

# Information Services for Dynamically Assembled Semantic Grids

Mehmet S. Aktas<sup>1,2</sup>, Geoffrey C. Fox<sup>1,2,3</sup>, Marlon Pierce<sup>1</sup>

<sup>1</sup> Community Grids Laboratory, Indiana University  
501 N. Morton Suite 224, Bloomington, IN 47404  
{maktas, gcf, mpierce}@cs.indiana.edu  
<http://www.communitygrids.iu.edu/index.html>  
<sup>2</sup> Computer Science Department, School of Informatics, Indiana University  
<sup>3</sup> Physics Department, College of Arts and Sciences, Indiana University

**Abstract.** The information management requirements in systems based on Web Service Architecture principles include both the management of large amounts of relatively static services and associated semantic information as well as the management of multiple dynamic regions (sessions or subgrids) where the semantic information is changing frequently. We design a hybrid Information Service supporting both the scalability of large amounts of relatively slowly varying data and a high performance rapidly updated Information Service for dynamic regions. We use the two Web Service standards: Universal Description, Discovery, and Integration (UDDI) and Web Services Context (WS-Context). We report initial results from a prototype that is applied to sensor and collaboration grids.

## 1 Introduction

E-Science Semantic Grids can often be thought of as dynamic collection of semantic subgrids where each subgrid is a collection of modest number of services that assembled for specific tasks such as forecasting earthquakes [1] or managing an audio/video collaboration session [3]. We term an actively interacting (collaborating) set of managed services as a *Gaggle* where services are put together for particular functionality. Semantic Grid may consist of several *Gaggles* each featuring intense local activity with less intense inter-gaggle interactions. Each *Gaggle* maintains most dynamic information which is the session related metadata generated as result of interactions among Grid/Web Services. *Gaggles* are also called as Grid Processes in the China National Grid. They are sessions in the field of collaboration. An infrastructure of the Semantic Grid is discussed in [2] where Grid Processes may be defined as cooperative processes that support the definition, management and integration of business processes. We also note that *Gaggles* may be composed from other “sub” *Gaggles* hierarchically.

Extensive metadata requirements of both the worldwide Grid and smaller sessions or “gaggles of grid services” that support local dynamic action may be investigated in diverse set of application domains such as sensor and collaboration grids. For example, workflow-style Geographical Information Systems (GIS) Grids such as the Pattern Informatics (PI) application [1] require information systems for storing both semi-static, stateless metadata and transitory metadata needed to describe distributed session state information. The PI application is an earthquake simulation and modeling code integrated with streaming data services as well as streaming map imaginary services for earthquake forecasting. Another example, collaborative streaming systems such as Global Multimedia Collaboration System (GlobalMMCS) [3] involve both large, mostly static information systems as well as much smaller, dynamic information systems. GlobalMMCS is a service-oriented collaboration system which integrates various services including videoconferencing, instant messaging and streaming, and is interoperable with multiple videoconferencing technologies.

Handling information requirements of these applications requires high performance, fault tolerant information systems. These information systems must be decentralized, relocate metadata to nearby locations of interested entities and provide efficient access, storage of the shared information, as the dynamic metadata needs to be delivered on tight time constraints within a *Gaggle*. Information Services support discovery and handling of services through metadata and are vital components of Grids [4].

### **1.1. Motivations**

We identify the following problems in Information Services supporting both traditional and Semantic Grids. First, Grid Information Services need to be able to support dynamically assembled service collections gathered at any one time to solve a particular problem at hand such as calculating damages from disruptions at the time of a crisis. Most of the traditional Grid Information Services [5-6] however are not built along this model. Second, Information Services should scale in numbers and geographical area. Most existing solutions [5-6] however have centralized components and do not address scalability and high performance issues. Third, Information Services need to be able to take into account user demand changes when making decisions on metadata access and storage. Fourth, Information Services need to be able to provide uniform interface for publishing and discovery of both dynamically generated and static information. Existing Grid Information Services however do not provide such capabilities. We therefore see this as an important area of investigation. This paper presents our design of an architecture and prototype to address the identified problems above. We describe a novel architecture for fault tolerant and high performance Information Services in order to manage distributed, dynamic session related metadata while providing consistent, uniform interface to both static and dynamic metadata.

## 1.2. Requirements

We design our architecture to meet the following requirements:

**Uniformity:** The types and update frequency of information may vary in both traditional and Semantic Grids. This requires a hybrid Information Service providing a uniform interface to dynamic/static metadata and supporting both the scalability of large amounts of relatively slowly varying information and a high performance rapidly updated Information Service for dynamic regions.

**Interoperability:** Information should be accessible by diverse set of consumer services through standard interfaces to increase usability. This requires leveraging existing Web Service standards for service discovery and communication to enable Information Services and consumer services to operate effectively together.

**Persistence:** Archival of session metadata may provide a metadata management system enabling session failure recovery or replay/playback capabilities for collaboration grids. This requires persistent metadata storage capability.

**Dynamism:** Dynamic metadata, i.e. rapidly updated and short-lived information need to be supported in both traditional and Semantic Grids. Furthermore, metadata need to be reallocated based on changing user demands and locations. This requires Information Services that can support metadata for dynamic regions and that can provide discovery of data-systems hosting the metadata under consideration in a dynamic fashion.

**Performance:** The update frequency on short-lived metadata may vary based on applications. Here, the system is required to support dynamic changes with a fine granularity time delay for the systems with a modest number of involved services (say, up to thousand services per session).

## 1.3. Contributions and Organization

The main contributions of this paper are two-fold. First, we present a novel architecture for a WS-Context [20] complaint metadata catalog service supporting distributed or centralized paradigms. We use an extended version of UDDI [21] for slowly varying metadata and present a uniform and consistent interface to both short-lived dynamic and slowly varying quasi-static metadata. We explore the application of context (session-related dynamic metadata) management in Grid systems to correlate activities in workflow-style applications, by providing a novel approach for management of widely distributed, shared session-related dynamic metadata. We investigate the problem of distributed session management in Grid applications, by providing an approach for distributed event (session metadata) management system enabling session failure recovery or replay/playback capabilities. We also address lack of search

capabilities in Grid Information Services, by providing uniform search interface to both interaction independent and conversation-based metadata enabling service discovery through events.

Our second contribution is the application of topic-based publish/subscribe methods to the problems of dynamic replication methodology to support dynamic metadata. We utilize a multi-publisher, multicast communication middleware and a topic-based publish/subscribe messaging system as a communication middleware to exchange messages between peers.

This paper is organized as follows. Section 2 reviews the state of art in existing information services and replica hosting environments. Section 3 reviews our design for information systems to support *Gaggles* paying particular attention to distributed data management aspects of the system. We discuss the status and the evaluation of our prototype in Section 4. In Section 5, we summarize and discuss future work.

## 2 Background

Most existing decentralized solutions to Information Services can be broadly categorized by the manner of in which decentralization is realized such as a) hierarchical, structured and b) unstructured, peer-to-peer (P2P). In structured architectures, components of the system are strictly controlled and may depend on each other for publishing and discovery of information. For an example, Globus Monitoring and Discovery System (MDS4) [5] has a hierarchical architecture where there is a single top-level Information Service that presents a uniform interface to clients to access data, while the data is collected by lower-level information providers. Another example is the structured P2P systems where the nodes in the systems are equally enabled and controlled and service information is disseminated to all nodes [7, 8].

Unstructured P2P architectures can be characterized as systems where there is lack of control on the capabilities of the system nodes and where there is no organizational structure. For an example, Relational Grid Monitoring Architecture (R-GMA) [6] presents a P2P architecture where consumers directly connect to information providers to retrieve the data without intermediary nodes. An extensive survey on Grid Information Services can be found at [9]. Architectures with pure decentralized storage models have focused on the concept of distributed hash tables (DHT) [7, 8]. DHT approach assumes possession of an identifier such as hash table that identifies the service that need to be discovered. Each node forwards the incoming query to a neighbor based on the calculations made on DHT. Although the DHT approach provides good performance on routing messages to corresponding nodes, it has various limitations such as primitive query capabilities on the database operations. Here, we design an architecture which can be defined as an unstructured P2P approach to P2P/Grid environment. We use multi-publisher message broadcasting through a topic-based publish/subscribe messaging system, which support access and storage decisions among distributed nodes.

Well-defined descriptions of resources, services and data constitute metadata. Metadata can be represented using varying metadata models such as XML Schema or Semantic Web languages (RDF, OWL, etc.). Here, we are mainly concerned with managing the metadata and delivering to clients, not with knowledge processing. We presume the metadata models to be application-specific and not defined by us. To this end, we are concentrating on distributed computing problems of managing metadata in the Semantic Grid. See Section 4 of this paper for more discussion.

An approach to solve the problem of locating services of interests is the UDDI Specifications [21] from OASIS (<http://www.oasis-open.org>). The UDDI is WS-I compatible and offers users a unified and systematic way to find service providers through a centralized registry of services. We identify the following limitations in UDDI Specifications. First, UDDI introduces keyword-based retrieval mechanism. It does not allow advanced metadata-oriented query capabilities on the registry. Second, UDDI does not take into account the volatile behavior of services. So, there may be stale data in registry entries. Third, UDDI does not support extensive metadata requirements of rich interacting systems. For instance, services may require an Information Service to publish and discover session metadata generated by one or more services as a result of their interactions. Fourth, since UDDI is domain-independent, it does not provide domain-specific query capabilities such as geo-spatial queries.

There have been some solutions introduced to provide better retrieval mechanism by extending existing UDDI Specifications. UDDI-M [10] and UDDIe [11] projects introduce the idea of associating metadata (name-value pairs) and lifetime with UDDI Registry. UDDI-M<sup>T</sup> [12-13] improves this approach in several ways such as improving the metadata representation from attribute name-value pairs into RDF triples to provide semantically rich service descriptions and relevant information. The Grimoires registry project (<http://twiki.grimoires.org/bin/view/Grimoires/WebHome>) extends the UDDI-M<sup>T</sup> to provide a registry which can support multiple service description models by taking into account robustness, efficiency and security issues. Another approach to leverage UDDI Specifications was introduced by METEOR-S [14] project which also utilizes semantic web languages when describing a service (such as data, functionality, quality of service and executions) in order to provide more expressiveness power and better service match-making process.

In our design, we too extend UDDI information model by providing an extension where we associate metadata with service descriptions. We use (name, value) pairs to describe characteristics of services similar to the UDDI-M and UDDIe projects. We expand on the capabilities that are supported by these projects, by providing domain-specific query capabilities. An example for domain-specific query capability could be XPATH queries on the auxiliary and domain-specific metadata files stored in the UDDI Registry. Another distinguishing aspect of our design is the support for session metadata. Our design supports not only quasi-static, stateless metadata, but also more extensive metadata requirements of interacting systems. UDDI-M<sup>T</sup> and METEOR-S projects are example projects that utilize semantic web languages to provide better service matchmaking in retrieval process. This research has been definitely investi-

gated [12-14] and so not covered in our design. We view dynamic and domain-specific metadata requirements of sensor/GIS and collaboration Grids as higher priority.

We use replication, a well-known and commonly used technique to improve the quality of metadata hosting environments, in our architecture. Sivasubramanian et al. [15] give an extensive survey on reviewing research efforts on designing and developing World Wide Web replica hosting environments, as does Robinovich in [16], paying particular attention to dynamic replication. As the nature of our target data is dynamic, we focus on data hosting systems that are handling with dynamic data. These systems can be discussed under following important design issues: a) distribution of client requests among data replicas b) selection of hosting environments for replica placement c) consistency enforcement.

**Distribution of client requests** is the problem of redirecting a client to the most appropriate replica server. Most existing solutions to this problem are based on DNS-Server such as in [17-18]. These solutions utilize a redirector/proxy server that obtains physical location of collection of data-systems hosting a replica of the requested data, and choose one to redirect client's request.

**Replica placement** is another issue that deals with selecting data hosting environments for replica placement and deciding how many replicas to have in the system. Existing solutions, that apply dynamic replication, monitor various properties of the system when making replica placement decisions [18-19]. For instance, Radar [18] replicates/migrates dynamic content based on changing client demands. Spread [19] considers the path between the data-system and client and makes decisions to replicate dynamic content on that path.

The **consistency enforcement issue** has to do with ensuring all replicas of the same data to be the same. Various techniques have been introduced in consistency management. For instance, the Akamai project [17] introduces versioning where a version number is encoded to document identifier, so that client would only fetch the updated data from the corresponding data hosting system. Radar [18] applies primary-copy approach where an update can be done only on the primary-copy of the data.

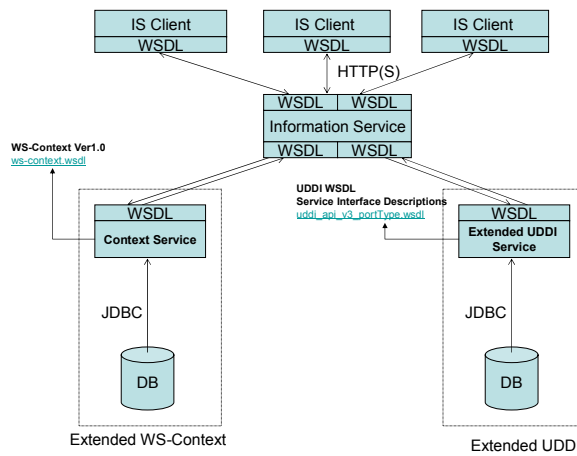
Our architecture differs from web replica hosting systems in the following ways. First, the intended use of our architecture is not to be a web-scale hosting environment. The scale of our target systems is in the order of a few dozen to at most a thousand entities participating in a session. Our target domains range from collaboration systems such as GlobalMMCS project to geographical information systems such as Pattern Informatics GIS-Grid. The participant entities of these systems might dynamically generate metadata during a session. Such metadata can be expected to be small in size and big in the volume depending on the Grid application. Second, existing solutions to dynamic replication assume all data-hosting servers to be ready and available for replica placement and ignore "dynamism" in the network topology. In reality, data-systems can fail anytime and may present volatile behavior. We use a

pure Peer-to-Peer approach, which is based on multi-publisher multicast mechanism, when distributing access and storage requests to data-systems.

### 3 Information Services

We have designed a novel architecture to Information Services presenting a uniform interface to support handling and discovery of not only quasi-static, stateless metadata, but also session related metadata. In order to be compatible with existing Grid/Web Service standards, we based the interface of our system on the WS-Context [20] and UDDI [21] Specifications. We have extended and integrated both specifications to provide uniform and consistent service interface to both dynamic and static metadata as illustrated in Figure 1.

Our approach is to utilize the existing state-of-art systems for handling and discovering static metadata and address the problems of distributed management of dynamic metadata. To do this, on receiving querying/publishing metadata requests, the system applies following steps to process service metadata. First, the system separates dynamic and static portions of the metadata. For instance, static metadata could be throughput or location of a service whereas dynamic metadata could be session identifier pointing to a workflow session in which the service is participating. Second, the system delegates the task of handling and discovery of static portion of the metadata to UDDI. As we research UDDI Specifications to integrate with our system, we have encountered various limitations in its capabilities which we address in a separate paper [22]. Third, the system itself provides handling and discovery using dynamic portions of the metadata in the metadata replica hosting environment.



**Fig.1.** Our design integrates both UDDI and WS-Context Metadata Services to provide uniform programming interface to service metadata.

The intended use of our approach is to support information in dynamically assembled Semantic Grids where “real-time” decisions are being made on which services to tie together in a dynamic workflow to solve a particular problem. One may think of WS-Context complaint Information Services as the metadata catalog for semantic metadata as in an RDF triple store. The semantic metadata expresses the relationships between resources and the applications that access the metadata catalog deduct further (inferred) information. In our design, the distinctive semantic richness comes from the highly dynamic architecture with metadata from more than two services (in contrast WS-Transfer, WS-Metadata Exchange Specifications that only easily get semantic enhancement from the two services that exchange metadata). We discuss various research issues in building Information Services for dynamically assembled Semantic Grids in the following section.

### **3.1. Fault Tolerant High Performance Information Services**

We have considered two application domains from sensor/GIS and collaboration grids to demonstrate the use of our system: GlobalMMCS and PI GIS-Grid. GlobalMMCS is a peer to peer collaboration environment where videoconferencing sessions can take place. Any number of widely distributed services can attend to a collaboration session. GlobalMMCS requires persistent archival of session metadata to provide replay/playback and session failure recovery capabilities. The PI GIS-Grid is a workflow-style Grid application which requires storage of transitory metadata needed to correlate activities of participant entities. Both application domains require a decentralized metadata hosting environment which can support both scalability (of large amounts of information) and performance requirements (of rapidly updated dynamic information). To this end, we identify two important research issues that need to be answered in our design: fault tolerance and high performance.

We use replication technique to provide fault tolerance and high performance which improves the quality of our data hosting environment. If one of the redundant storage elements goes down, it automatically consults remaining elements to restore itself. The replication technique can also lead into high performance by reducing the time between a client issuing a request and receiving the corresponding response. As the nature of our data is very dynamic, we use dynamic data replication technique, where data replicas may be created, deleted, or migrated among hosting data-systems based on changing user demands [16]. Two important aspects of dynamic replication are access and storage algorithms.

### **3.2. Access Algorithm**

The access algorithm distributes client requests to appropriate replica hosting data-systems. Our model is based on pure Peer-to-Peer approach where each node can probe all other nodes in the network to look up metadata. A primary role of the access algorithm is the discovery of one or more data-systems hosting the requested metadata. This discovery process consists of two steps: data-system discovery and access.



The first step concerns with selection of data-systems that can answer the client requests. The second step is to inform the data-system that is most appropriate for handling the request. In the first step, to find metadata, a node sends a probe message to all other nodes through a software multicast mechanism; target data-systems that host the metadata matching the probe send a response directly to requestor node. Here, response message consists of information regarding how well the data-system can handle this query. For instance, such information may include proximity information between the client and the data-system. On receiving response messages, the requestor node chooses the most appropriate data-system that can handle the request. In the second step, the requestor node sends the client request to the chosen data-system particularly asking to handle the request.

### **3.3. Storage Algorithm**

Storage algorithm selects data-systems for replica placement and decides how many replicas to have in the system. In our design, storage decisions are made autonomously at each node without any knowledge of other replicas of the same metadata. The storage decision is made based on the client requests served by that node. Storage process consists of two separate steps such as metadata placement and metadata creation. The first step has to do with selection of data-systems that should hold the replica and the second step has to do with metadata replica creation. In the first step, each node (data-system) runs the storage algorithm which defines client request thresholds for replica creation and deletion. If a metadata entry is in high demand which is above a pre-defined threshold, then the metadata is replicated. If a metadata entry is in low demand which is below a pre-defined threshold, it will be deleted. To replicate metadata, a node sends a “storage” message to all other nodes through a software multicast mechanism; target data-systems, that have available space, send a respond to directly requestor node. Here, the response message consists of various decision metrics such as client proximity information. On receiving the response messages, replica placement algorithm chooses the most appropriate data-system to replicate the metadata. In the second step, the requestor node sends a replica creation message directly to the chosen data-system asking to store a replica of metadata in consideration. This process creates a dynamic metadata storage in which metadata is moved based on changing client demands.

### **3.4. Multi-publisher Multicasting Communication Middleware**

An importing aspect of our system is that we utilize software multicasting capability which is an important communication medium supporting the ability to send out access and storage requests to the all nodes of the system. Any node can publish and subscribe to topics which in turn create a multi-publisher multicast broker network as communication middleware. Here, the publisher does not need to know the location and identities of receivers. It publishes a message to a topic to which all nodes subscribe. We use NaradaBrokering’s (NB) [23] publish/subscribe mechanism as a communication middleware for message exchanges between peers.

### 3.5. System Components

Our proposed architecture consists of various modules such as Query and Publishing, Expeditor, Access, Storage and Sequencer Modules. Architectural design of our system is illustrated in Figure 2.

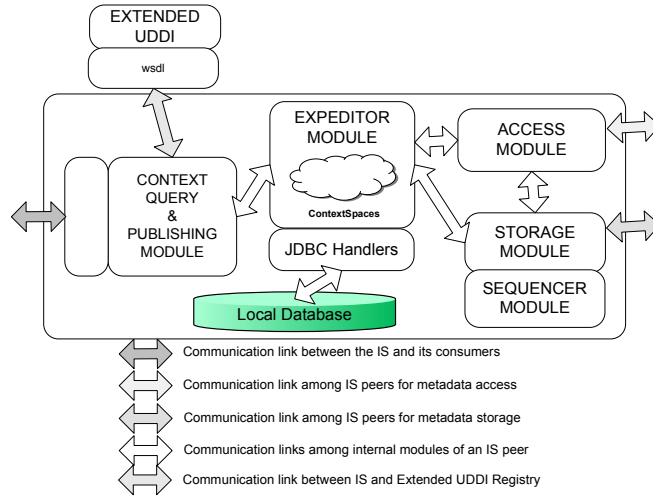


Fig.2. Architecture of an Information Service running on each peer

**3.5.1. Context Query and Publishing Modules:** These modules receive client requests through a uniform service interface for publishing/discovering dynamic and static metadata. The client query/publishing requests are processed and dynamic metadata parts of the queries are extracted. Then, the request is forwarded to Expeditor Module to find the results. Likewise, static metadata portion of the requests is relayed to external UDDI Service to publish/discover services through static metadata.

**3.5.2. Expeditor Module:** This is a generalized caching mechanism. Each node has a particular expediter. One consults the expediter to find how to get (or set) information about a dataset in an optimal fashion. The expediter is roughly equivalent to replica catalog in classic Grids. Expeditor forms a built-in memory and it maintains Context metadata objects in Context Spaces. A Context Space is an implementation of Tuple-Spaces concept [24]. Context Spaces allow us to apply space based programming to provide mutual exclusive access, associative lookup and persistence.

**3.5.3. Access Module:** This module runs the access algorithm mentioned above. It support request distribution by publishing messages to topics in NB network. It also receives messages (in respond to client request) coming from other peers and forward these query messages to Expeditor Module. The Access Module locates the nodes that are closest in terms of network distance with lowest load balance from the node re-

requesting access to the communal node in question. It also takes into account the load balance of each responding data-system when choosing the right data-system.

**3.5.4. Storage Module:** This module runs the storage algorithm. It interacts with the Expediter Module and applies the storage algorithm to all local Context metadata. If the metadata is decided to be replicated, then the storage module advertises this replication by multicasting it to available peers through NB publish/subscribe mechanism. The storage module also interacts with the Sequencer module in order to label each incoming metadata with a time stamp.

**3.5.5. Sequencer Module:** This module ensures that an order is imposed on actions/events that take place in a session. The Sequencer Module interacts with the Storage Module and labels each metadata which will be replicated in this replicated metadata hosting environment. The Sequencer Module interacts with Network Time Protocol (NTP) clients to achieve synchronized timestamps among the distributed nodes.

When receiving a query, the Query Module first processes the query and extracts the dynamic metadata portion of the query. Then, the Query Module forwards the query to Expediter, where the Expediter Module checks whether the requested data is in Context Spaces. If the Expediter Module can not find the result in Context Space or if the requested metadata is expired, then the query is forwarded to the JDBC Handler to query the data in local database. If the query asks for external metadata, then the Expediter will forward the query to Access Module, where the Access Module multicast a probe message to available Information Services through NB and communicates with the Information Services that are the original data sources for this query. The query is responded by an Information Service which may be the best qualified Information Service is to handle this query.

## 4 System Status and Evaluation

*Extended UDDI XML Metadata Services:* We have implemented extended UDDI XML Metadata Services [25] handling and discovery of static metadata based on the WS-I standard Uniform Description, Discovery, and Integration (UDDI) Specifications. We base our implementation on jUDDI (version 0.9r3), a free, open source, and java implementation of the specification. (More at <http://www.juddi.org>). jUDDI has been architected to act as the UDDI front-end on top of existing databases.

In our design, we only use a portion of the jUDDI library as UDDI-front end in order to implement extended version of UDDI XML API. We have discarded jUDDI servlet-based architecture and implemented Grid/Web Services interfaces as front access to UDDI Registries. We have enhanced jUDDI in the following ways. First, we expanded on UDDI XML Data Structure and implemented extensions to UDDI XML API to associate metadata with service entries. Second, we implemented a leasing capability. This solves a problem with UDDI repositories: information can become

outdated, so we automatically clean up entries by assigning them an expiration date. Leases on metadata may be extended. Third, we implemented GIS-specific taxonomies to describe Open GIS Consortium (OGC) compatible services such as Web Feature Services and their capabilities files. The “capabilities.xml” file is (in effect) the standard metadata description of OGC services. Finally, we implemented a more general purpose extension to the UDDI data model that allows us to insert arbitrary XML metadata into the repository. This may be searched using XPATH queries, a standard way for searching XML documents (<http://www.w3.org/TR/xpath>). This allows us to support other XML-based metadata descriptions developed for other classes of services besides GIS. The Web Services Resource Framework (WSRF), a Globus/IBM-led effort, is an important example.

Our approach allows users to insert both user-defined and arbitrary metadata into the UDDI XML metadata repository. Our work on UDDI is for a specific type of metadata: semi-static and context-free. UDDI is appropriate for data that is long-lived (i.e. should be true for months or years) and that is independent of the client interaction (i.e. all clients issuing the same requests get the same responses). We discuss the parts of our architecture that supports dynamic information in the short-lived service collections in the following section.

*WS-Context complaint XML Metadata Services:* We have implemented a centralized version of WS-Context complaint XML Metadata Services [25] handling discovery of dynamic, session related metadata. Here, session related metadata is short-lived and dependent on the client [26]. The WS-Context metadata service keeps track of context information shared between multiple participants in Web Service interactions. The context here has information such as unique ID and shared data. It allows a collection of action to take place for a common outcome. We utilize WS-Context Specification to maintain user profiles and preferences, application specific metadata, information regarding sessions and state of entities in these sessions. Each session is started by the coordinator of an activity. The coordinator service publishes the session metadata to Information Service and gets a unique identifier in return. The uniqueness of the session-id is ensured by the Information Service. Sessions can obviously be composed from other “sub” sessions hierarchically. Here, each session is associated with the participant services of that session. Dynamic session information, i.e. context, travels within the SOAP header blocks among the participant entities within the same activity. Our implementations of UDDI and WS-Context Metadata Services do not use XML databases but for efficiency convert the XML to SQL and store in MySQL database.

*Evaluation:* We have done preliminary testing [1] on the centralized version of the WS-Context XML Metadata Service’s primary operations which are GetContext and SetContext. Three measurement sets were made using a 50 byte string for GetContext. Each of the three sets consisted of 100 individual measurements. We also performed 3 sets of 100 measurements on the SetContext method. In average, we measure ~116 ms for GetContext and ~125 ms for SetContext functions to be performed. Both of these measurements are internal timings to process requests. We note that

these were subject to very large variations. We conclude from this that the actual internal processing time for small metadata pieces is typically smaller than the network invocation time and does not create an actual overhead. A more extensive evaluation regimen is being developed to determine the limits of our implementation.

## 5 Conclusions and Future Work

In this paper, we have identified an important gap in Information Services for Grids that is lack of support for dynamic information in dynamically assembled Semantic Grids. We have presented an architecture that addresses key issues of managing dynamic metadata such as a) providing an efficient metadata access and storage methodology by taking into account changes in user demands and b) providing a P2P approach for access/storage request distribution among the peers of the system to capture the dynamic behavior both in metadata and the network topology. We have discussed status of our implementation and report initial performance results from a prototype that is applied to sensor and collaboration grids.

Work remains to further develop a distributed metadata hosting environment by employing novel dynamic replication techniques and to evaluate the system as whole through extensive performance tests.

*Acknowledgement:* This work is supported by the Advanced Information Systems Technology Program of NASA's Earth-Sun System Technology Office.

## References

1. Galip Aydin, Mehmet S. Aktas, Geoffrey C. Fox, Harshwardhan Gadgil, Marlon Pierce, Ahmet Sayar. SERVOGrid Complexity Computational Environments(CCE) Integrated Performance Analysis, Accepted as poster and short paper in Grid2005, Seattle, USA
2. H. Zhuge. Semantic Grid: Scientific Issues, Infrastructure, and Methodology, Communications of the ACM. 48 (4) (2005)117-119.
3. Wenjun Wu, Geoffrey Fox, Hasan Bulut, Ahmet Uyar, Harun Altay "Design and Implementation of A Collaboration Web-services system", Journal of Neural, Parallel & Scientific Computations (NPSC), Volume 12, 2004.
4. B. Plale, P. Dinda, and G. Von Laszewski. Key Concepts and Services of a Grid Information Service. In Proceedings of the 15th International Conference on Parallel and Distributed Computing Systems (PDCS 2002), 2002.
5. Monitoring & Discovery System (MDS4) Web Site is available at <http://www.globus.org-/toolkit/mds>
6. A. Cooke, A.Gray, L. Ma, W. Nutt, J. Magowan, P. Taylor, R. Byrom, L. Field, S. Hicks, and J. Leake. R-GMA: An Information Integration System for Grid Monitoring. Proceedings of the 11th International Conference on Cooperative Information Systems, 2003.
7. Ratnasamy, Sylvia et al. A Scalable Content-Addressable Network. Proc. ACM SIGCOMM, pp 161-172, August 2001.

8. Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, Hari Balakrishnan. Chord: A Scalable Peer-to-Peer Lookup Protocol for Internet Applications. *IEEE/ACM Trans. on Networking*, 11 (1): 17-32, February 2003.
9. Serafeim Zanolis and Rizos Sakellariou. A Taxonomy of Grid Monitoring Systems. *Future Generation Computer Systems*, 21(1), January 2005, pp. 163--188.
10. V. Dialani. UDDI-M Version 1.0 API Specification. University of Southampton – UK. 02.
11. Ali ShaikhAli, Omer Rana, Rashid Al-Ali and David W. Walker. UDDIe: An Extended Registry for Web Services. *Proceedings of the Service Oriented Computing: Models, Architectures and Applications, SAINT-2003 IEEE Computer Society Press. Orlando Florida, USA, January 2003.*
12. Simon Miles, Juri Papay, Terry Payne, Keith Decker, and Luc Moreau. Towards a Protocol for the Attachment of Semantic Descriptions to Grid Services. In *The Second European across Grids Conference, Nicosia, Cyprus, pages 10, January 2004.*
13. Miles, S., Papay, J., Dialani, V., Luck, M., Decker, K., Payne, T., and Moreau, L. Personalized Grid Service Discovery. *Nineteenth Annual UK Performance Engineering Workshop (UKPEW'03), University of Warwick, Coventry, England, 2003.*
14. Verma, K., Sivashanmugam, K., Sheth, A., Patil, A., Oundhakar, S. and Miller, J. METEOR–S WSDI: A Scalable P2P Infrastructure of Registries for Semantic Publication and Discovery of Web Services, *Journal of Information Technology and Management.*
15. Sivansubramanian S., Szymaniak M., Pierre G., Steen M.V. Replication for Web Hosting Systems. *ACM Computing Surveys*. Vol. 6, No. 3, September 2004, pp. 291-334.
16. M. Rabinovich. *Issues in Web Content Replication. Bulletin of the IEEE Computer Society Technical Committee on Data Engineering, 1998.*
17. Dilley, J., Maggs, B., Parikh, J., Prokop, H., Sitaraman, R., and Wehl, B. Globally distributed content delivery. *IEEE Internet Computing* 6, 5 (Sept.), 50-58. 2002
18. M. Rabinovich, I. Rabinovich, R. Rajaraman, and A. Aggarwal. A Dynamic Object Replication and Migration Protocol for an Internet Hosting Service. *Proc. 19th Int'l Conf. Distributed Computing Systems*, pp. 101-113, June 1999.
19. P. Rodriguez, and S. Sibal. SPREAD: Scalable Platform for Reliable and Efficient Automated Distribution Computer Networks, vol. 33, nos. 1-6, pp. 33-49, June 2000.
20. Bunting, B., Chapman, M., Hurley, 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](http://www.arjuna.com/library/specs/ws_caf_1-0/WS-CTX.pdf)
21. Bellwood, T., Clement, L., and von Riegen, C. UDDI Version 3.0.1: UDDI Spec Technical Committee Specification. Available from <http://uddi.org/pubs/uddi-v3.0.1-20031014.htm>.
22. Mehmet S. Aktas, Galip Aydin, Geoffrey C. Fox, Harshwardhan Gadgil, Marlon Pierce, Ahmet Sayar, *Information Services for Grid/Web Service Oriented Architecture (SOA) Based Geospatial Applications, Technical Report, June, 2005*
23. 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.* See also: <http://www.naradabrokering.org>
24. N. Carriero and D. Gelernter. Linda in Context. *Commun. ACM*, 32(4): 444-458, 1989.
25. Extended UDDI and Fault Tolerant and High Performance Context Service Research is available at <http://www.opengrids.org>
26. Mehmet S. Aktas, Geoffrey C. Fox and Marlon Pierce. Managing Dynamic Metadata as Context. *The 2005 Istanbul International Computational Science and Engineering Conference (ICCSE2005), Istanbul, Turkey.*