

Patterns-based Evaluation of Open Source BPM Systems: The Cases of jBPM, OpenWFE, and Enhydra Shark

Petia Wohed¹, Nick Russell², Arthur H.M. ter Hofstede³,
Birger Andersson¹, and Wil M.P. van der Aalst^{2,3}

¹ Stockholm University/KTH, Stockholm, Sweden
{petia,ba}@dsv.su.se

² Eindhoven University of Technology, Eindhoven, The Netherlands
{n.c.russell,w.m.p.v.d.aalst}@tue.nl

³ Queensland University of Technology, Brisbane, Australia
a.terhofstede@qut.edu.au

Abstract. In keeping with the proliferation of free software development initiatives and the increased interest in the business process management domain, many open source workflow and business process management systems have appeared during the last few years and are now under active development. This upsurge gives rise to two important questions: What are the capabilities of these systems? and How do they compare to each other and to their closed source counterparts? In other words: *What is the state-of-the-art in the area?*. To gain an insight into these questions, we have conducted an in-depth analysis of three of the major open source workflow management systems – jBPM, OpenWFE and Enhydra Shark, the results of which are reported here. This analysis is based on the workflow patterns framework and provides a continuation of the series of evaluations performed using the same framework on closed source systems, business process modelling languages, and web-service composition standards. The results from evaluations of the three open source systems are compared with each other and also with the results from evaluations of three representative closed source systems: Staffware, WebSphere MQ and Oracle BPEL PM. The overall conclusion is that open source systems are targeted more toward developers rather than business analysts. They generally provide less support for the patterns than closed source systems, particularly with respect to the resource perspective, i.e. the various ways in which work is distributed amongst business users and managed through to completion.

Keywords: jBPM, OpenWFE, Enhydra Shark, Workflow Management Systems (WFMS), Business Process Management (BPM), open source software, workflow patterns.

1 Introduction

Workflow management systems (WFMSs) are software systems which support the execution of business processes in an organization. They take care of the distribution of work (and associated data) to the right people at the right time. They are configured based on (often graphical) process models. Similar to database management systems, WFMSs are also generic systems, but instead of supporting the storage and retrieval of data, they focus on the distribution of work to resources and provide facilities for monitoring its progress and managing it to completion.

Early work in the workflow domain during the 1970's focussed on the development of office information systems which aimed to streamline and automate the functioning of the modern office environment, however the immature technology of the time (e.g. limited computing accessibility, simplistic user interface capabilities, minimal network availability, etc), hampered their success. The technological advances made over the past two decades (particularly with respect to the widespread availability of the PC, graphical user interfaces and networking technology), have resulted in greater numbers of WFMSs coming to market, and an increased interest in business process management more generally.

Mindful of these developments and also of the increasing popularity of open source software, it is opportune to investigate the capabilities of open source WFMSs. A number of such systems have been released over the past six years (for specific details see [13, 24], each of which contains more than 30 offerings) and this raises the questions of how mature these systems are and whether they pose any serious competition to commercial systems in the area.

In order to answer these questions, in this paper *we aim to provide an insight into the state-of-the-art in open source workflow management systems*. In order to do so, we examine three of the major systems in the area [11], namely jBPM, OpenWFE, and Enhydra Shark. As an evaluation framework we use the workflow patterns [43], a collection of patterns which provide a taxonomy of generic, recurring concepts and constructs relevant to the workflow and business process management area. By systematically investigating the ability of each of the selected systems to support the patterns, we are able to build up a comprehensive picture of the scope and the suitability of the system for business process modelling and management.

The workflow patterns are particularly suited to evaluating the capabilities of WFMSs. In several earlier research initiatives, they have been successfully used to evaluate a number of proprietary systems [33] (e.g. Staffware⁴, IBM's WebSphere MQ, Pallas Athena's FLOWer, COSA, Sun's iPlanet, Oracle BPEL Process Manager), as well as to examine business process modelling languages (e.g. UML (Unified Modelling Language) Activity Diagrams [41, 40], BPMN (Business Process Modelling Notation) [42] and XPD (XML Process Definition Language) [1]), and standards for web-services composition (e.g. BPEL4WS (Business Process Execution Language for Web Services) [7], BPML (Business Process Modelling Language) [2]). The work presented here can be seen as a continuation of these evaluations. Using the same framework not only facilitates a thorough analysis of each of the systems based on a comprehensive and widely utilized evaluation framework, but also allows for broad comparison of the results with those obtained from earlier assessments.

Our choice of the Workflow Patterns as an analysis framework is also motivated by the following factors. First, it is a well accepted framework which has been widely used both for the selection of WFMSs (e.g. by UWV, the organization responsible for handling health and unemployment insurance in the Netherlands, the Dutch Justice Department, ArboNed, etc.) as well as for self-evaluations by individual vendors of the capabilities of their products (e.g. COSA, FLOWer, Staffware, IBM, etc.). Second, the workflow patterns have had proven impact on the workflow technology industry. They have triggered extensions to existing workflow management systems (e.g. FLOWer 3.0, Staffware Process Suite, Pectra Technology Inc's tool) and inspired the development of new ones (e.g. Zebra, Alphaflow)⁵. Third, the framework operates at a sufficiently detailed level such that it provides a comprehensive basis for assessing the capabilities of business process modelling languages and systems.

The evaluation discussed in this paper was undertaken as follows. Solutions for each of the patterns were sought in the three tools under consideration. Where successfully identified, they were implemented and tested. The initial results were summarized and each of the system vendors/developers was invited to provide feedback on their accuracy. On the basis of these responses, a final set of results were agreed upon. These results are presented here.

The main contributions of this work are:

- Identification of the capabilities and limitations of the evaluated systems, thus providing input for their future development.
- Explanation of how each of the patterns is captured in the systems, thus providing a source book of reusable knowledge for process designers that encounter the patterns when using one of these systems.
- Independent and unbiased evaluation of the capabilities of jBPM, OpenWFE and Enhydra Shark.
- Comparison of the capabilities of the open source systems with three representative closed source systems: Staffware, WebSphere MQ and Oracle BPEL PM.

The paper is structured as follows. In Section 2 we motivate our choice of open source systems. Section 3 gives an overview of the related work in the area and discusses our choice of the workflow patterns as an analysis framework in more detail. Section 4 briefly introduces the offerings under consideration: jBPM, OpenWFE, and Enhydra Shark. Sections 5, 6, and 7 present the results of the control-flow, data and resource patterns evaluations. Section 8 summarizes the results and compares them with the results of evaluations of three closed source systems: Staffware, WebSphere MQ and Oracle BPEL PM.

2 Choice of open source systems

There are a wide variety of open source workflow and business process management (BPM) offerings available. Based on a comprehensive literature survey, jBPM [16], OpenWFE [27] and Enhydra Shark [8] were selected for our evaluation. All of these systems are under active development, they are regularly updated and their web sites provide an indication of a sizeable and energetic user community. jBPM is part of JBOSS (a commercial company), Enhydra Shark supports XPD v1.0 (the standard proposed by the WFMC) and OpenWFE is an active project on SourceForge, labelled as "Production/Stable" with more than 100,000 downloads. Note that given the authors' close involvement with the YAWL Initiative, we refrain from commenting on the open source WFMS YAWL [3] in this paper.

To introduce readers to the area, we have identified the major open source WFMSs in Figure 1. The figure includes open source systems distributed through the well known source code repositories: SourceForge, RubyForge and ObjectWebForge⁶. It shows projects that were selected based on the following criteria: (i) contain the keyword "BPM" or "workflow" in their project name or description (which gave more than 460 projects, of which more than 430 were

⁴ Staffware was purchased by Tibco and has been incorporated into their iProcess Suite offering.

⁵ For specific details please, refer to www.workflowpatterns.com.

⁶ The Tigris.org, Bounty Source, BerliOS, Java Forge, and GNU Savannah code repositories were also searched but no results fulfilled the selection criteria for inclusion in the figure.

on SourceForge); (ii) describe projects on business process management and workflow system (i.e. projects focusing on document management, content management, groupware, wikis and not purely on business process and workflow management were discarded); (iii) are active projects, i.e. projects containing at least one file upload during the last two years (since June 2006); and (iv) have more than 20,000 downloads registered. Since download statistics are only publicly available for the projects on SourceForge, for projects fulfilling criteria (i)-(iii) which were distributed through RubyForge or ObjectWebForge, these statistics were requested from their developers. In Figure 1, the selected projects/open source offerings are shown in terms of thousands of downloads (the bars coloured in dark grey identify statistics that were retrieved from developers). The results are organized chronologically with respect to the projects' registration date in the relevant code repository.

Furthermore, three of the systems in the figure namely Taverna, ebills, and eLawOffice (marked with an asterisk and illustrated with a light-shaded bar), are domain specific systems (i.e. Taverna is a scientific WFMS, ebills is a system for supporting electronic payment workflows, and eLawOffice is tailored to support workflows in law firms) and therefore not particularly interesting for our study which aims to analyse generic business process management systems (i.e., WFMSs). As can be seen, the WFMSs in Figure 1 can be divided into three groups: systems with more than 100,000 downloads, i.e. jBPM (which clearly stands out), openWFE and Enhydra Shark; followed by Bonita and YAWL with approximately 70,000 and 60,000 downloads respectively; and systems with around 20,000-30,000 downloads. For the purposes of our evaluation of open source offerings, we restrict our examination to the three leading tools in this area: jBPM, OpenWFE and Enhydra Shark.

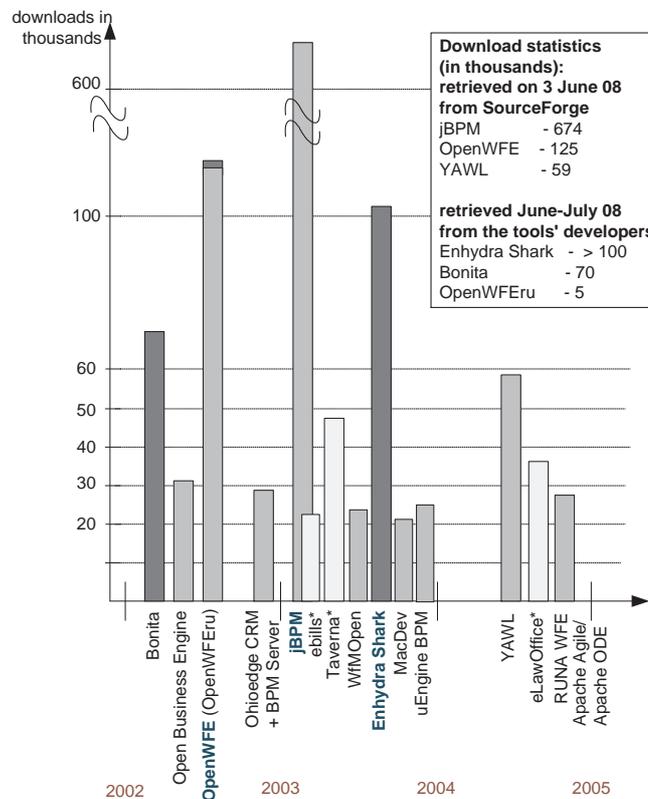


Figure 1 Open source workflow systems - history: the x-axis denotes the year in which the first version of the system was made available (e.g., via SourceForge) and the y-axis shows the total number of downloads

3 Related work

To our knowledge, there are three benchmarking frameworks in existence that have been used for evaluating and comparing modelling languages: the workflow patterns framework [43], Bunge, Wand and Weber's (BWW) framework [38], and the Semiotic Quality Framework [21].

The development of the *workflow patterns framework* started as a bottom-up, comparative analysis of workflow management software. Undertaken during 2000 and 2001, this analysis included the evaluation of 15 WFMSs, with a

specific focus on their underlying modelling languages [4]. The goal was to examine the expressive power of the underlying languages and identify similarities and differences between them. During this research effort, 20 control-flow patterns were experientially derived [4]. These patterns identify recurring, generic behaviours with respect to a specific aspect of process control. They range from simple characteristics to complex scenarios and provide a taxonomy for the control-flow perspective of processes.

Gradually the framework was extended to address the data and the resource perspectives [12] of workflows. While the control-flow perspective focuses on issues associated with execution ordering and behaviour of activities within a process, the data perspective deals with the issues of data representation and utilization in a WFMS. The resource perspective further complements the approach by describing the various ways in which work is distributed amongst and managed by the resources associated with a business process. In 2005, a set of 43 resource patterns [34] and a set of 40 data patterns [36] were added to the framework. Subsequent to this work, the initial set of 20 control-flow patterns was revisited and extended to 43 patterns in order to further clarify the operation of the original patterns and deal with omissions from the earlier set. In order to avoid any potential for ambiguous interpretation, the revised set of control-flow patterns was also formalized using coloured Petri nets [35]⁷.

The *Bunge, Wand and Weber (BWW) framework* [38], refers to Wand and Weber's tailoring and application of Mario Bunge's upper level ontology [5] within the information systems domain. According to the ontology classification in [10], the BWW framework would be classified as a top-level ontology for information systems. In contrast the workflow patterns framework would be classified as a *fine-grained* top-level ontology for *workflow and business process management systems*. The BWW framework was initially used for analysis and comparison of conceptual modelling languages, among them UML [26]. Later it was used for the analysis of process modelling languages e.g. Petri Nets [32], EPCs [32], IDEF3 [32], and BPMN [31, 32], as well as for candidate standards in the BPM domain (e.g. ebXML Business Process Specification Schema (BPSS), the Business Process Modelling Language (BPML), the Business Process Execution Language for Web Services (BPEL4WS) and the Web Service Choreography Interface (WSCI) [9]). However, as the BWW framework lacks conceptual structures for evaluating features essential to the workflow domain, such as the various types of splits and joins, iteration and cancellation constructs, and forms of concurrency restriction, we are critical of some of the conclusions drawn from these evaluations (e.g. "ebXML BPSS was a 'clear winner' in terms of ontological expressiveness" [9], p.288). Furthermore, the application of the framework as a theoretical foundation for conceptual modelling has also been a subject for criticism [44].

The *Semiotic Quality Framework* [21] for evaluating the quality of conceptual modelling languages is an extension of Lindland, Sindre and Solvberg's framework [22] for evaluating the quality of conceptual models. It is positioned by Wahl and Sindre to be at an even higher level of abstraction than the BWW framework (see [37], p.10). It outlines five aspects for language quality: domain appropriateness; participant's language knowledge appropriateness; knowledge externalizability appropriateness; comprehensibility appropriateness; and technical actor interpretation appropriateness. It includes aspects such as "easiness to learn and use a language" and relies on the workflow patterns for evaluation of the suitability aspects of a process modelling languages. It has been used for the evaluation of BPMN [37] and UML [20]. It is important to note that only a limited number of languages have been analysed using this framework (i.e. UML and BPMN) and, to our knowledge, it has not been used for the evaluation of actual WFMSs - a point which also applies to the BWW framework.

Due to the level of generality of the BWW framework and the Semiotic Quality Framework and the criticism associated with the BWW framework, the choice of evaluation framework for our study falls naturally on the workflow patterns framework. It should be noted that the workflow patterns framework deals with language and system suitability (i.e. modelling capabilities) for business process modelling and management. It does not cover aspects such as support, user friendliness, training effort, analysis, simulation and reporting capabilities etc. While important, we consider that evaluations focusing on these aspects require more subjective forms of assessment which fall outside of the scope of our work. They do, however, suggest suitable directions for future research in the area.

4 The Evaluated Systems

In this section we give a brief overview of the evaluated systems: JBoss jBPM version 3.1.4; OpenWFE version 1.7.3; and Enhydra Shark version 2.0. All three systems are written in Java and are extensible, enabling additional language features to be incorporated where required.

⁷ There also is a formalization of the initial 20 control-flow patterns in π -calculus [30]. However the purpose of this work was mainly to show that the patterns could be captured in π -calculus rather than providing a formal semantics for them or as an aid to identify new patterns.

4.1 JBoss jBPM

JBoss jBPM [16] is a WFMS and release version 3.1.4 was examined during the patterns-based evaluation. The system comprises the following main components:

- A workflow engine called JBoss JBPM core component (also referred to as the core process engine) which takes care of the execution of process instances.
- A process definition tool called JBoss jBPM Graphical Process Designer (GPD). It is a plugin to Eclipse, which provides support for defining processes in jPDL both in a graphical format and in XML format. jPDL (JBPM Process Definition Language) is the process language utilized by the system⁸.
- JBoss jBPM console web application which has two functions. It is a web based workflow client whereby users can initiate and execute processes. It is also an administration and monitoring tool where users can observe and intervene in currently executing process instances.
- JBoss jBPM identity component (which as stated in Chapter 11.11 in [14], will be developed in the future), which will take care of the definition of organizational information, such as users, groups and roles to which different tasks can be assigned. Currently the definition of this information is done through standard SQL insert statements directly into the workflow database.

A process definition in jPDL consists of a number of *nodes* and *transitions* between these nodes. In addition, the concept of *Swimlanes* [14], borrowed from UML, is used as a mechanism for assigning tasks to users. As mentioned earlier, users are defined in the workflow database using SQL insert statements. The most important nodes in jPDL are the *Start* and *End* nodes providing explicit denotation of the beginning and the completion of a process, the *State* node for modelling wait states, the *Task Node* for capturing work packages, and the routing nodes such as the *Fork*, *Join* and *Decision* for performing parallel execution, synchronization, and choice. A Task Node consists of one or more *Task* definitions. Tasks are atomic pieces of work either conducted by a human resource or an application. Graphically only Task Nodes are shown and not the tasks that they consist of. Furthermore, two constructs for hierarchical process decomposition are offered, namely the *Super State* and *Process State*. A Super State is used to organize a group of related nodes into a unit. A Process State is used for triggering the invocation of subprocesses which are defined independently from and outside of the process they are invoked from.

There is also a special node type called *Node* through which facilities are provided for the introduction of custom code to define desired new behaviour. This construct can be used both for invoking external applications, as well as for adding new control-flow behaviour to a model. It is important to note that a Node is responsible for the further propagation of task instances passing through it.

An alternative way to introduce customized behaviour is through *Actions*. Actions are defined through Java classes and specify new functionality, which is hidden in transitions or nodes and meant to implement behaviour important from a technical, but not from a business perspective (in the specification, an update of a database is mentioned as an example of an action). In contrast to Nodes, Actions can not influence control-flow and hence they are not responsible for any token propagation in a model. They are executed when the transition to which they are attached is executed, or the node for which they are defined reaches the state specified as the trigger (which is one of the following: node-enter, node-leave, before-signal and after-signal). As Actions and Nodes require coding, they are not considered in this evaluation. (Note that our evaluation focuses on suitability, i.e., the question is not whether something is possible but whether it is directly supported without additional programming in a conventional/general language like Java. Without this restriction any programming language would support any workflow pattern.)

Listing 1 shows an example of a jPDL specification represented in XML and Figure 2 shows its graphical representation. In this example there is one swimlane 'ernie' which contains a user 'ernie' who is responsible for the whole process where first tasks *A* and *B* of task1-node execute in parallel and when both are completed task *C* of task2-node is started. Note that the task nodes are shown graphically without their tasks, which in this particular example hides the parallelism present between tasks *A* and *B* in the flow. In the remainder of this paper closing tags in XML examples will be omitted where it improves readability. In addition, definitions of Swimlanes, Start state, and End state (i.e. lines 1-10 from Listing 1) will also be omitted so that examples can focus on those lines of code relevant to understanding the pattern solutions presented.

⁸ Processes can also be defined as Java objects or as records in the jBPM database (as stated in [14], Chapter 7 Persistence). However, these alternative notations, are outside the scope of our evaluation.

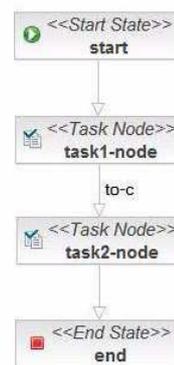
Listing 1

```

1 <process-definition
2   xmlns="urn:jboss.org:jpd1-3.1" name="Synchronization">
3   <swimlane name="ernie">
4     <assignment expression="user(ernie)"></assignment>
5   </swimlane>
6   <start-state name="start">
7     <task name="start-task" swimlane="ernie"></task>
8     <transition name="" to="task1-node"></transition>
9   </start-state>
10  <end-state name="end"></end-state>
11  <task-node name="task1-node">
12    <task name="A" swimlane="ernie"></task>
13    <task name="B" swimlane="ernie"></task>
14    <transition name="to-c" to="task2-node"></transition>
15  </task-node>
16  <task-node name="task2-node">
17    <task name="C" swimlane="ernie"></task>
18    <transition name="" to="end"></transition>
19  </task-node>
20 </process-definition>

```

Figure 2 JBoss JBPM - graphical representation of the model in Listing 1



For this evaluation, version 3.1.4 of the tool was analysed (i.e. jbpm-starters-kit-3.1.4 package from 25 January 2007). During late June a transition to an update of the tool was considered (i.e. jbpm-jpdl-suite-3.2.1 package from 25 June 2007). However, the state of the Workflow Client provided with this installation was considered immature. For instance every user was able to execute every task irrespective of whether the process definition assigned specific users to the various tasks. For this reason we decided to continue with the earlier but more stable version 3.1.4 of the tool.

4.2 OpenWFE

OpenWFE [27] is a WFMS of which version 1.7.3 was considered for the purposes of the patterns-based evaluation. The system comprises the following main components:

- An Engine routing work items to users in accordance with a process definition. Users are either human participants or automated agents. Automated agents are implemented in a language like Java or Python and defined as users in a workflow through a corresponding APRE, i.e. Automatic Participant Runtime Environment (Java classes are defined through Java APRE and Python script is defined through Pyya APRE).
- A web-based workflow design environment called Droflo which can be used to graphically define process models, which are then translated to OpenWFE's process definition language in XML format. OpenWFE uses its own process definition language. Due to the lack of documentation of the graphical notation, in this evaluation the XML format was used exclusively.
- A Webclient (also called webappserver) providing the user-interface to the Engine and to the worklists to which the engine distributes the work.
- A web-interface for resource management, called UMAN (for User MANagement). Due to lack of documentation on UMAN, all resource management during our evaluation was performed through direct modification of the files storing user information.

The development of OpenWFE has now migrated from Java to Ruby with the new implementation referred to as OpenWFEru. However, during our testing of OpenWFEru, at the end of June 2007, a client for OpenWFEru similar to that for OpenWFE was not yet available. As the lack of such a client would have considerably slowed our work, and the client is considered as an integral part of a WFMS, we decided to continue the work with the older Java version of the tool, i.e. OpenWFE 1.7.3.

Through the Webclient users access their worklists. The client allows them to view work items offered to them (and possibly to other users), choose a work item in order to execute it, and to report the completion of a work item. During process execution users can see a graphical representation of the process with the task they are working on highlighted. This graphical presentation is automatically generated by the system.

When logged in, users, also referred to as principals, can switch between the Launch view, where they can initiate processes for which they have the proper authorization and the Stores view, where they can access their work lists. Users can have certain privileges (read, write, delegate) for stores, which themselves may contain a number of participants. At design time work is assigned to participants which means that at runtime this work can be performed by any user that has write-access to the store of which this participant is a member. E.g. a work item <participant ref="role-alpha"> can be executed by user Alice if Alice has write-access to a store of which role-alpha is a member. The engine thus routes work according to the process logic, user privileges and participant membership of stores. Listing 2 shows an XML excerpt illustrating how this membership relation is defined. Listing 3 shows the definition of user Alice with the privileges assigned to her to read, write and delegate work items from Store.alpha.

The main routing constructs used in OpenWFE are *sequence* for sequential execution, *concurrency* for parallel execution, various dedicated constructs for *loops* (e.g. while, until), *iterator*, and the *concurrent iterator* to deal with multiple concurrently executing instances of a task. Attributes for these constructs allow further refinement of their behaviour. The language is block structured which imposes some restrictions on the specification of arbitrary forms of synchronization and iteration (for a general discussion on this issue see e.g. [19]).

Listing 2 Fragment of worklist-configuration.xml

```

1 <application-configuration
2   name="openwfe-worklist">
3   <service
4     name="Store.alpha"
5     class="[...].SimpleWorkItemStore">
6     <param>
7       <param-name>participants
8     </param-name>
9     <param-value>role-alpha
10    </param-value>
11  </param>
12  ...
13 </service>

```

Listing 3 Fragment from passwd.xml

```

1 <passwd> ...
2 <principal name="alice"
3   class="[...].BasicPrincipal"
4   password="+99-124+86+60">
5   <grant name="store.alpha"/>
6 </principal>
7 ...
8 <grant name="store.alpha"
9   codebase="file:./jars/r/*">
10  <permission
11    name="Store.alpha"
12    class="[...]StorePermission"
13    rights="read, write, delegate"/>
14 </grant>
15 ...

```

Listing 4 A process definition

```

1 <process-definition
2   name="cfpl Sequence"
3   revision="1.0">
4 <description>1 Sequence</description>
5 <sequence>
6   <participant ref="role-alpha"
7     description="design"/>
8   <participant ref="role-bravo"
9     description="implement"/>
10  <participant ref="role-charly"
11    description="deploy"/>
12 </sequence>
13 </process-definition>

```

A process can be decomposed into *subprocesses*. A subprocess can be defined as part of a main process or it can be specified independently. In the latter case it can be invoked by multiple processes.

In Listing 4 the specification of a simple process containing a sequence of activities is shown. In a process specification work items are assigned to participants (`<participant ref="role-alpha"/>`) and routing constructs, such as `<sequence>` in the example, are used to control the order in which these items can be performed. The description attribute (e.g. `description="design"`) is optional and it aids in understanding what function the task performs. Unfortunately, the description cannot be seen by users viewing their work list, it can only be seen when a work item is started. This makes it difficult for users to gain a detailed understanding of the work being offered to them.

4.3 Enhydra Shark

Enhydra Shark is a Java workflow engine offering from Together Teamlösungen and ObjectWeb [8]. The system contains the following components:

- a workflow engine called Shark;
- a Process Definition Tool called Together Workflow Editor (TWE) or JaWE;
- a Management Console, TWS Admin, which is a workflow client as well as an administration and monitoring tool;
- a set of ToolAgents (e.g. MailToolAgent, SOAPToolAgent, JavaScriptToolAgent, RuntimeApplicationToolAgent), which provide interfaces for different software, e.g. Javascript Interpreter.

The main products considered in this review were the engine distributed through Shark TWS-community-2.0-1 (TWS - Together Workflow Server) and the editor JaWE TWE-community-2.2-1 (TWE - Together Workflow Editor). The evaluation was done using the TWS Admin client.

Together Workflow Editor (TWE) [23] uses XPD 1.0 [39] as a process language. The main unit of organization in XPD 1.0 is a *package*. A package contains one or several *process* definitions and can contain global *data* entities common for process definitions within the package. A process consists of *activities* and *transitions* between them. Activities are assigned to *participants*. An activity may be performed *manually* or *automatically* and the participant may be either a human resource or an *application*. When an activity is completed according to the specification, the

thread of control advances to the next activity. *Conditions* that constrain the transitions can be specified. When multiple threads enter/exit an activity, a *transition restriction* is specified for the activity, indicating whether it is an AND or an XOR *Join/Split*.

A process can be hierarchically structured through the use of *Block activities*. A block activity contains an *Activity set* which is a set of subprocesses following the same syntax and semantic rules as for processes. Furthermore, an activity can invoke a process which is specified independently. There is a special kind of activity called a *Routing activity*, which corresponds to a “dummy” activity used only for routing the control-flow within a process. In addition to the XML process definition, Enhydra Shark also supports a graphical notation that was developed in-house.

5 Control-flow Patterns Evaluation

The workflow control-flow patterns (WCPs) focus on the ordering of activities within a process. They are divided into the following groups: *Basic Control-flow Patterns* capturing elementary aspects of control-flow; *Advanced Branching and Synchronization Patterns* describing more complex branching and synchronization scenarios; *Iteration Patterns* describing various ways in which task repetition may be specified; *Termination Patterns* addressing the issue of when the execution of a workflow is considered to be finished; *Multiple Instances Patterns* delineating situations where there are multiple threads of execution in a workflow which relate to the same activity; *State-based Patterns* reflecting situations which are most easily modelled in workflow languages that support the notion of state; *Cancellation Patterns* categorizing the various cancellation scenarios that may be relevant for a workflow specification; and *Trigger Patterns* identifying the different triggering mechanisms appearing in a process context.

The evaluation results for each of these groups are discussed below. Every subsection concludes with a table summarizing the results for the three offerings using a three point scale where ‘+’ corresponds to direct support for a pattern, ‘+/-’ indicates that direct support is present but in a limited way, and ‘-’ shows that the pattern is not directly supported. The evaluation criteria for assigning ‘+’ or ‘+/-’ are specified in [35].

5.1 Basic control-flow patterns

All three offerings provide direct support for the basic control-flow patterns. Listings 5-7 illustrate how these patterns can be captured in OpenWFE. The constructs <sequence>, <concurrency> and <case> implement the required behaviours for the Sequence, Parallel Split and Exclusive Choice patterns respectively. As the language is block structured, the closing tags </concurrency> and </case> provide the semantics for the Synchronization and Simple Merge patterns.

Enhydra Shark captures Sequences through a *transition* between activities (see Figure 3 and Listing 8). A Parallel Split (Listing 9) is captured through an *activity* node with the *transition restriction* <Split Type=“AND”>. Synchronization (Listing 10) is captured in the same way, but the transition restriction specified is <Join Type=“AND”>. Exclusive Choice and Simple Merge are captured as Parallel Split and Synchronization, respectively, but with XOR (instead of AND) specified as their transition restriction. For the Exclusive Choice *transition conditions* are also specified on the transitions leaving the XOR activity node. Graphically, transitions with transition conditions specified on them are shown with a bold arrow. It is possible to specify “Otherwise” as a transition condition for a transition (which will evaluate to true if the preceding conditions do not): this is graphically denoted by an orange coloured arrow.

Listings 11-16 show how the basic patterns can be captured in jBPM. In Figure 4, the corresponding graphical representations for jBPM are shown. A Sequence is represented by a *transition* between two *task-nodes*, the Parallel Split pattern is captured using a *fork node* with a number of outgoing transitions (Listing 12). An alternative solution is to specify the *tasks* to be run in parallel within the same task-node (Listing 13)⁹. Synchronization is captured through a *join node* with a number of incoming transitions (Listing 14).

⁹ Note that the graphical representation for the second solution for the Parallel Split pattern is similar to the graphical representation of the Sequence pattern, which may lead to confusion for some users.

Listing 5 WCP-1 Sequence in OpenWFE

```

1 <sequence>
2   <participant ref="role-alpha"
3     description="design"/>
4   <participant ref="role-bravo"
5     description="implement"/>
6   <participant ref="role-charly"
7     description="deploy"/>
8 </sequence>

```

Listing 6 WCP-2 Parallel Split and WCP-3 Synchronization in OpenWFE

```

1 <concurrency>
2   <participant ref="role-alpha"
3     description="implement"/>
4   <participant ref="role-bravo"
5     description="test"/>
6   <participant ref="role-charly"
7     description="document"/>
8 </concurrency>

```

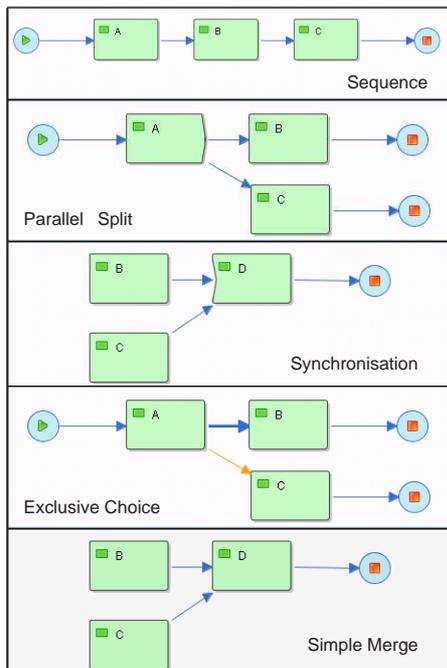
Listing 7 WCP-4 Exclusive Choice and WCP-5 Simple Merge in OpenWFE

```

1 <set field="alternative" type="string"
2   value="select task: 1, 2 or 3" />
3 <participant ref="role-alpha"
4   description="select task"/>
5 <case>
6   <equals field-value="alternative"
7     other-value="1"/>
8   <participant ref="role-alpha"
9     description="deploy"/>
10  <equals field-value="alternative"
11    other-value="2"/>
12  <participant ref="role-bravo"
13    description="correct"/>
14  <equals field-value="alternative"
15    other-value="3"/>
16  <participant ref="refactor"
17    description="task3"/>
18 </case>

```

Figure 3 Basic Patterns in Enhydra Shark



Listing 8 WCP-1 Sequence in Enhydra Shark

```

1 <Activity Id="A" Name="A">..</Activity>
2 <Activity Id="B" Name="B">..</Activity>
3 ...
4 <Transition From="A" Id="Tr1".. To="B"/>

```

Listing 9 WCP-2 Parallel Split in Enhydra Shark

```

1 <Activity Id="A" Name="A"> ..
2   <TransitionRestriction>
3     <Split Type="AND"> ..
4       <TransitionRef Id="Tr1"/>
5       <TransitionRef Id="Tr2"/>
6     ..
7 </Activity>
8 <Activity Id="B" Name="B"></Activity>
9 <Activity Id="C" Name="C"></Activity>
10 ...
11 <Transition From="A" Id="Tr1".. To="B"/>
12 <Transition From="A" Id="Tr2".. To="C"/>

```

Listing 10 WCP-3 Synchronization in Enhydra Shark

```

1 <Activity Id="B" Name="B">..</Activity>
2 <Activity Id="C" Name="C">..</Activity>
3 <Activity Id="D" Name="D"> ..
4   <TransitionRestriction>
5     <Join Type="AND"/>
6   ..
7 <Transition From="B" Id="Tr3" .. To="D"/>
8 <Transition From="C" Id="Tr4" .. To="D"/>

```

The Exclusive Choice pattern is captured by specifying *conditions* on the transitions leaving a *decision node* (Listing 15). The first expression that evaluates to true results in the continuation of the flow of control in the transition with which the expression is associated. The Simple Merge pattern is captured by a task-node which joins the transitions emanating from the set of task-nodes from which the thread(s) of control should be merged (Listing 16). Table 1 summarizes the support of the evaluated tools for the basic control-flow patterns.

Listing 11 WCP-1 Sequence in jBPM

```

1 <task-node name="task1">
2   <task name="task1a" swimlane="ernie">
3     <transition name="" to="task2">
4   </task-node>
5 <task-node name="task2">
6   <task name="task2a" swimlane="ernie">
7     <transition name="" to="end1">
8 </task-node>

```

Listing 12 WCP-2 Parallel Split in jBPM

```

1 <fork name="AND-split">
2   <transition name="trA" to="task1A">
3   <transition name="trB" to="task1B">
4 </fork>
5 <task-node name="task1A">
6   <task name="task1A" swimlane="ernie">
7   <transition name="" to="end1">
8 </task-node>
9 <task-node name="task1B">
10  <task name="task1B" swimlane="ernie">
11  <transition name="" to="end1">
12 </task-node>

```

Listing 13 WCP-2 Parallel Split in jBPM, alt2

```

1 <task-node name="task1-node">
2   <task name="task1A" swimlane="ernie">
3   <task name="task1B" swimlane="ernie">
4   <transition name="" to="end1">
5 </task-node>

```

Listing 14 WCP-3 Synchronization in jBPM

```

1 <fork name="fork1">
2   <transition name="tr1" to="task1A-node">
3   <transition name="tr2" to="task1B-node">
4 </fork>
5 <task-node name="task1A-node">
6   <task name="task1A" swimlane="ernie">
7     <transition name="" to="join1">
8 </task-node>
9 <task-node name="task1B-node">
10  <task name="task1B" swimlane="ernie">
11  <transition name="" to="join1">
12 </task-node>
13 <join name="join1">
14   <transition name="" to="task2-node">
15 </join>

```

Listing 15 WCP-4 Exclusive Choice in jBPM

```

1 <decision name="xor">
2   <transition name="urgent" to="task-b">
3     <condition>scenario==1</condition>
4   </transition>
5   <transition name="dont care" to="task-c">
6     <condition>scenario==2</condition>
7   </transition>
8   <transition name="forget about it" to="task-d">
9     <condition>scenario!=1
10    and scenario!=2</condition>
11 </transition>
12 </decision>

```

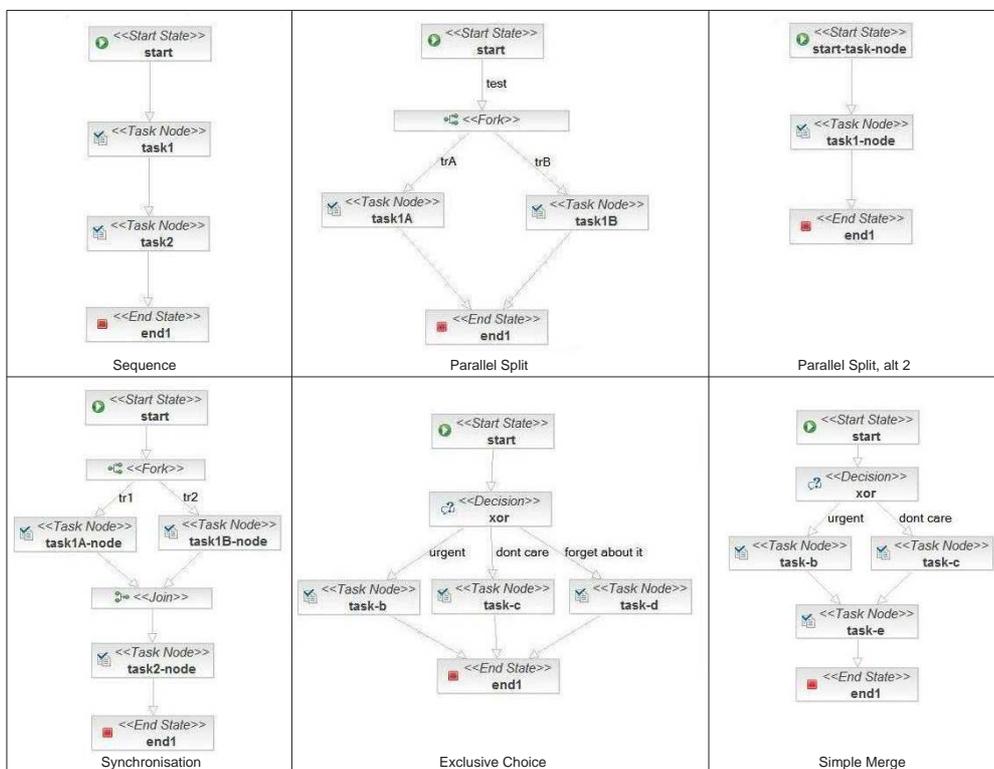


Figure 4 Basic Patterns in jBPM

Basic Control-flow Patterns	1	2	3
1. Sequence	+	+	+
2. Parallel Split	+	+	+
3. Synchronization	+	+	+
4. Exclusive Choice	+	+	+
5. Simple Merge	+	+	+

Table 1. Support for the basic control-flow patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

Listing 16 WCP-5 Simple Merge in jBPM

```

1 <decision name="xor">
2   <transition name="urgent" to="task-b">
3     <condition expression=.../>
4   </transition>
5   <transition name="dont care" to="task-c">
6     <condition expression=.../>
7   </transition>
8 </decision>
9 <task-node name="task-b">
10   <task name="task-b" swimlane="ernie">
11     <transition name="" to="task-e">
12   </task-node>
13 <task-node name="task-c">
14   <task name="task-c" swimlane="ernie">
15     <transition name="" to="task-e">
16   </task-node>
17 <task-node name="task-e">
18   <task name="task-e" swimlane="ernie">
19     <transition name="" to="end1">
20 </task-node>

```

Listing 17 WCP-8 Multi Merge in jBPM

```

1 <fork name="fork1">
2   <transition name="tr1" to="task1A-node">
3   <transition name="tr2" to="task1B-node">
4 </fork>
5 <task-node name="task1A-node">
6   <task name="task1A" swimlane="ernie">
7     <transition name="" to="task2-node">
8 </task-node>
9 <task-node name="task1B-node">
10   <task name="task1B" swimlane="ernie">
11     <transition name="" to="task2-node">
12 </task-node>
13 <task-node name="task2-node">
14   <task name="task2" swimlane="ernie">
15     <transition name="" to="end1">
16 </task-node>

```

5.2 Advanced branching and synchronization patterns

The support for the patterns in this group is limited. jBPM provides direct support for the Multiple Merge and the Generalized AND-Join patterns and partial support for the Thread Merge and the Thread Split patterns (it is rated a ‘-’ for the Structured Synchronizing Merge as the context condition of this pattern requires the presence of a corresponding Multiple Choice which is not supported by jBPM). If, in the bottom right-hand model of Figure 4, the decision node is replaced by a fork node, the model will capture the semantics of the Multiple Merge pattern. The jPDL definition for this pattern is given in Listing 17. Note that as a consequence of the use of the fork, it is guaranteed that the last node, i.e. “task2-node” will be executed twice, which distinguishes this pattern from the Simple Merge pattern. The Thread Merge and Thread Split can be achieved programmatically in Java by defining two new Nodes implementing the corresponding behaviour (according to the evaluation criteria of these patterns, the use of programmatic extension for achieving their behaviour rates as partial support).

Support for the Generalized AND-join in jBPM is demonstrated by the left-hand model in Figure 5. The join node in this model is initiated twice, correctly capturing the Generalized AND-join semantics. In this model task-C acts as a Multiple Merge and is executed twice (as both incoming branches are enabled separately) resulting in two concurrent threads of execution at the second fork node (each of which results in a thread of execution in both of the outgoing branches from the Fork node). The following join node can only execute when it has received a thread of control on each of its incoming branches. As there are two threads of control flowing down each of the incoming branches to the join node (as a consequence of the fork node firing twice), it is possible that more than one thread of control will be received on one input branch to the join before a thread of control is received on the other and the join is able to fire. Where this occurs, the join keeps track of the number of execution threads received on each incoming branch and it only executes when one has been received on each incoming branch. Any additional control threads are retained for future firings, hence the join will ultimately execute twice and two execution threads will be passed onto task-F.

The right-hand model in Figure 5 illustrates the scenario with which jBPM’s support for the Local Synchronizing Merge was evaluated. The join node in this model was intended to act as a local synchronising merge by keeping track of whether additional execution threads can still be received on input branches that have not yet been enabled. For example, after execution of the nodes taskA, taskB and task-d, the join must wait as it is possible that task-b will be executed. If task-a is chosen rather than task-b, the join can continue and trigger the execution of task-bd, if not, it must await the completion of task-b and then start the execution of task-bd. However, while syntactically correct, this process stalls when only one of the tasks task-b and task-d is selected and subsequently executed as the Join node continues to wait for the completion of both.

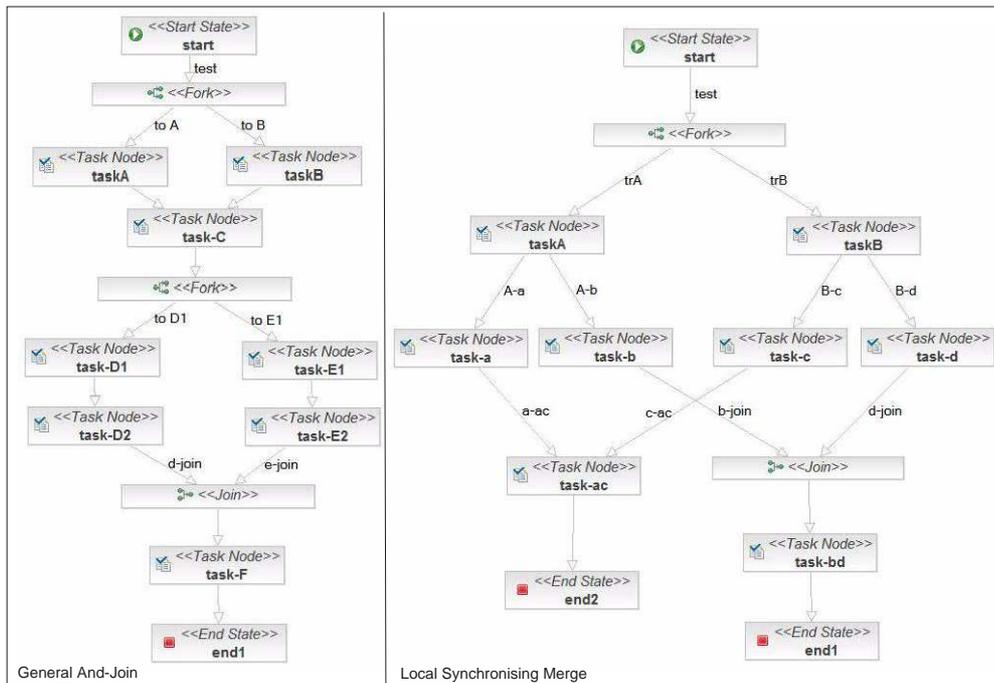


Figure 5 Patterns WCP-33 and WCP-37 in jBPM

Listing 18 WCP-9 Structured Discriminator in OpenWFE

```

1 <concurrency
2   count="1"
3   remaining="forget">
4     <participant ref="role-alpha"
5       description="mail contact"/>
6     <participant ref="role-bravo"
7       description="email contact"/>
8     <participant ref="role-charly"
9       description="sms contact"/>
10 </concurrency>

```

Listing 19 WCP-32 Cancelling Partial Join in OpenWFE

```

1 <concurrency
2   count="2"
3   remaining="cancel">
4     <participant ref="role-alpha"
5       description="mail contact"/>
6     <participant ref="role-bravo"
7       description="email contact"/>
8     <participant ref="role-charly"
9       description="sms contact"/>
10 </concurrency>

```

From the advanced branching and synchronization patterns, OpenWFE provides direct support for the Structured and Cancelling variants of the Discriminator and Partial Join patterns. These behaviours are achieved through the use of the <concurrency> construct with attribute settings that allow for the specification of the number of incoming branches that must receive an execution thread in order for the join construct to execute (i.e. the *count* attribute) and also for the specification of what should happen to the other (incoming) branches once the join has executed (the *remaining* attribute which can be set to ‘forget’ or to ‘cancel’ remaining execution threads received on these incoming branches). Listing 18 illustrates OpenWFE’s support for the Structured Discriminator pattern. At runtime three work items are started in parallel, the first one to complete (indicated by the *count*="1" setting) triggers the remainder of the flow (not shown) and the remaining work items are allowed to complete although this does not influence the flow in any way (due to the *remaining*="forget" setting). By changing the settings for the attributes ‘count’ and ‘remaining’ the other forms of Discriminator and Partial Join patterns are also easily captured. Listing 19 shows an example for the Cancelling Partial Join.

OpenWFE also provides partial support for the Multiple Choice pattern. This is achieved through a collection of if-then clauses executing concurrently, where several of the if-conditions may evaluate to true. Due to the structured nature of workflow specifications in OpenWFE there is no support for patterns such as the Multiple Merge and the Generalized AND-join. As with jBPM, the Structured Synchronizing merge is not considered to be supported in OpenWFE due to the absence of a dedicated Multiple Choice construct. The Local Synchronizing Merge is partially supported as the pattern can only be captured in structured scenarios¹⁰, while general scenarios such as the one presented for jBPM (in the right-hand model in Figure 5) cannot be handled in OpenWFE.

In this pattern category, Enhydra Shark only provides support for the Multiple Choice pattern, which similar to OpenWFE, is captured by defining overlapping conditions on transitions emanating from an AND-split activity (see

¹⁰ An example of such a scenario is presented in the animation for the pattern available at www.workflowpatterns.com/patterns/control/new/wcp37_animation.php

Figure 6 and Listing 20¹¹). Enhydra Shark also proposes the model in Figure 7a, as a solution for the Multiple Merge pattern. In this model activity A is a multi-choice activity after which one or several of the activities B, C, and/or D can be started. At completion each one of these activities asynchronously invokes an instance of the subprocess S1. However, this is not considered to fully capture the semantics for the Multiple Choice pattern, as general scenarios like the one shown in Figure 7b can not be captured in this manner. The same reasoning applies for OpenWFE (which also provides the possibility to invoke a subflow to run asynchronously to the main process).

Table 2 summarizes the results for the advanced branching and synchronization patterns.

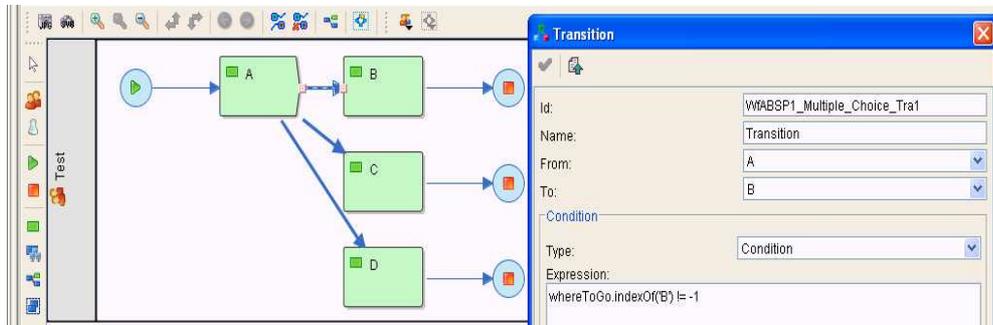


Figure 6 WCP-6 Multiple Choice in Enhydra Shark

Listing 20 WCP-6 Multiple Choice in Enhydra Shark

```

1 <Activity Id="A" Name="A">
2   <TransitionRestrictions>
3     <TransitionRestriction>
4       <Split Type="AND">
5         <TransitionRefs>
6           <TransitionRef Id="Tr1"/>
7           <TransitionRef Id="Tr2"/>
8           <TransitionRef Id="Tr3"/>
9         </TransitionRefs>
10        </Split>
11      </TransitionRestriction>
12    </TransitionRestrictions>
13    <ExtendedAttributes>
14      <ExtendedAttribute Name="VariableToProcess_
15        UPDATE" Value="X"/>
16    </ExtendedAttributes>
17  </Activity>
18 <Activity Id="B" Name="B"></Activity>
19 <Activity Id="C" Name="C"></Activity>
20 <Activity Id="D" Name="D"></Activity>
21 ...
22 <Transition From="A" Id="Tr1" .. To="B">
23   <Condition Type="CONDITION">X.indexOf('B') != -1
24 </Condition>
25 </Transition>
26 <Transition From="A" Id="Tr2" .. To="C">
27   <Condition Type="CONDITION">X.indexOf('C') != -1
28 </Condition>
29 </Transition>
30 <Transition From="A" Id="Tr3" .. To="D">
31   <Condition Type="CONDITION">X.indexOf('D') != -1
32 </Condition>
33 </Transition>
34 ...

```

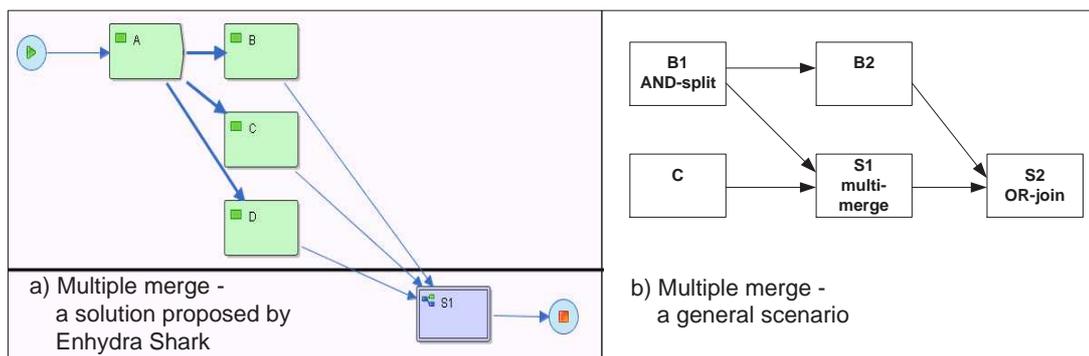


Figure 7 WCP-7 Multiple Merge proposed in an example contained in Enhydra Shark's default installation

¹¹ For space reasons the transition Id “WfABSP1_Multiple_Choice_Tra1” in Figure 6 is abbreviated to “Tr1” in Listing 20 and the variable “whereToGo” is replaced with “X”.

Advanced Branching and Synchronization	1	2	3
6. Multiple Choice	-	+/-	+
7. Structured Synchronizing Merge	-	-	-
8. Multiple Merge	+	-	-
9. Structured Discriminator	-	+	-
28. Blocking Discriminator	-	-	-
29. Cancelling Discriminator	-	+	-
30. Structured Partial Join	-	+	-
31. Blocking Partial Join	-	-	-
32. Cancelling Partial Join	-	+	-
33. Generalized AND-Join	+	-	-
37. Local Synchronizing Merge	-	+/-	-
38. General Synchronizing Merge	-	-	-
41. Thread Merge	+/-	-	-
42. Thread Split	+/-	-	-

Table 2. Support for the advanced branching and synchronization patterns in 1-JBoss jBPM 3.1.4, 2-OpenWFE 1.7.3, and 3-Enhydra Shark 2.0

5.3 Iteration patterns

In terms of the iteration patterns, jBPM directly supports the Arbitrary Cycles pattern, but does not have dedicated constructs for structured loops nor does it support recursion¹². jBPM's support for the Arbitrary Cycles pattern is shown in Listing 21 and Figure 8. The arbitrary cycle shown has two entry points (one through task-node B, the other through task-node E) and two exit points (one after task-node E and one after task-node G).

OpenWFE provides direct support for the Structured Loop and Recursion patterns. Both the while-do and repeat-until variants of the Structured Loop pattern are supported by the <loop> construct. A <loop> construct contains a *loop condition* specified through the use of <while> or <until> clauses and an *iteration body* enclosed within a <sequence> clause. The loop condition (i.e. the <while> and <until> parts) can appear before or after the iteration body (i.e. the <sequence> specification) which determines whether the loop is of type while-do or repeat-until (see listings 22-25 for illustrations of these solutions).

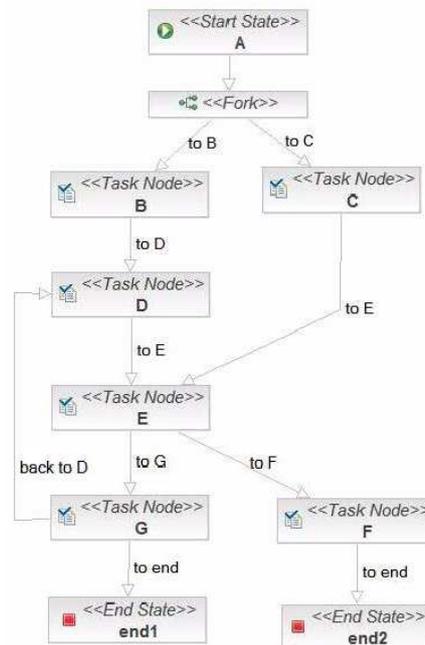
Listing 21 WCP-10 Arbitrary Cycles in jBPM

```

1 <start-state name="A">
2   <task name="A" swimlane="ernie">
3     <transition name="" to="fork1">
4   </start-state>
5 <end-state name="end1">
6 <task-node name="B">
7   <task name="B" swimlane="ernie">
8     <transition name="to D" to="D">
9   </task-node>
10 <task-node name="D">
11   <task name="D" swimlane="ernie">
12     <transition name="to E" to="E">
13   </task-node>
14 <task-node name="E">
15   <task name="E" swimlane="ernie">
16     <transition name="to G" to="G">
17     <transition name="to F" to="F">
18   </task-node>
19 <task-node name="G">
20   <task name="G" swimlane="ernie">
21     <transition name="to end" to="end1">
22     <transition name="back to D" to="D">
23   </task-node>
24 <fork name="fork1">
25   <transition name="to B" to="B">
26   <transition name="to C" to="C">
27 </fork>
28 <task-node name="C">
29   <task name="C" swimlane="ernie">
30     <transition name="to E" to="E">
31   </task-node>
32 <end-state name="end2">
33 <task-node name="F">
34   <task name="F" swimlane="ernie">
35     <transition name="to end" to="end2">
36 </task-node>

```

Figure 8 WCP-10 Arbitrary Cycles in jBPM, graphical representation of the model in Listing 21



¹² In the jBPM specification [15], Chapter 9.6 Superstates, it is indicated that superstates can be nested recursively. However, in the graphical process designer (GPD) it is currently not possible to specify the components of a superstate, hence we do not consider recursion to be supported.

Listing 22 WCP-21 while-do loop in OpenWFE, alt1

```
1 <loop>
2   <until> condition</until>
3   <sequence>iteration body</sequence>
4 </loop>
```

Listing 23 WCP-21 while-do loop in OpenWFE, alt2

```
1 <loop>
2   <while> condition </while>
3   <sequence>iteration body</sequence>
4 </loop>
```

Listing 24 WCP-21 repeat-until loop in OpenWFE, alt1

```
1 <loop>
2   <sequence>iteration body</sequence>
3   <until>condition</until>
4 </loop>
```

Listing 25 WCP-21 repeat-until loop in OpenWFE, alt2

```
1 <loop>
2   <sequence> iteration body</sequence>
3   <while> condition </while>
4 </loop>
```

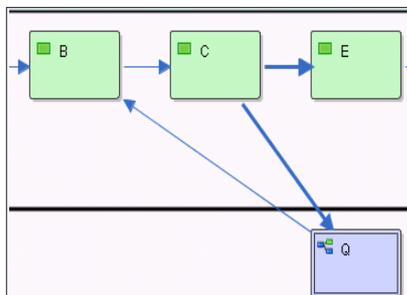
Furthermore, the `<iterator>` construct can be used to assemble loops. The `<iterator>` implements the behaviour commonly associated with for-loops from classical programming languages, where an index variable takes values from a specified range on a sequential basis and the body of the loop is executed once for each possible value. Listing 27 shows an example of an `<iterator>` construct in OpenWFE¹³.

The cursor concept in OpenWFE allows for the thread of control in a process instance to be moved back and forth. Consider for example Listing 28, at runtime the three activities can be executed in sequence, but at any point in time, users with sufficient privileges to execute these work items can advance the thread of control or move it back one or more steps in the flow and can even choose to exit the flow before the sequence is finished. The commands “back”, “skip”, “break” (or “cancel”), and “rewind” (moving back to the beginning of the cursor), are used for these purposes. These commands can either be specified at runtime in a field `__cursor_command__` (which appears by default in the user interface during the execution of a task specified within a cursor), or at design time as exemplified in Listing 29. It is possible to limit the use of these commands by specifying that they are disallowed in the definition of a given cursor, e.g. `<cursor disallow="rewind, break">`. Consequently, the cursor command can also be used for specifying arbitrary cycles with multiple exit points, thus fulfilling the criteria for full support of this pattern. Note, however that the cursor construct can not be used for capturing a scenario such as the one illustrated for jBPM in Figure 8 where the cycle can not only be exited at different points but also additional threads of control can enter the cycle at various places. In the new Ruby version of the tool, a *jump* construct is introduced [29] which seems to act as a goto operator and thus enables the modelling of arbitrary cycles with multiple entry points.

In terms of the iteration patterns, Enhydra Shark, similar to jBPM, directly supports the Arbitrary Cycles pattern (allowing cycles both with multiple entry and exit points), but does not have dedicated constructs for structured loops. Enhydra Shark also supports recursion. A workflow can invoke itself by means of subflow. For an example of this, see Figure 9 and Listing 26.

Table 3 summarizes the results for the support of the iteration patterns.

Figure 9 WCP-22 Recursion in Enhydra Shark, graphical representation of the model in Listing 26



Listing 26 WCP-22 Recursion in Enhydra Shark

```
1 <WorkflowProcess .. Id="cycle" ..>
2   ..
3   <Activity Id="B" Name="B">
4     .. <Join Type="XOR"/> .. </Activity>
5   <Activity Id="C" Name="C">
6     .. <Split Type="XOR"> .. </Split>
7     .. </Activity>
8   <Activity Id="E" Name="E"></Activity>
9   <Activity Id="Q">
10    <Implementation>
11      <SubFlow Id="cycle"/>
12    </Implementation>
13  </Activity>
14  ..
15 </WorkflowProcess>
```

¹³ With some small modifications, this solution is taken from Chapter 7 Workflow patterns, Pattern 14, in the OpenWFE specification [25].

Listing 27 WCP-21 For-loop in OpenWFE

```

1 <iterator
2   on-value="alpha,bravo,charly"
3   to-field="rolename">
4   <sequence>
5     <participant ref="role-#{f:rolename}"/>
6   </sequence>
7 </iterator>

```

Listing 28 The cursor construct in OpenWFE

```

1 <cursor>
2   <participant ref="role-alpha" description=
3     "drug experiment on animals"/>
4   <participant ref="role-bravo" description=
5     "small human sample experiment"/>
6   <participant ref="role-charly" description=
7     "large human sample experiment"/>
8 </cursor>

```

Listing 29 The cursor construct in OpenWFE, design time established behaviour

```

1 <cursor>
2   <set field="mark" type="integer"/>
3   <participant ref="role-alpha"
4     description="exam"/>
5   <if>
6     <greater-than field-value="mark"
7       other-value="5"/>
8     <break />
9   </if>
10  <participant ref="role-bravo"
11    description="study"/>
12  <participant ref="role-alpha"
13    description="complement"/>
14  <if>
15    <lesser-than field-value="mark"
16      other-value="3"/>
17    <back step="3"/>
18  </if>
19 </cursor>

```

Iteration Patterns	1	2	3
10. Arbitrary Cycles	+	+	+
21. Structured Loop	-	+	-
22. Recursion	-	+	+

Table 3. Support for the iteration patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

5.4 Termination patterns

All three offerings support the Implicit Termination pattern, i.e. a process instance completes when there are no remaining work items left to be done (i.e. there is no need to specify a single explicit end node). While jBPM notifies the user executing the last work item about the completion of a process, Enhydra Shark and OpenWFE do not. Currently, OpenWFE only identifies the completion of process instances in the log of execution history. In Enhydra Shark, a completed task is displayed in the Process List as having a “closed.completed” status. Table 4 summarizes these evaluation results.

Termination Patterns	1	2	3
11. Implicit Termination	+	+	+
43. Explicit Termination	-	-	-

Table 4. Support for the termination patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

5.5 Multiple instances patterns

There is virtually no support for Multiple Instance (MI) tasks in jBPM or Enhydra Shark. Direct support is only provided for the Multiple Instances without Synchronization pattern as new instances of tasks can be spawned off through the use of a loop. This is illustrated for jBPM with the model shown in Figure 10 where as part of the loop based on the *NewInstance?* task-node, new instances of the task-node MI-task are spawned off. In every iteration of such a loop an instance of a separate subprocess is invoked to run asynchronously with the main-process.

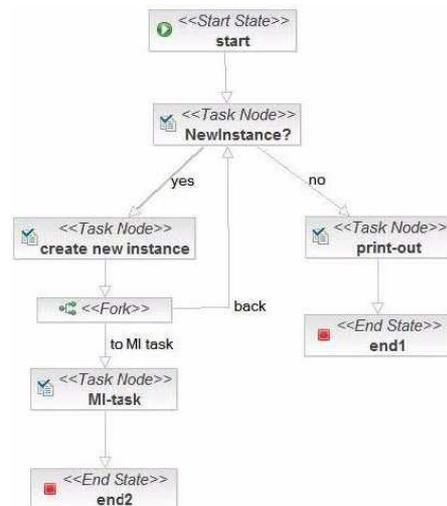
Similarly to jBPM and Enhydra Shark, the Multiple Instances without Synchronization pattern is achieved in OpenWFE by launching new instances as subprocesses through a loop. Setting the *forget* attribute to “true” ensures that synchronization over the created instances will not take place. This solution is illustrated in Listing 30. In this example three instances of the subprocess “run test” are spawned-off through the loop. The number ‘3’ of instances to be created is decided at design time and is hard coded in the loop condition.

In OpenWFE the Multiple Instances with a priori Design Time Knowledge and Multiple Instances with a priori Runtime Knowledge patterns are captured through the `<concurrent-iterator>` construct through which a number of instances (determined during design or runtime) are spawned off. At completion, the instances are synchronized. Listing 31 shows how three instances of a task “test” can be created that run concurrently. Each of these instances will have a unique number in its *index* field (1, 2 or 3 in this example) which is one of the items in the “on-value” list. At completion, the instances will be synchronized (through `</concurrent-iterator>`). Listing 32 shows how multiple

instances can be created on the basis of the values in a list, called “iteration_list” in this example. These values will be retrieved during the execution of the task “distribute tests to subjects”, i.e. at runtime but before the commencement of the multiple instances task “run test”. In “run test” a work item is created for every item in the iteration-list. The variable “user” acts as an index variable on the basis of which each work item is created. At completion the created work items are synchronized (through </concurrent-iterator>) and the task “summarize results” is enabled. Note that the work items of the multiple instance task “run test” are distributed to the participants included in the “iteration_list” (which is passed to distinct instances through the variable “user”, as illustrated in lines 8-9 in the listing). This is an example of the Deferred Allocation pattern from the Resource Perspective which will be discussed in more detail in Section 7.

The Multiple Instances without a Priori Runtime Knowledge pattern (i.e. WCP-15), which is one of the more advanced multiple instances patterns, is not supported in any of the offerings. The distinguishing feature of this pattern is that new instances of the multiple instances task can be created after the task has commenced. The concurrent-iterator construct in OpenWFE can only be used when the number of instances to be created is known before the commencement of the multiple instance task. The repeated invocation of a subprocess within a loop with the attribute forget=“true”¹⁴ does not synchronize the created instances and hence does not capture the full semantics of the pattern.

Figure 10 WCP-12 MI without Synchronization in jBPM



Listing 31 WCP-13 MI with a priori Design Time Knowledge in OpenWFE

```

1 <concurrent-iterator on-value="1, 2, 3"
2   to-field="index">
3   <participant ref="role-charly"
4     description="test"/>
5 </concurrent-iterator>
6 <participant ref="role-alpha"
7   description="summarize results"/>
  
```

Listing 32 WCP-14 MI with a priori Runtime Knowledge in OpenWFE

```

1 <sequence>
2   <set field="iteration_list"
3     type="StringMapAttribute"
4     value="type users: alpha,bravo.."/>
5   <participant ref="role-alpha"
6     description="distribute tests to subjects"/>
7   <concurrent-iterator on-field-value="iteration_list"
8     to-field="user">
9     <participant ref="role-${f:user}"
10      description="run test"/>
11 </concurrent-iterator>
12 <participant ref="role-alpha"
13   description="summarize results"/>
14 </sequence>
  
```

Listing 30 WCP-12 MI without Synchronization in OpenWFE

```

1 <sequence> <participant ref="role-a"
2   description="prepare first set of tests"/>
3 <loop>
4   <sequence>
5     <subprocess ref="run test" forget="true"/>
6     <inc variable="idx0" />
7   </sequence>
8 <while>
9   <lesser-than variable-value="idx0"
10    other-value="3" />
11 </while>
12 </loop>
13 <participant ref="role-bravo"
14   description="report on the commencement
15   of the tests"/>
16 </sequence>
17 <process-definition name="run test">
18   <sequence>
19     <participant ref="role-charly"
20       description="test"/>
21   </sequence>
22 </process-definition>
  
```

Listing 33 WCP-34 Static Partial Join for MI in OpenWFE

```

1 <sequence>
2   <set field="nr_of_inst"
3     value="1,2,3" />
4   <set field="nr_to_be_completed"
5     type="Integer" value="" />
6   <participant ref="role-alpha"
7     description="set up test"/>
8   <concurrent-iterator
9     on-value="{f:nr_of_inst}"
10    to-field="index"
11    count="{f:nr_to_be_completed}"
12    remaining="forget"
13    merge="last"
14    merge-type="mix">
15     <participant ref="role-bravo"
16       description="run test"/>
17 </concurrent-iterator>
18 <participant ref="role-alpha"
19   description="summarize test results"/>
20 </sequence>
  
```

Listing 33 shows how the *count* attribute of the <concurrent-iterator> construct can be used to achieve a partial join (note that in this listing a variable called “nr_to_be_completed” is used at runtime to determine the number of threads that need to be synchronized and merged). The Static Partial Join for Multiple Instances pattern is supported by setting the *remaining* attribute to “forget” so that any threads remaining after the completion of the multiple instance task are not cancelled. The Cancelling Partial Join for Multiple Instances pattern is implemented by setting

¹⁴ This is the solution proposed by OpenWFE in [25] (Chapter 7, Workflow Patterns, Pattern 15)

the *remaining* attribute to “cancel”. There is no dedicated support for cancelling (i.e. WCP-26) or force-completing multiple instance tasks (WCP-27) in OpenWFE, nor for dynamically creating new instances of a multiple instance task after it has started and for disabling that ability during execution (WCP-36). These results are summarized in Table 5.

Multiple Instances (MI) Patterns	1	2	3
12. MI without Synchronization	+	+	+
13. MI with a priori Design Time Knowledge	-	+	-
14. MI with a priori Runtime Knowledge	-	+	-
15. MI without a priori Runtime Knowledge	-	-	-
27. Complete MI Activity	-	-	-
34. Static Partial Join for MI	-	+	-
35. Static Cancellation Partial Join for MI	-	+	-
36. Dynamic Partial Join for MI	-	-	-

Table 5. Support for the multiple instances patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

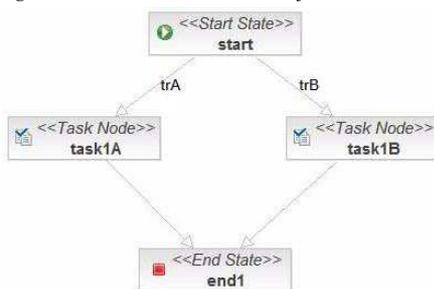
5.6 State-based patterns

Of the State-based patterns, jBPM only provides support for the Deferred Choice pattern. This is illustrated in Figure 11 and in Listing 34. Surprisingly, this pattern is not supported in OpenWFE. According to the self-evaluation of OpenWFE ([25], Chapter 7 Workflow Patterns, WP16) the code in Listing 18 with the *remaining* attribute set to “cancel” is proposed as a solution for the deferred choice. However, as discussed earlier, this is a solution for the Cancelling Discriminator and not the Deferred Choice pattern as all branches are started and the first to complete cancels the others whereas in a deferred choice the work items associated with the first of the tasks in each of the alternate branches associated with the deferred choice are offered and when the first task of one of these branches is chosen the others are withdrawn.

The `<interleaved>` construct in OpenWFE provides direct support for Interleaved Routing, as illustrated in Listing 35. At runtime work items corresponding to the three tasks are executed in arbitrary order and not concurrently. This construct however provides only partial support for Interleaved Parallel Routing as sequences of tasks (that need to be preserved) cannot be interrupted and have to be completed first before other activities (or other sequences of activities) can take their turn. In terms of Listing 36 this means that a work item corresponding to the “take subsidiary subject” task cannot be executed in between work items of the tasks “take methodology unit” and “write thesis”. The Critical Section pattern is not supported in OpenWFE as the sub-threads constituting the critical sections would have to be placed in one interleaved block thus limiting the use of the pattern in unstructured flows. Also the Milestone pattern is not supported in OpenWFE.

Enhydra Shark does not provide any support for State-based patterns. The results for this group of patterns are summarized in Table 6.

Figure 11 WCP-16 Deferred Choice in jBPM



Listing 35 WCP-40 Interleaved Routing in OpenWFE

```

1 <interleaved>
2 <participant ref="role-alpha"
3   description="take exam A"/>
4 <participant ref="role-alpha"
5   description="take exam B"/>
6 <participant ref="role-alpha"
7   description="take exam C"/>
8 </interleaved>
  
```

Listing 34 WCP-16 Deferred Choice in jPDL

```

1 <start-state name="start">
2   <task name="start-task" swimlane="ernie">
3     <transition name="trA" to="task1A">
4       <transition name="trB" to="task1B">
5     </start-state>
6 <end-state name="end1">
7 <task-node name="task1A">
8   <task name="task1A" swimlane="ernie">
9     <transition name="" to="end1">
10 </task-node>
11 <task-node name="task1B">
12   <task name="task1B" swimlane="ernie">
13     <transition name="" to="end1">
14 </task-node>
  
```

Listing 36 WCP-17 Interleaved Parallel Routing in OpenWFE

```

1 <interleaved>
2 <participant ref="role-alpha"
3   description="take subsidiary subject"/>
4 <sequence>
5   <participant ref="role-alpha"
6     description="take methodology unit"/>
7   <participant ref="role-alpha"
8     description="write thesis"/>
9 </sequence>
10 </interleaved>
  
```

State-Based Patterns	1	2	3
16. Deferred Choice	+	-	-
39. Critical Section	-	-	-
17. Interleaved Parallel Routing	-	+/-	-
40. Interleaved Routing	-	+	-
18. Milestone	-	-	-

Table 6. Support for the state-based patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

5.7 Cancellation patterns

Of the cancellation patterns only Cancel Case is supported in OpenWFE. Cancellation is specified at design time with the expression `<cancel-process/>`. However, any subprocesses invoked and running when the `<cancel-process/>` construct is reached do not get cancelled. Therefore, this pattern is only considered to be partially supported. In practice, cancellation of an activity or a case can be done by a workflow administrator through the command-line control program, but as this form of cancellation is an error-handling measure rather than a normal process activity, it is not considered to constitute support for the cancellation patterns.

In jBPM a task or a process can be forcibly ended at runtime by a user (who need not necessarily be an administrator). A forcibly ended task remains in the execution log, while a forcibly ended process is removed from the execution log. Therefore the Cancel Activity pattern is considered to be supported, while the Cancel Case pattern is not considered to be supported.

Enhydra Shark supports case cancellation at runtime. A case can be cancelled through the administration console by selecting “Terminate” for a specific instance. The case’s status is logged as “closed.terminated”. None of the other cancellation patterns are supported in OpenWFE, jBPM or Enhydra Shark.

The support for the cancellation patterns is summarized in Table 7.

Cancellation Patterns	1	2	3
19. Cancel Activity	+	-	-
20. Cancel Case	-	+/-	+
25. Cancel Region	-	-	-
26. Cancel MI Activity	-	-	-

Table 7. Support for the cancellation patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

5.8 Trigger patterns

The trigger patterns distinguish between transient and persistent triggers. Transient Triggers (WCP-23) are lost if not acted upon immediately (which implies that they need to be waited on at a particular place in a process) whereas Persistent Triggers (WCP-24) are retained by the process until they can be acted upon. Persistent Triggers are either buffered until the task which they target is enabled or they are able to trigger tasks that are not dependent on the completion of any preceding tasks.

Every task, i.e., `<participant>` definition in OpenWFE can be timed out (see Listing 38, where an instance of task1 will be skipped if not performed within 10 minutes). The default setting for timeout is one week. It is defined through the engine variable `..time_out..` and can be changed either by directly modifying its value in the configuration file or by (re)setting it from a process. To prevent a task from being timed out, its `timeout` attribute has to be set to “no”. While the timeout construct is convenient for defining back-up schemes that prevent process instances from stalling midway through their execution, they are intended more as a management tool than a real trigger mechanism.

With the `<sleep>` construct, the execution of a task can be delayed (see Listing 39, which is taken from the OpenWFE specification [28]). The time for the delay is either relative (`sleep for`) or absolute (`sleep until`). The expiration of the time interval defined for the `<sleep>` construct triggers the continuation of the flow. When the delay specified is relative, every trigger (time expiration) is mapped to a specific process instance. When the delay is absolute (and defined using global engine variables) its expiration may result in several process instances being triggered. The definition of a `<sleep>` construct at a particular point in the process, ensures that the trigger is expected and acted upon immediately when received. Through the use of this construct OpenWFE can provide support for transient triggers.

Listing 37 The when construct in OpenWFE - not working

```

1 <concurrency>
2   <when>
3     <equals variable-value="/triggerA"
4       other-value="13" />
5     <participant ref="role-alpha"
6       description="taskA"/>
7   </when>
8   <sequence>
9     <participant ref="role-bravo"
10      description="task1"/>
11     <set variable="/triggerA" value="13"/>
12     <participant ref="role-bravo"
13      description="task2"/>
14   </sequence>
15 </concurrency>

```

Listing 38 Timeout in OpenWFE

```

1 <sequence>
2   <participant ref="role-bravo"
3     description="task1" timeout="10m"/>
4   <participant ref="role-bravo"
5     description="task2"/>
6 </sequence>

```

Listing 39 The sleep construct in OpenWFE

```

1 <participant ref="role-bravo" />
2 <sleep for="10m" />
3 <participant ref="role-alpha" />
4 <set variable="date"
5   value="%c:tadd('%c:now()', '3d')" />
6 <sleep until="$date" />
7 <participant ref="role-bravo" />

```

No support for persistent triggers was identified in OpenWFE. The persistent trigger behaviour cannot be simulated using the <when> construct within a <concurrency> flow, as exemplified in Listing 37. The <when> construct operates as an asynchronous <if>, which repeatedly evaluates the specified condition (with the frequency defined in the configuration files) and when true executes the associated action. The surrounding <concurrency> expression is used to keep the <when> expression active during the execution of the remaining part of the process (which in this example is defined as a sequence). In the example an internal process variable “/triggerA” is monitored through the when-condition and a change in its value effectively acts as a trigger. The time interval to the next when-condition re-evaluation effectively causes a delay in the delivery of the trigger. For capturing external triggers, engine variables (e.g. “//triggerA”) could be used. Note that this is potentially problematic as no distinction can be made between different process instances. However, the most serious difficulty with this solution is that, as a consequence of continually re-evaluating the when-construct, the process continues to execute even when all intended process activity has been completed.

Transient triggers are supported in jBPM through the notion of State. When a process instance reaches a state on its execution path, it waits for a signal in order to leave the state and continue execution. No support for persistent triggers was observed in jBPM.

Enhydra Shark indicates that there is some support for transient triggers through the specification of a Deadline on an activity. However, the deadline functionality did not work in our tests. Another typical example of the use of transient triggers is to capture the receipt of an email. A MailToolAgent implements the connectivity needed for sending and receiving e-mail. However, in the evaluated version of the tool, the configuration needed for invoking the MailToolAgent did not work. Due to the problems with the deadline functionality and MailToolAgent configuration Enhydra Shark is not considered to support transient triggers. Nor does it support persistent triggers.

Table 8 summarizes the results for the trigger patterns.

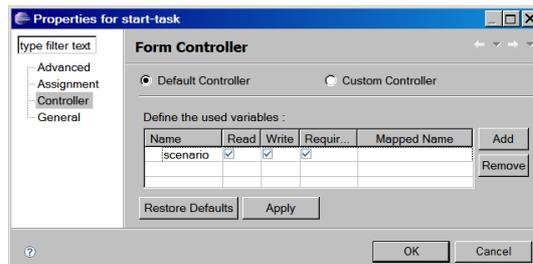
Trigger Patterns	1	2	3
23. Transient Trigger	+	+	-
24. Persistent Trigger	-	-	-

Table 8. Support for the trigger patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

6 Data Patterns Evaluation

The workflow data patterns (WDPs) describe the various means of data representation and utilization in the context of a workflow system. They are divided into five main groups: *data visibility patterns* characterizing the various ways in which data elements can be defined and utilized; *internal* and *external data interaction patterns* dealing with the various ways in which data elements can be passed between components within a process instance and also with the operating environment; *data transfer patterns* focusing on the way in which data elements are actually transferred between one process element and another; and *data routing patterns* capturing the various ways in which data elements can interact with other perspectives and influence the overall execution of the process. The evaluation results for each of these groups of patterns are discussed in turn in the following sections. However, first the approach that each of these three systems takes to the representation of data within workflow processes is introduced.

Figure 12 Graphical interface of controllers in jBPM



Listing 40 Variables in jBPM

```

1 <task name="start-task" swimlane="ernie">
2   <controller>
3     <variable name="scenario"
4       access="read,write,required">
5   </controller>
6 </task>

```

In jBPM, by default variables have a scope that encompasses the entire workflow process. They are created as part of the definition of an individual task via the *controller* construct. The intended function of a controller is to define the mapping between task and process variables. Currently only a one-to-one mapping is supported. (If a more complex mapping needs to be represented, a specific TaskControllerHandler capturing it would need to be implemented first.) Controllers not only deal with the mapping of process and task data, but they are also the only place where variables (both task and processes) can be defined. Figure 12 shows the graphical interface of a controller and the definition of a variable “scenario”. Listing 40 shows the corresponding XML representation. Mapped Name(s) define the task variable name, where it is different from the global variable naming (defined under Name). If specified, Mapped Name will be the name displayed to the user at runtime. The settings Read, Write and Required define whether the value in the task variable will be read from the corresponding process variable (and if the process variable has not yet been created, it is created), written back to it (if so, any previous value is overwritten) and whether the user input of a value for the variable is compulsory¹⁵.

In OpenWFE data handling is realized through variables and fields. *Fields* hold data values, which are populated and/or accessed by end users, while *variables* hold data values which are used for internal processing and are not visible to or accessible by end users. Data can be transferred between fields and variables at any time during process execution except during the execution of tasks. After being defined, a field is visible in all subsequent tasks unless its visibility is limited through the use of filters. In contrast, defined variables have a lifespan corresponding to that of the process instance. In the documentation [25], a distinction is made between local (subprocess) variables (which are denoted by the name of the variable) e.g. *varname*, process variables, e.g. */varname* (denoted by a slash preceding the name), and engine variables, e.g. *//varname* (which are denoted by two slashes preceding the name). Operations on data elements (other than user input) can not be performed inside a (atomic) task, instead data elements must be manipulated when they are passed between the tasks.

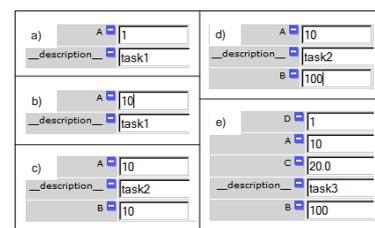
Listing 41 Variables and Fields in OpenWFE

```

1 <sequence>
2   <set field="A" type="integer" value="1" />
3   <set variable="/x" value="empty" />
4   <set variable="//x" value="{f:A}" />
5   <participant description="task1" ref="role-alpha" />
6   <set variable="/x" value="{f:A}" />
7   <set field="B" type="integer" value="{x}" />
8   <participant description="task2" ref="role-bravo" />
9   <set field="C"
10     value="{call:sum('/x','{f:A}')} " />
11   <set field="D" value="{//x}" />
12   <participant description="task3" ref="role-alpha" />
13 </sequence>

```

Figure 13 Fields interfaces in OpenWFE



Listing 41 illustrates variable and field definitions in a process and Figures 13a-e show how they appear in the user interface when the process is run. In Listing 41, field A and the variables /x and //x are defined prior to the commencement of ‘task1’. Field A is assigned the value ‘1’ as a default, /x is assigned the value ‘empty’ and //x is assigned the value of field A, (which is ‘1’). Figure 13a provides a screenshot of part of the user interface which indicates the values of data elements at the commencement of ‘task1’. Note that only field A (and not the variables) appear in this and that this field is writable by the user. Figure 13b shows a change in the value in A to 10 as a result of a change made by an end user. When ‘task1’ completes execution, the process continues to ‘task2’ but before it commences, the value of variable /x is assigned the value of field A, which is now 10 (see line 6 in the listing) and

¹⁵ During our test, variables with the setting Required (such as the variable scenario) were found to accept empty values, i.e. the Required setting did not work as expected.

field B is created with its default value set to that of variable x (see line 7). Note that variable x has not been defined. In practice, OpenWFE does not differentiate between subprocess and process scope (the moment a subprocess variable /x is defined, a corresponding process variable x is also created) which means that x is assigned the value of /x, hence field B is assigned the value '10'. At the commencement of 'task2' the values for both field A and B are as shown in Figure 13c. Both fields are writable by the end user and in Figure 13d the value of field 'B' is changed to 100 during the execution of the task. After the completion of 'task2', field C is defined (line 9 in the listing) and its initial value derived from the values of variable x and field A (line 10), which illustrates simple data manipulation. Finally, field D is defined and set to the value of the variable //x (line 11 in the listing), which is displayed at the commencement of 'task3' (see Figure 13e).

Enhydra Shark uses variables to store data relevant to a process. A variable is declared to be one of several types, e.g. Array or Basic. The default variable type is Basic with subtypes such as String or Integer (see Figure 14, left). New types of variables can be formed by either combining predefined types or restricting predefined ones. Also at design time variables may be given default values.

In Enhydra Shark a package aggregates one or several workflow processes. A variable may be defined at the process level or at the package level (see Figure 14, right hand side) using the Datafield tag. A datafield positioned at the process level defines a variable to be accessible by all instances of a particular process, whereas a datafield positioned at the package level defines a variable to be accessible by all process instances in the package. A consequence of this is that a variable defined at the package level may be used and reused by several processes. A package level variable, however, is not a global variable in the usual sense. When instantiating a process a copy of the variable is made accessible for the process instance. This copy resides entirely within the process instance and is not visible outside it. A package level variable cannot be used as a means for communicating between process instances.

At design time, the execution of an activity can be defined to involve reading or writing the value of a variable. For example, a transition condition from one activity to the next may be dependent on a specific value of a variable. A user executing the activity at run time will then be confronted with an input box to change the value and enable the transition. Variables can also be defined for manipulation during an automatic procedure. For instance, a variable may be incremented to keep count of iterations in a loop. The administration console allows the inspection of variables at any time during process execution.

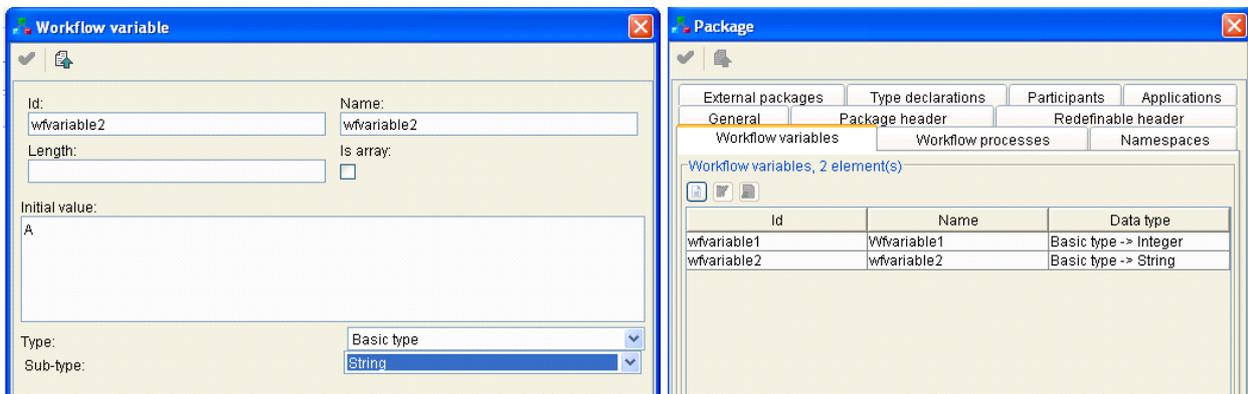


Figure 14 Variables in Enhydra Shark

6.1 Data visibility patterns

As previously indicated, in jBPM the most common approach to characterising data elements is via process variables hence, the Case Data pattern (WDP-5) is directly supported. However, owing to the fact that task variables are mapped to process variables on a one-to-one basis and the inherent difficulty of defining task variables that do not strictly correspond to existing process variables, the Task Data pattern (WDP-1) is not considered to be fully supported.

jBPM provides the notion of Super State which is intended to be used for grouping different nodes into logical units and hence serves as a mechanism for hierarchical structuring of processes. A corresponding symbol is also present in the graphical editor. However, it is not currently possible to graphically define the components which make up a Super State. In addition, the notion of Process State is introduced as a means of facilitating the invocation of subprocesses which are described in terms of distinct process fragments to the main process definition. However the use of this construct is problematic and leads to deployment failure. Hence, block structures are not implemented and there is no support for the Block Data pattern (WDP-2). The task-node construct can accommodate several tasks and

can potentially be considered as a partial implementation of the scope concept. However, no data elements can be defined for a task-node, thus the Scope Data pattern (WDP-3) is not supported.

As indicated in Section 5, jPDL lacks any notion of multiple instance tasks. If such tasks are required as part of a process model, a corresponding implementation of their behaviour would need to be completed first. jBPM provides the ability to embed such an extension within a node-type Node construct however the need for additional programmatic effort to achieve this outcome means that the Multiple Instance Data pattern (WDP-4) is not considered to be directly supported. Nor are any of the other visibility patterns supported.

In OpenWFE, all data elements are globally accessible throughout a process instance and, once defined, all fields are visible (unless restricted by the use of filters) to all subsequent tasks in a process instance. Hence OpenWFE directly supports the Case Data pattern (WDP-5) and consequently the Task Data pattern (WDP-1) is not supported. The presence of engine variables that can be accessed at any point in any process (e.g. *//varname*) confirms that Global Data (WDP-7) is supported. When a subprocess definition resides in the same file as the main process, a subprocess variable (e.g. */varname*) is also visible and modifiable from the main process. However local binding can be enforced through the use of a dot as prefix (i.e. *.varname*), hence OpenWFE also supports the Block Data pattern (WDP-2).

Filters can be used to specify read and write access restrictions on certain fields. Listing 42 shows an example of the definition of a filter called ‘student_view’ and its use with a specific task (i.e. task ‘register’). The setting `type="closed"` (on line 12) implies that any field not explicitly specified within the filter will not be visible when the filter is applied. The setting of attributes `add` and `delete` to “false” implies that additional fields can not be added or erased at runtime (which is not the case for the default setting). Four attributes are defined before the application of the filter to task ‘register’ and during the execution of the task only the fields ‘name’ and ‘address’ will be displayed (in addition to the `__description__` attribute of the task) while the attributes ‘result1’ and ‘result2’ will be filtered out. In particular the field ‘end_result’, which is included in the filter, is not displayed during the execution of ‘register’, as it has not been defined before the task.

Listing 42 Filters in OpenWFE

```

1 <sequence>
2   <set field="name" type="string" value="" />
3   <set field="address" type="string" value="" />
4   <set field="result1" type="integer" value="0" />
5   <set field="result2" type="integer" value="0" />
6   <participant description="register"
7     ref="role-student" filter="student_view" />
8   ...
9 </sequence>
10
11 <filter-definition name="student_view"
12   type="closed" add="false" remove="false">
13   <field regex="name" permissions="rw" />
14   <field regex="address" permissions="rw" />
15   <field regex="end_result" permissions="r" />
16   <field regex="__description__" permissions="r" />
17 </filter-definition>

```

Listing 43 A Web-Service Operation defined as a Participant in OpenWFE

```

1 <participant name="ws-weather">
2   ...
3   <param>
4     <param-name>wsdl</param-name>
5     <param-value>http://www.webservicex.net/
6       globalweather.asmx?WSDL</param-value>
7   </param>
8   <param>
9     <param-name>operationName</param-name>
10    <param-value>GetWeather</param-value>
11  </param>
12  <param>
13    <param-name>signature</param-name>
14    <param-value>city, country</param-value>
15  </param>
16  <param>
17    <param-name>return_to_field</param-name>
18    <param-value>weather</param-value>
19  </param>
20 </participant>

```

The use of filters for limiting access to fields qualifies as partial support for the Scope Data pattern (WDP-3) as the filter overwrites the default privileges for the data element and as such restricts rather than extends the access rights associated with it. As the number of data elements associated with a process instance tends to increase during its execution, the use of filters for restricting data visibility can become quite complex.

OpenWFE supports the Multiple Instance Data pattern (WDP-4) as is clear from Listing 32 in the previous section. At runtime each task instance that is created has its own value for the variable ‘user’. The Environment Data pattern (WDP-8) is also supported and external data in the operating environment can be accessed via web-services or external applications. In OpenWFE these are specified as participants in the participant-map configuration file. Listing 43 provides an example where the operation GetWeather from the Web-Service Global Weather is defined as an OpenWFE participant¹⁶, which means that it can, just like any other participant in the system, be invoked (i.e. have work distributed to it) from anywhere in a process specification. The *return_to_field* parameter specifies the location to which the environment data is transferred.

Enhydra Shark too offers process scoped variables, hence it supports the Case Data visibility pattern (WDP-5). It does not directly support Task Data (WDP-1), as data elements are defined at the package level or the process level. However, as it is possible to in a task define Javascript procedures, which can contain local data (such as loop counters), the support for Task Data is considered to be partial. (Similarly, through the Node construct in jBPM, task

¹⁶ This example is a part of the default OpenWFE 1.7.3 installation.

data can be defined and utilized. As this is achieved programmatically, the support for Task Data in jBPM is rated as partial.) Furthermore, Enhydra Shark supports Block Data visibility (WDP-2), through the notion of subprocess variables which are local to the corresponding subprocesses. In Enhydra Shark it is possible to define block activities. A block activity is an activity that encapsulates a set of activities into an embedded subflow. Data in a block activity can be shared with the containing process through parameter passing. That is, parameters of the block are mapped to variables in the calling process.

Through the notion of subprocess, Enhydra Shark supports shared implementation of tasks, hence it supports Multiple Instance Data (WDP-4). Global Data (WDP-7) is not supported. While variables are available at the package level (a package contains one or several workflows), these variables contain Case Data, i.e., their data values are not shared between several cases. Hence the use of these variables primarily serves as a means of reducing the number of variable definitions in a package.

Enhydra Shark supports Environment Data (WDP-8) through the use of Tool Agents. A Tool Agent provides an interface to a data source or an application in the operational environment. One tool agent, e.g. JavaScriptToolAgent (providing connectivity to the local Javascript interpreter), is available with the default installation in the evaluated product; hence this pattern is considered to be supported. Since coding in Java is needed, the support is considered as partial.

Table 9 summarizes the support of jBPM, OpenWFE and Enhydra Shark for the data visibility patterns.

Data Visibility	1	2	3
1. Task Data	+/-	-	+/-
2. Block Data	-	+	+
3. Scope Data	-	+/-	-
4. Multiple Instance Data	-	+	+
5. Case Data	+	+	+
6. Folder Data	-	-	-
7. Global Data	-	+	-
8. Environment Data	+/-	+	+/-

Table 9. Support for the data visibility patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

6.2 Data interaction patterns - internal

In jBPM, variables have process scope, hence the Data Interaction – Task to Task pattern (WDP-9) is supported through the global shared data approach. As there is no support for multiple instance tasks in jBPM and the fact that all attempts to implement task decomposition failed, we conclude that there is no support for data interaction patterns involving multiple instances, nor for data interaction to and from subworkflows. Achieving a similar effect through programmatic extensions is not considered to constitute partial support according to the evaluation criteria for these patterns. Support for the Data Interaction – Case-to-Case pattern (WDP-14) in jBPM is rated as partial as it can only be indirectly achieved through programmatic interaction with an external database, for example, by associating an action with a transition (as suggested in Section 9.5 of the documentation [14]).

OpenWFE provides a broader range of facilities for internal data interaction. Through the use of global variables, the Data Interaction – Task to Task pattern (WDP-9) is supported. Global variables are also utilized for data passing to and from blocks/subprocesses (WDP-10, WDP-11). Subprocesses are defined either in the same file as the main process from which they are invoked (see line 11 and lines 16-21 in Listing 44) or in a separate file (thus allowing different processes to invoke them as a part of their execution), see Listing 45 and line 12 in Listing 44. Unless limited through the use of filters, all fields (e.g. field ‘a’ and field ‘next’) become automatically visible to every subprocess, both internal and external, that is invoked after their definition. Furthermore, a variable (e.g. variable ‘x’) is accessible directly (i.e. both readable and writable) from any subprocess by using its name preceded by a slash (see lines 6-7 in Listing 45). When the subprocess is defined in the same file as the main process, the slash preceding the name of the variable can be omitted (as is done on line 18 in Listing 44). In addition, OpenWFE also supports data interaction with a subprocess through the use of parameters. For an example of this, see lines 11 and 19 in Listing 44 where the value of field ‘next’ is passed to the subprocess ‘internalB’ through the parameter ‘pr’.

Listing 44 Internal Data Interaction in OpenWFE

```

1 <process-definition name="dp10-11" revision="1">
2   <description>Data transfer-main</description>
3   <sequence>
4     <set field="a" type="integer" value="" />
5     <set variable="x" type="integer" value="2"/>
6     <set field="next" type="string" value="role-?" />
7     <participant description="task A" ref="role-alpha"/>
8     <if>
9       <lesser-than value="{call:mul(' ${f:a}' ' ${v:x}')}"/>
10        other-value="10" />
11        <subprocess ref="internalB" pr="{f:next}"/>
12        <subprocess ref="http://localhost:7079/sub.xml"/>
13      </if>
14      <participant description="task C" ref="role-alpha"/>
15    </sequence>
16    <process-definition name="internalB">
17      <sequence>
18        <set variable="x" value="{call:mul(' 2' ' ${v:x}')}"/>
19        <participant ref="{pr}" description="task B"/>
20      </sequence>
21    </process-definition>
22 </process-definition>

```

Listing 45 External subprocess in OpenWFE - file sub.xml in the workflow-definitions directory

```

1 <process-definition name="externalB"
2   revision="1">
3   <description>dp10-11 - subprocess
4   </description>
5   <sequence>
6     <set variable="/x" value="{call:
7       mul('4' ' ${v:/x}')}"/>
8     <participant ref="role-student"
9       description="task B-ext"/>
10  </sequence>
11 </process-definition>

```

Specific data elements can be passed to the instances of a multiple instance task, i.e. WDP-12 (as illustrated earlier in listings 31-33 in Section 5 where the concurrent-iterator construct was used to initiate multiple concurrent instances of a task), but aggregation of the data produced by these instances (WDP-13) is not supported. Instead, the data elements of one of the instances of the multiple instance task are taken and used in subsequent tasks. The attributes *merge* and *merge-type* can be used to define how the data from the different instances should be merged. Possible values for the attribute *merge* are ‘first’, ‘last’, ‘highest’ and ‘lowest’ and for the attribute *merge-type*, the values ‘mix’ and ‘override’ are defined. The attribute setting *merge*=‘last’ and *merge-type*=‘mix’ (as shown on lines 13 and 14 in Listing 33) implies that the data values for the last of the instances to complete will be taken and if any data values are missing they will be filled using the values from the previously completed instances.

The Data Interaction – Case to Case pattern (WDP-14) is only partially supported in OpenWFE as data is passed between cases using global variables (see the definition of the workflow variable //x on line 4 in Listing 41 for an example of this). This is considered to constitute partial support for the pattern as it results in the wider communication of a data element from one case to *all* other cases, and not to one specific case as is the intention of the pattern.

Enhydra Shark supports the Data Interaction - Task to Task pattern through global (process scoped) shared variables. It also supports data interaction between a Block Task and its Subprocess Decomposition. Global shared variables are used for data interaction to and from a Block Activity which is embedded in the main process, while parameter passing is used for communication to and from subflows that are defined outside of the main process. Data interaction to and from multiple instance tasks is not supported as the concept of task data is not supported. The Data Interaction - Case to Case pattern is partially supported through interaction with an external database. Similar to jBPM and OpenWFE, the more general behaviour of passing data from one case to all other cases is captured but not data passing from one case to another specific case.

Table 10 summarizes the results for the internal data interaction patterns.

Data Interaction - Internal	1	2	3
9. Task to Task	+	+	+
10. Block Task to Subprocess Decomposition	-	+	+
11. Subprocess Decomposition to Block Task	-	+	+
12. to Multiple Instance Task	-	+	-
13. from Multiple Instance Task	-	-	-
14. Case to Case	+/-	+/-	+/-

Table 10. Support for the internal data interaction patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

6.3 Data interaction patterns - external

OpenWFE supports both the Data Interaction – Task to Environment – Pull and Data Interaction – Environment to Task – Push patterns (i.e. WDP-15 and WDP-16). Listing 47 shows an example of a process capable of sending email to name@company.com, with the content of the field=“_subject_” in it (i.e. the Data Interaction – Task to Environment – Push-oriented pattern (WDP-15)). The mail is sent by a task called “send mail” executed by participant “notif-alpha”, the latter being an artificial participant configured in the participant-map file as such (an excerpt of

which is shown in Listing 46). Listing 48 demonstrates invocation of a web-service (and consequently demonstrates support for the Data Interaction – Task to Environment – Pull Oriented pattern (WDP-16)). To be able to invoke an external service, a corresponding participant needs to be defined in the participant-map configuration file. Listing 43 shows how the participant *ws-weather* corresponding to the operation *GetWeather* service *globalweather* is defined. The values retrieved from the fields *city* and *country* (as defined in lines 2 and 3 in Listing 48) during the execution of task “request-service” are sent to the *ws-weather* participant (line 6 Listing 48) when invoking the *GetWeather* operation and the answer received is assigned to the field *weather* (as defined in line 18 in Listing 43) during the execution of the task “display result” (lines 7 and 8 in Listing 48). Although OpenWFE does not provide support for task data (and instead supports case data), the data interaction is initiated at the task level rather than the case level, and hence is bound to a specific place/task in the process. Consequently the Data Interaction – Case to Environment – Push and Pull oriented (WDP-19 and WDP-20) patterns are not considered to be supported and neither are the Data Interaction – Process to Environment – Push-oriented and Pull-oriented (WDP-23 and WDP-24) patterns.

Although it is possible to use a relational database management systems for persistence and OpenWFE is distributed with code for MySQL, PostgreSQL, MS-SQL and DB2 connectivity, this functionality is not considered to provide support for data transfer between a process and external sources. It is really intended as a means of achieving data persistence in a WFMS, hence such a DBMS is considered to be an integral part of the WFMS (rather than of the environment).

Similar to OpenWFE, external data interaction in jBPM is realized at task (and transition) level, rather than at process level. The Data Interaction – Task to Environment – Push pattern (WDP-15) is most easily realized through the use of Actions, which are defined on transitions or at the entrance or the exit of nodes. In [6] pp 134-140, the connection of a jBPM task to an external database and the writing of data to that database is exemplified, by actually coding an ActionHandler Java class which enables the required data transfer. In a similar way, through the use of Nodes which can stop and wait for the receipt of data which they are able to retrieve from external sources, the Data Interaction – Environment to Task – Pull pattern (WDP-16) can also be implemented. Hence patterns WDP-15 and WDP-16 are considered to be supported. However, due to the extent of Java coding needed for achieving the required data interaction, the extent of the support for these patterns is considered to be partial. None of the other patterns in this category are supported by jBPM.

From this group of patterns, Enhydra Shark only offers support for the Task to Environment - Push and Environment to Task - Pull-Oriented patterns. A ToolAgent provides the ability for Enhydra Shark to interface with different software applications. For instance the SOAPToolAgent can be used for invoking web-services to which data can be sent and from which data can be received. The invocation of a tool agent is specified through ExtendedAttributes.

Table 11 summarizes the results for the external data interaction patterns.

Listing 46 WDP-15 in OpenWFE - excerpt from a participant configuration

```

1 <participant name="notif-alpha"
2 class="[...].MailNotifiedParticipant">
3   <!-- email notification parameter -->
4   <param>
5     <param-name>recipient-field</param-name>
6     <param-value>to</param-value>
7   </param>
8   <!-- parameters with their default values -->
9   <param>
10    <param-name>smtp-server</param-name>
11    <param-value>mail.company.com</param-value>
12  </param>
13  <param>
14    <param-name>smtp-port</param-name>
15    <param-value>25</param-value>
16  </param>
17  ...

```

Listing 47 WDP-15 in OpenWFE - excerpt from a process definition

```

1 <sequence>
2   <set field="to" value="name@company.com" />
3   <set field="__subject__" value="test" />
4   <participant description="send mail"
5     ref="notif-alpha" />
6 </sequence>

```

Listing 48 WDP-16 in OpenWFE - excerpt from a process definition

```

1 <sequence>
2   <set field="city" value="Stockholm" />
3   <set field="country" value="Sweden" />
4   <participant ref="role-bravo"
5     description="request service" />
6   <participant ref="ws-weather" />
7   <participant ref="role-bravo"
8     description="display result" />
9 </sequence>

```

6.4 Data transfer patterns

Data is transferred in jBPM using a Copy in/Copy out strategy (i.e. WDP-29). When dealing with concurrent tasks, this implies that at creation every work item receives a copy of the data values in the global process variables. Any updates to these values by a concurrent work item are only visible to that work item. At completion the values for each work item are copied back (if so specified) to the global process variables, overwriting the existing values for these variables. This means that the last work item to complete will be the last one to copy its values back to the process variables. To avoid potential problems with lost updates, the process designer needs to be careful when granting write access to variables for tasks that run concurrently and update global data.

Data Interaction - External	1	2	3
15. Task to Environment – Push-Oriented	+/-	+	+
16. Environment to Task – Pull-Oriented	+/-	+	+
17. Environment to Task – Push-Oriented	-	-	-
18. Task to Environment – Pull-Oriented	-	-	-
19. Case to Environment – Push-Oriented	-	-	-
20. Environment to Case – Pull-Oriented	-	-	-
21. Environment to Case – Push-Oriented	-	-	-
22. Case to Environment – Pull-Oriented	-	-	-
23. Process to Environment – Push-Oriented	-	-	-
24. Environment to Process – Pull-Oriented	-	-	-
25. Environment to Process – Push-Oriented	-	-	-
26. Process to Environment – Pull-Oriented	-	-	-

Table 11. Support for the external data interaction patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

Scripts can be used to manipulate data elements before a task is started and after it has completed, i.e. at the initiation and completion of a task node (see lines 2-3 and 12-13 in Listing 49)¹⁷. The event types “node-enter” and “node-leave” indicate when the script will be run. A script consists of an expression, describing what needs to be done, and a set of variables, specifying the input data to the script (if a variable definition is omitted, all process variables will be loaded into the script upon initiation). When a task-node contains a single task, the node-enter and node-leave scripts provide direct support for the two data transformation patterns WDP-32 and WDP-33.

Listing 49 Data transformation in jBPM

```

1 <task-node name="task1">
2   <event type="node-enter">
3     <script>
4       <expression> z = x + " " + y;
5       <variable name="A" access="read" mapped-name="x">
6       <variable name="B" access="read" mapped-name="y">
7       <variable name="C" access="write" mapped-name="z">
8     </script>
9   </event>
10  <task name="task1"> ...
11  <transition name="" to="task2">
12  <event type="node-leave">
13    <script>
14      <expression> c = c + " " + c;
15      <variable name="C" access="read,write"
16        mapped-name="c">
17    </script>
18  </event>
19 </task-node>

```

In terms of the Data Transfer patterns, OpenWFE provides direct support for the Data Transfer by Reference With Lock pattern (WDP-31). As all data elements are global for a process, each of them resides in a location accessible by all process components. Concurrency restrictions are enforced such that when a task is started by a resource, the corresponding work item (and associated data elements) is locked and the remaining resources with execution privileges for the task receive temporary read-only access to it. When dealing with concurrent tasks or multiple instance tasks, OpenWFE utilizes a Copy in/Copy out strategy (i.e. WDP-29). At creation, every work item receives a copy of the data values in the global variables. Any updates to these values by a given work item remain invisible outside that work item. At completion the values associated with the work item are copied back in accordance with the strategy defined in the synchronization construct, which has the potential to lead to lost update problems. Data transfer to and from multiple instances tasks is treated similarly (as discussed in the second paragraph of Section 6.2). To avoid the potential for lost updates, filters can be used on concurrent tasks to prevent simultaneous access to data elements. However, filters can not be used to solve this issue for multiple instances tasks.

Data transformations on input and output data (patterns WDP-32 and WDP-33 respectively) are supported. As previously mentioned, data transformations can be performed at any time between two tasks in a process. Examples of such transformations are illustrated by lines 6-7 in Listing 45 (where the value of the process variable /x is multiplied by four before the commencement of task “task B-ext”) and lines 9 and 10 in Listing 41 (where the value of the field C is calculated based on input from the preceding task “task2”).

Similar to jBPM and OpenWFE, Enhydra Shark supports transformations on input and output data (patterns WDP-32 and WDP-33) and Copy In/Copy Out (WDP-29) patterns. However the copy out operation is not applied when concurrent work items are operating on the same data, which may lead to updates being lost. When invoking subprocesses or block activities in Enhydra Shark, it is also possible to transfer data using parameters, hence there is support for the Data Transfer by Value – Incoming and Data Transfer by Value – Outgoing patterns (WDP-27 and WDP-28). The support is considered to be partial as these strategies are not applicable for data transfer between all activities and only apply where the data transfer is between block activities and subprocesses.

Table 12 summarizes the results for the data transfer patterns.

¹⁷ jPDL also supports scripts at the creation, start, assignment and end of a task. However, the graphical interface does not provide support for the definition of such scripts. Moreover, the task scripts we defined in the XML files specifying the processes did not execute at runtime.

Data Transfer	1	2	3
28. by Value – Incoming	–	–	+/-
28. by Value – Outgoing	–	–	+/-
29. Copy In/Copy Out	+	+	+
30. by Reference – Unlocked	–	–	–
31. by Reference – Locked	–	+	–
32. Data Transformation – Input	+	+	+
33. Data Transformation – Output	+	+	+

Table 12. Support for the data transfer in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

6.5 Data-based routing

There is no support for the Task Precondition - Data Existence (WDP-34) or Data Value (WDP-35) patterns in jBPM or Enhydra Shark, as there is no means to specify preconditions. OpenWFE supports both these patterns. The Task Precondition - Data Value pattern is illustrated through the if-statement on lines 8-13 in Listing 44, where if the sum of values in field ‘a’ and variable ‘x’ is less than ten, the internal subprocess is invoked otherwise an external subprocess is invoked. The Task Precondition – Data Existence pattern is supported in a similar way, except that the condition associated with the <if> statement contains an expression <defined variable-value=“variable-name”> testing whether the variable *variable-name* has been defined or not.

Enhydra Shark also partially supports the Task Postcondition – Data Value pattern (WDP-37). The support is partial because only the data types (but not their values) are checked. Furthermore, Enhydra Shark does not support the Task Postcondition – Data Existence pattern (WDP-36). Conditions like “varName != null” and “varName == null” can be specified, but whilst the first of these conditions always evaluates to true, the second one always evaluates to false because if a value is not entered, Enhydra Shark puts an empty string in the variable (varName in this example). If the variable varName is of type integer, an error message stating “Incorrect type!” is displayed.

In jBPM, the Task Postcondition – Data Existence pattern (WDP-36) is intended to be supported through the attribute Required for variables. However, it was found during testing that not providing a value for a variable which was marked as required did not prevent the user from registering the task as completed¹⁸. Hence there is no support for postconditions checking whether certain parameters have received a value, nor is there any support for the Task Postcondition – Data Value pattern (WDP-37). No support for the two postcondition patterns was observed in OpenWFE.

The Event-based Task Trigger pattern (WDP-38) is an extension of the Transient and Persistent Trigger patterns (CFP-23 and CFP-24), describing triggers in which data (and not just signals) is transferred. Furthermore, the Data-based Task Trigger pattern (WDP-39) focuses on the situation where the trigger is not an external event (e.g. mail arrival) but the change of an internal data condition.

Support for Data-based Task Triggers was not observed in OpenWFE. The pattern is currently not supported in jBPM as stated in jBPM [17], nor does Enhydra Shark support it. None of the offerings supports Event-based Task Triggers. In jBPM the State node can only be used for receiving signals as triggers where there is no data attached to the signals. In the version of Enhydra Shark evaluated, the ability of the MailToolAgent to receive triggers in the form of mail items was tested, but it did not work.

The support for the Data-based Routing pattern (WDP40) is rated as partial in both jBPM and OpenWFE as there is not full support for the multi-choice variant of the pattern. In contrast, Enhydra Shark fully supports the Data-based Routing pattern as it supports both the exclusive and multiple choice constructs.

The support for this group of patterns is summarized in Table 13.

Data-based Routing	1	2	3
34. Task Precondition – Data Existence	–	+	–
35. Task Precondition – Data Value	–	+	–
36. Task Postcondition – Data Existence	–	–	–
37. Task Postcondition – Data Value	–	–	+/-
38. Event-based Task Trigger	–	–	–
39. Data-based Task Trigger	–	–	–
40. Data-based Routing	+/-	+/-	+

Table 13. Support for the data based routing patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

¹⁸ The discussion on whether it is a bug in jBPM or a failure in JSF which assigns an empty string to variables when no value has been entered for them is ongoing. For details see [18].

7 Resource Patterns Evaluation

The workflow resource patterns (WRPs) focus on the manner in which work is distributed amongst the resources in an organization and managed through to completion. In WFMSs this translates to the way that task instances (i.e. work items) are distributed among the participants (i.e. the users) associated with a process. Figure 15a (reprinted from [34]) provides a general state transition diagram for the lifecycle of a work item. The prefixes S and R in the names of the transitions are abbreviations for *System* and *Resource* and indicate the initiator of a specific transition. A work item is first created by the system, which means that the preconditions for the enablement of the associated task have been met. At this point in the process, an instance of the task (also known as a work item) is created. This work item is either offered to a single resource or to a group of resources, or directly allocated to a specific resource. When it is offered, the offer is non-binding, i.e. a resource can choose to accept or reject it, and the work item remains on offer until it is finally selected by one of the resources to which it is offered. Once selected by a resource, the work item is allocated to that resource on a binding basis (i.e. the resource is expected to complete the work item at some point in the future). Subsequent to being allocated, a work item can be executed, (i.e. started and completed) at a time chosen by the resource. During execution the work item can be suspended and resumed. The execution may also lead to an unsuccessful completion and as a result the work item ends up in a failed state. Some workflow engines allow a work item to be started immediately after it has been offered, thus omitting the allocation step (and hence expediting work item commencement).

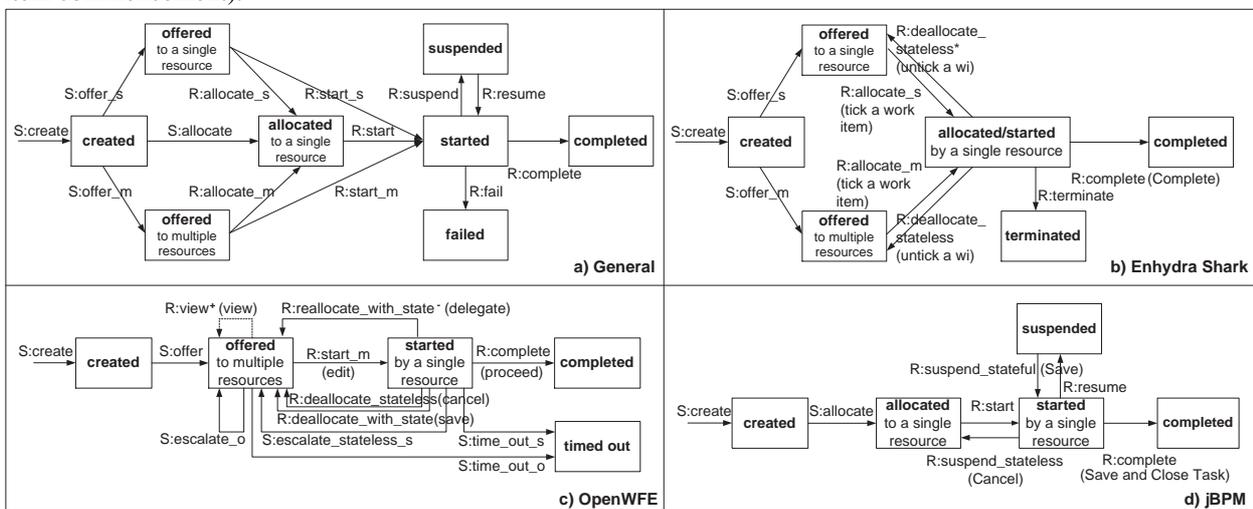


Figure 15 The lifecycle of a work item (the transition names given in parentheses are the corresponding commands in the user interfaces for these tools)

The resource patterns are grouped with reference to the work item lifecycle and the patterns are divided into the following categories: *creation patterns*, *push patterns*, *pull patterns*, *detour patterns*, *auto-start patterns*, *visibility patterns*, and *multiple resources patterns*. Before discussing how the three offerings support the patterns in each of these groups, it is worth mentioning that the notion of *offered* is missing in jBPM, the notion of *allocated* is missing in OpenWFE and the *allocated* and *started* states are merged in Enhydra Shark. Hence their lifecycle diagrams, presented in Figures 15b-d, deviate from the one shown in Figure 15a. These diagrams will be further explained throughout this section.

7.1 Creation patterns

The creation patterns describe the various ways in which a work item is handled after creation and prior to it being started. They are specified at design time and rely on the existence of an associated organizational model where resources can uniquely be identified and their roles and relationships within the organization are captured. In the work item lifecycle diagram this category of patterns correspond to the S:create transition. The major patterns belonging to this category are: *Direct Allocation* - which captures the ability to assign a task to a specific resource; *Role-based Allocation* - which depicts the ability to assign a task to a specific role (such that any resource playing this role can potentially execute it); *Deferred Allocation* - which allows the decision regarding how a task will be distributed to one or more resources to be deferred to runtime; *Separation of Duties* - (also known as the Four Eyes Principle) which ensures a task cannot be distributed to a resource that completed a nominated earlier task; *Retain Familiar* - which

captures the ability to allocate a work item to the same resource that completed a preceding work item in the same case; and *Automatic Execution* - which describes the ability of a task to be executed without it needing to be distributed to a human resource. Descriptions of the remaining creation patterns can be found in [34].

In jBPM, when a work item is created, it is directly assigned to a specific resource and included in the work list for that resource (which constitutes support for the Direct Allocation pattern). The use of assignment expressions in jPDL for the purposes of defining how tasks are distributed to users is illustrated in Listing 50.

Listing 50 WRP-1 Direct Allocation in jBPM

```
1 <task-node name="concurrency A">
2   <task name="task A1">
3     <assignment expression="user(ernie)">
4   </task>
5   <task name="task A2">
6     <assignment expression="user(bert)">
7   </task>
8 </task-node>
```

Listing 51 WRP-2 Role-based Allocation in jBPM

```
1 <task-node name="concurrency A">
2   <task name="task A1">
3     <assignment expression="group(programmers)">
4   </task>
5   <task name="task A2">
6     <assignment expression="group(testers)">
7   </task>
8 </task-node>
```

Listing 52 WRP-2 Role-based Allocation in OpenWFE

```
1 <concurrency>
2   <participant description="task A1"
3     ref="role-programmer">
4   <participant description="task A2"
5     ref="role-tester">
6 </concurrency>
```

Listing 53 WRP-1 Direct Allocation in OpenWFE - not supported

```
1 <sequence>
2   <participant description="task A1"
3     ref="user-ernie">
4   <participant description="task A2"
5     ref="user-ert">
6 </sequence>
```

In OpenWFE, when a work item is created, it is offered to a role (thus corresponding to the Role-based Allocation pattern). Listing 52 provides an example illustrating how this is to be specified. The work items distributed to a role are added to the worklist for a store based on the specification in the worklist-configuration file (refer back to Listing 2 for further details). These work items can then be executed by any user with read and write permissions for that store (see Listing 3 for an example of how this is specified in the passwd configuration file).

The Direct Allocation pattern is not supported in OpenWFE, i.e. it is not possible to assign work items directly to users as proposed in Listing 53. The flow halts in this listing at runtime because the users allocated to the tasks within it are not recognized and the corresponding work items do not get distributed to them. Direct allocation can be achieved by giving only one user access to a given store, but as this solution is really a special case of role-based allocation it is not considered to constitute support for the Direct Allocation pattern. Furthermore, as there is no means of defining an organizational structure in OpenWFE, i.e. there is no support for specifying relationships between users and roles, or any more general forms of organizational groupings and hierarchy, the Organizational Distribution pattern is also not considered to be supported.

Theoretically, jBPM also supports the ability to offer a work item to a group of resources. jPDL contains the primitives “user”, “group”, “membership” and “role” and provides the ability to allocate tasks to whole groups, thus potentially supporting the Role-based Allocation pattern (see Listing 51). Furthermore there are language primitives for distributing work to different roles played by the participants belonging to a specific group, and these could conceivably be used for supporting the Organizational Distribution pattern (see Listing 54). However, while the necessary language primitives are envisaged, the web-console’s current implementation (distributed with the jBPM 3.1.4 package¹⁹) does not provide any support for them. In particular, any selection by a resource of a work item offered to a group (of which it is a member) is not supported (i.e. the system does not provide support for transition R:allocate_m in Figure 15a). For the moment, an expression such as the one in Listing 51 leads to a stalled work item which is only visible in the database managing persistence, but which does not appear in any of the available worklists. For this reason, Role-based Allocation is not considered to be supported. Interestingly, the resource allocation in Listing 54 works when there is only one resource playing the specified role. In this case the work item is put in its worklist. If more than one resource plays the same role however, the flow stalls. Hence the Organizational Distribution pattern is also considered not to be supported.

¹⁹ The console in the latest release of the tool utilized during these evaluations, i.e. jbpml-jpdl 3.2.1, did not provide support for resource management. Any resource was able to start and complete any task independently of the process definitions in jPDL. For that reason we did not evaluate it further, preferring to stay with the older but more stable version of the tool.

Listing 54 WRP-10 Organizational Distribution in jPDL

```

1 <task-node name="concurrency A">
2   <task name="task A1">
3     <assignment expression="group(employees)
4       -->role(programmer)"> </task>
5   <task name="task A2">
6     <assignment expression="group(employees)
7       -->role(tester)"> </task>
8 </task-node>

```

Listing 55 CFP-14 MI and WRP-3 Deferred Allocation in OpenWFE

```

1 <sequence>
2   <set field="iteration_list"
3     type="StringMapAttribute"
4     value="users?: alpha,bravo.."/>
5   <participant ref="role-alpha"
6     description="distribute tests to subjects"/>
7   <concurrent-iterator
8     on-field-value="iteration_list"
9     to-field="user">
10    <participant ref="role-${f:user}"
11      description="run test"/>
12  </concurrent-iterator>
13  <participant ref="role-alpha"
14    description="summarize results"/>
15 </sequence>

```

Listing 56 WRP-3 Deferred Allocation and WRP-7 Retain Familiar in jBPM

```

1 <start-state name="start">
2   <task name="start task">
3     <assignment expression="user(ernie)">
4     <controller>
5       <variable name="empl"
6         access="read,write,required">
7     </controller>
8   </task>
9   <transition name="" to="task1">
10 </start-state>
11 <task-node name="task1">
12   <task name="task1">
13     <assignment actor-id="empl">
14   </task>
15   <transition name="" to="task2">
16 </task-node>
17 <task-node name="task2">
18   <task name="task2">
19     <assignment expression="previous">
20   </task>
21   ...
22 </task-node>

```

Enhydra Shark has very limited support for the creation patterns. It supports Direct, Role-based, and Deferred Allocation. A Participant in the workflow model can be mapped both to a specific User and to a Group of users defined in the process engine. Deferred allocation is implemented by using variables. A Performer variable can be specified for an activity and its value set at runtime.

Deferred Allocation is also supported by jBPM and OpenWFE. Similarly to Shark, it is implemented using variables which receive their values at runtime before the commencement of the relevant task. The assignment of “task1” on line 13 in Listing 56 shows how variable “empl”, populated by user ernie during the execution of “start-task”, is used for Deferred Allocation in jBPM. For OpenWFE a similar example is shown in Listing 32 (for convenience this listing is reprinted as Listing 55), where the solution for one of the Multiple Instance patterns was presented.

The Retain Familiar pattern is only supported in jBPM. The resource assignment for “task2” on line 19 in Listing 56 where the expression attribute is set to “previous” shows how this is implemented. In practice, using this assignment expression for every task in a process would result in all tasks in a case being executed by the resource who started the case. Note that this is considered as repeated use of the Retain Familiar pattern, rather than support for the Case Handling pattern, as every task needs to have an individual assignment expression (as opposed to a whole case, when created, being allocated to one resource). Consequently this solution is not considered to provide support for the Case Handling pattern.

The Automatic Execution pattern, which allows tasks to be executed by applications and is common to most workflow systems is supported by all three offerings. Table 14 summarizes the support of the tools for the creation patterns.

Creation Patterns	1	2	3
1. Direct Allocation	+	-	+
2. Role-Based Allocation	-	+	+
3. Deferred Allocation	+	+	+
4. Authorization	-	-	-
5. Separation of Duties	-	-	-
6. Case Handling	-	-	-
7. Retain Familiar	+	-	-
8. Capability-based Allocation	-	-	-
9. History-based Allocation	-	-	-
10. Organizational Allocation	-	-	-
11. Automatic Execution	+	+	+

Table 14. Support for the creation patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

7.2 Push patterns

The push patterns describe the different ways in which a work item is offered or allocated to resources by the system. These include indirect offerings where a work item is advertised on a shared work list, as well as direct offering and

allocation schemes where a work item is distributed to the work list of a specific resource. There are three groups of patterns belonging to this category. The first group captures the way in which the distribution is done. The three patterns in this group are: *Distribution by Offer-Single Resource*, *Distribution by Offer-Multiple Resources* and *Allocation-Single Resource* which correspond directly to the transitions S:offer_s, S:offer_m and S:allocate in Figure 15a. The second group captures the different means by which a resource is singled out from a set of possible candidates and allocated a work item. The patterns in this group are *Random Allocation*, *Round Robin Allocation* and *Shortest Queue* allocation. The third group focuses on the timing of the allocation process and the availability of a work item. It distinguishes three patterns: *Early Distribution*, *Distribution on Enablement* and *Late Distribution*.

Common to all three offerings is the characteristic that when a work item is created, it is immediately distributed, thus illustrating support for the Distribution on Enablement pattern. Furthermore, jBPM currently supports the Distribution by Allocation to a Single Resource pattern²⁰, OpenWFE supports the Distribution by Offer to Multiple Resources pattern and Enhydra Shark supports the Distribution by Offer to Single Resources and Distribution by Offer to Multiple Resources patterns. None of the allocation patterns (random, round robin or shortest queue) are supported, as these patterns are only relevant when a system selects a resource from a number of possible candidates and this does not occur in any of the offerings. Table 15 summarizes these results.

Push Patterns	1	2	3
12. Distribution by Offer-Single Resource	-	-	+
13. Distribution by Offer-Multiple Resource	-	+	+
14. Distribution by Allocation-Single Resource	+	-	-
15. Random Allocation	-	-	-
16. Round Robin Allocation	-	-	-
17. Shortest Queue	-	-	-
18. Early Distribution	-	-	-
19. Distribution on Enablement	+	+	+
20. Late Distribution	-	-	-

Table 15. Support for the push patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

7.3 Pull patterns

The pull patterns capture situations where resources that have been offered work items commit themselves to execute them at some future time. The situation where a resource simultaneously selects and commences the execution of an offered work item is also included in this category.

There are six pull patterns: *Resource-Initiated Allocation* which captures the situation where a resource commits to executing an offered work item at some future time (illustrated by transitions R:allocated_s and R:allocated_m in Figure 15a); *Resource-Initiated Execution - Allocated Work Item* where a resource commences a work item already allocated to it (depicted by the transition R:start in the lifecycle diagram); *Resource-Initiated Execution - Offered Work Item* where a resource simultaneously commits to executing and starts an offered work item (illustrated by the transitions R:start_s and R:start_m in Figure 15a); *System/Resource-Determined Work Queue Content* which describe the ability of the system/resource to control the content and the ordering of the work items in a work list; and *Selection Autonomy* which describes the ability of a resource to select the execution order for the work items in their work list.

In jBPM, a work item that is allocated to a resource can be started by the resource selecting the work item from their work list. This indicates support for the Resource-Initiated Execution – Allocated Work Item pattern. If several work items are available in a work list, the order of their execution is determined by the resource which also indicates support for the Selection Autonomy pattern. None of the other pull patterns are supported in jBPM.

In OpenWFE, a work item offered to a shared group list can be simultaneously selected and started by a resource in the group. Thus the Resource-Initiated Execution – Offered Work Item pattern is directly supported. The order in which work items are executed is determined by individual resources based on the contents of their work list hence the Selection Autonomy pattern is also supported. None of the other pull patterns are supported in OpenWFE.

In Enhydra Shark, a work item offered to a user appears in their work list; a work item offered to a group, appears in the work lists of every member of the group. Then the work item can be simultaneously selected and started by a resource (in the group). When selected/started by a resource, a work item that was offered to a group, disappears from the work lists of the remaining group members. Thus the Resource-Initiated Execution – Offered Work Item pattern is supported. If several work items are available in a work list, the order of their execution is determined by the resource, which indicates support for the Selection Autonomy pattern.

²⁰ In the latest tested release i.e. 3.2.1, jBPM seems to be preparing support for Distribution by Offer - Multiple Resources and the topic is frequently discussed on their forum, however such support was not available at the time of this evaluation.

Table 16 summarizes the support of the evaluated tools for the pull patterns.

Pull Patterns	1	2	3
21. Resource-Initiated Allocation	-	-	-
22. Resource-Initiated Execution – Allocated Work Item	+	-	-
23. Resource-Initiated Execution – Offered Work Item	-	+	+
24. System-Determined Work List Management	-	-	-
25. Resource-Determined Work List Management	-	-	-
26. Selection Autonomy	+	+	+

Table 16. Support for the pull patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

7.4 Detour patterns

Detour patterns refer to situations where the normal process by which work items are distributed to resources and managed through to completion is interrupted, resulting in a varied sequence of states and resource assignments for such work items. The interruption can be instigated by a resource or by the system. One such pattern captured in Figure 15a is *Suspension/Resumption*, which denotes the possibility of a resource to temporarily suspend the execution of a work item and to resume it at a later time. To avoid cluttering the diagram with atypical state transitions, the remainder of the detour patterns are not shown graphically. They are: *Delegation* which denotes the situation where a resource allocates a work item from their work list to another resource (in Figure 15a this would be captured by a transition from the “allocated” state back to itself); *Escalation* which describes the situation where the system proactively attempts to progress a work item that has stalled, by offering or allocating it to another resource or group of resources (in Figure 15a this would be represented by a number of possible transitions, each one going from the “started”, “allocated”, or “offered” states back to one of these states); *Deallocation* which denotes the situation where a resource relinquishes a work item which is currently allocated to them and make it available for (re-)distribution to other resources; *Reallocation* which captures the situation where a resource (re-)allocates a work item that they have started working on to another resource. In contrast to the Delegation and Deallocation patterns, Reallocation addresses work items that have been started. Also, in contrast to the Escalation pattern, Reallocation is initiated by a resource and not by the system. A reallocation is either *Stateful* (i.e. it preserves the state of a work item captured by the data associated with it) or *Stateless*. In the work item lifecycle, these patterns correspond to transitions from the “started” state back to the “allocated” state; *Skip* describes the scenario where a resource omits the execution of a work item assigned to them. This corresponds to a transition from the “allocated” state to the “completed” state; *Redo* which provides the possibility for a resource to execute a work item that has already been completed, thus introducing a backward transition from the “completed” to the “started” state; and *Pre-Do* which denotes the ability of a resource to execute a work item ahead of the current execution point in a case (a possibility that does not have a corresponding representation in the workitem lifecycle diagram).

Once a work item has been started in jBPM, it can be completed, or it can also be suspended hence demonstrating support for the Suspend/Resumption pattern, which is also shown in the lifecycle in Figure 15d.

In Enhydra Shark a user can “release” a task they selected and initiated (by un-checking it in the work list). If it is a work item that was initially offered to a group of users this will result in re-offering the work item to the group. Hence the Deallocation pattern is supported, although the semantics deviate slightly for the pattern. As the state “allocated” is merged with the state “started” (see Figure 15b), a work item needs to be started in order to be deallocated. Therefore a distinction can be made between stateful and stateless deallocations. The deallocation supported in Enhydra Shark is stateless.

In OpenWFE, a started work item is completed by selecting the “proceed” option from the task execution window. Once it is started, a work item can be reallocated to another group of resources which is achieved by selecting the “delegate” option. As the task must have been started in order to select the “delegate” option, the behaviour does not correspond to the Delegation pattern, but instead to the Reallocation pattern. This form of reallocation is stateful hence the Stateful Reallocation pattern is directly supported. Reallocation of a work item can only be initiated by users who possess the “delegate” right (defined in the passwd configuration file, recall Listing 3) in the store for which they are executing the work item in question. This is indicated in Figure 15c through the use of the suffix ‘-’ for the R:allocated_with_state transition. Delegation is not supported as a work item can only be forwarded to another resource once it has been started. In addition, a started work item can also be deallocated and selecting the “cancel” option results in a stateless deallocation with the work item being sent back to the group work list. Selecting the “save” option results in a stateful deallocation. Note that this is distinct from the Suspension/Resumption pattern as the interrupted work item does not remain in the work list of the resource who started executing it, but is sent back

to the group work list and could potentially be continued by any resource that is a member of the same group. Note also that the semantics deviate slightly from those for the Deallocation pattern. As the state “allocated” is missing in OpenWFE, a work item needs to be started in order to be deallocated. This implies that a distinction can be made between stateful and stateless deallocations.

Enhydra Shark does not support Stateful or Stateless Reallocation. The Reassign button in the Work List Console is currently not enabled (in the open source community version of the tool). It does not support Suspension/Resumption of a work item. In the evaluated version of the tool a “Suspension” and a “Resumption” button for a work item are displayed but not enabled for use. In Enhydra Shark a whole case can be suspended and resumed. In such a situation all work items available for the case are suspended (and resumed). They still appear in the work list(s), but it is not possible to allocate/start them (when one tries to do this, an incorrect error message is displayed, i.e. “the work item is allocated to another resource”).

OpenWFE supports the Escalation pattern, although the fact that the change of resource for the work item is initiated by the system means that it is not supported in the user interface of the tool. Escalation is captured via the “timeout” attribute through which the time duration for a task is limited to a specified value. For instance, in Listing 57 the duration of “task1” is limited to 30 seconds. After a task is timed out (which may occur both when the task is waiting to be started or after it has actually been started) the flow either continues to the next task (which in the lifecycle diagram is shown by the S:time_out_s and S:time_out_o transitions) or the timed-out work item is (re-) distributed to another group work list (captured in the work item lifecycle diagram through the S:escalate_stateless_s and S:escalate_o transitions). The new resource allocation, as is the case in the listing, needs to be explicitly defined using an if-then statement specifying the identity of the new resource assigned to the task. As indicated by the name S:escalate_stateless_s the escalation is stateless.

Enhydra Shark does not support Escalation. The notion of deadline is present for limiting the duration of an activity, but during runtime the deadline was not found to expire. The example provided with the installation package did not work either.

Listing 57 WRP-28 Escallation in OpenWFE

```

1 <sequence>
2   <set field="__subject__" value="task1" />
3   <set field="first name" value="name" />
4   <participant ref="role-bravo"
5     description="task1" timeout="30s"/>
6   <if>
7     <equals field-value="__timed_out__"
8       other-value="true" />
9     <participant ref="role-alpha"
10      description="task1" />
11   </if>
12   <participant ref="role-charly"
13     description="task2"/>
14 </sequence>

```

The existence of the <cursor> construct may at a first glance give the impression that the behaviour of the Skip pattern can be captured in OpenWFE. Tasks that are included within the <cursor> and </cursor> tags, may at runtime be given the value “skip” due to the __cursor_command__ field appearing in their user interface. However, this does not influence the execution of the current work item, which still needs to be completed by pressing the “proceed” button. As with any other completed work item, this work item is recorded in the execution log as having successfully completed with all the changes in the data made during its execution being preserved. In the __cursor_command__ field, the value ‘skip 2’ can be given in order to skip the subsequent task. However even this action does not provide support for the Skip pattern, because the work item for the subsequent task is not created, which is a prerequisite for the Skip pattern.

The Redo pattern is supported in OpenWFE through the concept of the <cursor>. An important issue associated with this pattern is that if a user goes back a number of steps and redoes a task, then they are forced to also redo any task that may be influenced by the redone task, i.e. the act of redoing a task should not break the integrity of the overall process model. This can be achieved by only allowing the “back” and “rewind” operations within a cursor (and disallowing the “skip”, “break” and “cancel” operations). However, when a process is rewound a number of steps, the task to be redone is offered not only to the resource who initiated the rewind, but to all members of the shared worklist, and it can therefore potentially be executed by anyone with sufficient privileges. Furthermore, the rewind command operates not at a task level but at the instruction level. This means that, for example, if there is a data-manipulation instruction between two tasks, it will be included as part of the rewind step. This necessitates very detailed knowledge of the process model by the end user when rewinding a process instance to a desired task, and requires that they are not only aware of the number of tasks that they will need to rewind, but also cognisant of any data manipulation operations and transfers defined in between tasks when determining the steps to be rewound. As a consequence of these peculiarities, support for the Redo pattern in OpenWFE is only considered to be partial.

The Pre-Do pattern is not supported in any of the tools. The pattern implies that a task can be performed in isolation before the thread of control has reached the task. This behaviour is not possible in any of the three tools. Table 17 summarizes the results for the detour patterns.

Detour Patterns	1	2	3
27. Delegation	-	-	-
28. Escalation	-	+	-
29. Deallocation	-	+	+
30. Stateful Reallocation	-	+	-
31. Stateless Reallocation	-	-	-
32. Suspension/Resumption	+	-	-
33. Skip	-	-	-
34. Redo	-	+/-	-
35. Pre-Do	-	-	-

Table 17. Support for the detour patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

7.5 Auto-start patterns

Auto-start patterns relate to situations where the execution of a work item is triggered by specific events in the lifecycle of the work item or the related process definition. Such events may include the creation or allocation of the work item, completion of another instance of the same work item or a work item that immediately precedes the one in question. There are four auto-start patterns: *Commencement on Allocation* (represented by the S:start transition in the work item lifecycle from the “allocated” to the “started” state) which describes the situation where a work item is started at the same time that it is allocated to a specific resource; *Commencement on Creation* which is similar to the Commencement on Allocation pattern in that it describes a transition triggered by the system, but in this case the commencement (and allocation) of the work item occurs directly after it is created. The rationale for this behaviour is that it expedites the throughput of work items in a process instance by cutting out some of the intermediate states (albeit with some resultant loss of flexibility in the deployment of the process). The remaining two patterns in this category capture the initiation of a work item based on the completion of a related work item. *Piled Execution* denotes the ability of the system to initiate the next work item for a given task (possibly in a different case) immediately after a resource has completed another instance of the same task. *Chained Execution* describes the ability of the system to automatically start the next work item in a case for a resource once they have completed the preceding work item. It differs from the Retain Familiar pattern in that the new work item is not only allocated to the same resource, but it is also started for them.

None of the auto-start patterns are supported in jBPM, OpenWFE or Enhydra Shark as all of them relate to situations where the system starts the execution of a work item ‘on behalf’ of a resource (see Table 18). In the evaluated offerings, only a resource can initiate the execution of a work item assigned to them. In Enhydra Shark, “Automatic” can be selected as a Start Mode for an activity (the default start mode is “Manual”). However, at runtime the work items of activities with this “Automatic” setting are not started automatically, but need to be started manually by selecting them from the corresponding work list.

Auto-start Patterns	1	2	3
36. Commencement on Creation	-	-	-
37. Commencement on Allocation	-	-	-
38. Piled Execution	-	-	-
39. Chained Execution	-	-	-

Table 18. Support for the auto-start patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

7.6 Visibility and multiple resources patterns

Unlike the previous groups of resource patterns, the *visibility patterns* define process characteristics that are not directly related to the work item lifecycle. Instead, they describe the openness of a system’s operation to scrutiny by process participants. There are two visibility patterns: *Configurable Unallocated* and *Allocated Work Item Visibility* which describe the ability to configure the visibility of unallocated and allocated work items to resources involved in a process.

OpenWFE provides limited support for both these patterns. When a user is defined (see Listing 3), the privileges they have to access different stores are also specified. Giving “read” privileges to a user for a store implies that they will be able to see both unallocated work items and also started (i.e. allocated) work items in that store. The fact that only users with read privileges can see the workitems in a store is depicted with the suffix ‘+’ in the R:view transition in Figure 15c²¹. This configuration cannot be performed by the user, but is performed by a workflow administrator.

²¹ Note also that this transition is slightly different than the other transitions in the lifecycle diagram: it does not imply any real change between the states that are its source and target. To show this difference a dotted arrow is used in the diagram.

It is interesting to note that it is not possible to configure the visibility so that only started or only unallocated items are visible. Hence the degree of support for the visibility patterns in OpenWFE is considered to be partial. There is no support for these patterns in jBPM or Enhydra Shark. Table 19 summarizes these results.

Finally, the *multiple resource patterns* describe situations where there is not a one-to-one correspondence between executing work items and resources (which is assumed to hold for the other resource patterns). There are two patterns in this group: *Simultaneous Execution* which describes the ability of a resource to execute several work items concurrently; and *Additional Resources* which describes the situation where more than one resource is assigned to the same work item. Neither of these patterns is supported by jBPM, OpenWFE or Enhydra Shark (see Table 19).

Visibility Patterns	1	2	3
40. Configurable Unallocated Work Item Visibility	-	+/-	-
41. Configurable Allocated Work Item Visibility	-	+/-	-
Multiple Resource Patterns			
42. Simultaneous Execution	-	-	-
43. Additional Resources	-	-	-

Table 19. Support for the visibility and multiple resource patterns in 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

8 Conclusion

8.1 Evaluation summary

This paper sought to investigate the state-of-the-art in open source workflow management software through a patterns-based analysis of three representative systems. Overall, one can conclude that the situation is promising and all three open source offerings that were examined provide a reasonably comprehensive range of functionality on which the automation of business processes can be based. This effectively demonstrates that the open-source community is capable of developing and delivering worthwhile contributions in the BPM domain.

Although the range of constructs supported by the three systems is somewhat limited, OpenWFE offers a considerably broader range of features than jBPM or Enhydra Shark. There are several potential areas for improvement and further development in all of the offerings. The evaluation results are compiled in tables 20-22 and summarized below. For comparison purposes tables 20-22 also incorporate the results from earlier patterns-based evaluations of Staffware²², IBM’s WebSphere MQ and Oracle BPEL Process Manager²³. See [43] for details about existing evaluations of closed source WFMSs.

For the control-flow perspective (Table 20), we can conclude:

- jBPM and Enhydra Shark support a relatively limited set of control-flow operators offering little support for patterns other than those related to basic control-flow;
- OpenWFE offers broader support for variants of the partial join (see column 2, patterns 30 and 32 in Table 20), some discriminator constructs (column 2, patterns 9 and 29 in Table 20) and also for controlled task concurrency, i.e. multiple instance tasks (column 2, patterns 12-14 and 34-35 in Table 20).

For the data perspective (Table 21):

- All three offerings support a limited range of data element bindings and rely heavily on case-level data elements. However, whilst simplistic, the data passing strategies employed in all three systems are reasonably effective and include consideration of important issues such as inline data manipulation when data elements are being passed;
- There are limited capabilities for handling external data interaction without utilising programmatic extensions (see columns 1,2,3 for patterns 17-26 in Table 21);
- The ability to influence control-flow from the data perspective is minimal. E.g. limited post-condition, trigger and data-based routing support is provided (see columns 1,2,3 for patterns 34-40 in Table 21);
- Although not evident in the tables, there are shortcomings when dealing with concurrency during data manipulation. Data is lost either because parallel updates on it are ignored (e.g. Enhydra Shark and to some degree OpenWFE), or because some of the updates are overwritten (e.g. jBPM and to some degree OpenWFE).

Basic Control-flow	A	B	C	1	2	3	Termination	A	B	C	1	2	3
1. Sequence	+	+	+	+	+	+	11. Implicit Termination	+	+	+	+	+	+
2. Parallel Split	+	+	+	+	+	+	43. Explicit Termination	-	-	-	-	-	-
3. Synchronization	+	+	+	+	+	+	Multiple Instances						
4. Exclusive Choice	+	+	+	+	+	+	12. MI without Synchronization	+	-	+	+	+	+
5. Simple Merge	+	+	+	+	+	+	13. MI with a pri. Design Time Knl	+	-	+	-	+	-
Advanced Synchronization							14. MI with a pri. Runtime Knl.	+	-	+	-	+	-
6. Multiple Choice	-	+	+	-	+/-	+	15. MI without a pri. Runtime Knl.	-	-	+/-	-	-	-
7. Str Synchronizing Merge	-	+	+	-	-	-	27. Complete MI Activity	-	-	-	-	-	-
8. Multiple Merge	-	-	-	+	-	-	34. Static Partial Join for MI	-	-	-	-	+	-
9. Structured Discriminator	-	-	-	-	+	-	35. Static Canc. Partial Join for MI	-	-	-	-	+	-
28. Blocking Discriminator	-	-	-	-	-	-	36. Dynamic Partial Join for MI	-	-	-	-	-	-
29. Cancelling Discriminator	-	-	-	-	+	-	State-Based						
30. Structured Partial Join	-	-	-	-	+	-	16. Deferred Choice	-	-	+	+	-	-
31. Blocking Partial Join	-	-	-	-	-	-	39. Critical Section	-	-	+	-	-	-
32. Cancelling Partial Join	-	-	-	-	+	-	17. Interleaved Parallel Routing	-	-	-	-	+/-	-
33. Generalized AND-Join	-	-	-	+	-	-	40. Interleaved Routing	-	-	-	-	+	-
37. Local Sync. Merge	-	+	+	-	+/-	-	18. Milestone	-	-	+/-	-	-	-
38. General Sync. Merge	-	-	-	-	-	-	Cancellation						
41. Thread Merge	-	-	+/-	+/-	-	-	19. Cancel Activity	+	-	+/-	+	-	-
42. Thread Split	-	-	+/-	+/-	-	-	20. Cancel Case	-	-	+	-	+/-	+
Iteration							25. Cancel Region	-	-	+/-	-	-	-
10. Arbitrary Cycles	+	-	-	+	+	+	26. Cancel MI Activity	+	-	+	-	-	-
21. Structured Loop	-	+	+	-	+	-	Trigger						
22. Recursion	+	+	-	-	+	+	23. Transient Trigger	+	-	-	+	+	-
							24. Persistent Trigger	-	-	+	-	-	-

Table 20. Support for the Control-flow Patterns in **A**–Staffware 10, **B**–WebSphere MQ 3.4, **C**–Oracle BPEL PM 10.1.2, **1**–JBOSS jBPM 3.1.4, **2**–OpenWFE 1.7.3, and **3**–Enhydra Shark 2.0

Data Visibility	A	B	C	1	2	3	Data Interaction-External (cont.)	A	B	C	1	2	3
1. Task Data	-	+/-	+/-	+/-	-	+/-	21. Env. to Case–Push	+/-	+/-	-	-	-	-
2. Block Data	+	+	-	-	+	+	22. Case to Env.–Pull	-	-	-	-	-	-
3. Scope Data	-	-	+	-	+/-	-	23. Process to Env.–Push	-	+/-	-	-	-	-
4. MI Data	+/-	+	+/-	-	+	+	24. Env. to Process–Pull	+/-	-	-	-	-	-
5. Case Data	+/-	+	+	+	+	+	25. Env. to Process–Push	-	+/-	-	-	-	-
6. Folder Data	-	-	-	-	-	-	26. Process to Env.–Pull	+	+	-	-	-	-
7. Global Data	+	+	+	-	+	-	Data Transfer						
8. Environment Data	+	+/-	+	+/-	+	+/-	27. by Value–Incoming	-	+	+	-	-	+/-
Data Interaction-Internal							28. by Value–Outgoing	-	+	+	-	-	+/-
9. Task to Task	+	+	+	+	+	+	29. Copy In/Copy Out	-	-	+	+	+	+
10. Block to Subpr. Dec.	+	+	-	-	+	+	30. by Reference–Unlocked	+	-	+	-	-	-
11. Subpr. Dec. to Block	+	+	-	-	+	+	31. by Reference–Locked	-	-	-	-	+	-
12. to MI Task	-	-	+/-	-	+	-	32. Data Transf.–Input	+/-	-	-	+	+	+
13. from MI Task	-	-	+/-	-	-	-	33. Data Transf.–Output	+/-	-	-	+	+	+
14. Case to Case	+/-	+/-	-	+/-	+/-	+/-	Data-based Routing						
Data Interaction-External							34. Task Precond.–Data Exist.	+	-	-	-	+	-
15. Task to Env.–Push	+	+/-	+	+/-	+	+	35. Task Precond.–Data Value	+	-	+	-	+	-
16. Env. to Task–Pull	+	+/-	+	+/-	+	+	36. Task Postcond.–Data Exist.	+/-	+	-	-	-	-
17. Env. to Task–Push	+/-	+/-	+	-	-	-	37. Task Postcond.–Data Val.	+/-	+	-	-	-	+/-
18. Task to Env.–Pull	+/-	+/-	+	-	-	-	38. Event-based Task Trigger	+	+/-	+	-	-	-
19. Case to Env.–Push	-	-	-	-	-	-	39. Data-based Task Trigger	-	-	-	-	-	-
20. Env. to Case–Pull	-	-	-	-	-	-	40. Data-based Routing	+/-	+	+	+/-	+/-	+

Table 21. Support for the Data Patterns in **A**–Staffware 9, **B**–WebSphere MQ 3.4, **C**–Oracle BPEL PM 10.1.2, **1**–JBOSS jBPM 3.1.4, **2**–OpenWFE 1.7.3, and **3**–Enhydra Shark 2.0

For the resource perspective (Table 22):

- Only simple notions of work distribution are supported (see columns 1,2,3 for patterns 15-17 in Table 22) and typically only one paradigm exists for work item routing in each offering e.g. push or pull (see columns 1,2,3 for patterns 12-14 and 21-23 in Table 22);
- There is no support for any form of work distribution based on organizational criteria, resource capabilities or execution history (see columns 1,2,3 for patterns 4-10 in Table 22);
- All three offerings provide relatively simple facilities for work item management e.g. there is limited ability to configure work lists at resource or system level (see columns 1,2,3 for patterns 40-41 in Table 22), no notion of concurrent work item execution (columns 1,2,3 for patterns 42-43) and no facilities for optimizing work item

²² Staffware has subsequently been bought by Tibco and used as the basis for Tibco’s i-Process Suite. In January 2008 Tibco’s i-Process Suite adopted BPMN as a graphical notation (supporting a subset of BPMN symbols) and XPDL 2.0 as a process definition language.

²³ As can be read from its name, Oracle BPEL Process Manager is based on BPEL4WS, i.e. Business Process Execution Language for Web Services language [7].

Creation Patterns	A	B	C	1	2	3	Pull Patterns, continuation	A	B	C	1	2	3
1. Direct Allocation	+	+	+	+	-	+	24. Sys.-Determ. WL Mng.	+	-	-	-	-	-
2. Role-Based Allocation	+	+	+	-	+	+	25. Rrs.-Determ. WL Mng.	+	+	+	-	+	-
3. Deferred Allocation	+	+	+	+	+	+	26. Selection Autonomy	+	+	+	+	+	+
4. Authorization	-	-	-	-	-	-	Detour Patterns						
5. Separation of Duties	-	+	-	-	-	-	27. Delegation	+	+	+	-	-	-
6. Case Handling	-	-	+	-	-	-	28. Escalation	+	+	+	-	+	-
7. Retain Familiar	-	+	+	+	-	-	29. Deallocation	-	-	+	-	+	+
8. Capability-based Alloc.	-	-	+	-	-	-	30. Stateful Reallocation	+/-	+	+	-	+	-
9. History-based Alloc.	-	-	+/-	-	-	-	31. Stateless Reallocation	-	-	-	-	-	-
10. Organizational Alloc.	+/-	+	+/-	-	-	-	32. Suspension/Resumption	+/-	+/-	+	+	-	-
11. Automatic Execution	+	-	+	+	+	+	33. Skip	-	+	+	-	-	-
Push Patterns							34. Redo	-	-	-	-	+/-	-
12. Distr. by Offer-Single Rsr.	-	-	+	-	-	+	35. Pre-Do	-	-	-	-	-	-
13. Distr. by Offer-Multiple Rsr.	+	+	+	-	+	+	Auto-start Patterns						
14. Distr. by Alloc.-Single Rsr.	+	+	+	+	-	-	36. Comm. on Creation	-	-	-	-	-	-
15. Random Allocation	-	-	+/-	-	-	-	37. Comm. on Allocation	-	+	-	-	-	-
16. Round Robin Alloc.	-	-	+/-	-	-	-	38. Piled Execution	-	-	-	-	-	-
17. Shortest Queue	-	-	+/-	-	-	-	39. Chained Execution	-	-	-	-	-	-
18. Early Distribution	-	-	-	-	-	-	Visibility Patterns						
19. Distribution on Enablement	+	+	+	+	+	+	40. Config. Unalloc. WI Vis.	-	-	-	-	+/-	-
20. Late Distribution	-	-	-	-	-	-	41. Config. Alloc. WI Vis.	-	-	-	-	+/-	-
Pull Patterns							Multiple Resource Patterns						
21. Rsr.-Init. Allocation	-	-	-	-	-	-	42. Simultaneous Execution	+	+	+	-	-	-
22. Rrs.-Init. Exec.-Alloc. WI	+	+	+	+	-	-	43. Additional Resources	-	-	+	-	-	-
23. Rsr.-Init. Exec.-Offered WI	+	+	+	-	+	+							

Table 22. Support for the Resource Patterns in A–Staffware 9, B–WebSphere MQ 3.4, C–Oracle BPEL PM 10.1.2, 1–JBoss jBPM 3.1.4, 2–OpenWFE 1.7.3, and 3–Enhydra Shark 2.0

throughput, e.g. automated work item commencement, chained execution (see columns 1,2,3 for patterns 36-39 in Table 22);

- One area where OpenWFE demonstrates noticeably better facilities is in terms of the range of detour patterns it supports e.g. escalation, deallocation, reallocation (see columns 1,2,3 for patterns 27-35 in Table 22).

8.2 Comparison to three closed source offerings

We used the workflow patterns framework not only because it is a powerful analysis tool, but also because many evaluations of proprietary tools have already been performed using it hence facilitating comparison with the results from earlier studies. If we compare the state-of-the-art in open source workflow systems (i.e. columns 1,2,3) to that in proprietary systems (i.e. columns A,B,C), the results in tables 20-22 show that none of the offerings stands out as being clearly superior to the others, although it can be argued that Oracle BPEL PM demonstrates a marginally wider range of features, whilst Enhydra Shark and jBPM clearly lag behind in terms of overall patterns support. Oracle BPEL PM and OpenWFE tend to demonstrate broader pattern support in their corresponding tool classes (i.e. open source vs proprietary), especially in the control-flow perspective. Moreover, it can also be observed that the proprietary tools are generally better equipped in the resource perspective and better able to support interaction with the external environment, whereas the open source systems essentially rely on their users having programming experience (e.g. Java) to achieve the required integration with other systems. In the data perspective jBPM, relying heavily on Java coding for data passing and manipulation, clearly lags behind the other offerings.

Overall one can conclude that the open source systems are geared more towards developers than business analysts. If one is proficient with Java, jBPM may be a good choice. However, choosing jBPM is less advisable if one is not a Java expert. Similarly, whilst OpenWFE has a powerful language for workflow specification in terms of its support for the workflow patterns, we postulate that it will be difficult to understand by non-programmers. Finally, Enhydra Shark's minimalistic support for the workflow patterns may require complicated work-arounds for capturing nontrivial business scenarios.

8.3 General reflections

We would like to note that there are two limitations associated with our study. First, while the patterns framework is a powerful analysis tool, it does not cover *all* aspects of a workflow or business process management system. Characteristics such as runtime performance, version management, model analysis, and administration facilities are some examples of features not examined by the framework. Each of these aspects is important and justifies in its own right a study similar to the one reported here. Secondly, the way in which the open source offerings were selected for evaluation also represents a possible limitation. While systematic, it omits a third category of offerings, namely

offerings that were initially developed as proprietary products which have subsequently been released as open source offerings, but are still only available through their vendors' websites, e.g. Intalio's BPMS.

Although not analysed by the patterns framework, some reflections made during our work concern documentation quality, development maturity (i.e. the degree to which an offering satisfies the promises made by its vendor), administration facilities, degree of openness, ease of installation and configuration, and availability of support. With respect to these factors, our experiences indicate that the open source workflow offerings examined leave considerable room for improvement.

OpenWFE lacks explanatory documentation for its graphical notation, the user management tool UMAN, and the command line administration tool called Control. This resulted in process models needing to be directly specified in XML, user management achieved by updating configuration files, and administration difficulties being faced when cancelling work items and process instances. Dealing with these issues necessitated a deep understanding of operational aspects of the tool.

jBPM's claim of fully supporting the original set of 20 control-flow patterns is rather overstated, as for many patterns the system relies on additional Java coding for their implementation. Furthermore, process decomposition defined using the jPDL language primitives Process State and Super State in the graphical editor did not work. These parts of the language were also poorly documented hence decomposition defined directly in the XML code was not possible. Finally, data entries defined as Required at design time could be left empty at runtime.

Enhydra Shark's vendor Together, in addition to providing an open source workflow engine, also provides a closed-source Together Workflow Server which incorporates some desirable functionality (e.g. administration and user functionality). It is also worthwhile mentioning a rather negative experience with Enhydra Shark's Process Engine: the recurrent appearance of a pop-up window containing vendor information which was displayed approximately every five minutes and blocked user interaction for around ten seconds.

In terms of administration support, all three tools fell somewhat short of expectations. An example is the area of user management. In jBPM, SQL statements are required to define users. Enhydra Shark requires special settings in one of the configuration files for groups defined through its user management tool. As mentioned, OpenWFE lack of documentation complicated user management.

Although support facilities were not evaluated, it is worth mentioning that the OpenWFE help forum provided rapid and generally useful assistance. In terms of interface design of the Process Engine and the Worklist handler, Enhydra Shark stood out. The installation of both OpenWFE and Enhydra Shark was straightforward. jBPM provides a convenient introductory web-tutorial.

Acknowledgement

We would like to thank John Mettraux for prompt and helpful responses on the OpenWFE help forum and Saša Bojanic for constructive and valuable feedback on Enhydra Shark. We would also like to thank the anonymous reviewers for providing valuable feedback on earlier versions of this paper.

References

1. W.M.P. van der Aalst. Patterns and XPD L: A Critical Evaluation of the XML Process Definition Language. BPM Center Report BPM-03-09, BPMcenter.org, 2003.
2. W.M.P. van der Aalst, M. Dumas, A.H.M. ter Hofstede, and P. Wohed. Pattern-Based Analysis of BPML (and WSCI). Technical Report BPM Center Report BPM-02-01, BPMcenter.org, 2002. Last accessed 22 July 08.
3. W.M.P. van der Aalst and A.H.M. ter Hofstede. YAWL: Yet Another Workflow Language. *Information Systems*, 30(4):245–275, 2005.
4. W.M.P. van der Aalst, A.H.M. ter Hofstede, B. Kiepuszewski, and A.P. Barros. Workflow patterns. *Distributed and Parallel Databases*, 14(1):5–51, July 2003.
5. M. Bunge. *Treatise on Basic Phylosphy Volume 3: Ontology I - The Furniture of the World*. Kluwer Academic Publishers, 1977.
6. M. Cumberlidge. *Business Process Management with JBoss jBPM: a Practical Guide for Business Analysts*. Packt Publishing, Birmingham, UK, 2007.
7. F. Curbera, Y. Golland, J. Klein, F. Leymann, D. Roller, S. Thatte, and S. Weerawarana. Business Process Execution Language for Web Services version 1.1. <http://dev2dev.bea.com/techtrack/BPEL4WS.jsp>.
8. Enhydra.org. Open Source Java XPDL Workflow. <http://www.enhydra.org/workflow/shark/index.html>. Last accessed 20 Nov 07.
9. P. Green, M. Rosemann, M. Indulska, and C. Manning. Candidate interoperability standards: An ontological overlap analysis. *Data Knowledge Engineering*, 62(2):274–291, 2007.

10. N. Guarino. Formal Ontology and Information Systems. In *Proc. of Formal Ontology in Information Systems (FOIS'98)*, pages 3–15. IOS Press, 1998.
11. P. Harmon. Exploring BPMS with Free or Open Source Products. BPTrends, Available at www.bptrends.com/publicationfiles/advisor200707311\%2Epdf, July 31 2007. Last accessed 27 Sep 07.
12. S. Jablonski and C. Bussler. *Workflow Management: Modeling Concepts, Architecture and Implementation*. Thomson Computer Press, London, UK, 1996.
13. Java-source.net. Open Source Workflow Engines in Java. Available at java-source.net/open-source/workflow-engines. Last accessed 27 Sep 07.
14. JBoss. JBoss jBPM 3.1 Workflow and BPM made practical. Available at docs.jboss.com/jbpm/v3.1/userguide/en/html/taskmanagement.html. Last accessed 27 Sep 07.
15. JBoss. JBoss jBPM 3.1 Workflow and BPM made practical, Chapter 9.6 Superstates. Available at docs.jboss.com/jbpm/v3.1/userguide/en/html/processmodelling.html. Last accessed 18 Dec 07.
16. JBoss. JBoss jBPM website. Available at www.jboss.com/products/jbpm. Last accessed 27 Sep 07.
17. JBoss jBPM forum. Data trigger task end. Available at www.jboss.com/index.html?module=bb&op=viewtopic&t=102283. Last accessed 27 Sep 07.
18. JBoss jBPM forum. Required variables. Available at www.jboss.com/index.html?module=bb&op=viewtopic&t=118219. Last accessed 27 Sep 07.
19. B. Kiepuszewski, A.H.M ter Hofstede, and C. Bussler. On Structured Workflow Modelling. In B. Wangler and L. Bergman, editors, *Proc. of the 12th Int. Conf. on Advanced Information Systems Engineering (CAiSE00)*, volume 1789 of LNCS, pages 431–445. Springer, 2000.
20. J. Krogstie. *Evaluating UML using a generic quality framework*, pages 1–22. IGI Publishing, Hershey, PA, USA, 2003.
21. J. Krogstie and A. Sølvsberg. Information System Engineering: Conceptual Modelling in a quality Perspective. Kompenium-forlaget, Trondheim, Norway, 2003.
22. O. I. Lindland, G. Sindre, and A. Sølvsberg. Understanding quality in conceptual modeling. *IEEE Software*, 11(2):42–49, 1994.
23. A. Madl et al.. Together Workflow Editor - User Guide. Together Teamlosungen EDV-Dienstleistungen GmbH, Available at <http://www.together.at/together/zznocms/twe/twedoc/twe.html>, 2007. Last accessed 05 Dec 07.
24. Manageability. Open Source Workflow Engines Written in Java. Available at www.manageability.org/blog/stuff/workflow_in_java. Last accessed 27 Sep 07.
25. J. Mettraux. The OpenWFE Manual: Open source Workflow Environment. Available at www.openwfe.org/manual/index.html. Last accessed 27 Sep 07.
26. A. L. Opdahl and B. Henderson-Sellers. Ontological Evaluation of the UML Using the BungeWandWeber Model. *Journal Software and Systems Modeling*, 1(1):43–67, 2002.
27. OpenWFE. OpenWFE website. Available at www.openwfe.org. Last accessed 27 Sep 07.
28. OpenWFE. OpenWFEru, Time Expressions, sleep. Available at http://www.openwfe.org/manual/ch06s08.html#expression_sleep. Last accessed 27 Sep 07.
29. OpenWFE. OpenWFEru, Workflow Patterns, Structural Patterns. Available at http://openwferu.rubyforge.org/patterns.html#pat_c. Last accessed 27 Sep 07.
30. F. Puhlmann and M. Weske. Using the Pi-Calculus for Formalizing Workflow Patterns. In W.M.P. van der Aalst et al., editor, *Proceedings from international conference on Business Process Management - BPM 2005*, volume 3649 of LNCS, pages 153–168. Springer-Verlag, 2005.
31. J.C. Recker, M. Indulska, M. Rosemann, and P. Green. How Good is BPMN Really? Insights from Theory and Practice. In Jan Ljungberg and Magnus Andersson, editors, *Proceedings 14th European Conference on Information Systems*, 2006.
32. M. Rosemann, J. Recker, M. Indulska, and P. Green. A Study of the Evolution of the Representational Capabilities of Process Modeling Grammars. In E. Dubois and K. Pohl, editors, *CAiSE*, volume 4001 of LNCS, pages 447–461. Springer, 2006.
33. N. Russell. *Foundations of Process-Aware Information Systems*. PhD Thesis, Queensland University of Technology, Brisbane, Australia, December 2007. Available at <http://www.yawlfoundation.org/documents/RussellThesisFinal.pdf>.
34. N. Russell, W.M.P. van der Aalst, A.H.M. ter Hofstede, and D. Edmond. Workflow Resource Patterns: Identification, Representation and Tool Support. In O. Pastor and J. Falcão e Cunha, editors, *Proceedings of the 17th Conference on Advanced Information Systems Engineering (CAiSE'05)*, volume 3520 of *Lecture Notes in Computer Science*, pages 216–232, Porto, Portugal, 2005. Springer.
35. N. Russell, A.H.M. ter Hofstede, W.M.P. van der Aalst, and N. Mulyar. Workflow Control-Flow Patterns: A Revised View. BPMcenter.org, BPM Center Report BPM-06-22, 2006. Available at <http://is.tm.tue.nl/staff/wvdaalst/BPMcenter/reports/2006/BPM-06-22.pdf>.
36. N. Russell, A.H.M. ter Hofstede, D. Edmond, and W.M.P. van der Aalst. Workflow Data Patterns: Identification, Representation and Tool Support. In L. Delcambre et al, editor, *Proc. of the 24th Int. Conf. on Conceptual Modeling (ER 2005)*, volume 3716 of LNCS, pages 353–368. Springer, 2005.
37. T. Wahl and G. Sindre. An Analytical Evaluation of BPMN Using a Semiotic Quality Framework. In T.A. Halpin, K. Siau, and J. Krogstie, editors, *Proceedings of the Workshop on Evaluating Modeling Methods for Systems Analysis and Design (EMMSAD'05), held in conjunction with the 17th Conference on Advanced Information Systems (CAiSE'05)*, pages 533–544. FEUP, Porto, Portugal, 2005.
38. Y. Wand and R. Weber. On the Ontological Expressiveness of Information Systems Analysis and Design Grammars. *Journal of Information Systems*, 3:217–237, 1993.

39. WfMC. Workflow Process Definition Interface - XML Process Definition Language, Version 1.0. Accessed Dec 2007 from http://www.wfmc.org/standards/docs/TC-1025_10_xpdl_102502.pdf, 2002.
40. P. Wohed, W.M.P. van der Aalst, M. Dumas, A.H.M. ter Hofstede, and N. Russell. Pattern-based Analysis of UML Activity Diagrams. BETA Working Paper Series, WP 129, Eindhoven University of Technology, Eindhoven, The Netherlands, 2004. <http://www.bpm.fit.qut.edu.au/projects/babel/docs/p242.pdf>.
41. P. Wohed, W.M.P. van der Aalst, M. Dumas, A.H.M. ter Hofstede, and N. Russell. Pattern-Based Analysis of the Control-Flow Perspective of UML Activity Diagrams. In L. Delcambre et al, editor, *Proc. of 24th Int. Conf. on Conceptual Modeling (ER05)*, volume 3716 of *LNCS*, pages 63–78. Springer Verlag, 2005.
42. P. Wohed, W.M.P. van der Aalst, M. Dumas, A.H.M. ter Hofstede, and N. Russell. On the Suitability of BPMN for Business Process Modelling. In Schahram Dustdar, José Luiz Fiadeiro, and Amit P. Sheth, editors, *Business Process Management*, volume 4102 of *Lecture Notes in Computer Science*, pages 161–176. Springer, 2006.
43. Workflow Patterns Initiative. Workflow Patterns - homepage. Available at www.workflowpatterns.com. Last accessed 27 Sep 07.
44. B. Wyssusek. On Ontological Foundations of Conceptual Modelling. *Scandinavian Journal of Information Systems*, 18(1):63–80, 2006.