

This item may be protected under Title 17 of the U.S. Copyright Law. It is made available by UMBC for non-commercial research and education. For permission to publish or reproduce, please contact the author.

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Facilitating E-Science Discovery Using Scientific Workflows on the Grid

Jianwu Wang^a, Prakashan Korambath^b, Seonah Kim^c, Scott Johnson^c, Kejian Jin^b, Daniel Crawl^a, Ilkay Altintas^a, Shava Smallen^a, Bill Labate^b, Kendall N. Houk^c

^a San Diego Supercomputer Center, UCSD, 9500 Gilman Drive, MC 0505, La Jolla, CA 92093, U.S.A.

^b Institute for Digital Research and Education, UCLA, 5308 Math Sciences, Los Angeles, CA 90095, U.S.A.

^c Department of Chemistry and Biochemistry, UCLA, Los Angeles, CA 90095, U.S.A.

Abstract: E-Science has been greatly enhanced from the developing capability and usability of cyberinfrastructure. This chapter explains how scientific workflow systems can facilitate e-Science discovery in Grid environments by providing features including resource consolidation, parallelism, provenance tracking, fault tolerance and workflow reuse. We first overview the core services to support e-Science discovery. To demonstrate how these services can be seamlessly assembled, a scientific workflow system, called Kepler, is integrated into the University of California Grid. This architecture is being applied to a computational enzyme design process, which is a formidable and collaborative problem in computational chemistry that challenges our knowledge of protein chemistry. Our implementation and experiments validate how the Kepler workflow system can make the scientific computation process automated, pipelined, efficient, extensible, stable, and easy-to-use.

1. Introduction

“E-Science is about global collaboration in key areas of science and the next generation of infrastructure that will enable it.”¹. Grid computing “coordinates resources that are not subject to centralized control by using standard, open, general-purpose protocols and interfaces, and deliver nontrivial qualities of service” [1]. For over a decade, Grid techniques have been successfully used to enable or facilitate domain scientists on their scientific computational problems by providing federated resources and services. Yet the software that creates and manages Grid

¹ John Taylor, Director General of Research Councils, Office of Science and Technology, UK

environments, such as the Globus toolkit², gLite³ and Unicore⁴, alone is not sufficient to manage the complex job control and data dependencies for many domain-specific problems. Such problems require combining more than one complex computational code into flexible and reusable computational scientific processes [2, 3]. Scientific workflow systems [4, 5] enable researchers to design computational experiments that span multiple distributed computational and analytical models, and in the process, store, access, transfer, and query information. This requires the integration of a variety of computational tools, including the domain-specific software, database programs as well as preparation, visualization, and analysis toolkits [2, 3].

In this chapter, we explain how scientific workflow systems can facilitate the e-Science discovery in Grid environments by providing features such as resource consolidation, parallelism, provenance tracking, fault tolerance and workflow reuse. The chapter is organized as follows. In Section 2, we summarize the core services needed for e-Science discovery. Section 3 demonstrates an assembly of these services by integrating the Kepler workflow system⁵ into the University of California Grid (UC Grid)⁶. In Section 4, an application for a theoretical enzyme design computation process is explained using the integrated architecture. We concluded the chapter in Section 5.

2. The Core Services to Support E-Science Discovery

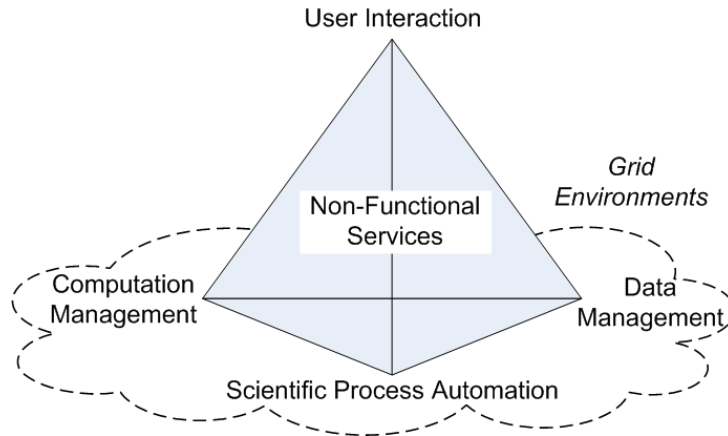


Fig 1. A core service classification to support e-Science discovery.

² <http://www.globus.org/toolkit/>

³ <http://glite.web.cern.ch/glite/>

⁴ <http://www.unicore.eu/>

⁵ <http://kepler-project.org>

⁶ <http://www.ucgrid.org/>

E-Science discoveries currently deal with increasingly growing scientific data and sophisticated computation. Usually these data and computation resources are distributed in geographically sparse locations and have a variety of usage modes. A typical computational experiment involves various tasks. By providing pipeline tools to connect the tasks and automate their execution, scientific process automation is increasingly important to help scientists easily and efficiently utilize the data and computation resources for their domain-specific scientific problems. This infrastructure should also enable scientists to interact with during the whole experiment lifecycle, such as triggering and adjusting the processes, monitoring their execution and viewing the resulting data. Further, non-functional services, such as security and failure recovery, are also important to ensure the whole process to be secure and fault tolerant. As shown in Fig. 1, computation management, data management, scientific process automation, user interaction and non-functional services are the core service categories to support e-Science discovery. Also these services are complementary to each other, and are often integrated in many e-Science projects.

In this section, we describe the main services in each category, discussing their purposes, challenges and the approaches and tools to enable them.

2.1. Computation Management

Over the past decade, Grid computation techniques have been successfully used to facilitate domain scientists with their scientific computational problems. Widely-used Grid software includes the Globus toolkit, gLite, Unicore, Nimrod/G⁷, and Condor-G⁸. More details on Grid computation management can be found in [6, 7].

2.1.1. Service-Oriented Computation

A Web service is “a software system designed to support interoperable machine-to-machine interaction over a network”⁹. Heterogeneous applications can be virtualized and easily interoperate with each other at Web service level by using the same interface definition approach. Original Web services, also called big Web services, are based on standards including XML, WSDL and SOAP. Another set of Web services, called RESTful Web services, are simple Web services implemented based on HTTP protocol and the principles of Representational State Transfer (REST) [8]. The Open Grid Services Architecture [9] defines uniform exposed service semantics for Grid components, called *Grid services*, such that Grid functionalities can be incorporated into a Web service framework.

⁷ <http://messagelab.monash.edu.au/NimrodG>

⁸ <http://www.cs.wisc.edu/condor/condorg/>

⁹ <http://www.w3.org/TR/ws-gloss/#webservice>

Through the introduction of Web and Grid services, a large number of computational resources in different scientific domains, e.g. bioinformatics, are becoming available. This presents several new challenges including service semantics, discovery, composition and orchestration [10].

2.1.2. Local Resource Management

A compute cluster resource is a collection of compute nodes connected through a private fast interconnect fabric, e.g., Infiniband, Myrinet, or even local area network (LAN), and operated by an organization such as a university. A local resource *manager* is an application that is aware of the resources in a cluster and provides an interface for users to access them. The job submission and execution on a cluster is usually managed through resource manager software, such as the Torque¹⁰, Sun Grid Engine (SGE, name changed to Oracle Grid Engine recently)¹¹ or Load Sharing Facility (LSF)¹². A resource *scheduler*, on the other hand, can simply allocate jobs in a first in first out (FIFO) basis, or follow some complex scheduling algorithms (e.g., preemption, backfilling, etc.) like the Maui¹³, MOAB¹⁴ or SGE scheduler. A resource manager can use multiple schedulers. For example, Torque can be integrated with either the Maui or MOAB.

The cluster owner sets policies on resource consumption such as which groups have access to it, how many resources can be allocated, whether they can run parallel jobs or only serial jobs, etc. This information is fed into the resource manager and is used when the scheduler executes jobs. A scheduler constantly monitors the cluster status and recalculates job priorities according to changes in the cluster environment.

Detailed techniques on local resource management can be found in [11, 12].

2.1.3. Resource Allocation

Resource allocation is the process of assigning resources associated with a Grid for a user application. It is a key service since there are usually many available local resources and also many user applications to be executed within one *Virtual Organization* (VO), where many different groups share resources through a collaborative effort.

One challenge here, called *resource scheduling*, is how to get a resource allocation result that satisfies user requirements, since resources in the Grid are not exclusive and may meet competing user requirements. It is already proved that the

¹⁰ <http://www.clusterresources.com/products/torque-resource-manager.php>

¹¹ <http://www.sun.com/software/sge/>

¹² <http://www.platform.com/workload-management>

¹³ <http://www.clusterresources.com/products/maui-cluster-scheduler.php>

¹⁴ <http://www.clusterresources.com/products/moab-cluster-suite.php>

complexity of a general scheduling problem is NP-Complete [13]. So many approximation and heuristic algorithms are proposed to achieve suboptimal scheduling in the Grid. Scheduling objectives are usually classified into *application centric* and *resource centric* [14]. The former one targets the performance of each individual application, such as makespan, economic cost, and quality of service (QoS). The latter one targets the performance of the resource, such as resource utilization and economic profit. Although it is active research area, the proposed solutions are not ready to be widely deployed in production environments and it is still common that users or communities provide their own simple resource allocation strategies using Grid interfaces to contact each local resource.

2.2. Data Management

Data management in the Grid, sometimes called Data Grid, can be seen as a specialization and extension of the Grid as an integrating infrastructure for distributed data management. This includes data acquisition, storage, sharing, transfer, archiving, etc. Representative Data Grid software includes the Globus toolkit, OGSA-DAI¹⁵, Storage Resource Broker (SRB)¹⁶ and its more recent version called integrated Rule Oriented Data System (iRODS)¹⁷. More comprehensive scientific data management and Data Grid surveys can be found in [15, 16, 17, 18, 19].

2.2.1. Data Acquisition

In general, scientific data may be created either from computation like scientific calculations and image processing, or through data collection instruments such as an astronomy telescope, earthquake-monitoring devices, and meteorology sensors, etc.

Once the experimental data is collected, it needs to be stored and transferred to the location where computing models can be run using that data to interpret the relationship or predict future events. The data often need to be shared among many researchers.

Data acquisition in large-scale observing systems, such as the National Ecological Observatory Network (NEON)¹⁸, is an emerging application area. These systems accommodate a broad spectrum of distributed sensors and continuously generate very large amount of data in real-time. Heterogeneous sensor integration [20] and data stream processing [21, 22] are two main challenges [23]. The details of these two problems are not in the scope of this chapter.

¹⁵ <http://www.ogsadai.org.uk/>

¹⁶ <http://www.sdsc.edu/srb/index.php>

¹⁷ <https://www.irods.org/>

¹⁸ <http://www.neoninc.org/>

2.2.2. Data Storage

Reliability, failover and Input/Output (I/O) throughput are critical factors for large datasets storage. Typical solutions include storing the data through RAID¹⁹ to achieve storage reliability by providing redundancy or distributed parallel file systems using metadata tables, such as Lustre²⁰ and PVFS²¹ to get higher I/O throughput.

One challenge in the Grid is how to provide a logical and simple view for researchers to access various types of geographically distributed data storage across a Grid environment. This is commonly handled by data storage abstraction techniques. For example, a logical data identifier, rather than its physical location, is provided to users to realize uniformed and easy data access. One tool that provides data abstractions is the SRB, which is a client-server middleware that provides a uniform interface for connecting to heterogeneous data resources over a network. The Replica Location Service (RLS)²² in the Globus toolkit also support data abstraction. Additionally, both SRB and RLS support data replica functionality to manage the multiple copies of the same data, which will get better response times for user applications by accessing data from locally “cached” data stores.

2.2.3. Data Transfer

A data transfer moves data between two physical locations. It is necessary to share data among the VO or realize better computation balance and performance. Challenges here include performance, security and fault-tolerance. There are many data transfer tools, e.g., FTP, scp, SRB, GridFTP, and others. FTP (File Transfer Protocol) is one of the universally available file transfer application, which functions over a network using TCP/IP based communication protocol such as the current Internet. scp (secure copy with OpenSSL encryption) is a simple shell command that allows users to copy files between systems quickly and securely. Using the above two tools do not need the expertise in Grid systems. GridFTP is built on top of FTP for use in Grid computing with the data encryption through the Globus Grid Security Infrastructure (GSI)²³. Additionally, GridFTP can provide third party transfer, parallel streams and fault tolerance. The SRB also provides strong security mechanisms supported by fine-grained access controls on data, and parallel data transfer operations.

2.2.4. Metadata

¹⁹ <http://en.wikipedia.org/wiki/RAID>

²⁰ <http://www.lustre.org/>

²¹ <http://www.pvfs.org/>

²² <http://www.globus.org/toolkit/data/rls/>

²³ <http://www.globus.org/security/overview.html>

Metadata is usually defined as “data about data”, which is regarded as “structured data about an object that supports functions associated with the designated object” [24]. Metadata is useful to understand, access and query the designated object. The metadata structures vary for different targets and usages [25]. Commonly used metadata categories in e-Science projects consist of dataset metadata (size, creator, format, access method, etc), application metadata (applicable operation system information, license, etc), resource metadata (node number, CPU speed, memory size, disk capacity, etc) and workflow metadata (creator, language, etc).

Semantics and ontology are more sophisticated techniques to describe and process metadata [26]. A domain ontology represents the particular meanings of terms as they apply to that domain. For example, the myGrid ontology helps the service discovery and composition in bioinformatics domain [27].

2.3. *Scientific Process Automation*

Scientific process automation provides pipeline tools to automate task execution. Scientific workflows are already regarded as the main solution to realize scientific process automation [28, 29]. Workflow is a higher-level “language” comparable to classic programming languages, such as scripting and object-oriented languages. The advantages of using workflow languages for scientific process include: 1) many workflow systems support intuitive process construction by “dragging and dropping” via a graphical user interface (GUI); 2) the components or sub-workflows in workflow are easy to share and reuse; 3) many workflow languages support parallelism intuitively; 4) workflow systems usually have built-in provenance support (see Section 2.3.4 for details); 5) some workflow systems are able to dynamically optimize process execution in Grid or other distributed execution environments.

Widely-used scientific workflow systems include Kepler, Pegasus²⁴, Taverna²⁵, Triana²⁶, ASKALON²⁷ and Swift²⁸. More detailed scientific workflow surveys can be found in [4, 5, 28, 29].

2.3.1. **Workflow Model**

There are different languages for representing scientific workflows [30, 31, 32], but they generally includes three types of components: tasks, control dependencies and data dependencies [33]. For example, in Fig. 2 the tasks T2 and T3 will be ex-

²⁴ <http://pegasus.isi.edu/>

²⁵ <http://www.taverna.org.uk>

²⁶ <http://www.trianacode.org>

²⁷ <http://www.dps.uibk.ac.at/projects/askalon/>

²⁸ <http://www.ci.uchicago.edu/swift/>

executed under different conditions. Additionally, T4 needs to get data from either T2 or T3 before its execution.




Fig.2. An example workflow composed of tasks and dependencies.

The tasks in the scientific computation process need to follow certain dependency logic to be executed. Usually the dependency logic can be described using *control flow*, *data flow*, or a *hybrid* of both. In control flows, or control-driven workflows, explicit control structures, including sequence, loop, condition and parallel, describe the dependency. In data flows, or data-driven workflows, data dependencies describe the relationships among tasks. Two tasks are only linked if the downstream task needs to consume data from the output of the upstream task. The hybrid method uses both control and data dependencies to enable powerful and easy logic description. Many scientific workflow systems, e.g., Kepler, Triana and Taverna, use hybrid methods.

2.3.2. Task Parallelism

Task parallelism occurs when the tasks in one workflow can execute in parallel, providing a good execution performance. The task parallelism patterns can be classified into three basic categories: *simple parallelism*, *data parallelism* and *pipeline parallelism* [34]. Simple parallelism happens when the tasks do not have data dependency in a data-driven workflow, or are in the same parallel control structure in a control-driven workflow. Data parallelism describes parallel execution of multiple tasks while different task processing independently on different parts of the same dataset. This employs the same principle as single instruction multiple data (SIMD) parallelism [35] in computer architecture. Pipeline parallelism describes a set of data are processed simultaneously among some sequential tasks, each task processing one or more data elements of the set. For example, the tasks in the Fig. 3 are executing simultaneously, each processing its own data.




Fig. 3. An example of pipeline parallelism in workflow.

2.3.3. Workflow Scheduling

Workflow scheduling maps the tasks of a process and its data to proper computational resources, in order to meet expected performance, such as minimal execution time. After scheduling, tasks in one workflow might be executed on many local resources in parallel on the Grid.

Workflow scheduling belongs to the application centric scheduling category of resource allocation (see Section 2.1.3), and focuses on the scheduling of process-based applications. Scheduling can be done either statically or automatically. Many heuristic scheduling approaches, such as genetic algorithm, are used to achieve suboptimal solutions [36, 37]. The local computation resource capacity, its advanced reservation capability and real-time status are usually needed to make decisions during the scheduling.

2.3.4. Provenance Management

Provenance plays a critical role in scientific workflows, since it allows or helps scientists to understand the origin of their results, to repeat their experiments, and to validate the processes that were used to derive data products. A major challenge of provenance management in Grid environments is how to efficiently store provenance information and easily query it in the future. There are three typical approaches to collect the execution information for data provenance: *centralized*, *decentralized* and *hybrid*. Centralized approach stores all data generated during a workflow's distributed execution in one centralized center. However, storing the data content of all distributed nodes in a single, centralized center is inefficient, especially when the dataset size is large. In a decentralized approach, provenance data is stored on the local node when it is generated (i.e., no data needs to move to a centralized center). While it is efficient to separate data storage to each distributed node locally, it becomes difficult to query and integrate this data in the future. In the hybrid approach, provenance data are stored locally in distributed nodes, and a centralized provenance catalog is employed to maintain the metadata and location information. After finding the needed data endpoint at the provenance catalog, users can get the data content from the corresponding nodes. In the hybrid approach, the burden for data transfer is reduced compared to the centralized

provenance system, and it is easier than the decentralized approach to do future provenance tracking. Note that one risk in both decentralized and hybrid approaches is that users may not be able to access the distributed nodes after the workflow execution, which is true when a resource is only open to some users for a limited time.

2.4. User Interaction

2.4.1. User Authentication

Grid credentials are commonly used to identify users to Grid resources. All certificate signing authorities (called CAs) need to have a policy to authenticate users before they issue credentials. The user requests a Grid certificate only once and then must be authorized to use Grid resources. In some organizations like UC Grid, all users can be positively identified as members of an organization using a Security Assertion Markup Language (SAML)²⁹ assertion from a Shibboleth Identity Provider (IdP)³⁰. Typically, IdPs will provide sufficient information to issue a Grid credential such as first name, last name, unique identifier, e-mail address, etc.

Some CAs let users keep credentials in their custody whereas other organizations maintain them in a credential management server such as MyProxy³¹. In the latter case, the credentials never leave the signing machine and only the short-lived credentials (called delegated proxy credentials) are provided to users. Users can always checkout credentials from those servers. The delegated proxy credentials usually have a short lifetime of less than eight hours and will be destroyed when they expire. Many CAs additionally have an annual renewal policy for their certificate, so when users are no longer associated with the original project, their certificates can be revoked. Their identities are published in a certificate revocation list, which is then distributed to all organizations where their certificates are trusted.

2.4.2. Portal

A Grid portal is a web server that provides web interfaces to Grid services such as job submission, file transfer, job status check, and resource information monitoring like cluster status, cluster workload, job queue information, etc. Some Grid portals provided generic capabilities that can be used by many types of users

²⁹ <http://saml.xml.org/>

³⁰ <http://shibboleth.internet2.edu/about.html>

³¹ <http://grid.ncsa.illinois.edu/myproxy/>

where as others, such as Science Gateways³², typically target a specific community of users. A Grid portal stores user information such as login identifier, the Grid resources the user is entitled to, the status of their submitted jobs, etc. When users login to a portal they are redirected to a credential management server, which allows the portal to authenticate Grid services on behalf of the user through a delegated proxy credential. Because a Grid portal holds some information about the role of the user, it can also authorize some Grid services such as application pool job submission.

One challenge here is to allow users to access multiple distributed resources and applications without separate authentications, which are usually supported by single sign-on techniques. GridShib³³ allows a portal to sign proxy certificates for the uniquely identifying subject assertion from a federated single sign-on service like Shibboleth³⁴. Shibboleth implements a federated identity standard based on SAML to provide identification and attributes of members in a federation. The primary purpose of Shibboleth is to avoid multiple passwords for multiple applications, interoperability within and across organizational boundaries, and enabling service providers to control access to their resources.

2.4.3. Job Monitoring

Users often like to know the overall load status and availability of resources. They also want to know the current execution status, e.g., which tasks are executing on which computers with which data. Therefore job monitoring services are also important for user interaction.

Some general job status information can be retrieved when the job is submitted to a Grid resource. For example, the Scheduler Event Generator (SEG) service in the Globus toolkit gets the status of jobs like pending, active, or done. The SEG service queries local cluster resource managers like SGE or PBS, and the job status information can be retrieved through the command-line or optionally pulled into a Grid portal or other GUI interfaces to display to users using a subscription API.

To view overall job statistics, typically cluster monitoring systems such as Ganglia³⁵ or Nagios³⁶ are deployed and display queue information, node information and the load on the cluster. A Grid monitoring service is also deployed to collect information from each resource's cluster monitoring tool, summarize it, and display it. This provides overall Grid job statistics that can be used by managers to ensure users' needs are being met. Such services are typically modeled after the Grid Monitoring Architecture [38], which was defined by Global Grid Forum. It

³² <https://www.teragrid.org/web/science-gateways/>

³³ <http://gridshib.globus.org/>

³⁴ <http://shibboleth.internet2.edu/>

³⁵ <http://ganglia.sourceforge.net/>

³⁶ <http://www.nagios.org/>

includes three main parts: 1) a *producer*, a process that produces events and implements at least one producer API; 2) a *consumer*, a process that receives events and implements at least one consumer API; and 3) a *registry*, a lookup service that allows producers to publish the event types and consumers to search them. Examples of Grid monitoring tools include the Globus Monitoring and Discovery Service (MDS).

2.4.4. Data Visualization

Data visualization is the creation of a visual representation of data, meaning "information which has been abstracted in some schematic form, including attributes or variables for the units of information" [39]. Through the presentation of visualized data, it is easy for scientists to study, analyze and communicate with one another. The primary reference model is called the filter-map-render pipeline [40]. The filter stage includes data selection, extraction and enrichment; the map stage applies a visualization algorithm to generate a viewable scene; and finally the render stage generates a series of images from the logical scene. To optimize the performance of data visualization, especially for large data set, many parallel visualization algorithms have been developed [41, 42, 43].

2.5. Non-Functional Services

We categorize a service to be non-functional if it is not usually explicitly used by users, but useful to ensure a certain property of the e-Science discovery, such as security. These services are often transparently provided in the whole lifecycle of user usage. These services are also orthogonal to the above services in that they can be integrated with each of them.

2.5.1. Security

Several security concerns in specific services have been discussed in Section 2.2.3 and 2.4.1. A key challenge of security management in Grid environments is that there is no centrally managed security system [44]. The GSI provides secure communication through authenticated services among resources and is widely used in many Grid systems and projects. It operates on the single sign-on concept meaning a single user credential will be valid among multiple organizations that trust the certificate signing authority of that credential. In GSI, every user needs to have a pair of X509 certificates. One is the private key that is used to encrypt the data, and the other is public key that is used to decrypt the data once it reaches the target resource. This kind of process is called asymmetric encryption because the keys that encrypts and decrypts are not the same. In practice, all Grid communica-

tions including GridFTP uses short-lived proxy certificates that allow the process to act on behalf of the user. The bearer of this proxy certificate has exactly the same capabilities as the original X-509 certificate holder until the proxy certificate expires. In a Grid infrastructure these proxy certificates are often leased from a certificate management server, such as MyProxy. The advantage is that a Grid user can securely store the long-lived private key in a secure machine and release only short-lived proxy credentials on a publically accessible certificate management server during Grid computing.

2.5.2. Failure Recovery

Due to the large number of resources in a Grid, there is a high probability of a single resource component failure during an application execution. Further, providing fault tolerance in a Grid environment can be difficult since resources are often distributed and under different administrative domains. Failures may occur in every level in the Grid architecture [45], which includes: 1) computational platforms, e.g., selected nodes for a job may run out of memory or disk space while the job is still running; 2) network, e.g., a local resource is temporarily unreachable during job execution; 3) middleware, e.g., a Grid service fails to start the job; 4) application services, e.g., domain-specific application gets an execution exception due to unexpected input; 5) workflow system, e.g., infinite loops during workflow execution; or 6) user, e.g., a user credential expires during workflow execution.

Many of these issues fall into the domain of system and network administrators, who must design infrastructure to provide redundant components. Here, we address only parts at the workflow level, where redundancy, retry and migration are the main fault tolerance policies. Using simultaneous execution on redundant resources, workflow execution will have a lower chance to fail. The workflow system can also retry a failed resource or service after a certain amount of time. In migration, a workflow system will restart on a different resource. For the latter two cases, checkpoint and intermediate data usually need to be recorded so that only a sub-workflow rather than the whole workflow will be re-executed.

3. Integrating the Kepler Workflow System in the University of California Grid

To demonstrate how the services in Section 2 can be seamlessly assembled, a scientific workflow system, called Kepler, is integrated into the UC Grid. The interoperation of these services and the characteristics of our integration will be discussed. Almost all the applications we use or discuss in this architecture are open source or open standard and can be easily implemented by other organizations.

3.1. University of California Grid

The UC Grid is a federation of all participating compute resources among UC campuses. Any organization that is willing to accept or follow the established policies and trust relations of UC Grid policy management authority can also be an autonomous member of UC Grid. Typically, the UC Grid resources are open to all faculty, staff and students who are members of any UC campus. The owners of the resources determine when and how much of those resources are allocated to the UC Grid resource pool.

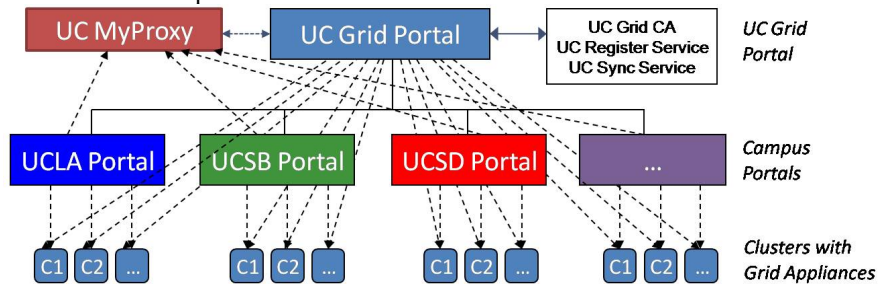


Fig.4. The architecture of University of California Grid.

As shown in Fig. 4, the UC Grid architecture mainly consists of four parts.

The UC Grid Campus Portal. This is the web interface, front-end to the UC Grid in a single campus. It provides the user interface and serves as a single point of access for users to all campus computing resources. It communicates with the Grid Appliances in a single campus. Because the users are known only at the campus level, all applications for a Grid account starts at the campus portal. The campus portal takes the request from the users and authenticates them through a Shibboleth based campus service, and if the authentication process is successful it sends a request to UC Grid Portal to sign an X-509 certificate for the user.

The UC Grid Portal. This is the web interface where users can login to access all the clusters in all ten campuses of the University of California. This is the super portal of all campus Grid portals and is the certificate signing authority for the entire UC Grid. Once the certificate is signed it is pushed to a MyProxy server that leases the short-lived credential every time a user login to the Grid portal. Any resource updates on campus portal will be updated on the UC Grid Portal immediately through a web service.

The UC MyProxy Server. The primary purpose of this server is to store user credentials and release them to UC Grid Portal when users login to the Portal.

Grid Appliance Nodes. These nodes are functionally equivalent to head nodes of a compute cluster where the UC Grid Portal software is deployed. These nodes need to be open only to the Grid portals and compute nodes inside the cluster. Additionally, these nodes need to be a job submission host to the local scheduler such as Torque, LSF or SGE.

The UC Grid is based on Globus toolkit 4.0.7, using its GridFTP, GRAM, MyProxy and MDS services. We have also implemented several services to meet our own requirements. Two important ones are the UC Sync Service and the UC Register Service. The UC Sync Service makes sure all database on both the campus portal and UC wide portal are synchronized so that they have updated user, cluster and application information. The UC Register Service is an automated process to authenticate, authorize and issue credentials for new users on the UC Grid Portal. For authentication purpose, it relies either on a Shibboleth based service or, in the absence of such services, make use of an ssh-based authentication mechanism. During the authorization process a cluster administrator has to verify the Unix user identifier of new user on a cluster.

The UC Grid Portal consists of application software that runs a portlet that is pluggable user-interface components in a portal to provide services such as Grid services. We employ Gridsphere³⁷ as our portlets container, which guarantees interoperability among portlets and portals through standard APIs.

The usage modes are quite different for users who already have access to some compute clusters and users who don't have access to any clusters. So users are classified into different user categories in the UC Grid: *cluster user* and *pool user*. Cluster users are those users with a Unix user identifier on one of the clusters who can access their resources directly without the Grid portal web interface. Pool users on the other hand are those users who do not have an account on any of the clusters or on the cluster where they want to run an application. The jobs submitted by pool users are run through guest accounts on all participating clusters when unused cycles are available on that cluster. Pool users can access the resources only by authorizing through their Grid credentials. Currently, pool users can only submit precompiled application jobs that are deployed in advance by the cluster administrator, as the Grid Portal does not have a mechanism to allow the pool users to upload their own binary files and guarantee the right run time architecture for that job. Some of the applications that pool users regularly use are Mathematica, Matlab, Q-Chem, NWChem, Amber, CPMD, etc. Typically, pool users need not worry about the target cluster as it is determined by the Grid Portal depending on the dynamic resource availability, and also the application availability on any of the clusters.

3.2. Kepler Scientific Workflow System

The Kepler project aims to produce an open source scientific workflow system that allows scientists to easily design and efficiently execute scientific workflows. Inherited from Ptolemy II³⁸, Kepler adopts the *actor-oriented modeling* [46] paradigm for scientific workflow design and execution. Each actor is designed to per-

³⁷ <http://www.gridsphere.org>

³⁸ <http://ptolemy.eecs.berkeley.edu/ptolemyII>

form a specific independent task that can be implemented as *atomic* or *composite*. Composite actors, or sub-workflows, are composed of atomic actors bundled together to perform complex operations. Actors in a workflow can contain ports to consume or produce data, called *tokens*, and communicate with other actors in the workflow through communication channels via *links*.

Another unique property inherited from Ptolemy II is that the order of execution of actors in the workflow is specified by an independent entity called *director*. The director defines how actors are executed and how they communicate with each other. Since the director is decoupled from the workflow structure, a user can easily change the computational model by replacing the director using the Kepler graphical user interface. As a consequence, a workflow can execute sequentially, e.g., using the Synchronous Data Flow (SDF) director, or in parallel, e.g., using the Process Network (PN) director [47].

Kepler provides an intuitive graphical user interface and an execution engine to help scientists to edit and manage scientific workflows and their execution. In the Kepler GUI, actors are dragged and dropped onto the canvas, where they can be customized, linked, and executed. Further, the execution engine can be separated from the user interface thereby enable the batch mode execution.

Currently, there are over 200 actors available in Kepler, which largely simplify the workflow composition. We will briefly describe the main distinctive actors that are used in this chapter.

Local Execution Actor. The *ExternalExecution* actor in Kepler is the wrapper for executing commands that run legacy codes in a workflow. Its purpose is to call the diverse employed external programs or shell-script wrappers.

Job Submission Actors. Kepler provides two sets of actors that can submit jobs to two typical distributed resources: Cluster and Grid. Each set has actors to be used for different job operations, e.g. create, submit, and status check.

Data Transfer Actors. There are multiple sets of data transfer actors in Kepler to support moving data from one location to another by different ways, e.g., FTP, GridFTP, SRB, scp, and others.

Fault Tolerance Actor. The fault tolerance actor [48], which is a composite actor that contains multiple sub-workflows, supports automatic exception catching and handling by re-trying the default sub-workflow or executing the alternative sub-workflows.

Besides the above actors, Kepler also provides the following characteristic capabilities.

Inherent and Implicit Parallelism. Kepler supports inherent and implicit parallelism since it adopts a dataflow modeling approach [49]. Actors in Kepler are independent from each other, and will be triggered once their input data are available. Since explicit parallel primitives, such as parallel-for, are not needed, workflow composition is greatly simplified. The workflow execution engine will parallelize actor execution automatically at runtime according to their input data availability.

Pipeline Parallelism. The execution of the tokens in the token set can be independent and parallel. Kepler supports pipeline parallelism by token streaming,

blocking and buffering techniques [34].

Provenance Tracking. Kepler provenance framework [50] supports collection and query on workflow structure and execution information in both local and distributed environments [51].

3.3. Integrated Architecture

By integrating Kepler into UC Grid, an overall architecture is presented in Fig. 5. Most of the services described in Section 2 are supported here.

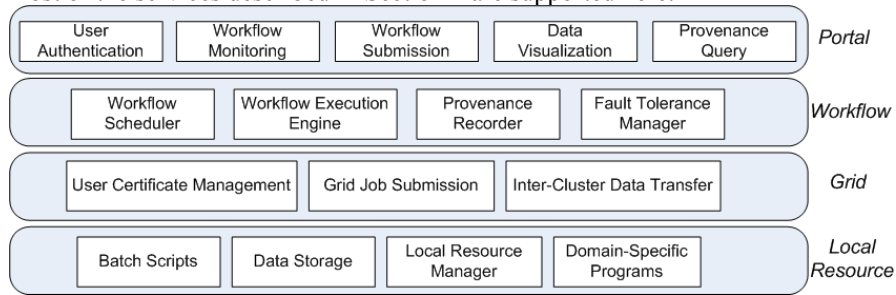


Fig. 5. A layered architecture to facilitate e-Science discovery.

3.3.1 Portal Layer

The portal layer mainly provides the User Interaction services described in Section 2.4 to enable users to interact with cyberinfrastructure and accomplish their e-Science discoveries.

User Authentication. In order for the Grid portal to execute Grid services on behalf of users, it must be provided with a delegated credential from a credential management server (MyProxy server is used here). When users login to the portal, they enter their username and corresponding MyProxy password. The portal will retrieve and store the short-lived delegation credential so that users can interact with Grid services. When the users log out, the delegated credentials are destroyed.

Workflow Submission. Users select a workflow by browsing the workflow repository in the Grid portal. The Grid portal then presents a web interface for users to configure the workflow parameters and upload the required input files that need to be staged on the target computing resources. Users are also allowed to upload their own workflows to the portal, but require authorization to avoid intentional or unintentional security problems. For example, malicious code might be inserted into a workflow by executing file delete commands in the Kepler local execution actor.

Once approved, the workflow is uploaded to the workflow repository and made available to other users through the Grid portal.

Workflow Monitoring. After a workflow is submitted through the Globus GRAM service for execution, the portal will monitor the progress and completion of the workflow execution by querying the Globus SEG service running on the target resource. Users can log in anytime to check the status, or the portal will send a notification email when the execution is done.

Data Visualization. Users can either download the output data to their local computer to visualize the data locally or chose one of the deployed application visualization services, e.g. Jmol³⁹, to visualize the data through the portal itself.

Provenance Query. During and after workflow execution, users can check provenance information via the query user interface in the portal. The query can utilize the provenance information from all previous workflow runs. For example, a user may want to understand how a parameter change influenced the results of one workflow, or how workflows are connected by a certain dataset.

3.3.2 Workflow Layer

The workflow layer provides the Scientific Process Automation services described in Section 2.3.

Workflow Scheduler. Currently, static workflow scheduling is supported by explicitly describing the scheduling policy in a separate workflow. Globus GRAM jobs are submitted through the workflow to initiate workflow task execution on resources. For workflow task execution on remote resources, data needs to be staged in before execution and staged out after execution. The capabilities of the resources must be known to achieve better overall load balancing. Sophisticated dynamic workflow scheduling capability is being studied as future work, such as automatic data staging and optimally splitting input data on multiple resources based on the resources' capability and real-time load status.

Workflow Execution Engine. Once an invocation request is received from the workflow scheduler along with the corresponding workflow specifications, the workflow execution engine will parse the specification, and execute the actors based on their dependencies and director configuration.

Provenance Recorder. During workflow execution, the provenance recorder listens to the execution messages and saves corresponding information generated by the workflow execution engine, such as an actor execution time and its input/output contents. It also saves a copy of the workflow for future verification and re-submission. The provenance will be stored locally and optionally merged into a centralized provenance database upon the workflow execution completion.

³⁹ <http://jmol.sourceforge.net/>

Fault Tolerance Manager. The fault tolerance manager will be triggered by exception messages generated from workflow execution engine, and will check whether the source of the exception is contained within the fault tolerance actor. If so, the alternative sub-workflows defined in the fault tolerance actor will be invoked with the corresponding configuration policy. Otherwise, the exception will be passed to the next higher fault tolerance actor, until the top level is reached, where the workflow execution failure message will be reported and workflow execution stopped.

3.3.3 Grid Layer

The Grid layer consolidates multiple resources and manages their computation and data resources, providing unified services for the portal and workflow layer. This layer is where the Globus toolkit software, such as GRAM, GridFTP and MyProxy Server, is located.

User Certificate Management. There should be at least one certificate signing authority, which stores the long-lived credentials for the user. The short-lived credentials are then pushed into a MyProxy server. The portal gets the delegated credential when users login to the portal. A workflow can also get the delegated credential through a MyProxy actor.

Grid Job Submission. A GRAM service will enable job submission on any accessible resources. A workflow execution can invoke multiple GRAM services on different resources to realize parallel computation.

Inter-Cluster Data Transfer. GridFTP permits direct transfer of files between the portal and the cluster resources or vice versa. Third party transfers can also be made between two clusters. A GridFTP actor is used to transfer data during workflow execution.

3.3.4 Local Resource Layer

The services provided by each compute cluster resource are located in the local resource layer. The compute nodes are not accessible directly from a portal. UC Grid communicates with the cluster through its Grid Appliance Node (see its details at Section 3.1) with both public and private interfaces.

Batch Scripts. The batch job script syntax and terminologies vary according to the type of scheduler. The Globus GRAM service executes a job manager service at the target to create the job submission script using information such as executable name, number of processors, memory requirement, etc. It is also useful to create ‘shims’ in the workflow [52], such as creating job scripts in accordance with the scheduler configuration of the host cluster.

Data Storage. Each cluster provides data storage service for their users up to a certain size limitation. Cluster users can store their data permanently on the clusters they can access, whereas pool users must save their data in portal provided storage servers. The data generated on the cluster by pool users must be downloaded; otherwise the data gets cleaned up periodically to make room for other pool users.

Local Resource Manager. A local resource manager is used to manage the job submission in a cluster. Several local resource managers are supported by the GRAM, such as SGE and PBS, which schedule jobs to the nodes in the cluster.

Domain-Specific Programs. Domain-specific programs are deployed to clusters to accomplish certain domain-specific computation problems. The applications widely used in UC Grid include Mathematica, Matlab, Q-Chem, NWChem, Amber, CPMD, etc.

4. Application in Computational Chemistry

In this section, we demonstrate how the services and integrated architecture described in Section 2 and 3 can facilitate e-Science discovery by applying them to a challenging application in computational Chemistry. The detailed information about the application can be found at [53].

4.1. Theoretical Enzyme Design Process

Enzymes are nature's protein catalysts that accelerate and regulate all metabolic reactions. To design new catalysts computationally and then to make them with molecular biological techniques will be a breakthrough in technology. An inside-out design approach has been developed [54]. In the process, quantum mechanical calculations give a theozyme [55], or theoretical enzyme, which is theoretical optimum catalyst. Numerous protein scaffolds are then screened to determine which can be used to display the side chains to mimic the geometry of the theozyme; this procedure generates a large library of potential catalysts that are evaluated for fidelity to the theozyme. The best of these are subjected to mutations to obtain a design that will both fold and catalyze the reaction. Typically, a theozyme with active sites is matched at least once per scaffold (226 protein scaffolds so far) to potentially accommodate the orientation of the model functional groups. The computation process needs to be repeated for many times with different theozymes and calculation options. The goal of this application is to accelerate and facilitate this important computation- and data-intensive process for chemists.

4.2. Conceptual Enzyme Design Workflow




Fig. 6. Conceptual workflow of enzyme design process.

As shown in Fig. 6, the conceptual enzyme design workflow takes quantum mechanical theozymes as inputs, and goes through three discrete steps before validation by experiments. The goal of this conceptual workflow is to standardize the enzyme design process through automation, and eliminate the need for unnecessary human interaction. Sequences of tasks in the enzyme design process are repeated using the same series of programs, which must be executed to design and evaluate an enzyme for a new chemical reaction. For example, a theozyme with active sites is matched at least once per scaffold to potentially accommodate the orientation of the model functional groups.

The entire enzyme design process can be performed independently for different scaffolds, and the total computation time for each scaffold can vary. The number of matches per theozyme is about 100-4,000 and computation time for a single scaffold on one CPU core is usually 1-3 hours. The number of enzyme designs generated by RosettaDesign per match is about 100-15,000 and computational time on one CPU core is usually 0.5-2 hours. One whole computation time required for all 226 scaffolds could take months on one single computer core and the total number of generated enzyme designs is about 7 million.

4.3. Executable Enzyme Design Workflow in Kepler

4.3.1. Workflow for Execution on the Local Cluster Resource

As shown in Fig. 7, a workflow is implemented to utilize the tens to hundreds of computing CPU cores in one cluster. The top-level workflow structure is the same as the conceptual workflow in the previous sub-section. Inside a composite actor, such as RosettaMatch, the sub-workflow dynamically creates job scripts according to a user's inputs and submits them to a job scheduler like SGE on the cluster using Kepler job actors. By distributing the jobs for all scaffolds through the job scheduler on the cluster, the jobs can be concurrently executed on many nodes on the cluster and the concurrency capability is limited only by the node capacity of the cluster.

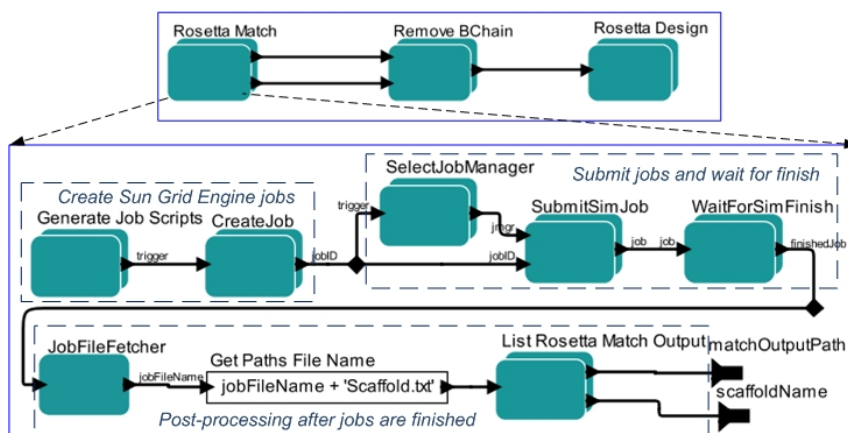


Fig.7. Kepler workflow for enzyme design processes on one cluster.

Many jobs will be executed through one composite actor, for example one job will be submitted in the RosettaMatch actor and 1-30 independent jobs in the RosettaDesign actor for each scaffold. Once a job is finished, it will trigger downstream actors, such as RemoveBChain actor, which post-process the output data from RosettaMatch. Kepler can naturally express the dependencies using connection links among actors and no explicit primitive or structure is needed to express the parallel. Using the PN director, all actors will be executed independently based on their input data availability, which can realize automatic parallel computation of these actors.

With pipeline parallel support in Kepler, the workflow input can be a vector of data, which greatly simplifies workflow construction. Each workflow execution will process many scaffolds, so we list the scaffold directory to create a vector of scaffolds and read it as input to the workflow rather than having many explicit parallel branches or loops on the whole workflow. Each scaffold does not need to wait for the completion of the RosettaMatch step for other scaffolds. It also fits that different scaffolds complete their execution within very different times. Additionally the sequence capacity can dynamically change during the workflow execution. For example, RosettaMatch can generate anywhere between 100-4,000 matches, for different scaffolds, and these matches generated as output files will dynamically be used as the input elements for the next steps in the workflow.

4.3.2. Workflow for Execution on the Grid

By adopting Globus as the Grid security and service infrastructure, the workflow shown in Fig. 8 is used to scheduling application execution among multiple cluster resources. For a local cluster, we execute the Kepler workflow shown in Fig. 7 through the Globus GRAM service deployed on the cluster. For remote clusters, besides the workflow execution, two extra tasks need to be added: 1) input data

needs to be staged in to the remote clusters in advance of the workflow execution, and 2) output data needs to be staged out from remote cluster back to the local cluster. We employ the GridFTP to do the data stage in and out. The computations on multiple clusters are independent of each other, so there is no control or data flow connections among the actors for these executions in the workflow, and the Kepler engine will run them in parallel. The workflows in Section 4.3.1 are easily reused in the two-level workflow structure. One challenge for this workflow is how to optimize the load balance on multiple clusters and the data stage in/out overhead.

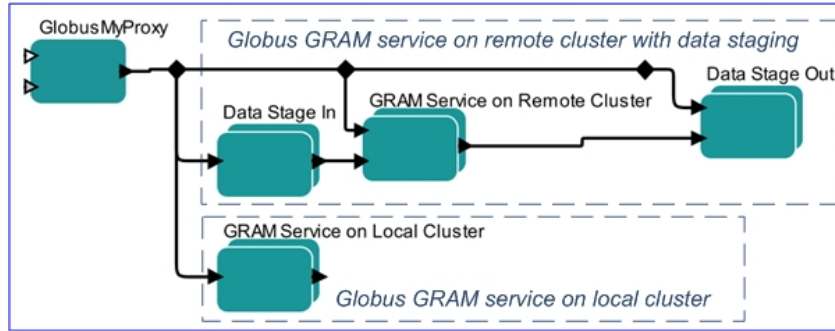


Fig. 8. Kepler workflow for enzyme design processes on multiple clusters. The GRAM service composite actor will invoke the workflow shown in Fig. 7 on the target cluster.

4.3.3. Provenance and Fault Tolerance Support in the Workflow

While using more clusters increases our computational capacity, it also increases our probability of failure. Although these exceptions happen only occasionally, the whole execution of the enzyme design workflow will crash without fault tolerance support.

To support fault tolerance at the workflow level, we adopt the Fault Tolerance actor for some sub-workflows. A simplified example is shown in Fig. 9, where the JobExecutionWithFT actor is a customized fault tolerance actor. There are two sub-workflows (shown in the bottom part of Fig. 9) in the actor, namely *default* and *clean-and-quit*, which will be invoked according to the configurations of the JobExecutionWithFT actor. The default sub-workflow submits a job, checks the job status, and throws an exception if the job fails. As specified in its configuration (shown in the right part of Fig.9), after catching the exception, the JobExecutionWithFT actor will re-execute the default sub-workflow after sleeping for 100 seconds. If the default sub-workflow still gets exceptions after three retries, the clean-and-quit sub-workflow, which cleans intermediate data and stops the whole execution, will be executed with the same input tokens.

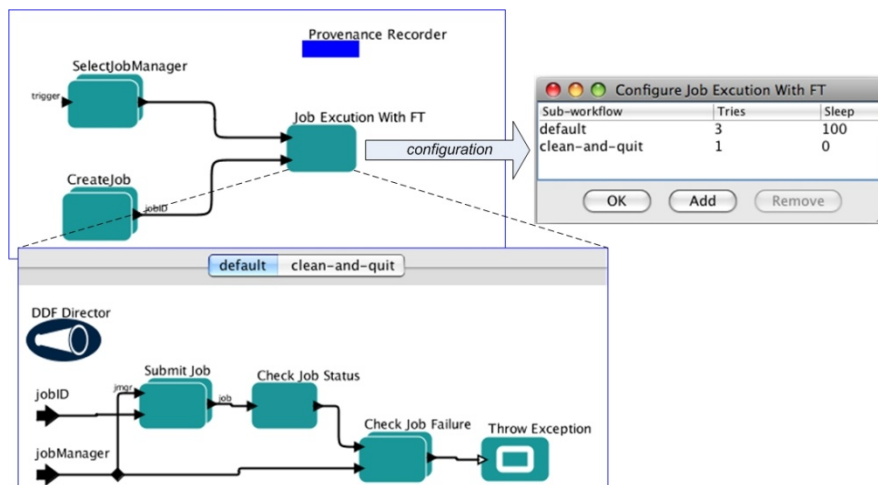


Fig.9. Fault tolerance and provenance support in Kepler workflow.

The exception handling logic can be easily expressed with the fault tolerance actor; no explicit loop and conditional switch is needed in the workflow. Further, the input data for sub-workflow retry does not need to be explicitly saved since they are automatically recorded by Provenance Recorder and will be fetched for re-execution.

Besides fault tolerance, Kepler provenance also supports collection and query on workflow structure and execution information in both local and distributed environments. Each enzyme design workflow execution will generate millions of designs and chemists may need the workflow to be executed for many times with different input models. Kepler provenance can help chemists to track the data efficiently in the future, such as querying which input model was used to generate one particular design.

4.4. Experiment

To measure the speedup capabilities of the workflows, we experimented on two clusters where Globus toolkit 4.0.7 and Sun Grid Engine job scheduler are deployed. Cluster 1 has 8 nodes, each with 24 GB of memory and two 2.6 GHz quad-core CPUs. Cluster 2 has 30 nodes, each with 4 GB of memory and two 2.2 GHz single-core CPUs.

Our first experiment executed the enzyme design workflow described in Section 4.3.1 on cluster 2 with different usable CPU core configurations. We also executed another workflow that only had the RosettaMatch part of the enzyme design workflow in order to determine the speedup difference. We ran the workflows with different inputs. As shown in Table 1, all these tests, no matter the

differences of workflow structures and inputs, have good scalability and speedup when the usable core number grows.

Table 1. Workflow execution with different usable CPU cores.

workflow	workflow input	output data	total job number	workflow execution time (unit: hour)		
				1 core	25 core	60 core
RosettaMatch	10 scaffold	0.29 GB	10	3.38	0.69	0.63
RosettaMatch	226 scaffold	27.96 GB	226	128.61	5.52	3.06
RosettaMatch + RemoveBChain + RosettaDesign	10 scaffold	10.92 GB	296	533.61	29.32	11.24

We also tested the enzyme design workflow in Section 4.3.1 and 4.3.2 to know the concurrence performance on these two clusters. From the execution data of the workflow execution only on cluster 1 and 2, we know cluster 1 is about twice as fast. So approximately twice as many inputs are distributed to cluster 1 and cluster 1 is set as the local cluster when the workflow is executed on the two clusters. The experiment data, shown in Table 2, demonstrates the good concurrence performance in the second and third test. The poor performance in the first test is because there are too few jobs compared to the number of CPU cores. We can also see the speedup ratios are not as good as those in the first experiment. The reasons are twofold: 1) it is hard to realize good load balance on multiple clusters since the execution time for different scaffold varies; 2) the data stage in and out phases for remote cluster may cause a big overhead if the transferred data is very large.

Table 2. Workflow execution with different clusters.

workflow	workflow input	total job number	workflow execution time (unit: hour)		
			cluster 1 (64 core)	cluster 2 (60 core)	cluster 1 and 2
RosettaMatch	10 scaffold	10	0.33	0.63	0.36
RosettaMatch	226 scaffold	226	1.52	3.06	1.33
RosettaMatch + RemoveBChain + RosettaDesign	10 scaffold	296	6.17	11.24	4.21

5. Conclusions

Increasingly e-Science discoveries are being made based on the enhanced capability and usability of cyberinfrastructure. In this chapter, we have summarized the

core services to support e-Science discovery in Grid environments. The five service dimensions, namely computation management, data management, scientific process automation, user interaction and non-functional services, are all important constituents and complementary to each other. To demonstrate how these services can be seamlessly assembled, we explained an integration of the Kepler workflow system with the UC Grid, and its application in computational chemistry. The implementation and experiments validate the capability of this integrated architecture to make a scientific computation process automated, pipelined, efficient, extensible, stable, and easy-to-use. We believe that, as the complexity and size of scientific problems grow larger, it is increasingly critical to leverage workflow logic and task distribution across federated computing resources to solve e-Science problems efficiently.

Acknowledgments: The authors would like to thank the rest of the Kepler and UC Grid community for their collaboration. We also like to explicitly acknowledge the contribution of Tajendra Vir Singh, Shao-Ching Huang, Sveta Mazurkova and Paul Weakliem during the UC Grid architecture design phase. This work was supported by NSF SDCI Award OCI-0722079 for Kepler/CORE, NSF CEO:P Award No. DBI 0619060 for REAP, DOE SciDac Award No. DE-FC02-07ER25811 for SDM Center, and UCGRID Project. We also thank the support to the Houk group from NIH-NIGMS and DARPA.

Reference

- [1]. Foster I (2002) What is the Grid? - a three point checklist. GRIDtoday, Vol. 1, No. 6. <http://www-fp.mcs.anl.gov/~foster/Articles/WhatIsTheGrid.pdf>
- [2] Sudholt W, Altintas I, Baldridge K (2006) Scientific Workflow Infrastructure for Computational Chemistry on the Grid. In Proceedings of the 1st Computational Chemistry and Its Applications Workshop at the 6th International Conference on Computational Science (ICCS 2006). LNCS 3993, pp. 69-76
- [3] Tiwari A, Sekhar AKT (2007) Workflow based framework for life science informatics. Computational Biology and Chemistry 31(5-6), pp. 305-319
- [4] Taylor I, Deelman E, Gannon D, Shields M (eds) (2007), Workflows for e-Science, pp. 376-394. Springer, New York, Secaucus, NJ, USA
- [5] Yu Y, Buyya R (2006) A Taxonomy of Workflow Management Systems for Grid Computing. J. Grid Computing, 2006 (3), pp.171-200
- [6]. Foster I, Kesselman C (eds) (2003) The Grid: Blueprint for a New Computing Infrastructure. Morgan Kaufmann Publishers, The Elsevier Series in Grid Computing, ISBN 1558609334, 2nd edition
- [7]. Berman F, Fox GC, Hey AJG (eds) (2003) Grid Computing: Making The Global Infrastructure a Reality. Wiley. ISBN 0-470-85319-0
- [8] Richardson L, Ruby S (2007) RESTful Web Services. O'Reilly Media, Inc., ISBN: 978-0-596-52926-0
- [9] Foster I, Kesselman C, Nick J, Tuecke S (2002) The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration. www.globus.org/research/papers/ogsa.pdf
- [10] Singh MP, Huhns MN (2005) Service-Oriented Computing: Semantics, Processes, Agents. John Wiley & Sons
- [11] Buyya R (ed.) (1999) High Performance Cluster Computing: Architectures and Systems. Volume 1, ISBN 0-13-013784-7, Prentice Hall, NJ, USA
- [12] Buyya R (ed.) (1999) High Performance Cluster Computing: Programming and Applications. Volume 2, ISBN 0-13-013785-5, Prentice Hall, NJ, USA

- [13] El-Rewini H, Lewis TG, Ali HH (1994) Task Scheduling in Parallel and Distributed Systems, ISBN: 0130992356, PTR Prentice Hall
- [14] Dong F, Akl SG (2006) Scheduling Algorithms for Grid Computing: State of the Art and Open Problems. Technical Report No. 2006-504, Queen's University, Canada, <http://www.cs.queensu.ca/TechReports/Reports/2006-504.pdf>
- [15] Gray J, Liu DT, Nieto-Santisteban M, Szalay A, DeWitt DJ, Heber G (2005) Scientific data management in the coming decade, ACM SIGMOD Record, v.34 n.4, p.34-41, doi://10.1145/1107499.1107503
- [16] Shoshani A, Rotem D (eds) (2009) Scientific Data Management: Challenges, Existing Technology, and Deployment, Computational Science Series. Chapman & Hall/CRC
- [17] Chervenak A, Foster I, Kesselman C, Salisbury C, Tuecke S (2000) The data Grid: Towards an architecture for the distributed management and analysis of large scientific datasets. Journal of Network and Computer Applications. Volume 23, Issue 3, July 2000, pp. 187-200. doi:10.1006/jnca.2000.0110
- [18] Moore RW, Jagatheesan A, Rajasekar A, Wan M, Schroeder W (2004) Data Grid Management Systems. Proceedings of the 21st IEEE/NASA Conference on Mass Storage Systems and Technologies (MSST)
- [19] Venugopal S, Buyya R, Ramamohanarao K (2006) A taxonomy of Data Grids for distributed data sharing, management, and processing. ACM Comput. Surv. 38(1)
- [20] Yick J, Mukherjee B, Ghosal D (2008) Wireless sensor network survey. Computer Networks, 52(12): 2292-2330, DOI: 10.1016/j.comnet.2008.04.002.
- [21] Fox G, Gadgil H, Pallickara S, Pierce M, Grossman RL, Gu Y, Hanley D, Hong X (2004) High Performance Data Streaming in Service Architecture. Technical Report. <http://www.hpsearch.org/documents/HighPerfDataStreaming.pdf>
- [22] Rajasekar A, Lu S, Moore R, Vernon F, Orcutt J, Lindquist K (2005) Accessing sensor data using meta data: a virtual object ring buffer framework. Proceedings of the 2nd Workshop on Data Management for Sensor Networks (DMSN 2005): 35-42
- [23] Tilak S, Hubbard P, Miller M, Fountain T (2007) The Ring Buffer Network Bus (RBNB) DataTurbine Streaming Data Middleware for Environmental Observing Systems. eScience 2007: 125-133
- [24] Greenberg J (2002) Metadata and the World Wide Web. The Encyclopedia of Library and Information Science, Vol.72: 224-261, Marcel Dekker, New York
- [25] Wittenburg P, Broeder D (2002) Metadata Overview and the Semantic Web, in Proceedings of the International Workshop on Resources and Tools in Field Linguistics
- [26] Davies J, Fensel D, van Harmelen F. (eds.) (2002) Towards the Semantic Web: Ontology-driven Knowledge Management. Wiley
- [27] Wolstencroft K, Alper P, Hull D, Wroe C, Lord PW, Stevens RD, Goble C (2007) The my-Grid Ontology: Bioinformatics Service Discovery. International Journal of Bioinformatics Research and Applications, 3(3):326-340
- [28] Ludäscher B, Altintas I, Bowers S, Cummings J, Critchlow T, Deelman E, Roure DD, Freire J, Goble C, Jones M, Klasky S, McPhillips T, Podhorszki N, Silva C, Taylor I, Vouk M (2009) Scientific Process Automation and Workflow Management. In Shoshani A, Rotem D (eds) Scientific Data Management: Challenges, Existing Technology, and Deployment, Computational Science Series. 476-508. Chapman & Hall/CRC
- [29] Deelman E, Gannon D, Shields MS, Taylor I (2009) Workflows and e-Science: An overview of workflow system features and capabilities. Future Generation Comp. Syst. 25(5): 528-540
- [30] Brooks C, Lee EA, Liu X, Neuendorffer S, Zhao Y, Zheng H (2007), Chapter 7: MoML, Heterogeneous Concurrent Modeling and Design in Java (Volume 1: Introduction to Ptolemy II), EECS Department, University of California, Berkeley, UCB/EECS-2007-7, <http://www.eecs.berkeley.edu/Pubs/TechRpts/2007/EECS-2007-7.html>
- [31] Scuff Language, Taverna 1.7.1 Manual, <http://www.myGrid.org.uk/usermanual1.7/>
- [32] SwiftScript Language Reference Manual. <http://www.ci.uchicago.edu/swift/guides/historical/languagespec.php>

- [33] Wang J, Altintas I, Berkley C, Gilbert L, Jones MB (2008) A High-Level Distributed Execution Framework for Scientific Workflows. In Proceedings of workshop SWBES08: Challenging Issues in Workflow Applications, 4th IEEE International Conference on e-Science (e-Science 2008) 634-639
- [34] Pautasso C, Alonso G (2006) Parallel Computing Patterns for Grid Workflows, In: Proc. Workshop on Workflows in Support of Large-Scale Science (WORKS06)
- [35] Flynn MJ (1972) Some Computer Organizations and Their Effectiveness. IEEE Trans. on Computers, C-21(9):948-960
- [36] Wicczorek M, Prodan R, Fahringer T (2005) Scheduling of scientific workflows in the ASKALON grid environment. SIGMOD Record 34(3): 56-62
- [37] Singh G, Kesselman C, Deelman E (2005) Optimizing Grid-Based Workflow Execution. J. Grid Comput. 3(3-4): 201-219
- [38] Tierney B, Aydt R, Gunter D, Smith W, Swamy M, Taylor V, Wolski R (2002) A Grid Monitoring Architecture. GWDPerf-16-3, Global Grid Forum <http://www.wdidec.lbl.gov/GGF-PERF/GMA-WG/papers/GWD-GP-16-3.pdf>
- [39] Friendly M (2009) Milestones in the history of thematic cartography, statistical graphics, and data visualization. Toronto, York University, <http://www.math.yorku.ca/SCS/Gallery/milestone/milestone.pdf>
- [40] Haber RB, McNabb DA (1990) Visualization Idioms: A Conceptual Model for Scientific Visualization Systems. IEEE Visualization in Scientific Computing:74-93
- [41] Singh JP, Gupta A, Levoy M. (1994) Parallel Visualization Algorithms: Performance and Architectural Implications, Computer, 27(7):45-55 doi:10.1109/2.299410
- [42] Ahrens J, Brislawn K, Martin K, Geveci B, Law CC, Papka M (2001) Large-scale data visualization using parallel data streaming. IEEE Comput. Graph. Appl., 21(4):34-41
- [43] Strengert M, Magallón M, Weiskopf D, Guthe S, Ertl T (2004) Hierarchical visualization and compression of large volume datasets using GPU clusters. Eurographics symposium on parallel graphics and visualization (EGPGV04), Eurographics Association: 41-48
- [44] Welch V, Siebenlist F, Foster I, Bresnahan J, Czajkowski K, Gawor J, Kesselman C, Meder S, Pearlman L, Tuecke S (2003) Security for grid services. Proceedings of the Twelfth International Symposium on High Performance Distributed Computing (HPDC-12). IEEE Press
- [45] Plankensteiner K, Prodan R, Fahringer T, Kertesz A, Kacsuk PK (2007). Fault-tolerant behavior in state-of-the-art grid workflow management systems. Technical Report. CoreGRID, <http://www.coregrid.net/mambo/images/stories/TechnicalReports/tr-0091.pdf>
- [46] Ludäscher B, Altintas I, Berkley C, Higgins D, Jaeger E, Jones M, Lee E, Tao J, Zhao Y (2005) Scientific workflow management and the Kepler system. Concurrency and Computation: Practice and Experience, 18 (10):1039-1065
- [47] Brooks C, Lee EA, Liu X, Neuendorffer S, Zhao Y, Zheng H (2007) Heterogeneous Concurrent Modeling and Design in Java (Volume 3: Ptolemy II Domains), EECS Department, University of California, Berkeley, UCB/EECS-2007-9, <http://www.eecs.berkeley.edu/Pubs/TechRpts/2007/EECS-2007-9.html>
- [48] Mouallem P, Crawl D, Altintas I, Vouk M, Yildiz U (2010). A Fault-Tolerance Architecture for Kepler-based Distributed Scientific Workflows. accepted by 22nd International Conference on Scientific and Statistical Database Management (SSDBM 2010)
- [49] Lee EA, Parks T (1995) Dataflow Process Networks. Proceedings of the IEEE, 83(5):773-799
- [50] Altintas I, Barney O, Jaeger-Frank E (2006) Provenance Collection Support in the Kepler Scientific Workflow System. Proceedings of International Provenance and Annotation Workshop (IPAW2006):118-132
- [51] Wang J, Altintas I, Hosseini PR, Barseghian D, Crawl D, Berkley C, Jones MB (2009) Accelerating Parameter Sweep Workflows by Utilizing Ad-hoc Network Computing Resources: an Ecological Example. In Proceedings of IEEE 2009 Third International Workshop on Scientific Workflows (SWF 2009) at Congress on Services (Services 2009) 267-274

- [52] Radetzki U, Leser U, Schulze-Rauschenbach SC, Zimmermann J, Lussem J, Bode T, Creemers AB (2006) Adapters, shims, and glue-service interoperability for in silico experiments. *Bioinformatics*, 22(9):1137–1143
- [53] Wang J, Korambath P, Kim S, Johnson S, Jin K, Crawl D, Altintas I, Smallen S, Labate B, Houk KN (2010) Theoretical Enzyme Design Using the Kepler Scientific Workflows on the Grid, accepted by 5th Workshop on Computational Chemistry and Its Applications (5th CCA) at International Conference on Computational Science (ICCS 2010)
- [54] Zanghellini A, Jiang L, Wollacott AM, Cheng G, Meiler J, Althoff EA, Röthlisberger D, Baker D (2006) New algorithms and an in silico benchmark for computational enzyme design. *Protein Sci.* 15(12): 2785-2794
- [55] Tantillo DJ, Chen J, Houk KN (1998) Theozymes and compuzymes: theoretical models for biological catalysis. *Curr Opin Chem Biol.* 2(6):743-50