

Application of Business Process Execution Language to Scientific Workflows

Asif Akram David Meredith Rob Allan

CCLRC e-Science Centre, CCLRC Daresbury Laboratory, Warrington, UK WA4 4AD
Email: {a.akram; d.j.meredith; r.j.allan}@dl.ac.uk

Abstract: This paper investigates the use of the Business Process Execution Language for Web services (BPEL4WS/BPEL) for managing scientific workflows. The complexity, unpredictability and inter-dependency of the components in a scientific workflow often demand great flexibility in a workflow-language in order to support; 1) exception handling, 2) recovery from uncertain situations, 3) user interactions to facilitate interactive steering and monitoring, 4) dynamism to adapt to the changing environment, 5) compensation handling to reverse the effects of previous activities that have been abandoned, and 6) flexibility to support dynamic selection of services at runtime and to support changing data requirements. These requirements are illustrated with examples taken from a real scientific workflow; the e-HTPX project for high throughput protein crystallography. In the context of addressing these requirements, the features of the BPEL4WS specification are discussed, which is widely regarded as the de-facto standard for orchestrating Web services for business workflows. An investigation and proposal for the use of the Web services Invocation Framework (WSIF) [38] to extend BPEL is also provided. In summary, by extending BPEL where necessary (in conjunction with standard BPEL features), workflows can be successfully adapted to fulfill all the requirements of complex scientific workflows.

Keywords: BPEL, scientific workflow, SOA, Web services, Web Services Invocation Framework.

1. Introduction

Modular software is designed to avoid failures in large systems, especially where there are complex user requirements. However, even in a modular system, inter-dependency of sub-components can make it difficult to meet changing user requirements without re-engineering code. A SOA is an architectural style whose goal is to achieve loose coupling among interacting software agents (services and clients). Agents can be composed in different ways to meet different requirements without having to re-write code. A SOA achieves loose coupling by employing two architectural constraints: 1) a small set of well-defined interfaces to all participating software agents and, 2) ensuring the interfaces are universally available to all providers and consumers. In simple terms a service is a function that is self-contained and immune to the context or state of other services. These services can communicate with each other, either through explicit messages (which are descriptive rather than instructive), or by a number of 'master' services that

coordinate or aggregate activities together. This is complementary to Object Oriented Programming. A SOA can also define a system that allows the on-demand binding of online resources as independent services. This can provide a more flexible, loosely coupled set of resources compared to traditional system architectures with the possibility to optimize for performance or functionality.

The role of a SOA it is not only limited to making software components available for on-demand consumption. SOA principles also influence the business logic of services by encouraging good design, i.e., by promoting modular and dispersed components that can be separately developed and maintained. In doing this, a SOA can help reduce IT infrastructure overhead. This also allows a greater emphasis to be placed on the core value of a service. In recent years, Web services have been established as a popular 'connection technology' for implementing a SOA. The interface required for Web services is described in a WSDL file (Web service Description Language [10]). Services exposed as Web services (such as those wrapping and exposing legacy codes) can be integrated into complex workflows that may span multiple domains and organizations.

The situation however, becomes even more complex when combining multiple geographically dispersed Web services together to form a single composite service in the form of a workflow (or sub-workflow). In this scenario, a composition language is (usually) required. There is growing interest in the use of Web services for scientific, parallel and distributed computing [5]. This interest is fuelled by the availability of different Web service toolkits that are tailored to Grid Computing, such as Globus Toolkit (GT) 3.x, GT4.x, Open Middleware Infrastructure Institute (OMII) and IBM Emerging Technologies Toolkit (ETK). Scientific applications and services that are deployed and executed using Service Oriented Grid Computing infrastructures on top of Web services are called Grid Services.

Integration of two or more services into a more complex workflow that is managed by a workflow engine is called 'Service Orchestration.' In this scenario, the workflow engine orchestrates the interactions between the partner services by acting as a broker or 'middle-man.' Service orchestration differs from 'Service Choreography' [22]. Choreography occurs in peer-to-peer style workflows where communications occur directly between partners (i.e., involving no orchestration by an intermediate workflow engine). Consequently, a set of agreed rules are required to define the roles of each peer in the collaboration and to determine which messages should be exchanged at particular

1 stages. This process is known as ‘Service choreography.’

2 With the popularity of Web services, several composition
3 languages used for service orchestration and choreography
4 have emerged in the last few years, such as BPML [36],
5 XLANG [26], WSFL [20], WSCDL [18], BPEL4WS [2],
6 GSFL [28] and WS Coordination [8]. BPEL4WS is generally
7 regarded as the de-facto standard for composing and
8 orchestrating workflows from Web services and will be
9 discussed in this paper. It is supported by commercial
10 vendors and by the open-source community who have
11 contributed good quality BPEL4WS editors and efficient
12 BPEL execution engines.

13 In Section 2, the requirements of scientific workflows are
14 discussed (illustrated with examples from a real scientific
15 process in e-HTPX). Section 3 outlines the main BPEL
16 architecture and built-in features provided within the
17 BPEL4WS. Section 4 provides a more in-depth analysis of
18 the BPEL4WS implementation in the context of addressing
19 the requirements of scientific workflows. Section 5 outlines
20 some of the limitations of BBEL.

22 2. Scientific Workflows

23
24 According to Kammer [17], “The unpredictability of business
25 processes requires that workflow systems support exception
26 handling with the ability to dynamically adapt to the
27 changing environment”. Traditional approaches to handling
28 this problem have fallen short, providing little support for
29 change, particularly once the process has begun execution.
30 Often, this unpredictability is even more apparent in fragile
31 or long-lasting scientific workflows where research evolution
32 often results in constantly changing requirements.
33 Consequently, this can directly impact; a) the changing
34 scientific functionality of a service, b) organizational
35 constraints and policies, and c) dependencies on underlying
36 data and legacy codes. As a result, a lack of interface
37 compatibility with semantically equivalent and competing
38 services can emerge. Dynamic scientific workflows often
39 need to be complemented with execution support at runtime,
40 such as dynamic scheduling, dynamic resource binding, and
41 infrastructure reconfiguration. These constraints demand
42 flexibility in a workflow language. Traditionally, workflows
43 were not solely designed to call different Web services in a
44 static or dynamic manner; rather, they are supposed to
45 provide; a) an execution environment for exception handling,
46 b) a compensation mechanism for partially completed or
47 failed jobs, c) dynamical invocation of services (adaptive in
48 nature to accommodate inconsistencies among services), and
49 d) graceful recovery from un-expected circumstances.

50 Scientific workflows can be more demanding in both their
51 nature and in their requirements than business processes. In
52 the following section, some of the requirements that may be
53 encountered in scientific workflows are outlined (some of
54 which are also common to business processes). The
55 requirements are hierarchical and co-dependent and include:
56 1) A Modular design; 2) Effective Exception handling (this
57 depends on a modular design); 3) Effective Compensation
58 (which should be possible from the exception handling code);
59 4) Workflow flexibility and adaptivity; 5) Support for long
60 running transactions and conversations with partner services;
61 and 6) Support for the management and monitoring of
62 workflow processes. Before discussing each of the
63

requirements, the e-HTPX project is briefly introduced.

2.1 Sample Application

The e-HTPX project [1, 31] is a distributed computing infrastructure designed for structural biologists to remotely plan, initiate and monitor protein crystallography experiments from their desktop computer. A number of key services for e-HTPX are being developed by leading UK e-Science, synchrotron radiation and protein manufacture laboratories (SRS Daresbury Laboratory [35]; The European Synchrotron Radiation Facility [32]; The Oxford Protein Production Facility [33]; York Structural Biology Laboratory [40]). The services being developed for the project define the ‘e-HTPX Pipeline,’ which is illustrated in Fig. 1, where:

Stage 1 – Protein Production

Stage 2 – Crystallization

Stage 3 – Data Collection

Stage 4 – Phasing (data processing)

Stage 5 – Solution of digital protein structure model

Stage 6 – Submission of protein model into public database.

The pipeline covers protein crystallization, delivery of a protein crystal to a synchrotron radiation facility, data collection using X-ray diffraction methods, and HPC data processing services. Remote access to these services is implemented by a collection of Web services. The data collection and data processing stages of the e-HTPX pipeline are particularly well suited to the application of workflow technologies. The following sections discuss some of the generic scientific workflow requirements. Real examples are illustrated from e-HTPX.

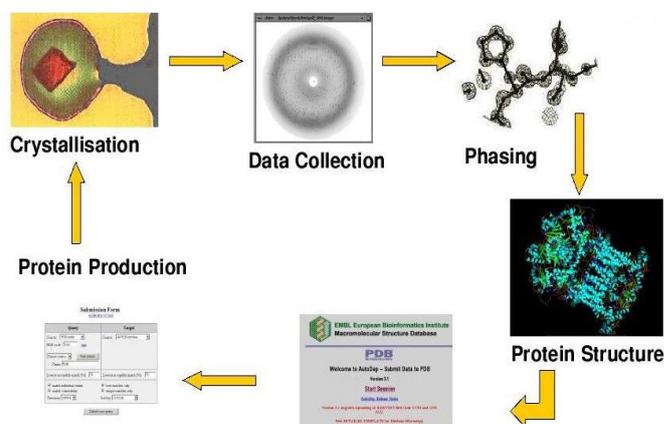


Fig. 1. e-HTPX Protein crystallography workflow.

2.2 Requirements for Scientific Workflows

2.2.1 Modular Design

The potential for complexity in a scientific workflow usually demands a thorough investigation of all the activities involved, including a detailed examination of inter-related dependencies and data. Failure of even one activity can hinder the successful execution of the whole workflow. Consequently, it is important to design modular workflows by grouping similar or related activities together and ensuring no side effects. Nutt [21] has discussed the modularization of workflows in detail in the context of: 1) Models of procedures, 2) Workflow modeling languages, 3)

1 Computation languages for modeling, 4) Workflow modeling
2 systems and, 5) Issues in workflow models. These modules
3 operate as ‘black boxes’ in the overall workflow, each having
4 their own variables, computational logic, dependency
5 constraints, security model, and event handlers. Workflow
6 systems aim to support the organization of work and require
7 that work activities are categorized and their relationships
8 described in some form of process model.

9 The components in a large scientific workflow are
10 arranged and connected around precise scientific
11 requirements, not on programming logistics. A large
12 repository of inter-connectable and re-usable components
13 opens a wide range of possibilities for processing and relating
14 diverse scientific data (the ‘Science-Factory’ concept) [34].
15 By modularizing workflows, module components can be
16 made sufficiently scalable and reliable to serve as repeatable
17 units. This helps automate the process of integrating diverse
18 resources and scientific services across faculties (and even
19 institutes). A workflow typically has a flexible n-tier design
20 that facilitates the independent upgrade of individual
21 components and tiers of the system. This helps ensure
22 workflow sustainability, especially for future initiatives [14].

23 With regard to e-HTPX, modularity has been integrated
24 from the outset and has been achieved by grouping inter-
25 dependent processes in the e-HTPX pipeline into distinct
26 Web service ‘Categories.’ Three Web service categories have
27 been devised; 1) Protein Crystallization, 2) Data Collection
28 (these services are tied to a particular synchrotron radiation
29 facility, e.g., services involved with shipping crystals to a
30 synchrotron and synchrotron specific data collection services),
31 3) Post-Data Collection Processing. The Web service
32 categories are effectively standalone and can operate in
33 isolation from each other. Each category is currently hosted
34 at a different institution (crystallization services hosted at
35 OPFF, data collection services hosted at specific synchrotron
36 – SRS and ESRF, and HPC data processing hosted at
37 Daresbury Laboratory). Modularizing related e-HTPX
38 processes into categories offers a significant advantage to the
39 user in that it is possible to consume services only for a single
40 category, rather than being tied to the entire e-HTPX pipeline.
41 In addition, modularization allows the service provider to
42 implement fine-grained authentication mechanisms that are
43 specific to the particular hosting institution (in conjunction to
44 the e-HTPX wide authentication mechanism). In doing this,
45 the institution can control access to their own resources (e.g.,
46 by facilitating special guest users/ home laboratory users).

48 2.2.2 Exception Handling

49
50 Even the best-planned workflows will most probably
51 encounter exceptions during their execution. Managing the
52 error involves handling the exception and rejoining the
53 normal workflow so that it may proceed correctly. The
54 occurrence of exceptions is a fundamental part of
55 organizational processes [25]. Exceptions can result from a
56 number of different sources such as inconsistent data [12],
57 divergence of tasks from the underlying workflow model [6],
58 unexpected contingencies [23] and un-modeled changes in
59 the environment [27]. A workflow model should provide
60 exception management with the ability to detect the
61 occurrence of exceptions, or even the possibility of
62 anticipating an exception thus improving the ability to react

and recover within the minimum amount of time.

A well-designed workflow provides exception handling at different levels throughout its execution. Exception handling can be categorized into three distinct types; 1) Global Exception Handlers, which are defined at the core workflow level, 2) Scoped Exception Handlers, which are defined at the modular level and are related to a set of similar and/ or inter-dependent activities, and 3) Inline Exception Handlers, which are tightly coupled to specific activities or tasks and are limited in their scope and re-usability. Bartherness and Wainer [4] have classified exceptions into two main categories; 1) events modeled by the system which deviate from the normal flow are called ‘expected exceptions’ (described as ‘variations’ by Saastamoinen [23]) and, 2) un-modeled events which cannot be handled by the actual activity/ task are called ‘unexpected exceptions.’ Global exception handlers are generic in nature and are used mainly to handle unexpected exceptions. Inline exception handlers are specific to the problem and process and therefore handle expected exceptions/ variations. Scoped exception handlers are hybrid in nature; they are generic but relate only to a set of activities and tasks within the module. Kammer [17] suggests that artifact recovery and exception handling should be localized if permissible rather than propagating the exception up through the exception hierarchy.

For e-HTPX, the modularity of the services allows each service category to handle its own scoped exception types (i.e., exceptions that are related to that specific service category). The most problematic exceptions for e-HTPX are those of the ‘Unexpected Global’ type. Unexpected errors occur mainly due to the physicality of the tasks that are invoked by a service and their unexpected failure (usually related to physical problems). This type of exception can occur during data collection if a sample-changing robot fails to operate correctly (causing a “robot collision”), or a crystal sample is damaged in transit to the synchrotron. Because these errors are physical in nature, they are difficult to forecast by the service and handle elegantly in the service software. Furthermore, because these types of error cannot be anticipated by the workflow-engine at design time, workflows are often poorly equipped to handle these types of exception. In such cases, it may be necessary to restart the workflow from a specific breakpoint and roll back to a point where all prior activities were successfully completed.

2.2.3 Compensation Mechanism

The goal of compensation is to reverse the effects of previous activities that have been abandoned. A workflow should call a compensating operation from an exception handler if exceptions are thrown (expected or unexpected - it must be noted however, that not every exception requires compensation). An effective compensating mechanism is more dependent on the participating services than on the workflow itself. Nevertheless, all situations requiring compensation should still be modeled for the workflow at design time in conjunction with exception handling. Compensation can be defined as ‘undoing steps in the business process that have already completed successfully,’ and is one of the most important concepts in business processes. Compensation behavior must be explicitly defined by the services to ensure reversal of a specific operation

1 when required. Workflows should be able to call those
2 compensating operations in order to release the specified
3 resource. Reversal is however not always possible.

4 Compensation handlers can be more important to business
5 processes where there are many different stakeholders with
6 financial interests compared to scientific workflows.
7 Nevertheless, scientific workflows often require
8 compensation, especially when considering long running
9 processes that make use of scarce resources (e.g., compute
10 resources, synchrotron radiation resources). For e-HTPX,
11 compensation is important for the data-processing stages of
12 the e-HTPX pipeline (e.g., the Bulk Molecular Replacement
13 service which utilizes HPC resources). Following submission
14 of a compute job, the remote user can monitor the results as
15 they are being generated. During the early stages of the
16 compute job, the user may wish to cancel because poor and/
17 or un-expected results are being generated (a function of data
18 quality). In this scenario, the compute resources can be
19 released as a result of compensation functionality.

20 21 **2.2.4 Flexibility and Adaptivity**

22
23 Scientific applications can often be highly dynamic and
24 subject to constant evolution. Consequently redesign and
25 optimization of an existing process often becomes
26 indispensable in a rapidly changing environment [15]. During
27 these redesigns, workflows have to be adjusted repeatedly. In
28 contrast, workflows modeled for traditional systems are often
29 static and are not designed to cope with a rapidly and
30 dramatically changing environment. In order to accommodate
31 rapidly changing and dynamic environments, workflows have
32 to be flexible and adaptive.

33 A workflow is flexible if it is subject to modification and
34 change, usually in terms of the data types involved and
35 service endpoint location. Flexibility is often required during
36 the early stages of a workflow implementation when the
37 services are immature. With regard to flexibility in data, this
38 can be achieved by applying generic 'base' or 'loose' data
39 types that may be extended at a later date to accommodate
40 new or more tightly constrained fields (in contrast to data
41 types that are 'narrow' or 'tightly' constrained to the
42 operation at the outset). With regard to service endpoint
43 location, most workflows integrate services having static
44 locations; they refer to single external process selected by the
45 developer at design time. This approach is appropriate for
46 highly constrained systems. However, in more flexible
47 workflows, the integration of additional and/ or re-located
48 services has to be accounted for. Consequently, potential
49 callouts and logic for selecting which partner service to
50 integrate has to be built into the business process. Flexibility
51 in the location of services can be achieved through WS-
52 Addressing [7], and by dynamically binding to services from
53 a service selection. Service endpoint location is usually
54 specified in an external configuration file or by user
55 interaction. The only assumption is that all the services have
56 'similar' interfaces, which may not always be the case.
57 Nevertheless, it may be possible to specify a different
58 workflow path for different services at the cost of increasing
59 workflow complexity (the technical details will be discussed
60 in next section).

61 Dynamic service selection from a service pool makes it
62 possible to schedule their execution dynamically at runtime.

For computationally extensive tasks, sub-tasks can be
allocated to different services executing in parallel. Flexible
workflow modeling addresses the following issues [21]: 1)
Too prescriptive. Often a workflow program over constrains
the solution by providing coordination or computation
constraints; 2) Too vague. The work specification does not
provide enough information to the workflow engine to enable
them to perform the desired work by invoking participating
services in required manner; 3) Exceptions. Exceptions occur
because the procedure is an incomplete specification of the
intended processing, because the intended procedure itself is
incomplete, or because the procedure is completely
inappropriate for the work [3]; 4) Changes. A workflow
procedure is specified, but the model must change frequently
due to changes in the nature of application/ experiments or
the computing environment; 5) Workflow as fiction. The
workflow specification assumes everything perfect and works
'by the book.' Assumptions according to an arbitrary
operational model are always too rigid and inefficient; 6)
Loss of Informal Information. When a manual procedure is
encoded as a workflow procedure, the encoding will almost
certainly fail to encode all the information flow associated
with manual processing; 7) Loss of informal processing. The
logical extension of ineffectiveness due to informal
information loss is ineffectiveness due to the loss of informal
work. In summary, modeling of workflows to accommodate
flexibility provides both efficient and adaptive workflows
since flexible workflows can adapt according to the
requirements.

Workflow flexibility is certainly a requirement for e-
HTPX service interaction. This is because the services often
undergo upgrades/ modifications in response to the advancing
nature of the research that underpins the service. A workflow
must be capable of selecting and consuming preferred
services from a service-pool at run-time. This is especially
relevant to the data collection services, which are duplicated
across multiple synchrotron facilities. Data flexibility is also
required. The e-HTPX Web services adopt a comprehensive
messaging model implemented in XML Schema. It is
anticipated that if the e-HTPX services are widely adopted at
multiple synchrotron facilities, the original 'parent' e-HTPX
XSD may be further refined to suit the particular needs of
that service implementation.

The adaptivity of a workflow is closely related to
flexibility. Workflow adaptivity can be static or dynamic
depending on whether changes and necessary adjustments
can be accommodated at design time or runtime respectively
(adaptivity and flexibility are closely related). The simplest
means in which workflow users adapt to change is by
responding to exceptions. These exceptions allow current
workflow implementations to break a pre-defined process in
an attempt to solve the problem. Workflow adaptivity is
further accommodated-for by adding support for extensibility.
For workflows, this often involves defining 'empty' activities
that act as placeholders for future extensions and adaptations.
These are also requirements of agile approaches to
developing distributed systems. The Workflow Management
Coalition (WfMC) [36], as part of its standards effort, has
created Interface Four, a mechanism that governs the hand-
off of work to other processes and messaging systems in an
attempt to mediate how exceptions will be handled.
According to the literature [36], the requirements of dynamic

1 adaptive workflows include; 1) Contingency Management
2 and Hand-off, which provide mechanisms for dealing with
3 and recovering from expected and unexpected divergence
4 from the intended process; 2) Partial Execution, which
5 supports creating and executing processes and process
6 fragments (modules) as they are needed, rather than requiring
7 the entire process be rigidly specified ahead of time; 3)
8 Dynamic Behavior in terms of both execution model and
9 object behaviors provides flexibility to modify workflow
10 paths and executed behaviors at run-time independent of
11 object data and, 4) Reflexivity, which allows a workflow
12 component to programmatically examine, analyze, create and
13 manipulate its own process and data.

15 2.2.5 Support for Long Running Processes / Transactions

16 Scientific workflows can be long running, often requiring
17 days (or even months) to complete. In e-HTPX, both
18 shipping crystals to a synchrotron facility and performing
19 data collection experiments occurs over a period of 2~3 days.
20 In the context of workflows, long running workflow
21 processes / instances are managed by the workflow engine.
22 Different workflow vendors provide workflow-engine
23 specific support to persist workflow state during power-outs,
24 hardware failure and software malfunction. For example,
25 Oracle BPEL Process Manager has a notion of 'hotspots.' A
26 hotspot is a marked stage in the workflow when the state of
27 the workflow instance is persisted.

28 In order to support long running conversations and
29 transactions between partners, both the workflow instance
30 and partner services have to maintain conversational state in
31 the message exchange. For workflows, this involves the
32 addition of process identifiers that are embedded and
33 exchanged between partners during a conversation (process
34 identifiers are termed 'correlation properties' in the context
35 of workflows). This concept is similar to maintaining state in
36 a http session.

39 2.2.6 Management of Workflow

40 Workflow management involves the addition of breakpoints,
41 addition of steering capabilities, and provision of monitoring
42 mechanisms. Workflow management is complicated and, in
43 practice, workflows can be managed in many different ways.
44 Effective management depends on the underlying structure of
45 the workflow engine and on the environment in which the
46 workflow executes. Nevertheless, there should still be
47 support for managing workflows in terms of specifications.
48 Workflow management is crucial for compute-intensive and
49 long-running workflows and services. The following points
50 address the main workflow management concepts:

51 *Breakpoint/ Checkpoint:* A breakpoint is an arbitrary location
52 in the workflow that divides different stages that are
53 crucial to the overall functioning of the workflow.
54 Persistence of state is useful at breakpoints for a number
55 of reasons; to enhance fault tolerance, to allow re-
56 execution with different experimental data sets/ bindings,
57 and to facilitate inspection/ editing of intermediate data
58 values.

59 *Steering:* Steering further generalizes breakpoints to
60 accommodate rollback, restart, and modification to the
61 processes at runtime. A good management mechanism

should permit the re-execution of a disrupted workflow,
permit workflow re-starts at arbitrary midpoints, and
allow rollback to a previous internal state (invoked by
computer or user). Kaiser [16] summarizes some of the
main steering concepts as: reset, re-schedule, restart,
undo, complete, abort, recover, ignore or jump operations
in the process interpreter.

Monitoring: Workflow monitoring is crucial for optimization
and analysis of internal state, especially with regard to
data flow, inspection of values, re-scheduling of tasks
and re-ordering of tasks for parallel execution.

3. BPEL4WS

The goal of the BPEL4WS specification is to provide a
notation for specifying business process behavior based on
Web services. BPEL4WS builds on top of the Web service
stack and is defined entirely in XML, combining various
technologies including WS-Reliable Messaging, WS-
Addressing [7], WS-Coordination [8] and WS-Transactions
[11] (BPEL however, is not strictly limited to only these
technologies). BPEL4WS also provides a rich vocabulary for
the description of business processes supporting two different
types of process declaration (Abstract and Executable
Processes).

3.1 Abstract Process

An 'Abstract Process' specifies the message(s) exchanged
between participating services without exposing the internal
details of the process. The objective of an abstract business
process in BPEL is therefore to specify only the externally
observable aspects of process behavior, also known as 'public
process behavior' without revealing their internal
implementation. Abstract business processes are not
executable and must be explicitly declared as 'abstract.' An
abstract process usually represents an executable process.
Abstract processes only provide the complete description of
external behavior that is relevant to a partner or several
partners involved in the workflow. There are two good
reasons to separate the public aspects of business process
behavior from internal or private aspects. 1) To hide the
internal complexity and decision making of the businesses
process from partners, 2) Separating public from private
processes provides the freedom to change private aspects of
the process implementation without affecting the observable
behavior. Observable behavior must clearly be described in a
platform independent manner and captures behavioral aspects
that may have cross-enterprise significance.

Abstract processes have restricted access to the syntax and
semantics available to executable BPEL processes. Although
BPEL abstract processes are not executable, they can
indirectly impose behavioral regulations upon private
processes executed by the BPEL orchestration server. For
example, a BPEL abstract process may impose sequencing
rules for invoked operations within the BPEL private process.
These rules can define inter-dependencies between
executable operations. The BPEL orchestration server should
throw an exception when any attempt is made to run these
operations out of order.

In summary, systems integration requires more than the
ability to conduct interactions by using standard protocols.

1 The full potential of Web services as an integration platform
2 will be achieved only when applications and business
3 processes are able to integrate their complex interactions
4 through standard integration models. The abstract process
5 provides this missing process integration model, which can
6 be tailored and customized in the executable process. Often
7 however, the use of abstract processes is overlooked when
8 designing workflows with BPEL.

10 3.2 Executable Process

11
12 An 'Executable Process' allows specification of the exact
13 details of an abstract process workflow. Executable processes
14 extend abstract processes and different executable processes
15 can extend the same abstract process. The BPEL specification
16 defines various building blocks, known as 'activities.' These
17 activities can be used to specify an executable process's
18 business logic to any level of complexity. The use of these
19 activities can be further controlled through corresponding
20 attributes, which gives the developer more flexibility. These
21 activities can be grouped into different categories based on
22 their usage (Basic, Structured and Special):

24 3.2.1 Basic Activities

25
26 *invoke*: Calls a Web service; which can be either one-way or
27 request-response.
28 *receive*: Receives an incoming Web service call.
29 *reply*: Sends a response to a received Web service call.
30 *assign*: Used for copying and manipulating data using XPath.
31 An assign activity has the sub element:
32

- *copy*: which itself wraps two elements 'from' the
33 source and 'to' the sink.

34 *throw*: Used to throw an exception when the process drifts
35 from its designated path.
36 *rethrow*: To forward a fault from inside a fault handler.
37 *empty*: A 'no operation' activity used as an extensibility
38 feature.
39 *wait*: Pauses the flow of a process for a specified period
40 (time-out) or until a particular time has passed (alarm).

42 3.2.2 Structured Activities

43
44 *sequence*: Executes one or more activities sequentially.
45 *flow*: Executes one or more activities in parallel.
46 *switch*: Executes one of several paths based on the value of a
47 condition.
48 *while*: Executes a specified activity as long as a condition is
49 true.
50 *RepeatUntil*: Repeats execution of contained activity until the
51 condition becomes true. (Feature of BPEL 2.0)
52 *ForEach*: Executes contained activity exactly N+1 times
53 where N equals the final value minus the start value.
54 (Feature of BPEL 2.0)
55 *pick*: Waits for a message to arrive; it can time out, raise an
56 alarm or throw an exception.
57 *variables*: Defines global variables used by the process or
58 scoped variables within the scope.

60 3.2.3 Special Activities

61
62 *validate*: To validate XML data stored in variables.
63

scope: Defines a new scope for variables, fault handlers, and
compensation handlers.

compensate: Invokes compensation on an inner scope that
has previously completed.

faultHandlers: Used to catch exceptions thrown by the
'throw' activity during the execution of the process.

terminate: Immediately terminates a business process
instance.

extensionActivity: A wrapper for domain specific language
extension.

exit: To immediately terminate execution of a business
process instance.

3.3 BPEL Process Execution

In BPEL, the main workflow is an executable process. The
BPEL engine can manage the lifecycle of multiple workflow
instances. The engine navigates through the workflow and
invokes required partner Web services. Partner services can
also invoke workflow instances. A Web service can
participate by implementing a role according to a service
'partner type.' This is achieved by referencing a port-type(s)
contained in a common WSDL interface that is supported by
all other partners. A service that implements this port-type
interface can be a partner in the process. The BPEL
specification also provides constructs to handle remote and
long-running transactions (LRT) and exceptions using a
'correlation-set' (also used by other related standards, e.g.,
WS-Coordination and WS-Transaction). The correlation-set
is a token in a Web service message that identifies the
specific workflow instance. A partner service must maintain
the same token in an ensuing request/response message,
which can be used by the workflow engine to route the
message to the correct workflow instance. This enables
features like exception handling and compensation in long
running activities. A concrete example of a correlation-set is
further discussed and illustrated in Section IV. In BPEL, the
<invoke> activity performs or invokes a single operation that
involves the exchange of input and output messages. The
<invoke> and <receive> activity both support correlation-sets.
The dependency between the workflow process and the
orchestrated services is established using partner 'link types.'

Within the BPEL workflow process itself, different
activities pass information among themselves using global
data variables. The BPEL specification had built-in support
for local transactions, exception handling and recursive
decomposition.

4. Using BPEL4WS to Implement Scientific Workflows

The richness of the BPEL specification stems from the range
of different activity types. Only a few of these activities have
been outlined in the previous section. Those presented
comprise the main building blocks of a workflow. In general,
the BPEL vocabulary is tailored more to the requirements of
business processes, which often have different requirements
to scientific workflows. Nevertheless, different BPEL
activities can still be used (sometimes in non-conventional
ways) to address the requirements of scientific workflows. In
this section we provide a more in-depth analysis of the
BPEL4WS implementation in the context of the requirements

of scientific workflows.

4.1 Modular Design

A modular design is the first requirement of a scalable and maintainable workflow. A modular design facilitates grouping similar or inter-related activities together to define complex algorithms. These modules can be atomic in nature fulfilling the requirements of ACID transactions (Atomicity, Consistency, Isolation, and Durability). An ACID transaction either successfully executes the whole module, or if an error occurs, the compensation handlers of the already completed sub-tasks are invoked. The BPEL specification provides variety of structured activities for implementing modularity in a workflow (sequence, flow, scope, while, switch). These structured activities can have any number of sub activities to model the logic of the overall workflow.

A 'scope' activity defines its own local variables, exceptions, compensation mechanisms and event handlers. A scope activity is independent of the main process and is used to model subordinate workflows within the main process. An example of a scope activity that can be defined within a BPEL script is shown in Fig. 2. A scope activity propagates its outcome to other structured activities within the workflow by updating global variables. Global variables are declared at the beginning of the BPEL workflow script and are accessible from any structured and non-structured activity.

```
<scope>
  <variables>
    <!-- Variables definitions local to
scope -->
  </variables>
  ....
  <faultHandlers>
    <!-- Fault handlers local to scope -->
  </faultHandlers>

  <compensationHandler>
    <!-- Compensation handlers discussed
later -->
  </compensationHandler>

  <eventHandlers>
    <!-- Event handlers are discussed later
in this article -->
  </eventHandlers>

  <!-- Activities structured or non-
structured -->
</scope>
```

Fig. 2. Layout of the BPEL4WS scope activity.

A 'sequence' activity allows the user to define a set of activities that will be executed sequentially. These activities could include invocation of partner services, exception handlers, conditional-statements or even empty activities.

The 'flow' activity is used to group different activities or tasks that are executed in parallel. There are different scenarios when the 'flow' activity is helpful, particularly when the outcome of multiple independent tasks is required for next stage and the sequential execution of these independent activities would be less efficient.

A 'while' activity is used to define an iterative activity that is performed until the specified Boolean condition evaluates to false. A 'while' activity is often useful in scientific

workflows involving parameter sweeps, where iterations are executed with different sets of parameters. The Boolean conditional logic of the while activity can involve complex logic if necessary using logical operators (e.g., NOT, AND, OR).

The last structured activity is the 'switch' activity, which expresses conditional behavior. It consists of one or more conditional branches defined by a 'case' activity, followed by an optional 'otherwise' element. The 'case' or 'otherwise' activity wraps the set of inter-related tasks to make a single module.

All the structured activities can have child structured-activities, resulting in extremely flexible workflows. Inter-related tasks can be further grouped together based on their dependency, requirements, and domain or common variables. For the most part however, the best approach to achieve effective modularity is to group sub-tasks based on the variables they share, as this isolates and confines the local variables within controlled boundaries. Sub modules within the module improve the maintenance and scalability of the workflow as they can be easily replaced as required to accommodate change in the business process, or run in parallel to improve the efficiency of the main workflow.

4.2 Exception Handling

The BPEL specification provides a set of in-built exceptions for different situations; 'createInstance' exceptions are thrown when the workflow script cannot be instantiated or loaded, 'initializePartnerRole' exceptions are thrown when partner services cannot be invoked according to the role in the script, 'messageExchange' exceptions are thrown when the request or response message format is not consistent with the WSDL interface and 'queryLanguage' exceptions are thrown when the response or request cannot be queried by sub elements during an assignment activity.

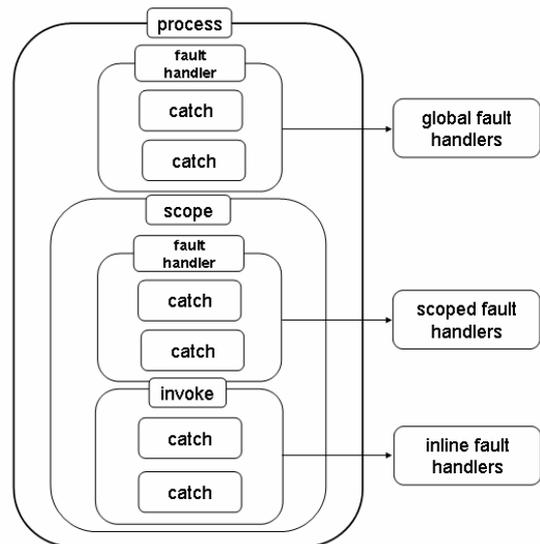


Fig. 3. Different exception handling mechanisms in the BPEL4WS.

BPEL also supports the different fault handler types that were outlined in the previous section, i.e., 'Global,' 'Local' and 'Inline' exception handlers. This provides the BPEL

1 workflow developer with an elegant recovery mechanism.
 2 Fig. 3 summarizes different level of fault handlers supported
 3 by the BPEL4WS specification. Fault handlers specified at
 4 different levels filter exceptions so that only generic
 5 exceptions are propagated up through the exception hierarchy.
 6 Fig. 4 shows an example of a global or process level
 7 exception handler, defined as a child of the root `<process>`
 8 element. Process exceptions include those caused by network
 9 and remote Web service errors. The process level fault
 10 handlers catch all un-caught exceptions thrown by all
 11 activities, including nested and non-nested tasks. The main
 12 role of the process level fault handlers is to handle the un-
 13 caught exceptions propagated from the modules. In the
 14 workflow script there can be several `<catch>` activities that
 15 handle specific faults and `<catchAll>` activities to handle all
 16 other faults. To specify which fault should be handled by a
 17 catch block, the developer must either specify a 'faultName'
 18 or a 'faultVariable.'

```

19 <process ...>
20 .....
21 <faultHandlers>
22   <catch faultName="ns:xxx" >
23     <!-- Perform an activity -->
24   </catch>
25   <catch faultVariable="anyVariable" >
26     <!-- Perform an activity -->
27   </catch>
28   ...
29   <!-- catchAll is optional -->
30   <catchAll>
31     <!-- Perform an activity -->
32   </catchAll>
33 </faultHandlers>
34 </process ...>

```

34 Fig. 4. Process (Global) level fault handlers in the BPEL4WS
 35 script. The process level fault handlers are children of the
 36 `<process>` element.

```

38 <process ...>
39 .....
40 <scope>
41   ....
42   <faultHandlers>
43     <catch faultName="ns:scope" >
44       <!-- Perform an activity -->
45     </catch>
46     <catch faultVariable="anyVariable" >
47       <!-- Perform an activity -->
48     </catch>
49     ...
50     <!-- catchAll is optional -->
51     <catchAll>
52       <!-- Perform an activity -->
53     </catchAll>
54   </faultHandlers>
55   ....
56   <!-- Different activities -->
57 </scope>
58 .....
59 </process ...>

```

56 Fig. 5. Scope level fault handlers in the BPEL4WS script.
 57 The scope level fault handlers are defined within the scope.

59 As discussed earlier, the 'scope' activity can define its own
 60 local variables, exception handlers and event handlers and
 61 nested sub-tasks. These scoped exceptions are declared
 62 within the 'scope' activity as shown in Fig. 5.

Fault handlers attached to a particular 'scope' handle faults
 of all the nested activities defined within the 'scope'. These
 exceptions are more specific to the business logic of the
 module in which they are defined. Faults that are not caught
 within the scope are re-thrown to the enclosing scope, until
 they are caught by an enclosing scope activity or are
 propagated to a global fault handler. Scopes in which faults
 have occurred are considered to have ended abnormally even
 if a fault handler has caught the fault successfully. In this
 case, execution control is passed to the enclosing scope,
 which can either take an alternative approach, such as
 invoking other semantically equivalent partner services, or
 can retry after delay.

The BPEL specification also provides inline exceptions
 that can be handled within the 'invoke' primitive activity. In
 general, an invoke primitive activity should define inline
 exception handlers rather than being wrapping within its own
 scope activity and defining scoped handlers. The syntax of an
 inline fault handler defined within a invoke activity is similar
 to that of a 'faultHandler' block (this is shown in Fig. 6).

```

<invoke ... >
  <catch faultName=" ns:xxx" >
    <!-- Perform an activity -->
  </catch>
  <catch faultVariable=" anyVariable " >
    <!-- Perform an activity -->
  </catch>
  ...
  <catchAll>
    <!-- Perform an activity -->
  </catchAll>
</invoke>

```

Fig. 6. Inline fault handlers in the BPEL4WS script defined
 within the invoke activity.

4.3 Compensation Handler

BPEL supports compensation of errors by defining
 compensation handlers, which are normally specific to a
 particular scope and are often related to 'Long-Running
 Transactions' (LRT). The compensating mechanism
 supported in the BPEL specification is similar to the fault
 handling approach, allowing three types of 'compensate
 activity' (process level, scope level and inline). The
 compensate activity uses the 'scope' attribute to identify
 which scope to compensate. Once any activity (e.g., 'invoke')
 has run successfully, its compensation handler awaits for a
 call to execute. It may be invoked explicitly, arising from a
 running 'compensate' activity, or implicitly due to a fault.
 Normally, compensation handlers are invoked from fault
 handlers. It is important to understand that compensation
 handlers are only called for successfully completed scopes
 and activities. If an error occurs in another scope / activity,
 the compensation handlers of the already completed scopes /
 activities are invoked to perform any necessary compensation
 (e.g., cleanup). It is also important to understand that in a
 BPEL script, a compensation handler can only be invoked by
 another compensation handler or from a fault occurrence that
 is defined in a higher scope or process. In other words, a
 compensation handler in a lower scope can only be invoked

by the occurrence of a higher (parent) scope fault or from a parent scope compensate activity.

Fig. 7 illustrates the compensation mechanism in the BPEL4WS process. The main process has two independent scopes, scopeA and scopeB. Both scopes have one invoke activity. The invokeA activity in the scopeA has inline compensation handler, whereas the invokeB activity in the scopeB has inline fault handler. The main process calls the activity invokeA (1) and after successful execution of invokeA, the process calls invokeB (2). The invokeB activity throws an exception and is caught by the inline fault handler (3). The inline fault handler re-throws the exception, which is caught by the fault handler in the parent activity (4) (in the main process). The global fault handler performs optional logging, notifies the user via server specific features and then calls the compensation handler of invokeA (5). The inline fault handler of invokeB cannot call the compensation handler of invokeA as both are at the same level in the workflow hierarchy.

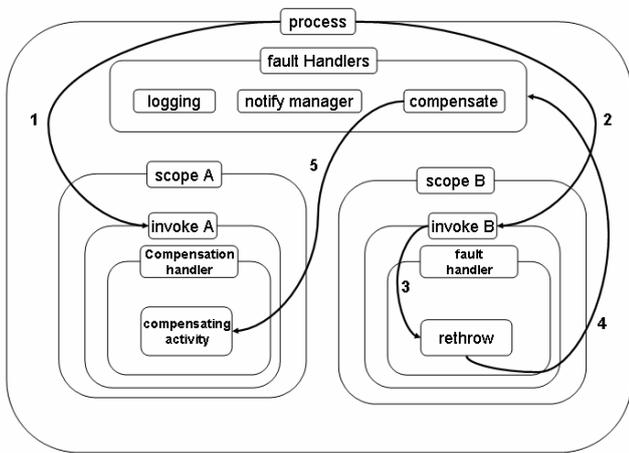


Fig. 7. Compensation mechanism in BPEL4WS. Refer to text for an explanation to the sequence of actions.

The compensation mechanism in a workflow should not be considered a remedy for deviations from the planned process unless the partner services can provide sufficient compensation. It is the responsibility of the partner service to implement these compensation strategies. It must be noted however, that often Web service designers concentrate only on the implementation of the core service with little consideration of compensation. For example, many scientific workflows are often long running, but offer no provision to release unwanted resources. The compensating features provided by BPEL are often under-utilized due to poor service design and over reliance on RPC style Web services. Any sort of compensation is only possible for non-blocking style Web services, implemented using the asynchronous Message Exchange Pattern (MEP).

4.4 Data Flexibility, Workflow Flexibility and Adaptivity

As previously discussed, scientific services are often subject to change, especially with regard to the data types and service-endpoint locations. Data flexibility is supported using 'generic' or 'loose' Web service data types. A 'loosely typed' Web service means the WSDL file does not contain a tightly

constrained XML schema in the type system to fully define the format of the messages. Instead the WSDL defines generic types to encapsulate data within messages (i.e., encoded Strings, CDATA sections, base64 encoded sections, SOAP attachments, xsd:anyType, xsd:any). The main advantage of loose types is that different data formats can be encapsulated without having to update the WSDL interface (e.g., different document types can be transmitted as SOAP attachments or different markup fragments such as XML can be encoded in strings). This allows Web service components to be replaced in a very flexible 'plug-and-play' fashion. It must be noted however, that the loosely typed approach requires extra negotiation between providers and consumers in order to agree on the format of the data that is encapsulated by the generic type. Consequently, an understanding of the WSDL alone is usually not sufficient to consume the service. Fig. 8 shows an example of a loose / generic data type. The generic XML data type 'any' is a wrapper for arbitrary XML fragments. Consequently, this type can be mapped to different implementations, for example, either as a javax.xml.soap.SOAPElement for JAX RPC Web services or as an org.apache.axis.message.MessageElement for Apache Axis Web services. This technique produces a greater workload for the Web service that retrieves the XML fragment from the message, but provides flexibility to replace optional values with default values rather than generating errors.

```
<complexType>
  <sequence> <any maxOccurs="1"/>
</sequence>
</complexType>
```

Fig. 8. The <xsd:any> loose data type.

Workflow flexibility provides dynamic selection of services at runtime from a pool of equivalent services, and can be implemented in different ways. A common practice is to use an exception handler to call an appropriate compensation handler through a switch statement in response to an exception resulting from an unavailable service. The compensation handler may then invoke an alternative service. However, this is a non-conventional approach, and may make the workflow hard to manage and maintain. This is illustrated in Fig. 9.

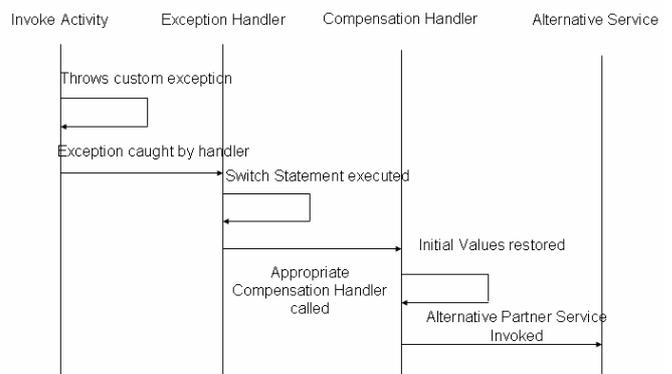


Fig. 9. A non-conventional approach to compensation handling in a flexible workflow. The logic used to call an alternate service may become overly-complex when using custom or non-conventional techniques.

1 A more conventional approach involves dynamically binding
 2 the partner Web services. Dynamic binding allows the
 3 workflow to invoke alternative services through
 4 configuration of inputs at run-time. This approach eliminates
 5 the need for static management of the partner services during
 6 the development phase. Dynamic binding assigns WS-
 7 Addressing [7] compliant endpoint references (EPR) to
 8 partner Web services (known as 'links') during execution of
 9 the workflow. An EPR can encapsulate an endpoint defined
 10 for any protocol (e.g., http, jms, https, smtp). All services
 11 replaceable through a dynamic binding must have the same
 12 (WSDL) interface. Fig. 10 shows an example of a BPEL4WS
 13 script that dynamically binds the partner services. The
 14 variable 'dynamicPartnerReference' can be updated at
 15 runtime based on user input, querying the online status of a
 16 Web service or from a configuration file.

17 Workflow adaptivity is mainly facilitated by the 'empty'
 18 activity, which means 'no operation.' An 'empty' activity can
 19 be used anywhere in the workflow as a 'place-holder' for
 20 extension, and can be replaced by any functional activity as
 21 required. For example, when a partner service provides a new
 22 compensation service, an 'empty' activity can be replaced by
 23 an actual 'invoke' activity to call the new service.

24 Another provision provided by the BPEL specification that
 25 facilitates workflow flexibility is the 'link' activity, which
 26 defines dependencies between different workflow activities.
 27 The 'link' activity defines the 'source' and 'target' activities.
 28 The target activity can be executed only after the source
 29 activity has successfully completed. A single target activity
 30 can have multiple source activities, in which case the target is
 31 executed only after successful completion of all source
 32 activities. Source and target activities can belong to different
 33 'structured' activities, and belong to different parallel
 34 executing blocks. Consequently, links can be used to create
 35 very flexible dependency chains if required. A workflow
 36 engine also manages the links between activities simplifying
 37 development. Links can have associated transition conditions
 38 and predicates that are evaluated at runtime and determine
 39 whether a link is followed or not.

41 4.5 Long Running Transactions / Processes

42 Workflow processes can occur over a long period of time
 43 requiring long running conversations and transactions
 44 between partners. Long running processes require the
 45 workflow engine to manage and correctly identify the specific
 46 workflow instance so that Web service messages can be
 47 routed correctly. To implement long running conversations
 48 and transactions, process identification tokens (identifiers)
 49 have to be embedded and maintained within Web service
 50 messages (either in the SOAP header or in the SOAP body).
 51 The tokens allow the BPEL engine to route the message to
 52 the correct workflow instance. Tokens can be created and
 53 embedded either by the invoking activity or by a responding
 54 activity. If a partner service is a regular Web service (i.e., not
 55 a partner BPEL workflow), then it must at least maintain the
 56 tokens in a message payload in order to maintain the long
 57 running transaction. For the most part, when considering
 58 regular Web services, the invoking workflow is responsible
 59 for creating and embedding the tokens. In BPEL, these tokens
 60 are termed 'correlation properties.' BPEL defines two WSDL
 61 extensions for correlation properties, the first for declaring
 62
 63

the correlation properties in the WSDL file and specifying its
 type (refer to Fig. 11a), and the second for declaring its
 location in a WSDL message (refer to Fig. 11b). Correlation
 properties defined in a WSDL message are termed 'property
 aliases.'

```
<!-- Importing namespace related to WSA -->
xmlns:wsa="http://schemas.xmlsoap.org/ws/2003
/03/addressing"

<!-- Declaring variable of data type EPR -->
<variable name="dynamicPartnerReference"
element="wsa:EndpointReference"/>

<!-- Updating the variable
dynamicPartnerReference -->
.....
<!-- Assigning dynamicPartnerReference to
partnerLink -->
<assign>
  <copy>
    <from variable="
dynamicPartnerReference "/>
    <to partnerLink="DummyService"/>
  </copy>
</assign>
```

Fig. 10. Dynamic binding of the partnerLink Web service in the BPEL4WS script.

```
a)
<bpws:property name="correlationData"
type="xsd:int"/>

b)
<bpws:propertyAlias
messageType="ns1:Callback"
part="jobNumber"
propertyName="ns1:correlationData"/>

<bpws:propertyAlias messageType="ns1:
updateVariable "
part="jobNumber"
propertyName="ns1:correlationData"/>
```

Fig. 11. WSDL extensions defined by the BPEL for the correlation properties.

Multiple correlation properties declared in the WSDL
 interface can be combined to produce a 'correlationSets'
 element that is defined in the BPEL script. A
 'correlationSets' element defines the identification token of a
 long running workflow instance. Fig. 12 shows an example
 of a 'correlationSets' element that is defined in a BPEL script.
 The correlation set is based on the property aliases defined in
 the WSDL.

```
<correlationSets>
<correlationSet name="myCorrelationSet"
properties="nal:correlationData"/>
</correlationSets>
```

Fig. 12. CorrelationSets in the BPEL script.

A correlation set that has been defined previously in a BPEL
 script can be later referenced by a long running activity.
 During process execution, either the invoking workflow or
 the partner service must create an instance of the correlation
 set to successfully correlate subsequent messages with the
 correct workflow instance. In Fig. 13, the invoke activity

1 initiates the correlation set which is reused by the receive
2 activity.

```
3  
4 <invoke name=" submitJob "  
5 partnerLink="jobProcessor"  
6 operation="anyLRP"  
7 inputVariable="serviceRequest ">  
8 <correlations>  
9 <correlation initiate="yes"  
10 set="myCorrelationSet"  
11 pattern="out"/>  
12 </correlations>  
13 </invoke>  
14 <receive name="jobResult"  
15 partnerLink="client"  
16 portType="client:BPPELProcess"  
17 operation="update"  
18 variable="updateVariable"  
19 createInstance="no">  
20 <correlations>  
21 <correlation initiate="no"  
22 set="myCorrelationSet"/>  
23 </correlations>  
24 </receive>
```

Fig. 13. Correlation set initiated in the invoke activity and maintained in the receive activity.

26 If a workflow engine successfully locates an existing
27 workflow instance identified by the correlation set, the
28 message can be routed to the correct instance. If an instance
29 cannot be located, a new instance may be created if the
30 createInstance attribute of the receive activity specifies a
31 value of 'yes' (refer to Fig. 13). If multiple instances qualify,
32 then a runtime exception is raised, indicating an ambiguity
33 caused by a correlation set that is non-specific.

4.6 Management of Workflow

37 The BPEL specification does not provide any direct
38 mechanism to support management and monitoring of a
39 workflow script. This next section outlines and proposes
40 several potential strategies for workflow management and
41 monitoring.

42 The <pick> activity. BPEL provides the <pick> activity,
43 through which the workflow can be made to wait for the
44 occurrence of selected events (manually selected if
45 necessary). Events can be message events, handled with the
46 <onMessage> activity, or chronologic alarm events, which
47 are handled with the <onAlarm> activity. For each event,
48 there is an activity or a set of activities that should be
49 performed. For example, a <pick> activity may specify
50 several <onMessage> elements and several <onAlarm>
51 elements. The <onMessage> activity waits for the messages
52 from the partner services or for the user interaction (user
53 involvement in the workflow is vendor specific). In summary,
54 the <pick> activity is only a limited management facility.

55 Wrap engine-specific classes with Web Services
56 Invocation Framework (WSIF) [38]. Different workflow
57 engines provide their own proprietary management
58 capabilities and user interaction functionalities. These
59 'engine-specific' management capabilities are generally not
60 portable. For example, the Oracle BPEL Process is a freely
61 available implementation of the BPEL specification and
62 provides its own API to interact with the workflow engine at

runtime. To achieve portability, custom and engine-specific
classes can be wrapped and exposed through WSIF, the Web
service Invocation Framework. WSIF exposes Plain Old Java
Objects (POJO), Enterprise Java Beans (EJB) and Java
Connectors (JCA) as Web services (and provides
accompanying WSDL interfaces). In the case of Oracle
BPEL Process, these classes are;

C:\OraBPPELPM_1\integration\orabpel\system\classes.

These classes can be exposed through the use of WSIF,
packaged together as jar files, and placed in the common
directory of a Web application container. Persistence of data
during different stages of workflow is often achieved by
exposing EJB's with WSIF.

Use checkpoints. In addition to wrapping engine-specific
classes with WSIF, it is also possible to (partly) increase
portability by extending engine-specific management
facilities with standard checkpoints in the workflow.

A workflow monitoring web service. It is relatively easy to
describe the state of a workflow in XML, especially since the
BPEL specification is defined in XML. Consequently, it is
easy to create a workflow monitoring Web service that can be
used by workflow clients to monitor or track the execution of
a workflow in a portable way. If a workflow state is available
in XML format then it can be exposed as global variables that
can be queried by the monitoring service. It is also possible to
retrieve a workflow snapshot at run time in XML. This can
be transformed as required using XSLT into user acceptable
formats. However, it must be stated that extracting an XML
snapshot of the current state of the executing workflow
requires some understanding of the underlying workflow
engine. Development of a custom monitoring service can be
especially important when using workflows to coordinate
large parameter sweeps. This is because it is very important
to be able to monitor a large number of workflow instances
and, at minimum, to be able to stop them when execution is
not proceeding correctly. If workflow monitoring can be
integrated with existing monitoring capabilities, it will help
to give a more complete picture of workflow execution
progress at any given moment [24].

A workflow history web service. A monitoring service can
be enhanced by a workflow execution history service, which
can be very helpful when attempting to determine what
previously occurred to a workflow instance.

5. Limitations of BPEL4WS

The following bullets (1 to 5) discuss the limitations of the
BPEL4WS specification. Different workflow engine vendors
have attempted to address these limitations by augmenting
their products with additional features. These features are
implemented either through specification extension, or by
providing custom APIs. Luckily, most of the workflow
engines are either open source or free to use, allowing the
workflow designer to select an appropriate product based on
their requirements.

Reusability of primitive and structured activities in BPEL
is limited. For example, it is not possible to re-execute an
activity that is defined earlier in the script by referring to it
later. In this scenario, the activity needs to be re-defined
where necessary, which makes the workflow script overly
complicated and unnecessarily large. The ActiveBPEL™
Designer [30] facilitates reuse of activities through the use of

1 'BPElets.' This custom feature allows the designer to create
 2 their own BPEL compositions (activities) that are likely to be
 3 reused later in the process or in other projects. A list of
 4 'BPElets' can be added to the BPEL Tools palette. This
 5 allows the workflow designer to drag a BPElet onto a
 6 diagramming canvas. BPEL Designer automatically
 7 generates and adds the underlying BPEL code to the
 8 workflow. Despite ease of use for the designer, there is still
 9 repetition of code as new code fragments are added for each
 10 BPElet.

11 It is not possible to trigger an <onMessage> or <onAlarm>
 12 event within the workflow for the purposes of dynamically
 13 modifying the workflow at runtime. This shortcoming can be
 14 addressed to some extent by throwing an exception and
 15 invoking a compensation handler. Alternatively, a link can be
 16 used to notify an activity when another activity discovers a
 17 condition of interest.

18 The BPEL specification does not explicitly support user
 19 interactions, which results in different levels of support for
 20 different workflow engines. As a result, vendor-specific code
 21 dependencies are often introduced. Support for user
 22 interaction varies in functionality and syntax. For example,
 23 Fig. 14 shows the injection of user interaction code in the
 24 'IBM WebSphere* Application Server Process
 25 Choreographer,' which has extended the 'invoke' activity
 26 [19].

```

  28 <bpel:invoke partnerLink="null"
  29     portType="abc:maintenancePT"
  30     operation="performMaintenance"
  31     inputVariable="ClientInformation"
  32     outputVariable="statusAndCost">
  33     <wpc:staff>
  34         <wpc:potentialOwner>
  35             <staff:membersOfRole
  36                 role="fieldWorker"/>
  37             </wpc:potentialOwner>
  38             <wpc:reader>
  39                 <staff:membersOfRole
  40                     role="manager"/>
  41             </wpc:reader>
  42         </wpc:staff>
  43     </bpel:invoke>
  
```

41 Fig. 14. Human interaction support in the IBM WebSphere*
 42 Application Server Process Choreographer via the null
 43 partnerLink and staff extension.

```

  45 xmlns:task="http://services.oracle.com/bpel/t
  46 ask"
  47 <partnerLink name="customTaskManager"
  48     partnerLinkType="task:TaskManager"
  49     partnerRole="TaskManager"
  50     myRole="TaskManagerRequester"
  51 />
  
```

52 Fig. 15. Human interaction support in the Oracle BPEL
 53 Server via the TaskManager.

54 Similarly the Oracle BPEL Process Manager supports user
 55 interaction through the 'Task Manager Service.' The Task
 56 Manager is a built-in service that enables the developer to
 57 include manual tasks in BPEL processes. Task Manager is an
 58 asynchronous service and provides two interfaces for
 59 interaction. Fig. 15 demonstrates the use of the Task Manager
 60 Service. A partner link called 'customTaskManager' is added
 61
 62
 63

which represents the Task Manager Service. The location of
 the Oracle BPEL Task Manager WSDL interface is:
 http://localhost:9700/orabel/default/TaskManager/TaskMan
 ager?wsdl. The remaining steps are similar to invoking
 operations on any Web service.

Recently a new project BPEL4People has been launched
 with the objective of standardizing the explicit inclusion of
 manual tasks in BPEL processes. The most important
 extensions introduced in BPEL4People are 'people activities'
 and 'people links'. A people activity is used to define user
 interactions; in other words, tasks that a user has to perform.
 For each people activity, the BPEL server must create work
 items and distribute them to users eligible to execute them.
 People activities can have input and output variables and can
 specify deadlines.

Often, there are situations when a developer would want to
inject Java code directly into a workflow script, either to
 implement simple actions or to re-use code that already
 models an activity. To do this, the developer has two standard
 options:

- Wrap local objects as local Web services and invoke the
 services from the workflow script. This approach can
 affect the performance of the overall workflow if a large
 number of Java objects are deployed as local Web
 services, especially if they are referenced frequently.
- Use Web Services Invocation Framework (WSIF) [32] to
 invoke local Java objects using the SOAP protocol. This
 approach requires learning the WSIF API and
 configuration with the workflow engine. Oracle BPEL
 Process Manager has built in support for WSIF.

The non standard option is to rely on proprietary solutions.
 The ActiveBPEL Engine [29] supports 'Custom Invoke
 Handlers' to invoke local or custom objects. Oracle BPEL
 Process Manager extends the BPEL specification through the
 use of the <bpelx:exec> activity. The <bpelx:exec> allows
 the user to embed Java code in the workflow script, similar to
 Java code in Java Server Pages (JSP). IBM and BEA also
 floated the concept of injecting Java code in the workflow
 script in the form of a new specification called 'BPEL for
 Java' (BPELJ) [5]. The BPELJ specification was not widely
 accepted by the Web service community and its support is
 limited to their own orchestration engines.

The BPEL4WS specification does not address Web
service message level security. For example, a partner service
 may require authentication credentials such as signed
 Security Assertion Markup Language (SAML) assertions [9]
 to be added to the SOAP message header. Consequently,
 different orchestration engines have extended the original
 BPEL specification to allow manipulation of the SOAP
 header. Oracle BPEL Process Manager has two extensions to
 either compose or consume SOAP headers in a workflow
 script ('inputHeaderVariable' and 'outputHeaderVariable').
 An alternative approach would be to use custom invoke
 handlers using APIs of choice to embed security credentials
 into the SOAP message.

5. Conclusion

Scientific workflows are similar to business processes in
 many respects, but their constantly changing and potentially
 long-running nature makes it difficult to anticipate all
 contingencies at design time. This produces a need for agile

1 monitoring and steering capabilities for scientific workflows.
 2 BPEL4WS is designed specifically to orchestrate Web
 3 services into workflows from the business perspective, but
 4 there is growing interest in the use of a Web services for
 5 scientific parallel and distributed computing [12]. This
 6 interest is fuelled by the availability of Web services that are
 7 tailored to Grid Computing in compliance with Web Services
 8 Resource Framework (WSRF) [39]. The emergence of
 9 BPEL4WS as the de-facto industry standard for Web service
 10 orchestration is significant for scientific domains. In house
 11 developed solutions are appealing in the short term as they
 12 may be well tailored and tuned to address specific scientific
 13 problems. However, their cross platform availability and
 14 acceptance across multiple institutes is limited. Relying on
 15 standards attracts a broader audience enabling greater
 16 contributions from all quarters of research. We conclude that
 17 BPEL4WS is currently the best candidate for orchestrating
 18 scientific and Grid services. The limitations of BPEL4WS
 19 can be overcome by supporting standard pattern-based
 20 technologies such as WSIF, J2EE and WS-* specifications.
 21 The use of WSIF especially facilitates portability and plug-in
 22 of required code into a workflow such as engine-specific
 23 support for workflow monitoring and management.
 24 Development of the WSIF based workflow plug-in is the
 25 natural and standard way to add missing components in the
 26 BPEL workflow. In most cases, careful application of
 27 BPEL4WS with a well-anticipated and investigated
 28 orchestration model can improve a scientific workflow. A
 29 desktop BPEL4WS execution engine gives a researcher
 30 control on the runtime configuration, monitoring and steering
 31 of a workflow with flexible execution, enhanced user
 32 interaction, and intelligent response to dynamic situations. A
 33 portal-based solution may be more appropriate for making
 34 pre-defined workflows available to a large number of end
 35 users.

37 Acknowledgements

38
 39 The authors acknowledge funding for this work from the
 40 EPSRC WOSE project: Workflow Optimization for Services
 41 in e-Science and the BBSRC/ DTI e-HTPX project: An e-
 42 Science Resources for High Throughput Protein
 43 Crystallography.

45 References

46
 47 [1] R Allan, R Keegan and D Meredith, e-HTPX – HPC,
 48 Grid and Web-Portal Technologies in High Throughput
 49 Protein Crystallography, All Hand Meeting, Nottingham,
 50 2004, pp. 187-193.
 51 [2] T Andrews, F Curbera, H Dholakia, Y Goland, J Klein, F
 52 Leymann, K Liu, D Roller, D Smith, S Thatte, I
 53 Trickovic and S Weerawarana, Business Process
 54 Execution Language for Web services version 1.1,
 55 (BPEL4WS), May 2003, [ftp://www6.software.
 56 ibm.com/software/developer/library/ws-bpel.pdf](ftp://www6.software.ibm.com/software/developer/library/ws-bpel.pdf)
 57 [3] C Argyris and D Schön, Organizational Learning: A
 58 Theory of Action Perspective Revista Espanola De
 59 Investigaciones Sociologicas, Issue 77/78, 1973, pp. 345-
 60 350.
 61 [4] P Barthelmeß and J Wainer, “Workflow systems: A few
 62 definitions and a few suggestions”, Conference on

Organizational Computing Systems, Milipitas, CA, 1995,
 pp. 138-147.
 [5] M Blow, Y Goland, M Kloppmann, F Leymann, G Pfau,
 D Roller and M Rowley BPELJ: BPEL for Java
 technology, March 2004, [http://www-128.ibm.com/
 developerworks/library/specification/ws-bpelj/](http://www-128.ibm.com/developerworks/library/specification/ws-bpelj/)
 [6] G Bolcer, Flexible and Customizable Workflow
 Execution on the WWW, Ph.D. Thesis, University of
 California, Irvine. September, 1998.
 [7] D Box, E Christensen, F Curbera, D Ferguson, Jeffrey
 Frey, M Hadley, C Kaler, D Langworthy, F Leymann, B
 Lovering, S Lucco, S Millet, N Mukhi, M Nottingham, D
 Orchard, J Shewchuk, E Sindambiwe, T Storey, S
 Weerawarana and S Winkler, Web services Addressing
 (WS-Addressing), August 2004, [http://www.w3.org/
 Submission/ws-addressing/](http://www.w3.org/Submission/ws-addressing/)
 [8] L F Cabrera, G Copeland, M Feingold, R W Freund, H T
 Freund, J Johnson, S Joyce, C Kaler, J Klein, D
 Langworthy, M Little, A Nadalin, E Newcomer, D
 Orchard, I Robinson, J Shewchuk and Tony Storey, Web
 Services Coordination (WS-Coordination) 1.0, August
 2005, [ftp://www6.software.ibm.com/software/developer/
 library/WS-Coordination.pdf](ftp://www6.software.ibm.com/software/developer/library/WS-Coordination.pdf)
 [9] S Cantor, J Kemp, R Philpott and E Maler. Security
 Assertion Markup Language (SAML) V2.0, 15 March
 2005, [http://www.oasis-open.org/committees/tc_
 home.php?wg_abbrev=security#samlv20](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security#samlv20)
 [10] E Christensen, F Curbera, G Meredith and S
 Weerawarana, Web services Description Language
 (WSDL) 1.1, March 2001, <http://www.w3.org/TR/wsdl>
 [11] W Cox, F Cabrera, G Copeland, T Freund, J Klein, T
 Storey and S Thatte, Web Services Transaction (WS-
 Transaction) 1.0, January 2004, [http://dev2dev.bea.com/
 pub/a/2004/01/ws-transaction.html](http://dev2dev.bea.com/pub/a/2004/01/ws-transaction.html)
 [12] G Cugola, E D Nitto, A Fuggetta and C Ghezzi, A
 framework for formalizing inconsistencies and deviations
 in human-centered systems. ACM Transactions on
 Software Engineering and Methodology (TOSEM), Vol.
 5, No. 3, 1996, pp. 191-230.
 [13] W Emmerich, B Butchart, L Chen and B Wassermann,
 Grid service orchestration using the business process
 execution language (BPEL). UCL-CS Research Note
 RN/05/07.
 [14] S Fitzgerald and M A Hutchins, How workflow
 automation improves employee exit management. Intel
 Corporation, Internal Report, January 2005,
[http://www.intel.com/it/digital-office/reducing-
 departure-costs.pdf](http://www.intel.com/it/digital-office/reducing-departure-costs.pdf)
 [15] Y Han, A Sheth, C Bussler, A taxonomy of adaptive
 workflow management. Computer Supported
 Cooperative Work (CSCW), Seattle, November 1998.
 [16] G E Kaiser, S E Dossick, W Jiang, J J Yang and S X Ye,
 WWW-based collaboration environments with
 distributed tool service. World Wide Web Journal,
 Baltzer Science Publishers, Vol. 1, No. 1, 1998, pp. 3-25.
 [17] P J Kammer, G Bolcer, R N. Taylor and M Bergnam,
 Techniques for supporting dynamic and adaptive
 workflow. Computer Supported Cooperative Work
 (CSCW), Vol. 9, 2000, pp. 269-292.
 [18] N Kavantzias, D Burdett, G Ritzinger, T Fletcher and Y
 Lafon, Web services Choreography Description
 Language Version 1.0, December 2004,

- 1 <http://www.w3.org/TR/2004/WD-ws-cdl-10-20040427/>
2 [19]R. Khalaf, A. Keller, and F. Leymann, Business
3 processes for Web services: Principles and applications.
4 IBM Systems Journal, Vol. 45, No 2, 2006.
5 [20]F Leymann, Web services Flow Language (WSFL)
6 Version 1.0, [http://www-306.ibm.com/software/
7 solutions/webservices/pdf/WSFL.pdf](http://www-306.ibm.com/software/solutions/webservices/pdf/WSFL.pdf)
8 [21]G J Nutt, The evolution towards flexible workflow
9 systems. Distributed System Engineering. Vol. 3, 1996,
10 pp. 276-294.
11 [22]C Peltz, Web services orchestration and choreography
12 IEEE Computer, Vol. 36, No. 10, 2003, pp. 46-52
13 [23]H Saastamoinen, On the Handling of Exceptions in
14 Information Systems, PhD thesis, University of
15 Jyväskylä, 1995.
16 [24]A Slomiski , On using BPEL extensibility to implement
17 OGSF and WSRF Grid workflows. Concurrency and
18 Computation: Practice and Experience, Vol. 18, No. 10,
19 2005, pp. 1229 - 1241.
20 [25]L A Suchman., Office procedure as practical action:
21 models of work and system design. ACM Transactions
22 on Information Systems, Vol. 1, No. 4, October 1983, pp.
23 320 -328.
24 [26]S Thatte, XLANG, [http://www.gotdotnet.com/team/
25 xml_wsspecs/xlang-c/default.htm](http://www.gotdotnet.com/team/xml_wsspecs/xlang-c/default.htm)
26 [27]W J Tolone, Introspect: a Meta-Level Specification
27 Framework for Dynamic Evolvable Collaboration
28 Support, Ph.D. Thesis. University of Illinois at Urbana-
29 Champaign, 1996.
30 [28]P Wagstrom, S Krishnan, and G v Laszewski, GSFL: An
31 OGSF Workflow Framework, [http://www-
32 unix.globus.org/cog/papers/gsfl-paper.pdf](http://www-unix.globus.org/cog/papers/gsfl-paper.pdf) , 2002
33 [29]ActiveBPEL Engine (2.0), <http://www.activebpel.org/>
34 [30]ActiveBPEL™ Designer, [http://www.active-
36 endpoints.com/products/activebpeldes/index.html](http://www.active-
35 endpoints.com/products/activebpeldes/index.html)
37 [31]e-HTPX - A Project for High Throughput Protein
38 Crystallography, <http://www.e-htpx.ac.uk>.
39 [32]ESRF – The European Synchrotron Radiation Facility,
40 <http://www.esrf.fr/>.
41 [33]OPPF – The Oxford Protein Production Facility,
42 <http://www.oppf.ox.ac.uk>
43 [34]Science Factory, Integrated Biological Computing
44 Environment, [http://www.science-factory.com/products/
46 workflow.html](http://www.science-factory.com/products/
45 workflow.html)
47 [35]SRS – Synchrotron Radiation Source, Daresbury
48 Laboratory, <http://www.srs.ac.uk/srs/>
49 [36]The Workflow Management Coalition,
50 <http://www.wfmc.org/>
51 [37]Web services Choreography Description Language
52 Version 1.0, [http://www.w3.org/TR/2004/WD-ws-cdl-
54 10-20040427/](http://www.w3.org/TR/2004/WD-ws-cdl-
53 10-20040427/)
55 [38]Web service Invocation Framework (WSIF),
56 <http://ws.apache.org/wsif/>
57 [39]Web Services Resource Framework (WSRF), 1.2,
58 <http://www.globus.org/wsrf/>
59 [40]YSBL – York Structural Biology Laboratory,
60 <http://www.ysbl.york.ac.uk/>

Unclear reference source info: #15

Author Bios

A Akram is currently doing research in the CCLRC e-Science Centre, Daresbury Laboratory. He is actively involved in evaluation of emerging Web services standards and composition languages. Asif is a Chemical Engineer, graduated from the University of Punjab, Pakistan. He has worked in the petro-chemical industry and has taught in different universities before moving to research field.

Dr D Meredith is member of the UK Grid Operations Support Centre for the National Grid Service and Grid Technology Group at CCLRC Daresbury Laboratory. David's main research areas include distributed computing and enterprise systems, with special application to the physical sciences. David's background is as a researcher in computational modeling of petroleum systems and Lithosphere dynamics.

Dr R Allan is head of Grid Technology Group at CCLRC Daresbury Laboratory, with an interest in distributed service development and deployment and user interfaces for research, mainly in the natural sciences. Rob's background is as a researcher in theoretical atomic physics doing large scale computational modeling. He has also developed numerical algorithms and parallel computing and Grid tools for the UK Collaborative Computational Projects and HEC communities, e.g., through CCP2, CCP6, HPCI and UKHEC initiatives. He has skills in FORTRAN, Perl and C/C++. Rob has a practical knowledge of Grid and Web services middleware and helped to set up CCLRC's e-Science Centre and the UK's Grid Support Centre and has acted as PI or co-PI in numerous high-profile research council and JISC grants. He has written several technical evaluations, including the DTI-funded Globus evaluation report (December 2001) and the JISC Sakai Evaluation report (December 2004). Rob is currently acting as e-Research representative on the Joint Framework Working Group (JISC UK and DEST Australia) and is chair of the Operations Board of the NW-GRID and chair of the Service Delivery Board of National Centre for e-Social Science.