

# GSFL: A Workflow Framework for Grid Services

Sriram Krishnan<sup>1,2</sup>, Patrick Wagstrom<sup>1,3</sup>, Gregor von Laszewski<sup>1</sup>

<sup>1</sup>Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439

<sup>2</sup>Indiana University, Bloomington, IN, 47405

<sup>3</sup>Illinois Institute of Technology, Chicago, IL, 60616

## Abstract

The *Open Grid Services Architecture* (OGSA) tries to address the challenge of integrating services spread across distributed, heterogenous, dynamic *virtual organizations*, using the concepts and technologies from both the Grid and Web service communities. The Web service community has realized that Web services can reach their full potential only if there exists a mechanism to describe the various interactions between the services and dynamically compose new services out of existing ones. This is true in the case of Grid services as well. In this paper, we analyze existing technologies that address *workflow* for Web services, and try to leverage them for Grid services, which have different needs from standard Web services. We discuss these special needs, and present the *Grid Services Flow Language* (GSFL), which addresses them for Grid services within the OGSA framework.

KEY WORDS: Grid; OGSA; OGSI; Grid Services; Web Services; Workflow

## 1 Introduction

The Web services approach is rapidly gaining momentum in the industry. The W3C [9] defines a Web service as a software application identified by a URI [13], whose interfaces and bindings are capable of being defined, described and discovered by XML artifacts, and which can support direct interactions with other software applications using XML based messages via internet-based protocols. A more general and descriptive definition can be found in [20], where a Web service is defined as a *platform and implementation independent* software component that can be:

- *Described* using a service description language
- *Published* to a registry of services
- *Discovered* through standard mechanisms
- *Invoked* through a declared API, usually over a network
- *Composed* with other services

The goal of Web services is interoperability. It follows from the definitions that a requestor can access Web services using *standard* mechanisms. In the ideal world, any requestor can interface with any application that claims to be a Web service irrespective of the language and the environment that either of them use. This makes the Web services approach appealing to modern enterprise and inter-organizational computing systems.

Grid computing involves the agglomeration of a variety of resources in dynamic, distributed *Virtual Organizations* [18]. *Grid technologies* are infrastructures which support the sharing and coordinated use of such diverse resources. Grid technologies [16] are currently seeing widespread adoption in the scientific computing community. Apart from the problems which are inherent to dealing with resources, viz. algorithms and problem-solving techniques, resource management, security, instrumentation and performance analysis, network infrastructure, etc., Grid technologies have to solve similar problems as the ones addressed by the Web services, viz. description, discovery, communication, remote invocation, etc. Recently, this fact has been recognized by the Grid community, resulting in the development of the Open Grid Services Architecture (OGSA) [17].

OGSA is the result of lessons learned while developing the Globus Toolkit [7], which has become the De-facto standard for Grid computing. OGSA uses the Web Services Description Language (WSDL) [14] to achieve self-describing and discoverable services. It defines a set of standard interfaces that a *Grid Service* may export, that enable features such as discovery, service lookup, lifetime management, notification, and credential management. However, in a similar vein to Web services technology, Grid services can realize their full potential only if there is a mechanism to dynamically compose new services out of agglomeration of existing ones. In order to do so, we not only have to describe the order in which these services and their methods execute, but also present a way in which such an agglomerate can export itself as a service. In this context, we define the term *workflow* as a set of rules that define the interactions between a set of services in order to be composed into a *meta-service*.

In this paper we show the current state of workflow languages for Web services and how they relate to Grid services. We point out some shortcomings of such languages in the context of the Grid. As a solution, we propose the Grid Services Flow Language (GSFL), which we describe in detail. We present the research problems posed and the approach we take towards them.

## 2 Technology Survey

The field of workflow languages for Web services has seen a lot of activity in the recent past. A variety of methods have been proposed as standards by the major Web services software providers. Since it is impossible to do a thorough survey of all the existing technologies due to the high rate of change in this field, we focus on some of the larger projects that are getting attention at this time.

### 2.1 Web Services Flow Language (WSFL)

WSFL [21], which is proposed by IBM, is one of the approaches to Web services workflow. WSFL describes the composition of Web services using a *flow model* and a *global model*. The flow model defines a series of *activities* that represent the operations of the composite Web service, and specifies the order in which these activities execute. It defines the flow of control and data between the various activities using *controlLinks* and *dataLinks* respectively. In most cases, the data flow will closely follow the control flow. However, WSFL is flexible enough to accommodate cases where this may not be true.

The global model defines how the activities of the composite Web service are mapped into the operations of the individual Web services, using what they call *plugLinks*. While the flow model describes the orchestration of the various activities in the workflow, the global model describes how these are implemented by the participating Web services. WSFL identifies the services participating in the workflow using a *locator* element, which supports the following bindings :

- *static*, where a reference to a WSDL or WSFL definition is provided.
- *local*, where the service implementation is local.
- *uddi*, where the service implementation is looked up using the UDDI [2] API.
- *mobility*, where the service provider is referenced in a message generated by some activity.
- *any*, where the service provider is not restricted by the flow model.

WSFL also supports lifecycle operations for the flow model of the composite Web services. It supports operations such as *spawn*, *call* (a blocking spawn), *suspend*, *resume*, *enquire*, and *terminate*. The advantages are WSFL are its logical consistency with WSDL, and the ability to define Web services which are recursively composed of other Web services.

Version 1.0 of WSFL was released by IBM in May 2001. However, not much has been written about WSFL since then. There is no popular implementation of WSFL available either, although there seem to be a few groups which are working on one [15].

## 2.2 XLANG: Web Services for Business Process Design

XLANG[27], a proposal by Microsoft Corporation, is a language that is used to model business processes as autonomous agents. In WSDL the unit of action is an operation, this operation can be either on a stateless service (such as a stock quote) or on a stateful (subservient) service where the interaction defines the beginning and end of the process. There is a third model, where business processes may be *autonomous agents*, such as a supply chain. In this chain, input and output messages occur in a defined order, deemed a *service process*. It is based on  $\pi$ -calculus theory for the connection and synchronizations of automata.

XLANG defines the following set of operations as extensions to the standard WSDL operations to assist in this modeling:

- *delays* allow a thread to stall for a specified time period, or until another condition is met
- *raise* is a method to raise exceptions for certain actions
- *process control* combines actions together with conditional and iterative statements
- *correlation* provides a method for declaring longer running conversations
- *transaction support* allows definition of rollback procedures if one action in the execution fails
- *contracts* create agglomerate services by facilitating one way bindings between ports

Like many of these technologies XLANG is still evolving. Currently it lacks methods to add services dynamically to a business process and also does not have a mechanism to export the methods of these services as part of the workflow. These will be addressed in future revisions of the language. However, little has been written about XLANG since its initial release in May 2001.

## 2.3 Web Services Conversation Language (WSCL)

WSCL [28] is a conversation language framework, proposed by the Hewlett-Packard Company, for modeling the sequencing of the interactions or operations of *one interface*. It fills the gap between mere interface definition languages that do not specify any choreography and more complex process or flow languages that describe complex global multi-party conversations and processes. The major elements of the WSCL specification are :

- *Document type descriptions*, which reference the types of documents that the service can accept or transmit, defined using XML Schemas [3].
- *Interactions*, which model the actions of a conversation between two participants.
- *Transitions*, which specify the ordering relationships between the interactions.
- *Conversations*, which list all the interactions and transitions that make up a conversation.

The conversations are the public interfaces supported by a service. They add additional semantics to the WSDL for the service by also specifying the possible ordering of the operations. However, WSCL does not address the recursive composition of Web services, which is what we aim to do for Grid services.

## 2.4 Other Related Work

The Web Services Choreography Interface (WSCI) [11] proposed by Sun, Intalio, SAP and BEA is a language aimed at application to application integration on a tighter level than that proposed by XLANG. However, WSCI was proposed in June 2002, and is still fairly new at the time of writing this paper. The Business Process Modeling Language (BPML) [10] is a meta-language for the modeling of business processes. BPML provides an abstracted execution model for collaborative and transactional business processes based on the concept of a transactional finite-state machine. DAGMan [6] is a meta-scheduler for Condor [5] that manages dependencies between jobs. Despite the fact that DAGMan does not deal with the workflow for Web services, the concept of using a directed acyclic graph to represent a set of programs where the inputs, outputs, and the execution are inter-dependent, can be applied to describe the dependencies between the Web services. The XCAT Application Factories [19] address workflow issues for Grid based components within the Common Component Architecture (CCA) [12] framework. XCAT allows components to be connected to each other dynamically, making it possible to build applications in ways not possible with the standard Web services model.

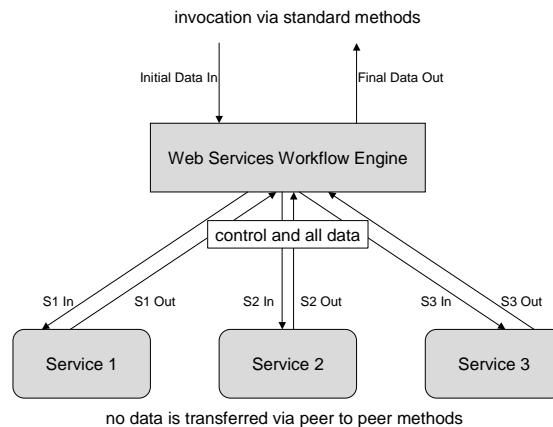


Figure 1: Web Services Workflow Model

### 3 Grid Workflow Requirements

As a result of the above survey and our analysis of Grid use cases, we have established a set of requirements for a workflow specification for the Grid. In this section, we describe these requirements, and how existing Web service technologies do not address all of them, despite providing invaluable techniques that we reuse.

Just like the Web service technologies aim to do, the Grid workflow specification should allow specific activities implemented by individual services to be exported as activities of the workflow. It should also allow the exported activities to trigger a chain of other activities. Current technologies such as WSFL address this issue effectively. Hence, we try to incorporate these features presented by WSFL into the Grid Services Flow Language. Furthermore, the activities exported in such a manner should also be described in the same manner as the service itself. In this sense, the specification should be rich enough to describe the workflow such that the WSDL for the workflow entity (henceforth referred to as a *workflow coordinator*) can be auto-generated from the specification. The workflow coordinator must be able to handle the methods that have been dynamically exported as a composition of the various activities of the workflow, in such a way that clients can access them using the same standard tools that they use to deal with the individual services. This is an important requirement for recursive composition of services.

As we observed, and as has been pointed out by [19], existing Web services define their workflow in such a way that the workflow engine has to intermedate at each step of the application sequence, as shown in Figure 1. This is because most of the current set of workflow technologies are designed for business to business communication, where there may only be a moderate level of data transmission across the Web services. As a result, the workflow engine does not end up being a real bottleneck. However, for Grid based services, exchanging large amounts of data is the norm. Having a central workflow engine relay the data between the services would be a bad idea in this case. The workflow specification needs to be able to allow communication between the services as depicted in Figure 2.

As we mention before, OGSA adds extensions to WSDL in order to address Grid-specific needs. It addresses communication between Grid services using *notificationSources* and *notificationSinks*, which allow services to carry out asynchronous delivery of messages between each other. It is needed that GSFL provide a mechanism to connect *notificationSources* and *notificationSinks*, thus obviating the need for the workflow engine to mediate at every step. Additionally, OGSA uses *Registries* and *Factories* for locating and creating Grid services, respectively. This has to be appropriately handled by GSFL.

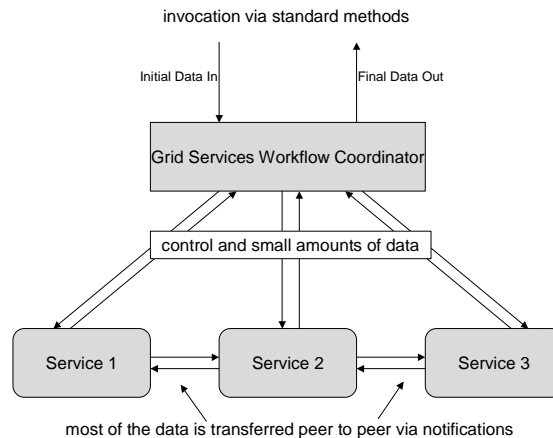


Figure 2: Grid Services Workflow Model

It is conceivable that certain Grid services in the workflow will not be executing while others are. This may be due to the fact that the services which need to execute earlier run for weeks or that the service which executes later needs data from the former to bootstrap itself. The Grid workflow specification should be able to handle this particular need. Additionally, it should also be able to handle instantiation of the individual Grid services on a per method or a per workflow instance basis. In case the Grid services are instantiated on a per workflow instance basis, certain activities exported by the workflow may not function because a certain Grid service may have run to completion, or may not have been instantiated yet. In such a case, certain semantics need to be added to the ordering of the exported activities, as suggested by WSCL.

In the following section, we describe the Grid Services Flow Language, and how we address the requirements specified.

## 4 GSFL Overview

The Grid Services Flow Language is an XML based language that allows the specification of workflow descriptions for Grid services in the OGSA framework. It has been defined using XML Schemas. A simplified architecture is shown in Figure 3. It has the following important features, which we expand in the following subsections.

- *Service Providers*, which are the list of services taking part in the workflow.
- *Activity Model*, which describes the list of important activities in the workflow.
- *Composition Model*, which describes the interactions between the individual services.
- *Lifecycle Model*, which describes the lifecycle for the various activities and the services which are part of the workflow.

### 4.1 Service Providers

All the services that are part of the workflow have to be specified in the list of *serviceProviders*. The service providers are identified through out the GSFL document by a unique *name*, which is specified as part of the definition. The definition also contains the *type* of the service provider, which is the type of the Grid service, as specified by its WSDL specification. Service providers can be located using the *locator* element, which allows looking up service providers

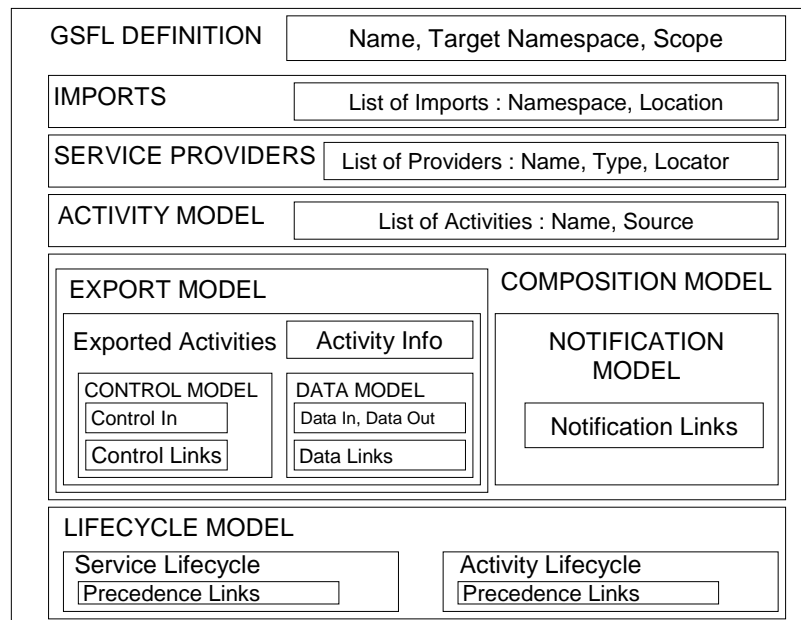


Figure 3: Architecture of the GSFL

in a number of ways. Services can be located *statically*, via a static URL which would point to an already running service. They can also be started up using *factories*, the handle to which would be available in the GSFL document. Services can also be looked up using *registries*.

## 4.2 Activity Model

The *activityModel* lists all the operations belonging to the individual service providers, that have any role to play in the workflow. It contains a list of activities, each of which has a *name* for identification purposes, and a *source*, which is a reference to an operation in a Web service defined by an *endPointType* element. The *endPointType* element contains the names of the operation, *portType*, *portName* and *serviceName* for a particular operation.

## 4.3 Composition Model

The *compositionModel* describes how the different Grid services are composed to form a new Grid service. It describes the control and data flow between various operations of the services, and also the direct communication between them in a peer-to-peer fashion. It consists of the following subparts.

### 4.3.1 Export Model

The *exportModel* contains the list of activities that have to be exported as operations of the workflow process. Any client can invoke these operations on the workflow instance using standard mechanisms. Since the workflow instance can also be viewed as a standard Grid service, it can be used recursively as part of another workflow process. For each activity exported, the control and data flow is described by the *controlModel* and the *dataModel* respectively.

The *controlModel* describes the chain of activities that are invoked when the exported activity is invoked by a client. Each *controlModel* element contains an attribute *controlIn* which references the first activity to be executed when the

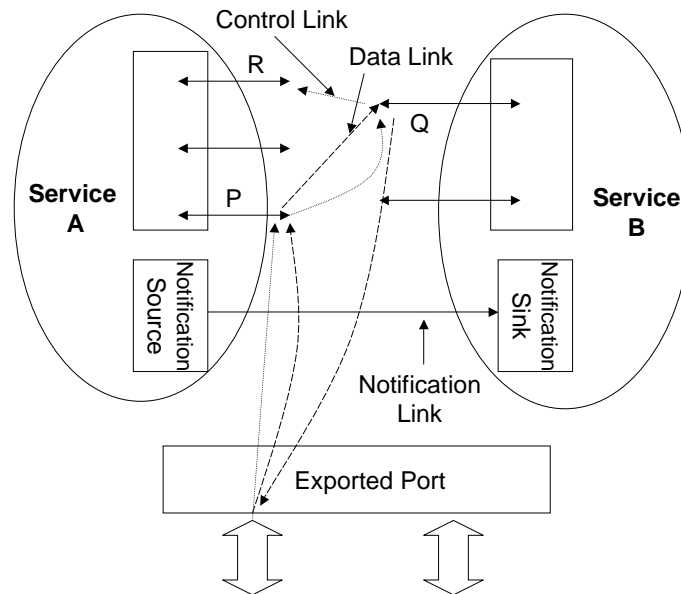


Figure 4: The Composition Model in action

exported activity is invoked. Each controlModel also contains a series of *controlLinks* which is a precedence list of all activities that need to be successively invoked, as part of this exported activity.

The *dataModel* describes the flow of data that occurs when an exported activity is invoked. This flow may not necessarily mirror the flow of control between the various activities. Each dataModel element contains an attribute *dataInTo* (shown as “Data In” in Figure 3) which signifies the activity which will receive the data provided as input to the exported activity. A *dataOutFrom* attribute (shown as “Data Out” in the figure) designates the activity from which the data is returned to the caller.

The GSFL document provides enough information in the dataModel and the controlModel to not only dynamically generate the WSDL for the activities exported, but also support the invocation of such exported activities dynamically, as we will explain in Section 5.

#### 4.3.2 Notification Model

The *notificationModel* provides the solution to the problem of the workflow engine mediating at every step of an activity. As mentioned before, OGSA services communicate with each other using notificationSources and notificationSinks. The notificationModel provides a mechanism to link such sources to sinks and vice versa, along with a particular *topic*, using *notificationLinks*. The services can now communicate large amounts of data among each other, without having the need to go through the workflow engine. Users can still use the control and data models to communicate control messages and small amounts of data between each other, but the use of notification messages is the recommended form of communication for large amounts of data.

The composition model is illustrated by a simple example shown in Figure 4. Two services, A and B, constitute the workflow. The notification model consists of a single notification link  $A \rightarrow B$ , which connects the notification source of service A to the notification sink of service B. The export model consists of a couple of activities that are exported,

one of which is shown in detail in the figure. One of the exported activities is implemented by operations  $P$  and  $R$  of service A, and operation  $Q$  of service B. The control model of the exported activity consists of control links  $P \rightarrow Q$  and  $Q \rightarrow R$ . Operation  $P$  of service A serves as the controlIn for the exported activity. The data model consists of a single data link  $P \rightarrow Q$ . This may possibly be because operation  $R$  may not need any data to be invoked. This is an example where the data links need not necessarily be the same as the control links. Operation  $P$  serves as the dataInTo, while operation  $Q$  serves as the dataOutFrom for the exported activity. Thus, invocation of the exported activity will trigger a set of operations described above, following the flow of control and data described by the different types of links.

## 4.4 Lifecycle Model

The *lifecycleModel* addresses the order in the services and the activities execute. The *serviceLifecycleModel* contains a list of precedence links describing the order in which the services are supposed to execute. Hence, all services need not be instantiated at start up, and can be started up once the preceding services have stopped executing.

The *lifecycleModel* uses the *scope* attribute for the workflow, which can be one of *session* or *application*. Session scope means that no state will be maintained between calls to the workflow engine. All calls to the engine are legal. Services will be instantiated for each call using the *serviceLifecycleModel*, and these services will be alive when calls to them are made.

Application scope means that state will be maintained in the workflow engine between calls. Services will be instantiated only once per instance of the workflow, using the *serviceLifecycleModel*. Hence, all calls to the workflow engine may not be valid, because the services that implement these activities may not be alive. Hence, we add an *activityLifecycleModel*, which describes the order in which they can be invoked. In other words, some activities can be invoked only if certain other activities have already been successfully invoked, e.g. a *checkout* operation in an online shopping system may be invoked only after one or more *buy* operations. Following the *activityLifecycleModel* will ensure that all the services will be alive when calls to them are made in the proper order.

We believe that we have addressed the requirements for workflow for Grid services using the above design. We now discuss some of the issues involved in the implementation of a GSFL engine.

## 5 Implementation Details

As a basis for the prototypical development of GSFL we have selected the OSGI Technology Preview [22]. This is a Java based implementation of the current Grid Services Specification [29] utilizing Apache Tomcat [23] and Apache Axis [24]. Tomcat is the servlet container that is used in the official Reference Implementation for the Java Servlet [26] and JavaServer Pages [25] technologies, while Axis is an implementation of the Simple Object Access Protocol (SOAP) [1] submission to W3C. In the following subsections, we describe a few important parts of the GSFL implementation.

### 5.1 GSFL Parsing

We use Castor [4] for parsing the GSFL schema, and generating Java bindings for the same. Castor is an open source data binding framework for Java, and it can be used to generate the source for Java classes representing an XML schema. It also provides methods to unmarshal XML documents conforming to a particular schema into corresponding Java objects, and vice versa. Using Castor, we auto-generate the Java bindings for our XML schema, thus obviating the need for us to write code to do the same. This significantly reduced our code development time.

Although Castor maps the GSFL schema into appropriate Java classes, it has no understanding of the semantics associated with the elements of the schema, e.g. Castor has no idea that the *Import* type in the GSFL schema has enough information (namespace & location) to import another GSFL document. Castor will only map a type such as *Import* to a Java class having the fields - namespace & location, with the appropriate getter and setter methods. Hence,



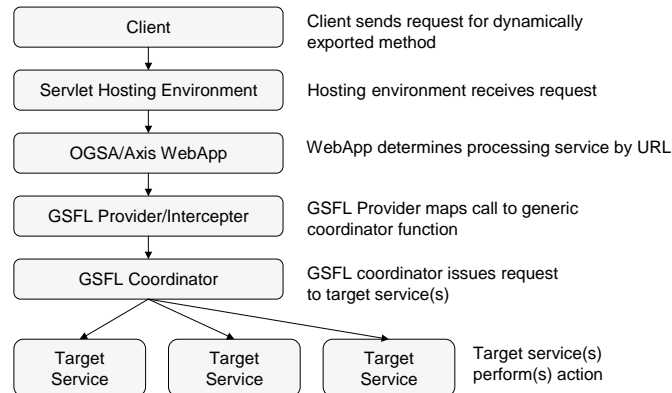


Figure 5: GSFL Control Flow

we had to add a wrapper around the Castor generated classes so that such situations can be handled. However, this was much easier than having to write all of the above code ourselves.

## 5.2 WSDL Auto-generation

Apart from using Castor for the Java bindings for the GSFL document, we use the WSDL4J [8] for dealing with the WSDL documents of the individual services. WSDL4J is a toolkit that allows the creation, representation, and manipulation of WSDL documents describing services.

As we described in Section 4, the GSFL document contains enough information about the activities that are exported. Since the WSDL for the services participating in the workflow is available, the *Types* for the input and output messages, and the *Messages* themselves can be easily found out for these activities. Information about the rest of the WSDL, such as *Operations*, *Ports*, etc. can be easily derived from the GSFL itself. The key in the auto-generation of the WSDL for the workflow was to represent enough information the GSFL schema.

## 5.3 GSFL Coordinator

The heart of a GSFL Workflow is the GSFLCoordinator service, this service creates virtual ports and services that map to processes internal to the workflow. These ports are virtual because they do not physically exist on the GSFLCoordinator. However, a client can invoke methods on these ports, as they get mapped to a set of other calls with respect to the GSFLCoordinator.

A diagram representing the usual flow of control is shown in Figure 5. A client wishing to execute a GSFL workflow first begins an instance of the GSFL coordinator via the standard OGSA factory methods. The coordinator instance is then sent the GSFL that represents the workflow. Upon receipt of the GSFL document, the coordinator dynamically generates a WSDL document with all of the newly exported operations included. This WSDL can then be used by clients wishing to execute the workflow, and the operations that are dynamically exported. When a client calls an operation that has been exported by the GSFL coordinator, a request is sent via the servlet container to the OGSA webapp (Web Application). The OGSA webapp utilizes the mappings in the `server-config.wsdd` to send the call the GSFLProvider/Interceptor, which has been implemented as an extension to the standard OGSA RPCURIProvider class, which is responsible for dispatching incoming requests to the correct service instance based on the URL of

the request. If the standard provider had been used, this invocation would have failed, since the operation does not physically exist on the GSFL Coordinator. However, the custom provider intercepts this call, and sends it to a generic marshaler function in the coordinator. Based on the information provided in the GSFL document, the coordinator then processes the request and maps the operation to a set of calls, which are in-fact implemented by the set of services participating in the workflow.

The provider has provisions to differentiate between calls to operations in dynamically exported ports, and the static ones. One of the static operations of interest is *generateWSDL*, which auto-generates the WSDL for the coordinator, inclusive of the dynamic ports, as described in Section 5.2.

## 6 GSFL Example

An example of a possible GSFL document for a filtered search service can be seen in Figure 6. The GSFL document shown contains the composition model shown in Figure 4, along with a few additions. It consists of three service providers - A, B, C. The activities that are part of the workflow, viz. P, Q, R, S, T, have been listed in the activity model. The export model consists of two exported activities - PQR and ST. The composition of PQR has been explained in Section 4.3. Activity ST is composed of two other activities - S and T. It consists of a control link  $S \rightarrow T$ , and a data link  $S \rightarrow T$ . In this case, the control and data flows are identical. Since the scope of the workflow is “session”, the lifecycle model of the workflow consists of only the service lifecycle. The figure shows that services A and B have to be active at the same time, while C can start up after A and B are done.

Despite the fact that the figure shows a toy example, it demonstrates the various constructs of the language, and the ease with which a workflow can be defined using individual services.

## 7 Conclusions

In this paper, we described the Grid Services Flow Language (GSFL), a workflow framework for Grid based services. We surveyed existing technology that addresses workflow for Web services, and investigated their applicability to Grid services. Existing Web services technologies provided a number of desirable features that we could reuse, however certain requirements such as peer-to-peer service interaction and complicated lifecycle management for the services were found lacking. We designed GSFL such that it addresses these requirements, while still being able to integrate the features provided by existing Web service technologies.

There still are broad areas of application and enhancement that are possible in the Grid Services Flow Language. This is still a work in progress, and the language will continue to evolve depending upon the requirements of the Grid community. Ideas for future features include the ability to handle and process exceptions, similar to those found in XLANG; automatic integration with a graphical workflow editor; and enhanced ordering of tasks in the workflow, possibly with constructs such as loops and switch statements.

## References

- [1] Simple Object Access Protocol (SOAP) 1.1. <http://www.w3.org/TR/SOAP>, May 2000.
- [2] UDDI Technical White Paper, September 2000. Universal Description, Discovery and Integration is for discovering web services.
- [3] XML Schema Part 0: Primer. <http://www.w3.org/TR/xmlschema-0/>, May 2001.
- [4] Castor. <http://castor.exolab.org/>, 2002.
- [5] Condor : High Throughput Computing. <http://www.cs.wisc.edu/condor/>, 2002.
- [6] DAGMan (Directed Acyclic Graph Manager). <http://www.cs.wisc.edu/condor/dagman/>, 2002.

```

<definitions name="Sample" scope="session" targetNamespace="http://www.mcs.anl.gov/gsf1">

<!-- List of Service Providers -->
<serviceProvider name="A" type="AServiceType">
  <locator type="factory" handle="http://localhost:8080/ogsa/services/AServiceFactory" />
</serviceProvider>

<serviceProvider name="B" type="BServiceType">
  <locator type="factory" handle="http://localhost:8080/ogsa/services/BServiceFactory" />
</serviceProvider>

<serviceProvider name="C" type="CServiceType">
  <locator type="factory" handle="http://localhost:8080/ogsa/services/CServiceFactory" />
</serviceProvider>

<!-- List of activities -->
<activityModel>
  <activity name="P">
    <source serviceName="A" portType="APort" portName="aPort" operation="P" />
  </activity>
  <activity name="Q">
    <source serviceName="B" portType="BPort" portName="bPort" operation="Q" />
  </activity>
  <activity name="R">
    <source serviceName="A" portType="APort" portName="aPort" operation="R" />
  </activity>
  <activity name="S">
    <source serviceName="A" portType="APort" portName="aPort" operation="S" />
  </activity>
  <activity name="T">
    <source serviceName="C" portType="CPort" portName="cPort" operation="T" />
  </activity>
</activityModel>

<!-- The composition model -->
<compositionModel>
  <exportModel>
    <exportedActivity>
      <exportedActivityInfo name="PQR" portType="exportPortType" />
      <controlModel controlIn="P">
        <controlLink label="link0" source="P" target="Q" />
        <controlLink label="link1" source="Q" target="R" />
      </controlModel>
      <dataModel dataInTo="P" dataOutFrom="Q">
        <dataLink label="link0" source="P" sink="Q" />
      </dataModel>
    </exportedActivity>
    <exportedActivity>
      <exportedActivityInfo name="ST" portType="exportPortType" />
      <controlModel controlIn="S">
        <controlLink label="link0" source="S" target="T" />
      </controlModel>
      <dataModel dataInTo="S" dataOutFrom="T">
        <dataLink label="link0" source="S" sink="T" />
      </dataModel>
    </exportedActivity>
  </exportModel>

  <notificationModel>
    <notificationLink label="link0" source="A" sink="B" topic="AB"/>
  </notificationModel>
</compositionModel>

<!-- Lifecycle for the services -->
<lifecycleModel>
  <serviceLifecycleModel>
    <precedenceLink label="link0">
      <parent>
        <element serviceName="A" />
        <element serviceName="B" />
      </parent>
      <child>
        <element serviceName="C" />
      </child>
    </precedenceLink>
  </serviceLifecycleModel>
</lifecycleModel>
</definitions>

```

Figure 6: A Sample GSFL Document

- [7] The Globus Project. <http://www.globus.org/>, June 2002.
- [8] The Web Services Description Language for Java Toolkit (WSDL4J). <http://www-124.ibm.com/developerworks/projects/wsd4j>, July 2002.
- [9] World Wide Web Consortium. <http://www.w3.org/>, June 2002.
- [10] Assaf Arkin. Business Process Modelling Language. <http://www.bpml.org/bmpi-downloads/BPML-SPEC-1.0.zip>, June 2002.
- [11] Assaf Arkin, Sid Askary, Scott Fordin, Wolfgang Jekeli, Kohsuke Kawaguchi, David Orchard, Stefano Pogliani, Karsten Riemer, Susan Struble, Pal Takacs-Nagy, Ivana Trickovic, and Sinisa Zimek. Web Services Choreography Interface. <http://www.sun.com/software/xml/developers/wsci/index.html>, June 2002. Version 1.0.
- [12] R. Armstrong, D. Gannon, A. Geist, K. Keahey, S. Kohn, L. McInnes, S. Parker, and B. Smolinski. Toward a Common Component Architecture for High-Performance Parallel Computing. In *Proceedings of High Performance Distributed Computing*, pages 115–124, Redondo Beach, California, 1999. CCA Forum, .
- [13] T. Berners-Lee, R. Fielding, and L. Masinter. Uniform Resource Identifiers (URI): Generic Syntax. <http://www.ietf.org/rfc/rfc2396.txt>, August 1998.
- [14] Erik Christensen, Francisco Curbera, Greg Meredith, and Sanjiva Weerawarana. Web Services Description Language. <http://www.w3.org/TR/wsdl>, March 2001. Revision 1.1 - March 15, 2002.
- [15] Distributed Systems Department Pervasive Collaborative Computing Environment Project (PCCE), LBL. PCCE Quarterly Reports. <http://www-itg.lbl.gov/Collaboratories/quarterly-reports.html>, April 2002.
- [16] I. Foster and C. Kesselman, editors. *The Grid: Blueprint for a Future Computing Infrastructure*. Morgan Kaufmann Publishers, July 1998.
- [17] Ian Foster, Carl Kesselman, Jeffrey Nick, and Steven Tuecke. The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration. <http://www.globus.org/research/papers/ogsa.pdf>, January 2002.
- [18] Ian Foster, Carl Kesselman, and Steve Tuecke. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *International Journal of Supercomputing Applications*, 15(3), 2002.
- [19] Dennis Gannon, Rachana Ananthkrishnan, Sriram Krishnan, Madhusudhan Govindaraju, Lavanya Ramakrishnan, and Aleksander Slominski. Grid Web Services and Application Factories. <http://www.extreme.indiana.edu/xcat/AppFactory.pdf>, June 2002.
- [20] Steve Graham, Simeon Simeonov, Toufic Boubez, Doug Davis, Glen Daniels, Yuichi Nakamura, and Ryo Neyama. *Building Web Services With Java*. SAMS, 2002.
- [21] Frank Leymann. Web Services Flow Language. <http://www-4.ibm.com/software/solutions/webservices/pdf/WSFL.pdf>, May 2001.
- [22] OGSi Technology Preview Release. <http://www.globus.org/ogsa/releases/TechPreview/>, June 2002.
- [23] Apache Jakarta Project. Apache Tomcat. <http://jakarta.apache.org/tomcat/>, June 2002.
- [24] Apache XML Project. Apache Axis. <http://xml.apache.org/axis/>, June 2002.
- [25] Sun Microsystems Inc. Java Server Pages. <http://java.sun.com/products/jsp>, 2002.
- [26] Sun Microsystems Inc. Java Servlet Technology. <http://java.sun.com/products/servlet/index.html>, 2002.
- [27] Satish Thatte. XLANG: Web services for Business Process Design. [http://www.gotdotnet.com/team/xml\\_wsspecs/xlang-c/default.htm](http://www.gotdotnet.com/team/xml_wsspecs/xlang-c/default.htm), 2001.
- [28] The Hewlett-Packard Company. Web Services Conversation Language (WSCL) 1.0. <http://www.w3.org/TR/wscl10/>, March 2002.
- [29] Steven Tuecke, Karl Czajkowski, Ian Foster, Jeffrey Frey, Steve Graham, and Carl Kesselman. Grid Service Specification (Draft 2). [http://www.gridforum.org/ogsi-wg/drafts/GS\\_Spec\\_draft02\\_2002-06-13.pdf](http://www.gridforum.org/ogsi-wg/drafts/GS_Spec_draft02_2002-06-13.pdf), June 2002.