize and control the transport and storage of goods within a warehouse. This is achieved through the definition and processing of complex transactions, including shipping, receiving, put-away, picking and issuing of goods. The objective of the warehouse management system is to provide a set of computerised procedures supporting all the aforementioned activities: from the handling of goods reception, storage and shipping, to the management of all the physical storage facilities.

Warehouse management often utilizes auto AIDC (Automatic Identification and Data) technology, such as bar-code scanners, mobile computers, wireless LANs and potentially RFID (Radio Frequency Identification) to efficiently monitor the flow of products. Current systems require that, once data has been collected, synchronization with a centralized system is performed. The centralized system is in charge of controlling all aspects of warehouse management. Pervasive flows offer the possibility to distribute the control logics and hence to improve flexibility and context awareness of the executed processes.

The (logical and physical) structure of a warehouse is described by Figure 1. Doors are the locations where the goods arrive at or leave the warehouse. Trucks drive up to the doors of a warehouse in order to unload or load goods there. Staging areas are used for interim storage of goods in the warehouse. These are located in close proximity to the doors associated to them. The storage area is organized in several zones corresponding to different storage types. Storage types are physical or logical subdivisions of a warehouse complex, characterized by its warehouse technique, the space used, its organizational form, or its function.

Warehouse management requires the execution of different procedures which refer to the different objects and human actors, including *good receipt*, *issuing* and *transfer*. Here we focus on the first procedure, which describes the three steps (see Figure 2) that have to be performed when the goods arrive at the warehouse: *delivery* (a truck has reached the warehouse and is docked at a door); *unload* (goods are unloaded and temporally stored in the staging area); and *put-away* (goods are moved to the storage area).

3 Background

Our work is strongly based on a proposal for the extension of WS-BPEL to deal with adaptation [16] and the process calculus Blite [13].

3.1 WS-BPEL and APFs

Similar to the well-known workflows, APFs consist of a set of activities and a corresponding execution order, which is specified using control elements such as sequence, choice or parallel operators. After a deep analysis and comparison of today's workflow standards using criteria such as industry impact, robustness aspect and extensibility [2], the ALLOW project has chosen WS-BPEL as a nucleus for an APF language.

A particular feature of APFs is that they are situated in the real world. This realizes the *pervasiveness* of the flows and is achieved in two ways. First, the flows are logically attached to physical entities (which can be either objects or humans) and move with them through different contexts. Secondly, they run on physical devices (e.g. PDAs, desktops). Today's workflow languages (e.g. WS-BPEL) provide no possibility to explicitly model the context model and constraints on the workflow environment. The pervasiveness of a WS-BPEL workflow can be specified only through ad-hoc interactions with external context-aware services. Clearly, this solution affects the readability and transparency of the pervasive aspect within the specified workflow. A first extension [9] that has been done aims at providing a modeling approach to annotate WS-BPEL processes with contextual constraints and an execution model to monitor those constraints during process execution.

Another important aspect of APFs is their *adaptiveness*. A flow is a dynamic entity that modifies its execution in order to adapt to changes in the execution environment. A key enabling factor for automated adaptation mechanisms is a convenient way of embedding the adaptation logic within the specification of a flow. Adaptation mechanisms cannot be limited to standard recovery constructs (e.g. fault/event/compensation handlers in WS-BPEL), but should also support the specification of flexible context-aware reactions to adaptation needs that can be used to handle run-time flow deviations without requiring a flow recovery/failure. The aim of the work in [16] is to present a set of primitives and principles that can support the encoding of context-aware run-time deviations and changes within a flow model in a secure (from an execution perspective) and convenient (from a modelling perspective) way. The authors propose a set of *built-in adaptation* modeling constructs that can be useful to add dynamicity and flexibility to flow models and for each construct they define the corresponding WS-BPEL extension. In particular, the proposed constructs are (i) conditional branches within flows with context conditions as guard conditions, (ii) context handlers that allow to automatically react to context conditions violation during the execution of the flow, and (iii) constructs that allow to specify a set of alternative scopes, each handling a specific execution context, and that allow to jump at run-time from one scope to another, whenever the context changes or the assumptions on the context turn out to be wrong.

<i>Basic activities</i>	$b ::= \text{inv } \ell^i \circ \bar{x} \mid \text{rcv } \ell^r \circ \bar{x} \mid x := e$ empty throw exit	invoke, receive, assign empty, throw, exit
<i>Structured activities</i>	$a ::= b \mid \text{if}(e)\{a_1\}\{a_2\} \mid \text{while}(e)\{a\}$ $a_1 ; a_2 \mid \sum_{j \in J} \text{rcv } \ell_j^r \circ_j \bar{x}_j ; a_j$ $a_1 \mid a_2 \mid [a \bullet a_f \star a_c]$	basic, conditional, iteration sequence, choice (with $ J > 1$) parallel, scope
<i>Start activities</i>	$r ::= \text{rcv } \ell^r \circ \bar{x} \mid \sum_{j \in J} \text{rcv } \ell_j^r \circ_j \bar{x}_j ; a_j$ $r ; a \mid r_1 \mid r_2 \mid [r \bullet a_f \star a_c]$	receive, choice sequence, parallel, scope
<i>Services</i>	$s ::= [r \bullet a_f] \mid \mu \vdash a \mid \mu \vdash a, s$	definition, instance, multiset
<i>Deployments</i>	$d ::= \{s\}_c \mid d_1 \parallel d_2$	deployment, composition

Table 1. Syntax of Blite

3.2 WS-BPEL and Blite

Blite [13] is a formal language for describing web service orchestrations. It has been designed as a formal model to capture the essentials of WS-BPEL, a de-facto standard for describing web services. Blite is a process calculus and as such it has a well-defined notion of syntax and operational semantics. The language includes features such as service definition and instantiation, typical flow constructs, communication primitives and failure handling and compensation mechanisms. We offer here a brief, intuitive overview of Blite and refer to [13] for a detailed presentation.

The syntax of the Blite language is summarised in table 1 (borrowed from [13]). *Basic activities* include variable assignments, flow success related operations (throw and exit) and communication primitives to send (inv) or receive (rcv) values from partner links. *Structured activities* organise basic activities in flows by using typical flow constructs such as branches, sequences, loops, fork&join and (input-prefixed) choices¹. In addition, scopes can be defined with appropriate failure and compensation activities. *Start activities* are structured activities starting with a choice of receive operations. The reason for this is that service definitions are inactive until they receive a request. Services already instanced, instead are represented by their memory μ (an assignment of values to variables) and their flow (a structured activity). Deployments (i.e. the system) are sets of correlated services.

4 A formal language for adaptable pervasive flows

This section presents our formalisation of adaptable pervasive flows in terms of a flow language that we call APFoL. We present here our language for adaptable pervasive flows, describing each ingredient and giving its encoding in Blite. It is worth mentioning that the encoding automatically equips our language with a formal semantics.

We start offering an informal presentation of our visual notation, which we plan to formalise in the future, possibly by means of a graphical encoding of our language.

¹ Called *pick* in [13] but *choice* here to avoid confusion with the pick flow construct of APFoL.

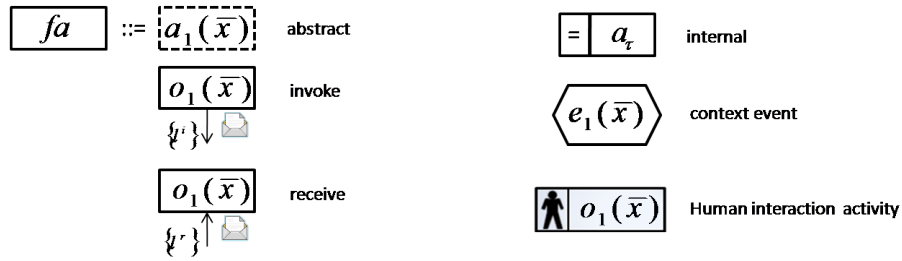


Fig. 3. Visual syntax for Flow Activities of Table 2.

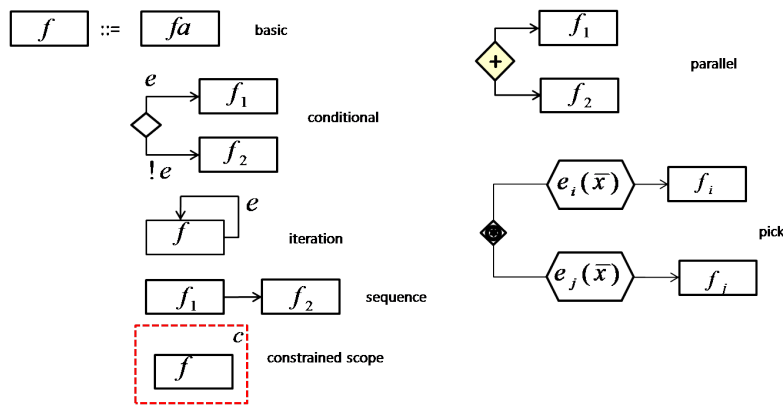


Fig. 4. Visual syntax for Flow Instances (Part-I) of Table 2.

Here we just present the informal visual notation as Figures 3, 4 and 5, that act as an illustration of the textual notation that we shall describe in detail.

We now present the textual notation of APFoL (summarised in Table 2), which basically extends the syntax of Blite (c.f. Table 1) with ad-hoc constructs for built-in adaptation mechanisms.

The main differences with respect to Blite’s syntax regard basic and structured activities. We call them called *basic and structured flows* in APFoL to avoid confusion and stick to the APF slang. They include some new constructs to model the relevant primitives of adaptable pervasive flows. Services and deployments remain identical but, again, we call them differently (*flows and flow systems*) to avoid confusion. Finally, we avoid presenting the productions for *start flows* for simplicity: they are a straightforward adaptation of those for *start activities* in the same way as flow instances are an adaptation of activity instances.

First, APFoL includes the same control flow constructs of Blite. In addition, even if not part of the primitive syntax summarised in Table 1, APFoL includes typical control flow constructs such as different forms of branching (e.g. switch) and looping (e.g. loop-exit) which are straightforwardly encoded in Blite.

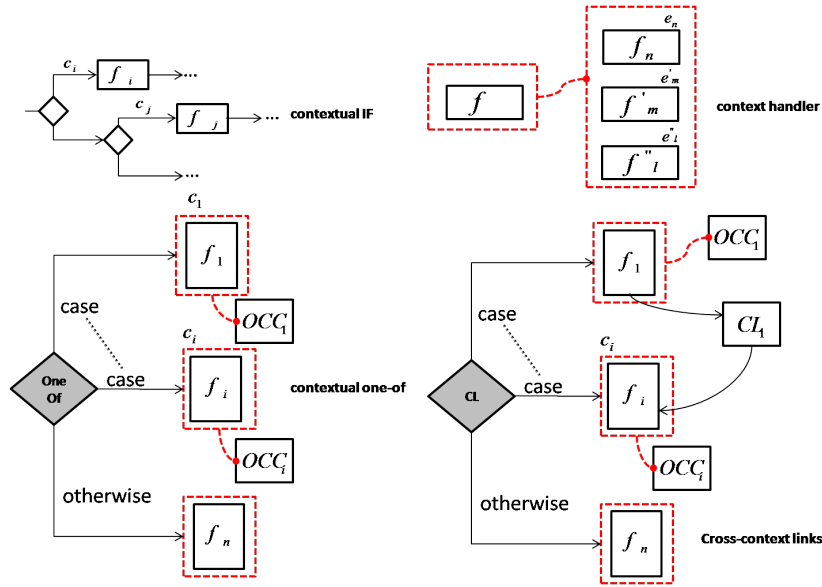


Fig. 5. Visual syntax for Flow Instances (Part-II) of Table 2.

An *abstract activity* $A(\bar{x})$ represents either a partial design-time specification of a flow model or an abbreviation of a complex activity. Abstract activities are modelled just as function symbols A of type fa (flow activity) or fr (start activity), thus $A(\bar{x}) : fa \cup fr$ in Table 2. For each such symbol we assume a definition to exist (for abbreviations) or to be given at run-time (for partial designs). Refining an abstract activity then means replacing the left-hand side of a definition by its right-hand side. Clearly, this is not a real extension of the language and is standard machinery of all algebraic specifications.

Communication activities allow for sending (resp. receiving) a message to (resp. from) another flow. Invoke and receive activities are synchronous. Thus we encode them as suggested in [13] by the authors of Blite, namely by a pair of receive and invoke actions. In the definition, id stands for the flow instance identity. A sending activity is thus encoded as the invocation of operation o at partner l , where the identity of the flow is passed to receive the response. Similarly, the receive activity expects to receive an invocation of operation o at himself (id) together with the invoker's identity q which is used to send the response.

Data manipulation activities are internal activities that change the value of local variables and do not interact with their environment. Data manipulation activities are modelled as structured activities whose component basic activities a_τ are all assignments. Note that the grammar for a_τ can be given but, for the sake of a clear presentation, prefer to avoid adding an ad-hoc syntactical category and its (rather redundant) productions.

Human interaction activities are activities that require an interaction with a human, e.g. displaying or getting information through a device. Human interaction activities are modelled as communication operations, since devices are represented by flows.

<i>Flow activities</i>	$fa ::= \text{sinv } l \circ \bar{x} \mid \text{srcv } l \circ \bar{x} \mid a_\tau$ $\langle e\bar{x} \rangle \mid A(\bar{x})$	invoke, receive, internal context event, abstract activity
<i>Flows instances</i>	$f ::= fa \mid \text{if}(e)\{f_1\}\{f_2\} \mid \text{while}(e)\{f\}$ $f_1 ; f_2 \mid f_1 \parallel f_2$ $\llbracket f \rrbracket_a^c \mid \text{pick}_{i \in J}(e_i \rightarrow f_i)$ $\text{clF}_{i \in J}(c_i \rightarrow f_i)$ $\llbracket f \rrbracket_{f_1}^{e_1} \llbracket f_n \rrbracket_{f'_1}^{e'_1} \llbracket f'_m \rrbracket_{f''_1}^{e''_1} \llbracket f''_l \rrbracket$ $\text{one - of} \llbracket f_1 \rrbracket_{r_1}^{c_1} \dots \llbracket f_n \rrbracket_{r_n}^{c_n}$	basic, conditional, iteration sequence, parallel constrained scope, pick contextual IF context handler contextual one-of
<i>Flows</i>	$ff ::= [\text{fr} \bullet f] \mid \mu \vdash f \mid \mu \vdash \dot{f}, ff$	definition, instance, multiset
<i>Flow system</i>	$fs ::= \{\text{ff}\}_c \mid fs_1 \parallel fs_2$	deployment, composition

Table 2. Syntax of APFoL

Sending activity	$\text{sinv } q \circ \bar{x} \stackrel{\text{def}}{=} \text{inv } \langle \text{id}, q \rangle \circ \bar{x} ; \text{rcv } \langle \text{id} \rangle \circ \text{ack}$
Receiving activity	$\text{srcv } q \circ \bar{x} \stackrel{\text{def}}{=} \text{rcv } \langle \text{id}, q \rangle \circ \bar{x} ; \text{inv } \langle q \rangle \circ \text{ack}$
Context event	$\langle e \bar{x} \rangle \stackrel{\text{def}}{=} \text{rcv } \text{ContextManager } \text{get}_e \bar{x}$
Internal event	$\langle e \rangle \stackrel{\text{def}}{=} \text{while}(e)\{\text{empty}\}$
Constrained scope	$\llbracket f \rrbracket_a^c \stackrel{\text{def}}{=} \langle c \rangle [\{f\}^{-c} \mid (-c); \text{throw}] \bullet a * \text{empty}$
Pick	$\text{pick}_{i \in J}(e_i \bar{x}_i \rightarrow a_i) \stackrel{\text{def}}{=} \sum_{i \in J} \langle e_i \bar{x}_i \rangle ; a_i$
Contextual IF	$\text{clF}_{i \in J}(c_i \rightarrow f_i) \stackrel{\text{def}}{=} \text{switch}_{i \in J} c_i \rightarrow \llbracket f_i \rrbracket_{\text{throw}}^{c_i}$

Table 3. Blite encoding of the main ingredients of APFoL.

Context events are a special type of activities for receiving events broadcasted by a particular entity called *Context Manager*. More precisely, during this activity the flow execution waits until the event is received. We model context events as particular receive operations. More precisely, we shall model context managers as services that broadcast their events e via replicated sending operations named get_e . Thus, the reception of the event is modelled as a reception operation with the context manager as partner, operation get_e and the corresponding tuple of values.

We shall also use a sort of *internal events*, denoted by $\langle e \rangle$ whose meaning is to wait until the expression (i.e. a condition or trigger) e is true.

A *constrained scope* $\llbracket f \rrbracket_a^c$ is a flow f enclosed into a scope with unique entry and exit points, a constraint c and adaptation a (triggered if the constraint is not valid). they are represented in our syntax by terms of the form $\llbracket f \rrbracket_a^c$, where f is the normal flow, c is the constraint and a is the adaptation to be performed in case the condition fails inside the scope. This is modelled in Blite by exploiting the failure mechanism. First, we wait until the condition is true. Then we open a Blite scope where we put a condition observer

$$\begin{aligned}
& \llbracket \text{main} \rrbracket \llbracket \overset{\text{fault}_1}{ff_1} \rrbracket \dots \llbracket \overset{\text{fault}_n}{ff_n} \rrbracket \llbracket \overset{\text{event}_1}{ef_1} \rrbracket \dots \llbracket \overset{\text{event}_m}{ef_m} \rrbracket \llbracket \overset{\text{block}_1}{bf_1} \rrbracket \dots \llbracket \overset{\text{block}_l}{bf_l} \rrbracket \stackrel{\text{def}}{=} \\
& \text{new } done; \\
& done := false; \\
& \llbracket \{main\}^{suspend} \rrbracket \\
& \quad | \text{fault}_1 \rightarrow throw \\
& \quad | \dots \\
& \quad | \text{fault}_n \rightarrow throw \\
& \quad | \text{event}_1 \rightarrow ef_1 \\
& \quad | \dots \\
& \quad | \text{event}_m \rightarrow ef_m \\
& \quad | \text{block}_1 \rightarrow bf_1; done := true \\
& \quad | \dots \\
& \quad | \text{block}_l \rightarrow bf_l; done := true \\
& \llbracket \overset{\text{fault}_1 \vee \dots \vee \text{fault}_{1n}}{\text{switch}\{\text{fault}_1 \rightarrow ff_1 \dots \text{fault}_n \rightarrow ff_n\}} \rrbracket \\
suspend & \stackrel{\text{def}}{=} (\neg done \wedge \text{block}_1) \vee \dots \vee (\neg done \wedge \text{block}_l)
\end{aligned}$$

Fig. 6. Context handler in Blite

flow in parallel with the normal flow conditioned to $\neg c^2$. Conditioning is necessary to avoid the normal flow to progress in case the context condition violated. Note that this semantics does not really interrupt the flow. The exception is raised and the failure code performs the adaptation activity, only when the observer is executed.

A *pick* is a branching point in the process where the alternatives are based on events, rather than the evaluation of expressions. More precisely, it is the receipt of a message or of a context event that determines which of the paths will be taken. Event-based decision is modelled by a non deterministic choice of activities preceded by the corresponding triggering event.

A *contextual IF* allows to define several flow fragments as possible branches in the execution of the flow. Each flow fragment has an associated *context condition*. We can define also one flow fragment without a context condition, which will encode the default behaviour. The operational semantics of this construct is similar to a traditional if: the first fragment for which the *context condition* holds will be selected and executed.

A *context handler* is a particular flow associated to a *main* flow or flow scope. It specifies an alternative flow (in the form of a contextual IF) to be applied if a corresponding scope condition is violated. There are various forms of conditions within the context handler flow.

- *Fault-triggering* suspend all active tasks in the main flow and execute the corresponding error-handler. If the main flow is a flow scope, the fault is propagated to the enclosing scope.

² This is done by inserting a condition (in form of an *internal event*) between each activity and can be easily defined in a recursive manner.

```

one – of  $\llbracket f_1 \rrbracket_{r_1}^{c_1} \dots \llbracket f_n \rrbracket_{r_n}^{c_n} \stackrel{\text{def}}{=}
\text{new } success;
\text{loop}
\text{switch}
\quad c_1 \rightarrow \llbracket f_1; success := true \rrbracket_{r_1}^{c_1}
\quad \dots
\quad c_n \rightarrow \llbracket f_n; success := true \rrbracket_{r_n}^{c_n}
\text{if}(success) \text{exit};$ 
```

Fig. 7. Contextual One-of in Blite

- *Event-triggering* conditions can be *non-blocking* (execution of the main flow proceeds normally and the corresponding flow is executed concurrently) or *blocking* (execution of the main flow suspends, then the corresponding flow is executed and finally the main flow is resumed).

The encoding of context handlers in Blite is defined in Figure 6. The construct consists of a main flow *main* and three sets of observers: fault-handlers, non-blocking event handlers, and blocking event handlers, whose triggers are respectively denoted by *ff*, *ef* and *bf*. The idea is as follows: the main flow (conditioned to *suspend*) is put in parallel with various observers one for each fault and event trigger. When a fault condition triggers, an error is raised and handled by the error handler which selects a fault and fires the corresponding flow. Non-blocking events trigger the corresponding flow. Blocking events perform similarly, but note that the *suspend* condition depends on the blocking event conditions: if one of them is true and the corresponding flow has not been performed, the main flow remains suspended until *done* becomes true³.

A *contextual one-of* consists of a set of alternative flow fragments, each of them associated to a contextual condition modeling the contextual assumption for that fragment, and a rollback flow that can be executed to undo the partial and unsuccessful work of the fragment. At run-time, the first flow fragment for which the contextual condition holds is chosen and executed. During the fragment execution, its context condition is monitored and, as soon as it is violated, the following actions are performed:

1. *stop execution*: all running activities within the fragment are stopped;
2. *undo partial work*: the roll-back flow associated to the current fragment is executed;
3. *context jump*: the first fragment for which the associated context holds is executed and its context condition is monitored.

Roll-back flows can throw fault/exceptions (e.g. to handle the fact that the work done within the fragment cannot be undone), and in this case the flow is terminated following

³ Note that in order to guarantee a unique *done* variable we declare it as *new* at the beginning of the encoding. This is not a feature of Blite but can be added straightforwardly by considering the local store μ as a stack of assignment sets instead of a plain set.

```

CL  $\llbracket f_1 \rrbracket_{r_1}^{c_1} \dots \llbracket f_n \rrbracket_{r_n}^{c_n} \stackrel{\text{def}}{=}
\text{new } success; \text{new } next\_flow; next\_flow := any;
\text{loop}
\text{switch}
\quad c_1 \wedge proceed(1) \rightarrow \llbracket f_1; success := true \rrbracket_{r_1}^{c_1}
\quad \dots
\quad c_n \wedge proceed(n) \rightarrow \llbracket f_n; success := true \rrbracket_{r_n}^{c_n}
\text{if}(success) \text{exit};$ 
```

Fig. 8. Contextual One-of with CLs in Blite

normal flow fault handling. If this is not the case, and the roll-back flow completes successfully, the main flow is considered successfully running.

The encoding in Blite is rather easy, a loop is used to guarantee that performing a rollback returns to the selection of one of the choices. The only way to exit a loop is to successfully finish one of main flows.

When using the *contextual one-of*, it may be the case that, when jumping from one execution context to another, we do not want to undo the work done or the complete flow rollback is not possible. The *cross-context link* (CL) is designed especially for this case. CLs connect two activities of different scopes within a *contextual one-of*. CLs allow adapting to a context change by jumping from a certain execution state of the current activity (*source* activity) to an execution activity (*target* activity) of another fragment suitable for the actual context. After the jump the flow instance must be in a consistent state. Therefore, a CL has an associated flow needed to prepare the flow to the jump. At runtime, if the contextual condition associated to the running scope turns out to be false, two possibilities are considered:

1. if there exists some context link leaving the active activity for which the context condition holds:
 - (a) the roll-back flow associated to the cross-context link is executed
 - (b) the monitoring for the new context condition is activated
 - (c) the flow execution is re-started from the target activity of the CL
2. otherwise the condition violation is handled as described for the standard *contextual one-of*.

The encoding in Blite is similar to the encoding of ordinary contextual one-of. The first difference is that the guard of each flow f_i is enriched with $proceed(i)$ which is an abbreviation for $next_flow = i \vee next_flow = any$. This serves to control which flow should be executed next. The second difference is that each compensation c must take the ad-hoc roll-back flows for cross-context jumps into consideration. With $|c|$ we denote the introduction of a choice that decides whether to apply the ordinary compensation c or the roll-back flow associated to the jump to the next flow.

Box Unloading($b, wa, iLoc, stgLoc, strLoc$)

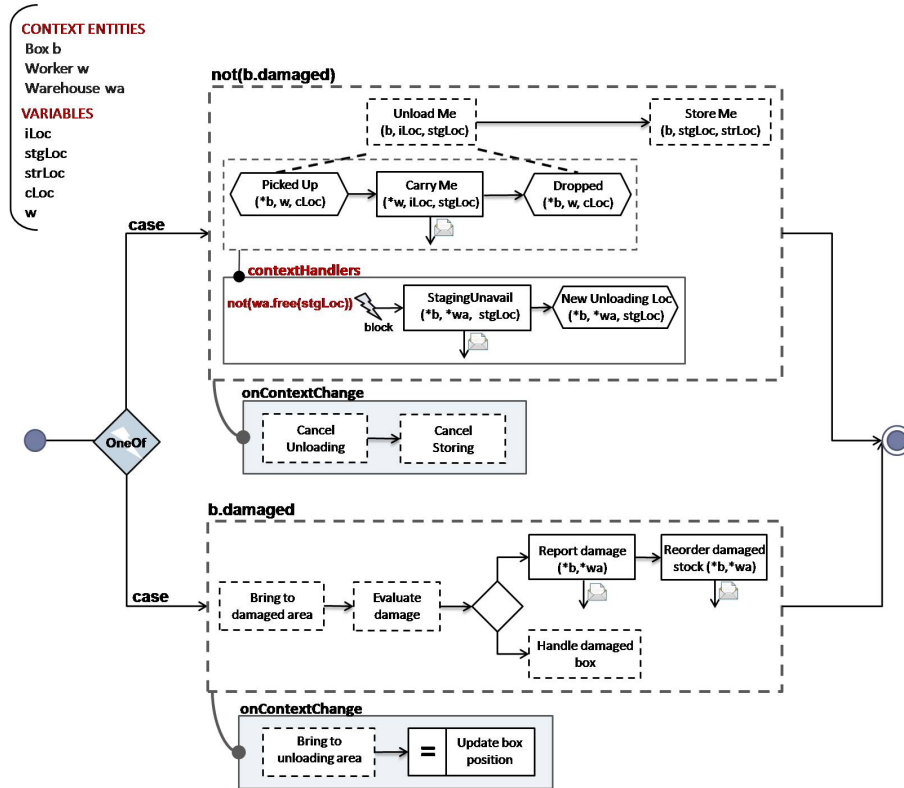


Fig. 9. The Box Unloading example

5 The Box Unloading example

In this section we present a complete example that summarizes most elements introduced before. For exemplification, we consider the box flow. A first problem that can occur here is that the box can be damaged. The damage may have occurred either before, during transportation, but it may also occur at any point while the box is being unloaded to the staging area, or moved to the storage area.

In Figure 9 we use the Contextual OneOf construct to model the handling of damaged boxes. In case the box is not damaged, the first flow scope is chosen and executed. If at any point the box gets damaged, the context condition *not(b.damaged)* is violated and the *onContextChange* flow is executed. That is, pending activities for unloading and/or storing are canceled (e.g. the reserved staging/storage location is made available for other boxes, the request for unloading/storing sent to workers are revoked). The specific activities to be performed clearly depend on the state of execution, due to this abstract activities are specified, namely *Cancel Unloading* and *Cancel Storing*

The APFoL code of this example is flow Box Unloading shown in Figure 5, while its WS-BPEL code is listed in [3].

6 Conclusion and Future Work

We have described a preliminary version of APFoL, a language for adaptable pervasive flows with formal support. More precisely, we have presented a language whose textual notation is based on Blite [13], a process calculus for WS-BPEL.

The formalisation of the language equips the language with a well-defined (and hence non-ambiguous) semantics. More precisely, the formalisation as a process calculus (Blite) facilitates the use of automated verification techniques. Some support for the analysis and verification of Blite specifications exists (an encoding from Blite into another calculus with tool support [4]), but we are working in a direct implementation of Blite semantics in the rewrite engine Maude [8], in order to exploit its generic built-in capabilities in form of analysis tools such as a model checker and a theorem prover. With such an implementation at hand, APFoL can be implemented as a derived rewrite theory in Maude.

Our approach has been illustrated with examples from the Warehouse case study of the Allow project [1].

We plan to develop a graph-based formalisation of our visual notation, possibly basing existing techniques for the graphical encoding of process calculi (e.g. [6]). The main goals of having a formal graph-based representation is formalise the relation between textual and visual notation and to enable the use of graph transformation techniques, and their corresponding tools.

In the future we would like to investigate the connection with apparently similar approaches in the data base community around the notion of *business artifacts* [17, 11], which are flows attached to physical objects moving through different contexts.

Acknowledgments

The authors would like to thank Roberto Bruni for sharing with us his knowledge on WS-BPEL formalisations and the authors of Blite (Rosario Pugliese, Francesco Tiezzi and Alessandro Lapadula) for providing us with useful material.

References

1. EU-FET project 213339 ALLOW. <http://www.allow-project.eu/>.
2. D3.1 Basic flow-model and language for Adaptable Pervasive Flows. ALLOW Project Deliverable, Nov 2008.
3. APFoL homepage. <http://www.antonibucchiarone.it/apfol.html>.
4. Blite: A formal account of WS-BPEL. <http://rap.dsi.unifi.it/blite/>.
5. OASIS WSBPEL Technical Committee. Web Services Business Process Execution Language, version 2.0. Available at <http://docs.oasis-open.org/wsbpel/2.0/wsbpel-v2.0>, 2007.
6. R. Bruni, F. Gadducci, and A. Lluch Lafuente. A graph syntax for processes and services. In *9th International Workshop on Web Services and Formal Methods (WS-FM'09)*, 2009.

7. M. Bundgaard, A. J. Glenstrup, T. T. Hildebrandt, E. Højsgaard, and H. Niss. Formalizing higher-order mobile embedded business processes with binding bigraphs. In D. Lea and G. Zavattaro, editors, *COORDINATION*, volume 5052 of *Lecture Notes in Computer Science*, pages 83–99. Springer, 2008.
8. M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. L. Talcott. *All About Maude — A High-Performance Logical Framework. How to Specify, Program and Verify Systems in Rewriting Logic*, volume 4350 of *Lecture Notes in Computer Science*. Springer, 2007.
9. H. Eberle, S. Fll, K. Herrmann, F. Leymann, A. Marconi, T. Unger, and H. Wolf. Enforcement from the inside: Improving quality of bussiness in process management. In *IEEE 7th International Conference on Web Services (ICWS'09)*, 2009. To appear.
10. K. Herrmann, K. Rothermel, G. Kortuem, and N. Dulay. Adaptable Pervasive Flows - An Emerging Technology for Pervasive Adaptation. In *Workshop on Pervasive Adaptation (PerAda)*. IEEE Computer Society, September 2008.
11. R. Hull. Artifact-centric business process models: Brief survey of research results and challenges. In *OTM '08: Proceedings of the OTM 2008 Confederated International Conferences, CoopIS, DOA, GADA, IS, and ODBASE 2008. Part II on On the Move to Meaningful Internet Systems*, pages 1152–1163, Berlin, Heidelberg, 2008. Springer-Verlag.
12. C. Laneve and G. Zavattaro. Web-pi at work. In *Proceedings of the 1st International Symposium on Trustworthy Global Computing (TGC'05)*, volume 3705 of *Lecture Notes in Computer Science*, pages 182–194. Springer, 2005.
13. A. Lapadula, R. Pugliese, and F. Tiezzi. A formal account of WS-BPEL. In *10th International Conference on Coordination Models and Languages (COORDINATION'08)*, volume 5052 of *Lecture Notes in Computer Science*, pages 199–215. Springer, 2008.
14. N. Lohmann. A feature-complete Petri net semantics for WS-BPEL 2.0. In *4th International Workshop on Web Services and Formal Methods (WS-FM'07)*, volume 4937 of *Lecture Notes in Computer Science*, pages 77–91. Springer, 2008.
15. R. Lucchi and M. Mazzara. A pi-calculus based semantics for WS-BPEL. *Journal of Logic and Algebraic Programming*, 70(1):96–118, January 2007.
16. A. Marconi, M. Pistore, A. Sirbu, H. Heberle, F. Leymann, and T. Unger. Enabling Adaptation of Pervasive Flows: Built-in Contextual Adaptation. Manuscript submitted to an international conference. <http://www.antoniobucchiarone.it/publicazioni/SOA-TR09.pdf>, June 2009.
17. A. Nigam and N. S. Caswell. Business artifacts: An approach to operational specification. *IBM Systems Journal*, 42(3):428–445, 2003.
18. C. Ouyang, E. Verbeek, W. M. P. van der Aalst, S. Breutel, M. Dumas, and A. H. M. ter Hofstede. Formal semantics and analysis of control flow in WS-BPEL. *Science of Computer Programming*, 67(2-3):162–198, 2007.
19. W. M. P. van der Aalst and K. M. van Hee. *Workflow Management: Models, Methods, and Systems*. MIT Press, 2002.