

VLab: Collaborative Grid Services and Portals to Support Computational Material Science

Mehmet A. Nacar, Mehmet S. Aktas, and Marlon Pierce

Community Grids Lab

Indiana University

Corresponding Author: mnacar@cs.indiana.edu

Zhenyu Lu and Gordon Erlebacher

School of Computational Science and Information Technology

Florida State University

Corresponding Authors: zhenyulu@cs.fsu.edu, erlebach@csit.fsu.edu

Dan Kigelman, Evan F. Bollig, Cesar De Silva, Benny Sowell, and David A. Yuen

Minnesota Supercomputer Institute

University of Minnesota

Corresponding Author: rao@msi.umn.edu

Abstract: We present the initial architecture and implementation of VLab, a Grid and Web Service-based system for enabling distributed and collaborative computational chemistry and material science applications for the study of planetary materials. The requirements of VLab include job preparation and submission, job monitoring, data storage and analysis, and distributed collaboration. These components are divided into client entry (input file creation, visualization of data, task requests) and backend services (storage, analysis, computation). Clients and services communicate through NaradaBrokering, a publish/subscribe Grid middleware system that identifies specific hardware information with topics rather than IP addresses. We describe three aspects of VLab in this paper: 1) managing user interfaces and input data with Java Beans and Java Server Faces; 2) integrating Java Server Faces with the Java CoG Kit; and 3) designing a middleware framework that supports collaboration. To prototype our collaboration and visualization infrastructure, we have developed a service that transforms a scalar data set into its wavelet representation. General adaptors are placed between the endpoints and NaradaBrokering, which serve to isolate the clients/services from the middleware. This permits client and service development independently of potential changes to the middleware.

1. Introduction: Grid Enabling Material Science Applications

The Virtual Laboratory for Earth and Planetary Materials (VLab) is a National Science Foundation-funded interdisciplinary research collaboration whose primary objective is to investigate planetary materials at extreme conditions based on ab-initio computational techniques to better understand the processes that create earth-like and other planetary objects. Such calculations typically involve hundreds or thousands of computer runs. These runs occur in stages, with complex interactions. They are often managed by several researchers. To address challenges in collaborative and distributed computing, VLab brings together a team that includes researchers in computational material science, geophysics, scientific visualization, Grid computing, and information technology. Additional information on VLab is available from [1].

Some of the many problems that VLab must address include the ability to create input files through portals, submit jobs, store and retrieve the job input and output data on demand, analyze and visualize the data, and store the data. These tasks must be possible in a distributed

environment and the flow of information must be accessible to multiple collaborating researchers, although they might not be co-located. An additional constraint on our system is that it must be robust, i.e., fault tolerant. When working in a complex multi-user environment, it is inevitable that some components will fail. However, these failures should not affect the work of an individual researcher. Thus, we have chosen to connect the users of the systems (referred to as clients) and the various tasks requested by the users (storage, visualization, analysis, job submission, etc.) as services using NaradaBrokering [2], a middleware system that builds in many of the required features.

In its initial phase, VLab follows a well-established pattern for building Grids: application codes on remote machines are accessed securely through Grid services through a browser portal. This follows the common three-tiered architecture [6]. A user interacts with a portal server through a Web browser. The portal server in turn connects to remote Grid services that manage resources on a backend system that provides the computing power for running the codes. The longer term research goal however is to go beyond these traditional approaches. The distinguishing feature of our research is the use of the publish/subscribe paradigm, which completely decouples the clients from the services. Users have no knowledge of the resources allocated to their requests, although they will have the capability to monitor task progress.

Many of VLab's workhorse simulation codes are included in the "Quantum Espresso" package developed by the Democritos group [3]. As a starting point towards developing an automated web service workflow, we consider PWSCF, the Plane Wave Self Consistent Field [4] code in the Espresso suite. PWSCF is a parallelized application, often submitted to supercomputing resources via a batch queuing system. Common Grid technologies such as WS-GRAM, Reliable File Transfer, GridFTP, and Grid security, from the Globus Toolkit [5], provide the means to interface external applications with several popular schedulers, such as LSF, Condor, and PBS. Several additional problems must be addressed and are discussed in this paper: a) managing user inputs as persistent, archived, hierarchical project metadata (Section 2); b) simplifying and monitoring complicated, multi-staged job submissions using Grid portal technology (Section 3); and c) integrating VLab applications with Grid messaging infrastructure [2] to virtualize resource usage, provide fault tolerance, and enable collaboration (Section 4).

2. Managing User Interfaces and Input Data

Job submission tasks are broken down into the following components: 1) provide a user front end to input the data to generate a PWSCF input file, 2) move the input file to a backend resource, usually the computer that will run PWSCF, and 3) run PWSCF. Grid Web portals [6] are used to manage the collection of processes that define the submission task. This is a classic Grid portal problem. As detailed in a follow-up issue to [6] (currently in preparation), most current Java-based portals are developed around the "portlet" approach [7]. Portlets provide a consistent framework by which Web applications may be packaged and deployed into standard-compliant portlet containers. A portlet is typically a single, mostly self-contained application, such as the set of Web forms, Java code, and third party jars needed to submit the PWSCF code. Portlets are deployed into portlet containers, which are responsible for general purpose tasks, such as handling user login, providing a layout manager that arranges the portlets on the user's display, determining the user's rights to access particular portlets, and remembering the user's customizations (such as page arrangements and skin colors). We follow this approach and adopt the GridSphere [8] portal container for VLab deployment. By using standard compliant portlets,

we may later adopt other containers (such as uPortal or Jetspeed2), and can share VLab portlets with other collaborators who may prefer these different containers.

Portlets may be developed using several different Java web technologies. For VLab, we decided to test portlet development with Java Server Faces (JSF) [16]. JSF is a model-view-controller style framework for building web applications that provides three very important advantages for all applications. First, Web developers do not need to explicitly manage HTTP request parameters. This eliminates the dependency of the backing Java code on specific parameter names in the HTML `<input>` tags. Second, the Web form's business logic is encapsulated in simple Java Beans. Each HTML `<input>` parameter in a web form is associated with a property field in the Java Bean that manages the page. Finally, the scope of the Java Bean (i.e. session, request, or application) is configurable by the developer and managed by JSF. Developers do not need to explicitly manage their variables' life cycle.

One develops JSF applications by developing web pages with HTML and JSF tag libraries. The code for processing the user input (the "business logic") is implemented using Java Beans, which are associated with JSF pages in a configuration file (`faces-config.xml` by default). A full description of JSF is beyond our current scope, so interested readers should consult [16]. However, the implications of the points to science portals are important. The immediate result is that we do not need to adopt HTML parameter naming conventions for our Web pages (and thus break our forms when we change names). More importantly, we can develop our web application code as Java Beans ("backing beans"), which can be shielded from the Java Servlet specification. This allows us to reuse Java Bean code in non-JSF applications, take advantage of XML bean serialization tools, develop simple standalone unit tests, and generally take advantage of Java Bean-based "Inversion of Control" [14] frameworks.

We have developed PWSCF input pages with JSF to collect the user input needed to create a PWSCF input page. A sample page is shown in Figure 1. Users must fill out two pages of forms to describe their problem and have a chance to preview and (if experts) edit the generated input file manually before submission. Users may also upload additional input data (atomic pseudo-potentials) from their desktop. The linked input pages and backing Java Bean code together constitute a portlet. One of the issues we address is the persistent preservation of user input data. The form of Figure 1 is tedious to fill out, and quite typically a user will want to make minor modifications to a particular job and resubmit it later. This is part of the larger problem of metadata management, which has been investigated by projects such as the Storage Resource Broker [17] and Scientific Annotation Middleware [18]. For VLab, we are evaluating the use WS-Context [9], a lightweight, Web Services based metadata system. A "context" is simply a URI-named collection of XML fragments that may be arranged in parent-child relationships [10, 11]. Context servers are normally used as lightweight metadata storage locations that can be accessed by multiple collaborating web services.

In our current work, the data collected from the user interface input form (Figure 1) is written into a unique context associated with that user session. This data is stored persistently using a MySQL database, although this implementation detail is not relevant to the PWSCF developer. Each user has a base context, which is subdivided into one child context per user session. These child contexts are used to store specific input parameter values for that particular job submission. These sessions may then later be browsed and the data recovered for subsequent job submission—the form in Figure 1 has its values repopulated

Although we may store and recover values one at a time from the context storage, we are developing a way to more easily store and recover entire pages using Java Bean serialization.

PWSCF Submission			
step 1 of 3			
General			
Calculation	vc-relax	Restart Mode	from scratch
Nstep	1	Max cpu seconds	3600
dt	50	pressure	0.0
N. atoms	2	N. types	2
E conv	1.0E-4	F conv	1.0D-2
iprint	50		
Cell			
IA1 (Å)	IB1 (Å)	IC1 (Å)	
5.0	5.0	5.0	
cos γ	cos β	cos α	
0.0	0.0	0.0	
Symmetry	cubic	Cell mass	0.0020
Cell vector coordinates relative to IA1:			
A	1.0	0.0	0.0
B	0.0	1.0	0.0
C	0.0	0.0	1.0
Electrons			
<input checked="" type="checkbox"/> spin polarized	# of species	2	
<input checked="" type="checkbox"/> LDA + U	# of species	2	
Occupations	smearing	gaussian	(for occupations = smearing)
	Smearing magnitude (Ry)	0.01	
ecutwfc	72.0	ecutrho	280.0
nbnd	38	nelec	64
conv_thr	1.0d-06		
K Points			
automatic			
	Mesh	Displacement	
4	4	4	
Reset Next			

Figure 1 PWSCF input forms are developed with JSF.

holder; in other words, it is a directory where contexts with similar properties are stored. Similar to an access control list in a UNIX file system, each session directory may have associated metadata, called “session directory metadata.” Session directory metadata describes the child and parent nodes of a session. This enables the system to track the associations between sessions. One can create a hierarchical session tree where each branch can be used as an information holder for contexts with similar characteristics. These contexts are labeled with URIs, which give structured names to tree elements. For example, “vlab://users/jdoe/session1” may refer to a session directory where contexts are stored and linked to a session name “session1” and user name “jdoe”. Upon receiving the system response to a request for session creation, the user can store the context associated to the unique session identifier assigned by the WS-Context Store. This enables the WS-Context store to be queried for contexts associated to a session under consideration. Each context is stored with unlimited lifetime as the WS-Context Store is being used as an archival data store.

3. Task Management

In the previous section we described the use of JSF to create input pages for collecting input data from the user. These input values are used to create an input file of the form expected by the PWSCF application. We are now ready to make contact with the Grid. Recall that we are using a three-tiered model for our system: the portal server manages clients to Grid services, which in turn provides access to backend computing resources. Our requirements at their simplest are a)

We are developing XML serialization of the entire input page using software from the Castor Project. This will allow us to serialize entire page contents, store them into the WS-Context server, and then un-serialize them to reconstruct the input form parameter values.

A serialized Java Bean object may be stored and queried in WS-Context XML metadata store using following programming logic. Following WS-Context specifications, a Java object may be considered to be a “context”, i.e., metadata associated with a session. When storing a context, we first create a session in WS-Context store.

Here, a session can be considered an information

transfer the PWSCF input file to the desired backend resource (i. e. one with PWSCF installed on it), b) invoke the PWSCF application, c) monitor the application, and d) access the data.

Many portals have provided these capabilities, and general purpose portlets for performing these tasks are available from the Open Grid Computing Environments (OGCE) project [19] and GridSphere [8]. Java-based Grid portals and portlets quite often are based on the Java CoG Kit [12], which provides a client programming interface for interacting with standard Grid services such as GRAM, GridFTP, MyProxy, and Condor. More recently, the Java CoG has been significantly redesigned to provide abstraction layers for common tasks (job submission and remote file operations). These abstraction layers mask the differences between different Globus Toolkit versions and also support alternative tools such as job submission with Condor. The redesigned Java CoG also provides a way to group these tasks into workflow graphs that can accomplish sequences of operations. This is more thoroughly reviewed in [13].

Although existing portlets may be adapted to handle VLab tasks such uploading input files to remote machines and invoking PWSCF, this adaptation still involves a lot of work and does not reuse code. One of our goals in this project, in collaboration with the OGCE, is to simplify Grid portlet development by using JSF tag libraries that wrap the Java CoG abstraction classes. This allows us to associate multiple actions with a single HTML button click. These actions can furthermore be grouped into composite tasks that correspond directly to Java COG workflow graphs described in [13].

JSF presents us with a problem, however. It only manages individual session beans, but a user may need to submit many independent jobs within a single session, each with its own bean. Also, we must link several beans into compositions of multiple grid tasks—even the simple PWSCF submission combines file transfer and job submission tasks into a single button click. The Task Manager and Task Graph Manager described in this section represent our current solution to these problems.

The *Task Manager* handles independent user requests, or tasks, from the portlet client in Grid services. The user request-generating objects are simply Java Bean class instances that wrap common Grid actions (launching remote commands, transferring data, performing remote file operations) using Java COG classes. We define a general-purpose interface called *GenericGridBean*, which specifies the required get/set methods of the implementing class. These are data fields such as “host name”, “toolkit provider”, and so forth. *GenericGridBean* implementations include *JobSubmitBean*, *FileTransferBean* and *FileOperationBean*. When a client invokes a particular type of action, it does so indirectly through the Task Manager Bean. The TaskManager is responsible for passing property values and calling action methods of task beans. Once a user request is caught, the task manager instantiates a task bean object and its event listener. It then persistently stores them to a storage service, which in our implementation is a WS-Context service. The Storage Service has methods that can access these bean instances with a unique key called “taskname”. A JSF validator guarantees that each “taskname” parameter is unique within the session scope.

The Task Manager is also responsible for monitoring task beans and managing their lifecycles. When a task is initially submitted, we store its property values and state in the Storage Service (Figure 2). Live objects correspond to the COG states “unsubmitted”, “submitted”, and “active”. When tasks enter “completed” or related states (“failed”, “canceled”), its submission, completion dates and output file(s) metadata are stored as well. This allows us to recover the submitted job’s properties (such as input file used and execution host) for later editing and resubmission. One

current drawback to this initial scheme is that live objects will be lost when the session expires. Globus Toolkit services provide persistence with clients through callback ID handles, but we must add this capability to the Java COG.

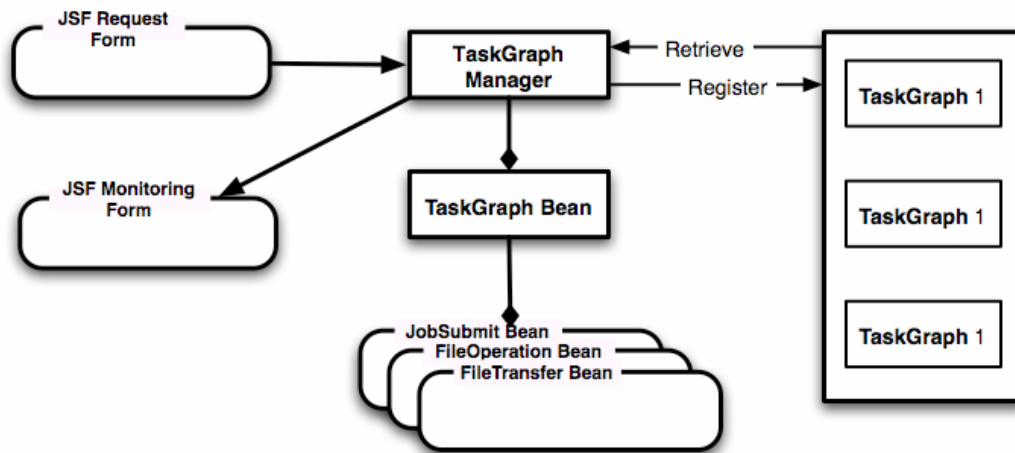


Figure 2 The Task Graph Manager manages Grid submissions with multiple steps.

Formatted: Caption, Space Before: 5 pt, After: 5 pt

Deleted: ¶

We are also investigating another solution to the persistence problem. One of the side effects of the JSF approach is that the Java Bean classes that encapsulate the Grid business logic may be developed independently of the JSF container. This will potentially allow us to develop a purely Web Service version of our system, in which our Task Beans run as standalone Web Services. This approach will simplify the management of user objects that are not directly tied to Tomcat session objects. To implement this, we are investigating the use of Inversion of Control pattern implementations like Spring [14]. Under this scenario, the Spring framework will take care of bean lifecycles and task properties will be injected into Spring through the use of web services.

In addition to managing multiple independent tasks, we must also often manage coupled tasks: a single button click may require two or more actions that depend on one another. As described in [13], the Java COG Kit provides the TaskGraph class to handle directed acyclic graph-like workflows composed of atomic tasks. As with the Task Manager previously, we have defined a *Task Graph Manager* class, which is a Java Bean wrapper around the COG TaskGraph class. The TaskGraph Manager handles multiple-step task submission, as depicted in Figure 2. The TaskGraph Manager coordinates user requests with TaskGraph backing beans. Each TaskGraph bean is itself composed of instances of the GenericGridBean implementation (JobSubmit, FileOperation, and FileTransfer Beans). We express the dependencies using JSF tag library extensions so that the JSF application developer can encode the composite task graph workflow out of reusable tags as shown in JSF snippet below. The TaskGraph Manager submits and monitors TaskGraphs through an action method when the user clicks the submit button. The TaskGraph Manager has monitoring and event handling mechanisms. The TaskGraph Manager redirects user inputs into the proper TaskGraph bean instance and manages its lifecycle. We may also use the Task Graph Manager to retrieve specific TaskGraph instances so that we may, for example, check the current status of a running job.

```
<o:taskGraph id="myGraph" binding="#{taskgraph.taskgr}" >
```

```

<o:task id="task1" method="#{task.create}" type="FileTransferInput" />
<o:task id="task2" method="#{task.create}" type="JobSubmit" />
<o:task id="task3" method="#{task.create}" type="FileTransferOutput" />
<o:taskAdd id="taskadd1" name="task1" depends="task2" />
<o:taskAdd id="taskadd1" name="task2" depends="task3" />
<o:contextStore id="context" type="taskgraph" />
</o:taskGraph>

```

JSF tags represent a multi-staged task submission. In this scenario, the task graph is composed of three tasks, and each task is dependent on previous one. The first task is to move the input file from a remote location that corresponds to a GridFTP file transfer call. The second task is to submit a job that is a GRAM job submission call. The third task is another GridFTP file transfer to stage output file(s) to a remote archival or file server. *taskAdd* tags define dependencies between tasks. The *contextStore* tag initiates the storage of taskgraph archival data onto the storage server.

4. Collaborative Web Services

The PWSCF application represents a straightforward Grid application. VLab will also need to address more challenging problems: determining best available backend resources for scheduling applications and developing collaborative web services that allow multiple clients to interact with the same service. For example, the Task Manager and TaskGraph Manager described in Section 3 specify specific host computers in the current implementation. Virtualizing the backend connection through proxy services is highly desirable for both load balancing and fault tolerance. In this section, we explain our use of message oriented middleware (NaradaBrokering) to investigate solutions to these problems.

The desired functionality ascribed to VLab demands a flexible environment that supports a variety of services related to computation, storage, visualization, database transactions, and processing. In addition, system scalability, expandability, and fault tolerance are essential components. We have chosen NaradaBrokering [2] as a middleware system that already incorporates many of these elements.

NaradaBrokering [2] is a Grid and Web Service-compatible middleware framework formed by a cooperating set of brokers whose role is to process messages sent to it, each with a topic tag, and route them to any subscribers to that topic. The use of a topic tag is the hallmark of publish/subscribe systems, which make it possible to make requests and receive replies without regard to the specific system processing the requests. Using NaradaBrokering, we have developed a prototypical network, geographically distributed to demonstrate the features of a future VLab system. The components of our architecture are illustrated in Figure 3.

We have developed a wavelet transform service for demonstration purposes. Large-scale datasets are stored on two servers in Minnesota (implemented) and Indiana (not yet implemented). A client can request the wavelet transform of a specific dataset, and have a specified number of wavelet coefficients transmitted to the client. We implement the client as an applet (which can itself be embedded in a portlet) for acceptable interactivity and image rendering. However, the approach is general and can be applied also to other Web applications (i.e. JSF portlets described in Section 2). In our wavelet application, the applet displays the coefficients as spheres centered at the location occupied by the center of the 3D wavelet. High performance graphics are obtained through the use of JOGL, an OpenGL API for Java.

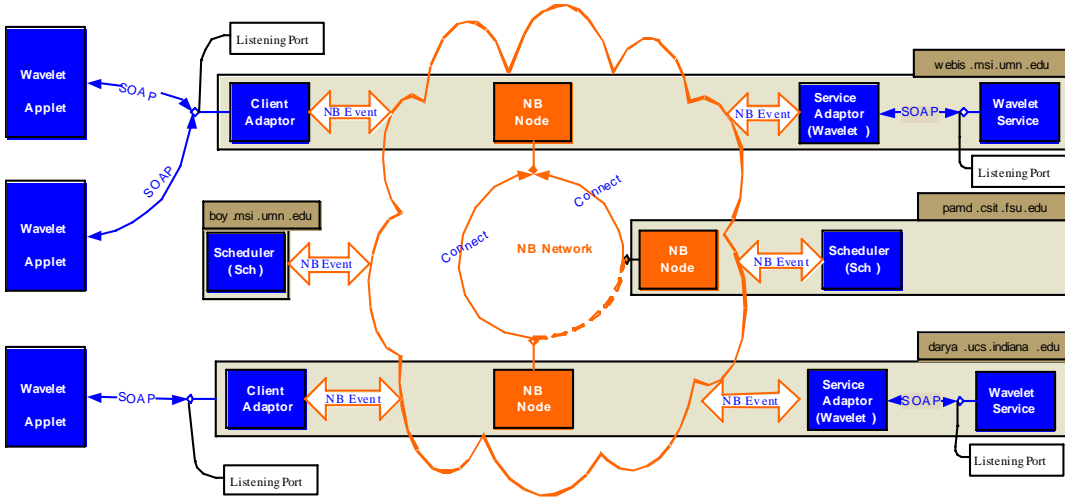


Figure 3 A network of three NaradaBrokering brokers is deployed across three sites. Attached to the network are two wavelet services (with service adaptors), two schedulers, and several client adaptors. Users access the wavelet service via collaborative applets.

Although we demonstrate our framework with a wavelet transformation service for visual impact, this service can be replaced with any other service that conforms to the WSDL standard. The system has several components that work together to isolate the user from the task of finding available services, finding the servers that support them, and distributing the workload. As shown in Figure 3, these components are the Broker Network, the schedulers, the services (with associated adaptors), the clients (with their adaptors), and the service registries (not shown). The *Broker Network* encompasses a collection of three NaradaBrokering nodes connected to one another. Any entity (client or service) publishing to any of these nodes will have its request propagated through the entire Broker Network towards one or several destinations. NaradaBrokering is responsible for efficient routing, guaranteed delivery, or in the event of problems, appropriate error feedback to the sender.

The *Schedulers* manage the task execution process. There are several schedulers to provide fault tolerance into the system, in the event that one or more servers that host the scheduler become unavailable. Different schedulers might also handle different types of requests. When a client makes a request, a scheduler is identified to handle the request and establish a connection between service and client. The scheduler is also responsible for splitting larger tasks into sub-tasks (not implemented), finding the appropriate services and submitting these tasks with due regard to task order. The *Services* are responsible for completion of requests initiated either by the clients or by the schedulers on their behalf. Data generated from these services are either returned to the clients or to the schedulers who then forward the data to the clients. Each *service adaptor* is an interface between a service (which has no knowledge of the NaradaBrokering middleware and communicates via the SOAP protocol), and NaradaBrokering, which understands messages labeled with an associated topic. Thus, the adaptors wrap and unwrap messages received by and sent to the services to make them conform to the NaradaBrokering protocol. To this end, the adaptors simply add a publishing topic, state whether message refers to a subscription or a publishing message, and optionally adds some headers to the message, such as IP addresses. *Services* are responsible for executing the tasks generated by the clients. *Clients*

may be applets, either standalone or within a portal environment, or portlets. They communicate with the middleware through the use of client adaptors, which play the same role with respect to clients as do the service adaptors with respect to services.

The key concept within our system is that each entity (any one of the system's components with the exception of the brokers), has a unique ID, and the entity subscribes to its own ID. Furthermore, each entity subscribes to a category ID related to the services it can provide. For example, a scheduler subscribes to the ID "sch", while the wavelet server might subscribe to the ID "wav". A client first sends a message with the topic "sch". Through the publish/subscribe mechanisms, all entities that subscribed to "sch" will respond, namely all the schedulers. Currently, the first scheduler to respond is chosen. If one scheduler is unavailable, the second scheduler is chosen (both messages reach the requesting client). Included in the return message from the scheduler, its unique ID is included, so that the client may connect directly to it. The client then sends its specific task request (in this case, the request to perform the wavelet transform on a specific file) to the scheduler. The message headers include the name of the desired service (in this case the service has topic "wav"). The scheduler sends a message out requesting a wavelet service that has subscribed to the topic "wav". Once found, the wavelet service returns a message with its own unique ID. The scheduler can now act as a proxy for the client, or link the client to the service directly. We note that the original client task has its own ID associated with it, and any other client that uses that same ID will share its display with the original client. The *registry* entities provide the mechanism through which additional clients can enquire as to existing client interactions (implemented). Whether or not this client can connect to the ongoing sessions is a function of the authentication model (not yet implemented).

5. Conclusions and Future Work

The work described above represents the first steps in the development of a robust, collaborative system to provide simple, yet flexible workflows for research scientists interested in conducting complex scientific computational tasks rather than worry about the intricacies of the underlying computational frameworks. To this end, we have addressed three important components of this framework: data entry, job submission, and backend services. Our work is guided by the principle of ease of use, fault tolerance, collaboration, reusable code, and persistent records. The use of the PWSCF code serves as a prototypical code, which forms a single cog in a more complex data entry, job submission, and data analysis cycle. More generally, there are several codes that must be submitted, often in large numbers, and there is often causality between results already generated and the codes to be submitted as a result of particular analyses. Implementation of a system that takes care of these dependencies (quasi-) automatically is an end objective of this project. To this end, we will focus our attention on several important components of the system.

Scheduling and Load Balancing: Although we include at least two entities of each type in our distributed system (2+ schedulers, 2+ wavelet services, etc.), there is as yet no attempt to take network and server load into account when choosing which units will perform the actual work. It is currently a first-come, first-chosen approach. We will evaluate existing work and enhance our schedulers and services to activate themselves based on a more realistic measure of instantaneous or extended load.

Collaboration: Collaboration is a natural attribute of our system. Two user tasks that subscribe to identical topics automatically receive the same information. We will investigate approaches to

achieve this collaboration both at the visual level (shared user interfaces and displays), with the possibility of multiple users controlling the input. Much work has been done in this area, albeit (to the authors' knowledge) not within the context of publish/subscribe middleware.

Workflow: Complex workflows are important within VLab. Recent research has shown how to implement strategies for specifying workflows across multiple services. This work will be integrated within our system to properly link input, job submission, analysis, feedback to the user, and finally, automatic (or semi-automatic) decisions regarding the next set of simulations to submit.

In future work, the Web Service Resource Framework [5] and particularly the WS-Notification specification family may be used to replace our pre-Web Service topic system. Support for WS-Notification is currently being developed in NaradaBrokering [15]. When this is available we will evaluate its use in our system.

This work is supported by the National Science Foundation's Information Technology Research (NSF grant ITR-0428774, 0427264, 0426867 VLab) and Middleware Initiative (NSF Grant 0330613) programs.

6. References

- [1] Virtual Laboratory for Earth and Planetary Materials Project Web Site: <http://www.VLab.msi.umn.edu/>.
- [2] Shrideep Pallickara and Geoffrey Fox *NaradaBrokering: A Distributed Middleware Framework and Architecture for Enabling Durable Peer-to-Peer Grids* in Proceedings of ACM/IFIP/USENIX International Middleware Conference Middleware-2003, Rio Janeiro, Brazil June 2003.
- [3] Democritos Scientific Software Web Site: <http://www.democritos.it/scientific.php>.
- [4] Plane-Wave Self Consistent Field Web Site: <http://www.pwscf.org/>.
- [5] I. Foster, *Globus Toolkit Version 4: Software for Service-Oriented Systems*. IFIP International Conference on Network and Parallel Computing, Springer-Verlag LNCS 3779, pp 2-13, 2005.
- [6] Hey, A. and Fox, G., eds. Concurrency and Computation: Practice and Experience, Vol. 14, No. 13-15 (2002). Special issue on Grid Computing Environments.
- [7] Abdelnur, A., Chien, E., and Hepper, S., (eds.) (2003), *Portlet Specification 1.0*. Available from <http://www.jcp.org/en/jsr/detail?id=168>.
- [8] The GridSphere Portal web site: <http://www.gridsphere.org>.
- [9] Bunting, B., Chapman, M., Hurlery, O., Little M., Mischinkinky, J., Newcomer, E., Webber J, and Swenson, K., *Web Services Context (WS-Context)*, available from http://www.arjuna.com/library/specs/ws_caf_1-0/WS-CTX.pdf.
- [10] Mehmet S. Aktas, Geoffrey Fox, and Marlon Pierce. *Managing Dynamic Metadata as Context*. June, 2005. Available from http://grids.ucs.indiana.edu/ptliupages/publications/maktas_iccse05.pdf.
- [11] Fault Tolerant High Performance Information Services Web Site: <http://www.opengrids.org>
- [12] Gregor von Laszewski, Ian Foster, Jarek Gawor, and Peter Lane, *A Java Commodity Grid Kit*, Concurrency and Computation: Practice and Experience, vol. 13, no. 8-9, pp. 643-662, 2001, <http://www.cogkit.org/>.
- [13] Kaizar Amin and Gregor von Laszewski. *Java Cog Kit Abstractions*. 2005. Available from http://www.cogkit.org/release/4_1_2/manual/abstractions.pdf.
- [14] Rod Johnson, *Expert One-On-One J2EE Design and Development*. Wrox, 2003.
- [15] Shrideep Pallickara, Harshawardhan Gadgil and Geoffrey Fox *On the Discovery of Topics in Distributed Publish/Subscribe systems*. In Proceedings of Grid 2005, Seattle, WA., Nov 2005.
- [16] Craig McClanahan, Ed Burns, and Roger Katain (eds). *Java Server Faces Specification 1.1*
- [17] [Rajasekar, 2003] Rajasekar, A., Wan, M., Moore, R., Schroeder, W., Kremenek, G., Jugatheesan, A., Cowart, C., Zhu, B., Chen, S.Y., and Olschaowsky, R. "Storage Resource Broker - Managing Distributed Data in a Grid," *Computer Society of India Journal, Special Issue on SAN, Vol. 33, No. 4, pp. 4254* Oct 2003. Other Storage Resource Broker publications available from <http://www.sdsc.edu/srb/Pappres/Pappres.html>.
- [18] Jens Schwidder, Tara Talbott, and James D Myers, *Bootstrapping to a Semantic Grid*, Proceedings of the *Semantic Infrastructure for Grid Computing Applications Workshop* at IEEE/ACM International Symposium on Cluster Computing and the Grid CCGRID 2005 May 9-12, 2005, Cardiff, UK.
- [19] The Open Grid Computing Environments Project web site: <http://www.colllab-ogce.org>.