An Ontology-based Grid Service for Multimedia Search

M. Amoretti, D. Bianchi, G. Conte, M. Reggiani, F. Zanichelli
Dipartimento di Ingegneria dell'Informazione, University of Parma
Parco Area delle Scienze, 181/A - 43100 Parma, Italy

{amoretti, bianchi, conte, reggiani, mczane}@ce.unipr.it

Abstract

Nowadays ICT technologies are deeply shaping the ways teaching and learning take place in the Universities. The effective exploitation of a potentially huge amount of instructional multimedia material calls for powerful search capabilities. The ideas and the tools developed in the framework of the Semantic Web can be very useful to build an enriched service for information discovery in next-generation e-learning systems.
In this paper we present an ontology-based Grid Service for searching multimedia contents within our Grid-based e-learning environment. To annotate the material we have used an ontology describing the nature of the multimedia resources.

1. Introduction

The development of next-generation e-learning systems will face very challenging tasks as ICT technology advancements give rise to higher user expectations. Ubiquitous network access, smart terminals, ever increasing computational, storage and broadband capacities (at feasible costs) have a great potential to transform the ways teaching and learning take place in the Universities and in other organizations. While current e-learning technologies are being deployed with promising results, we believe that dramatic improvements in terms of system scalability, extensibility, content adaptation and transparency to users can be obtained by service-oriented Grid Computing.

Grid Computing has emerged as a new paradigm in distributed computing with the goal of providing access to resources and services of different administrative domains in a transparent, seamless and secure way [Foster et al., 2002]. Recent significant activities coordinated by the Global Grid Forum (GGF) [GGF] are currently defining a service-oriented framework based on the concept of Grid Service, *i.e.* the *Open Grid Service Infrastructure (OGSI)* [Tuecke et al., 2003]. A major goal of this approach is the possibility to simplify the design of geographically distributed services that may interact with each other. Moreover, this environment integrates the Grid Security Infrastructure (GSI) framework, allowing the authentication of user through a trust manager.

We are investigating a campus-wide modular architecture for multimedia content access and distribution. Using Grid technologies, we are developing an extensible service for uniformly searching and accessing multimedia contents available at several local video servers. In this paper we describe an ontology-based engine which increases the flexibility of the search service, allowing to perform extended queries to semantically enriched data, not only related to the available multimedia objects but also to other entities of the academic domain.

The remaining of the paper is organized as follows. Section 2 describes the state of the art in Semantic Web and ontologies. Section 3 briefly recalls the design issues of the multimedia architecture we are deploying, and illustrates recent improvements on the prototype. Sections 4 and 5 describe the ontology-based search engine we integrated in our architecture. Section 6 provides references to other research and experimentation work in Semantic Web and service-oriented multimedia systems. Finally, section 7 provides some concluding remarks and discusses future work.

2. Ontologies

Metadata are useful to describe the content and the features of information entities and apply to many fields like digital libraries, multimedia collections, repositories of learning materials. Metadata were introduced to overcome the failure of the search engines based on keywords and boolean query (and/or of keywords) in document retrieval. In fact, systems based only on keywords cannot interpret correctly the meaning of terms, due to the lack of reference to a contest.

One of the first and well known example of metadata used to describe documents is the one proposed by The Dublin Core Metadata Initiative, that introduced a limited set of well defined relations like ”creator”, ”title” etc. Nevertheless, to add structure to metadata some higher level of representations is needed. Classification systems and controlled dictionaries or terminologies (in the form of taxonomies and Thesauri) are normally used to add semantics to the data. The problem to add semantics to data is very general and was also faced by the proposal of the *Semantic Web* [Berners Lee et al., 2001], whose aim was to make the Web understandable not only by humans but also by machines.

The Semantic Web gives rise to a number of technologies to enrich the web data. The first of these technologies is the Resource Description Framework (RDF) that allows the representation of the metadata relative to a resource not depending on the resource itself. To represent classifications, terminologies and, in general, domain knowledge the use of ontology is needed. Many languages to represent ontology were developed. The W3C has proposed the Web Ontology Language (OWL) as a new standard. Also some query languages were proposed to access data in the ontologies, mainly based on description logic.

In details, an ontology is an explicit specification of some domain. It is a formal and declarative representation which includes the vocabulary (or names) for referring to the terms in that subject area and the logical statements that describe what the terms are, how they are related to each other. Ontologies therefore provide a vocabulary for representing and communicating knowledge about some domain and a set of relationships that hold among the terms in that vocabulary. Ontology are used:

1. to give defined semantics to terms and to a representation of a domain;
2. to share and exchange knowledge about a domain and to guarantee interoperability;
3. to add flexibility to the search of content.

The ideas and the tools developed in the framework of the Semantic Web are useful to build a service for information discovery. In the following, we present a Grid Service to discover multimedia material in an educational contest for a University.

3. A Multimedia Architecture Based on Grid Services

In the following we introduce a set of objectives and requirements we are taking into account in the design of an architecture for campus-wide multimedia content access and distribution [Amoretti et al. 2004a]. Several complex, often interacting issues contribute to the effectiveness of such a rich service. In our view the main ones are the following:

- functional integration of search and delivery capabilities;
- representation level at which search is performed (*i.e.* text- or knowledge-based);
- ability to integrate pre-existing multimedia services hiding their heterogeneity;
- security management (authentication, authorization, accounting, etc.);
- ability to interact with resource managers to request resource allocations;
- conformance to standards to allow interoperability and easier integration into other systems.

We are developing a prototype using the *Globus Toolkit 3* [Globus], an open source implementation of the Open Grid Services Architecture (OGSA) [Foster et al., 2002], which is a framework based on OGSI. Service-orientation has been adopted because of the need to operate on a wide range of resources such as data, storage, network, software applications, using standard high-level descriptions. For such purposes, Grid Services appear more suitable than other middleware technologies. CORBA [CORBA], e.g., is becoming a standard solution for embedded systems, for which it provides useful services (such as Notification and Concurrency), and Real-Time features. Nevertheless CORBA, which requires that resources are mapped to distributed objects, does not provide service-oriented design handles that could simplify virtualization, *i.e.* resource access across multiple heterogeneous platforms with uniform service semantics.

Our prototype includes a *Search Service* and several *Local Search Services* (one for each content provider), as illustrated in Figure 1. The user (in our e-learning application a student interested in a recorded video lecture) accesses to the Search Service, which keeps track of the available multimedia servers and activates Local Search Services to process user queries. The search results (e.g. title, format, length, localization of a video lesson) are collected by the Search Service and returned to the user, which can then select the desired object to be requested to the particular multimedia server. Finally, a streaming session is established between the remote server and the user host.

While a university campus might host a large number of multimedia servers, many of these could be temporarily off-line or mismanaged. Thus, monitoring the state of multimedia servers is required to ensure that the search service will deal only with operational resources. A simple idea would be to include a polling functionality in the Search Service, periodically asking each resource about its aliveness. This solution can be inefficient as it increases the workload of the Search Service and reduces the efficiency of the whole system. Therefore, we chose instead a push model: each resource periodically notifies its aliveness, updating a data structure of the Search Service, that contains information about the state (active or inactive) of each multimedia server. If the update from a resource is not received within a time deadline, the Search Service attempts at querying the resource and updates accordingly the list of multimedia servers. In the following we describe in more detail an ontology-based implementation of the Local Search Service.

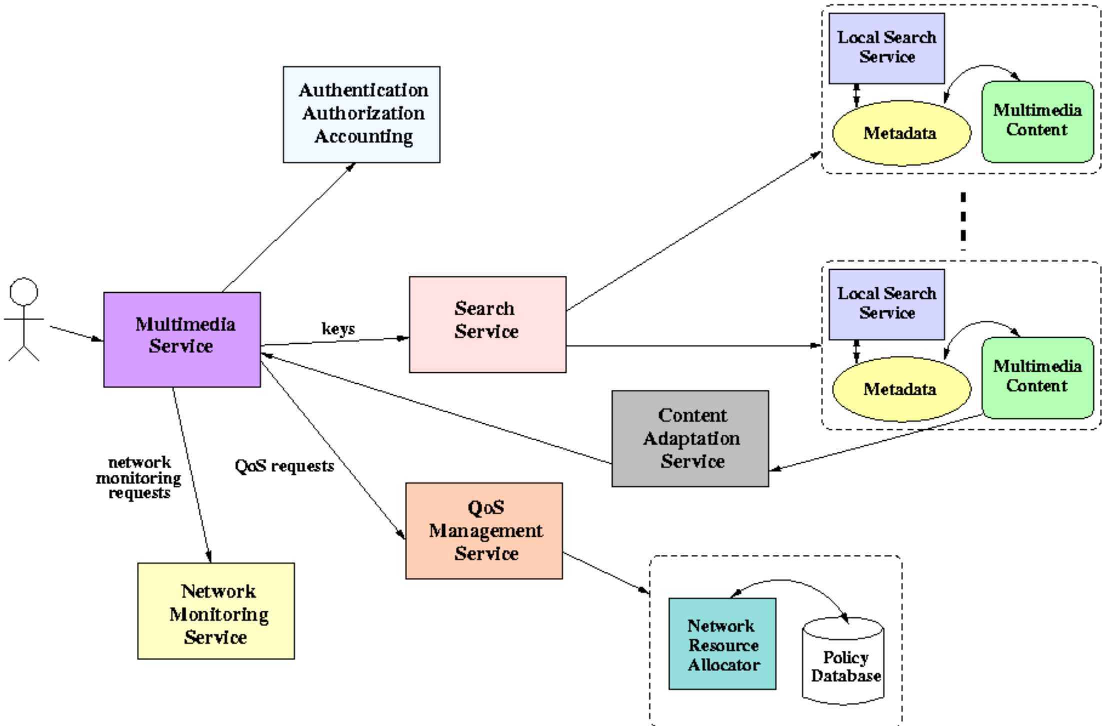

Figure 1: High-level organization of the service-oriented multimedia distribution architecture.

4. The Ontology-based Local Search Service

The heterogeneity and the actual location of the content repositories are hidden to the user, *i.e.* multimedia servers could greatly differ in terms of metadata, search capabilities and management policies. Because of that, it could be unfeasible to make the Search Service consistent and interoperable with all the remote multimedia servers.

Since the results to the queries must be presented to the user in a homogeneous format, we adopted the solution of deploying, on each multimedia server, a Local Search Service empowered with mechanisms to access the specific information data structure of the local server and to return the list of matching multimedia objects to the Search Service.

For any user query, an instance of Local Search Service is spawned on each active multimedia server. The Search Service does not connect to an existing instance. Instead, it first connects to the Local Search Service Factory, requests the creation of a new instance, uses it, and then destroys it. This kind of behaviour distinguishes Grid Services from Web Services, and in this specific context allows different queries to be processed concurrently, while subsequent refinements of a query are evaluated within the same Local Search Service instance.

The Local Search Service exposes a WSDL interface which defines the following methods:

- searchByAuthor(name, surname)
- searchByProfessor(name, surname)

- searchByKey(keyword)
- searchByTheme(themeName)
- searchByCourse(courseName)
- browse()

The latter returns the complete list of available multimedia objects, while the formers return a partial list including only specific content.

If the Local Search Service is ontology-based, it also exposes the following methods:

- loadOntology(), which initializes the ontology model;
- resetModel(), which finalizes the ontology model;
- searchMore(), which extends the search by relaxing some constraints.

The searchMore() functionality strongly relies on the Grid Service feature of statefulness. Once a search is performed with traditional mechanisms (searchBy..), the service instance caches the results which can be subsequently refined. The search engine uses the relations between the ontology individuals to extend the search domain and to return more results, increasing the probability of matching the query of the user. A Web Service would not be able to perform refined searches without continuous input submission from the client application, which should be able to store partial results.

The interaction between the search engine and the OWL ontology (which is described in the next section) is realized with *Jena* [Jena], a Java framework for building semantic applications. There are many ways of describing an ontology, and a variety of opinions as to what kinds of definition should go in one. In practice, the contents of an ontology are largely driven by the kinds of application it will be used to support. Jena does not take a particular view on the minimal or necessary components of an ontology. Rather, it tries to support a variety of common styles of use. In particular, Jena provides a programmatic environment for RDF, RDFS and OWL, including a rule-based inference engine.

Starting from the ontology, Jena creates a model which depends on both the asserted statements in the underlying RDF graph, and the statements that can be inferred by the reasoner being used, if any (in the specific case of our service, the reasoner uses the OWL rules).

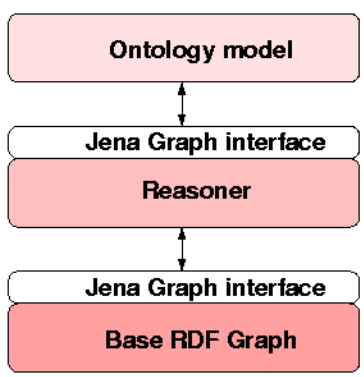

Figure 2: Ontology model construction in Jena.

As illustrated in figure 2, the asserted statements are held in the base graph. This presents a simple internal interface, Graph. The reasoner, or inference engine, can use the contents of the base graph and the semantic rules of the language, to show a more complete set of statements, *i.e.* including those that are entailed by the base assertions. This is also presented via the Graph interface, so the model works only with that interface. This allows us to build models with no reasoner, or with one of a variety of different reasoners, without changing the ontology model.

5. The Multimedia Ontology

The main concepts of our ontology are documents, as well as other entities of the academic domain such as people (teachers and students), courses and a taxonomy of the subjects of the taught disciplines.

People participating in the material production and authoring are described by the Person class and its subclasses, as illustrated in figure 3. They can be producers or publishers of documents. People can be teachers of a course or they can be involved in course activities. To extend the search capabilities, some relations between persons are introduced. People may collaborate in teaching a course or in a research project. We argue that materials produced by people with common activities and interests have a high probability to be correlated.

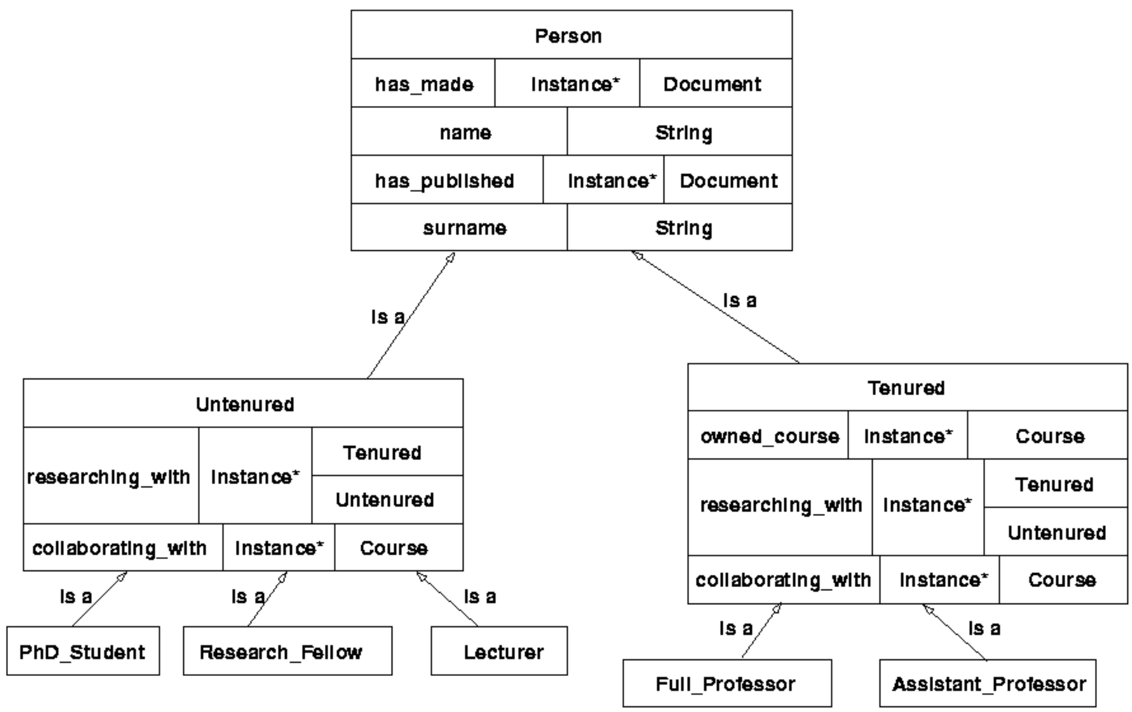

Figure 3: The Person class and its subclasses.

To annotate the material we use the Document class, which describes the nature of the searchable resources (figure 4). The latter can be in audio, video, text, image or multimedia format. Multimedia documents are composed of one or more documents of the other kinds (as happens in a SMILE content package). All documents have attributes like title, author, publisher, etc.

Each document is also described by a list of keywords. The content is described by topics (class Theme) that are derived from the taxonomy of the subject matter. As illustrated in figure 5, these topics are organized in a hierarchy that describes in detail their mutual relation. For each theme there is a link to the father node (the more general topic) and other attributes are used to specify the documents related to the topic and the courses that teach the topic. Also the courses, with their attributes (people involved in teaching, pertinent documents) are represented.

The concepts represented in the ontology allow the search of the multimedia repository using a number of different criteria. Material may be accessed by author, course, teacher, keyword or referencing a topic in the subject taxonomy. More flexible searches may be performed navigating the ontology. If the user has searched for a particular topic, then the system can look also in the siblings nodes and in the upper level of the taxonomy to find related topics.

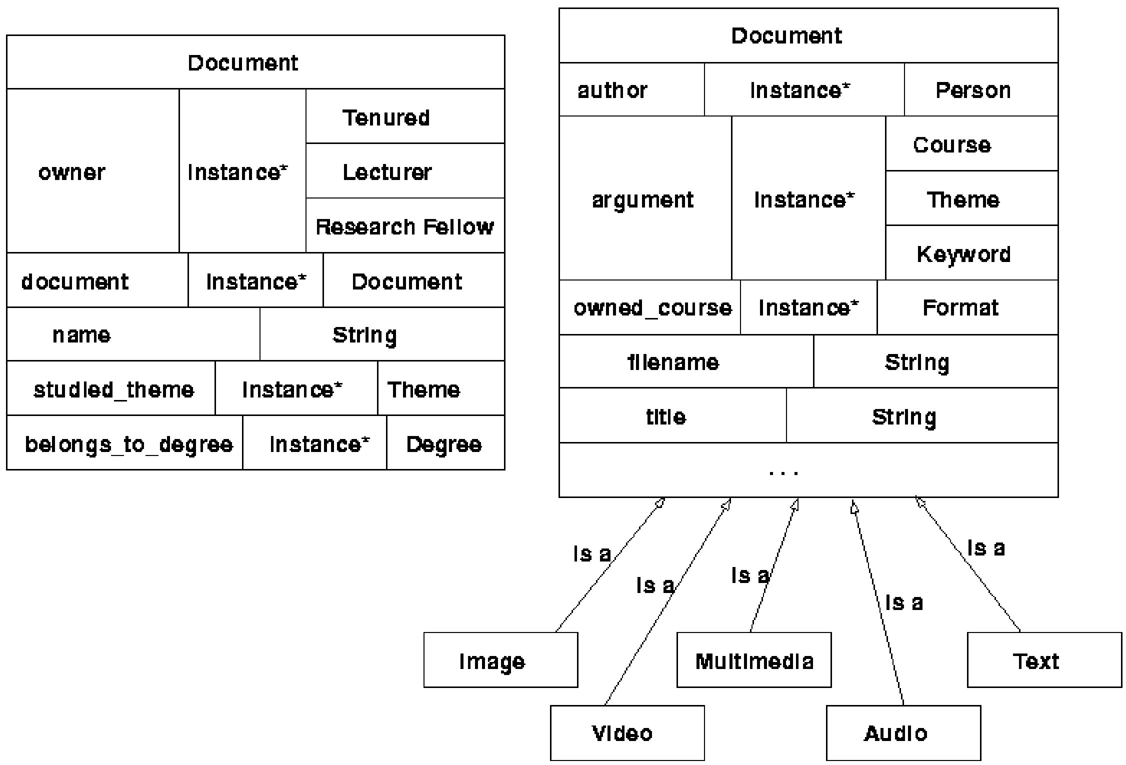

Fig. 4: The Document class and its elements.

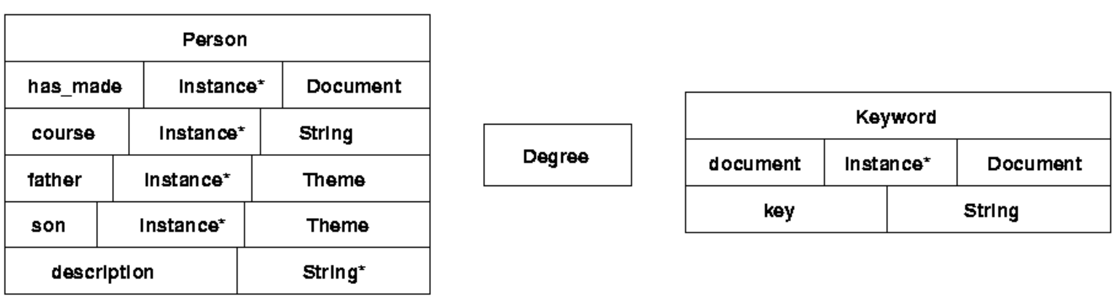

Fig. 5: The Theme and Keyword classes.

Other extensions are possible. Search by author may be extended seeking for person teaching in the same course or collaborating in the same research group. These search extensions require to maintain a context, which is built using the results of the previous user queries. The Grid technology helps giving a stateful service for the client instance. The discovery engine is realized using the RDQL language implemented in the Jena toolkit to search the RDF triples representing the OWL ontology. We have experimented the discovery service building a taxonomy for the Operating Systems subject matter (figure 6). People involved in the course were also represented. The multimedia material consists of a number of multimedia lessons that were introduced as instances of the ontology.

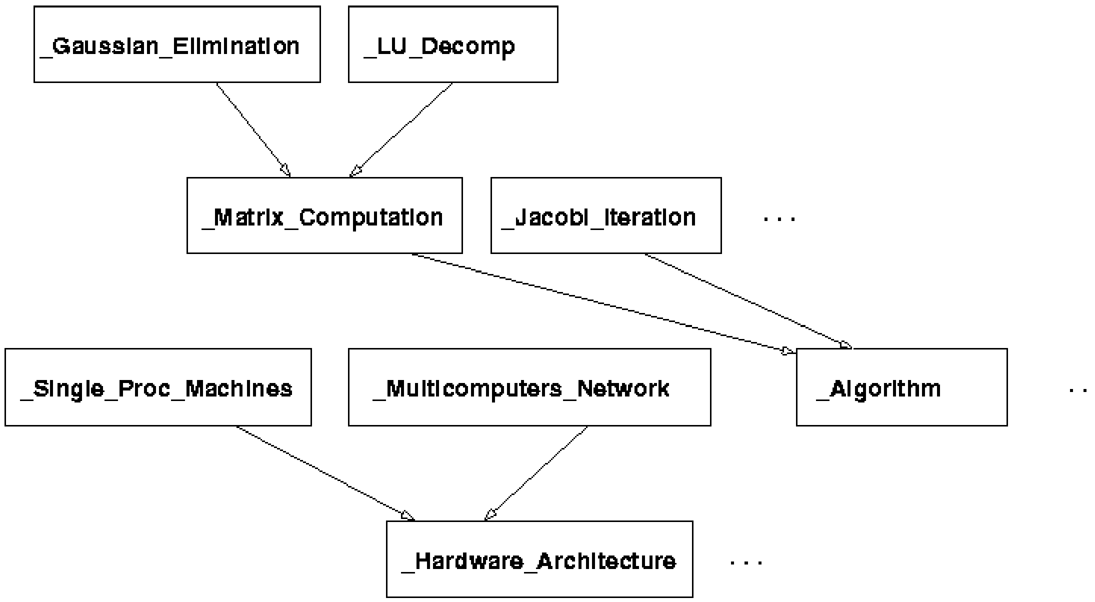

Fig. 6: The Operating Systems taxonomy.

6. Related Work

The novelty of our approach to the community support problem is the Grid Service-oriented design. In literature there are few papers which relate e-learning and multimedia distribution to the Grid context from the OGSI [Tuecke et al., 2003] point of view. And even if service-orientation is described as the state of the art for community support systems, the lack of prototype implementations is evident.

OntoEdu [Cui et al., 2004b] is proposed as a flexible architecture for e-learning based on several new technologies, such as ubiquitous computing, ontology engineering, Semantic Web and Grid Computing. The core of OntoEdu is the educational ontology, which is divided in five parts: user adaptation, automatic composition, education ontology, service model and content model. These concepts are realized using a grid-based design, but Grid Services do not refer to the OGSI specification and the educational ontology appears to be implemented outside the service grid. In practice, the educational ontology is queried by an educational server to find the services which are suitable for the users' needs.

An interesting framework to support semantic description of services, even if not implemented in a working prototype, is described in [Cui et al., 2004a]. The service classification defines Core Services (Access, Communication, Information, Query) and Appended Services (e.g. Sign-On, Locate, View, Update). The set of proposed services can be deployed in alignment to the OGSA [Foster et al., 2002] guidelines but also extended to provide additional functionality that is missing from today's Grids but required by an e-learning network (*i.e.* P2P support).

An e-learning Grid architecture is illustrated in [Pankratius and Gottfried, 2003]. Users access to a Learning Management System (LMS), which transfers the fabric layer applet of the Grid as mobile code which can be started locally on the users' machines. This enables communication with the Grid and the users' machines themselves to become resources in the Grid. Basic elements of the LMS are the Grid Learning Objects (GLOBs), which are associated to units of study and endow conventional e-learning content and a Grid Application Layer to access the Grid. Each GLOB contains an overview of the lesson, metadata, which are used to find the GLOB, several Reusable Information Objects (RIO), which represent the content of the unit of study, and a summary. The feasibility of these concepts is shown with the outline of a possible implementation of a Grid application for photo-realistic visualization that can be used for e-learning courses in medicine.

There is an effort to introduce a standardization for multimedia-related metadata. The MPEG-7 [MPEG7] descriptions of content may include the creation and production processes, the information related to the usage of the content, the storage format, and encoding information. Structural information on spatial, temporal or spatio-temporal components of the content (scene cuts, segmentation in regions, region motion tracking) are also included. Moreover, low level features (such as color, texture etc.) and high level features (objects, events) are described, along with information about how to browse the content in an efficient way and about the interaction with the content (user preferences, usage history). MPEG-21 [MPEG21], instead, aims at defining a normative open framework for multimedia delivery and consumption to be used used by all the players in the delivery and consumption chain. The goal of MPEG-21 is to support users to exchange, access, consume, trade and otherwise manipulate digital items in an efficient, transparent and interoperable way.

In the context of multimedia ontologies, [Hunter, 2003] proposes a top level ontology capable of integrating a number of specific standard metadata dictionaries like MPEG-7 (which describes multimedia content in details) and MPEG-21 (which concerns the rights associated to multimedia content). One of the motivation of such integration is to enable a single search interface across heterogeneous metadata descriptions and content in distributed archives.

7. Conclusions and Future Work

In this paper, we presented an ontology-based search service for multimedia content access and distribution. The service has been implemented as a Grid Service and easily integrated in the campus-wide e-learning architecture we are developing. The possibility to realize extended searches to semantically enriched data (searching for similar topics or for collaborating people) gives to the content discovery process a degree of flexibility not available in conventional systems.

The ontology we have proposed may be extended adding features to describe with more details the media content. Currently the media is described as a whole, while an indexing of segments may improve the search capabilities allowing a student to find a well defined topic. It would be useful to describe also some features of the user to personalize the search (user model, user history). This may be implemented by borrowing some metadata from MPEG-7. Moreover, a better description of the multimedia material production may be important, but a full MPEG-21 specification seems too detailed for a use in a closed environment like a University Campus.

In general, all the services of the proposed multimedia system are suitable to be integrated in the *SP^{2}A* project [Amoretti et al., 2004b], a service-oriented architecture which take advantage of the strengths of two technologies at different levels, by relying on a P2P network as a decentralized carrier of light/simple service advertisements and on Grid service query and invocation mechanisms which exploit the robust and secure OGSI technologies. The project is still in an early phase of development and several open issues will be investigated in our future research. In particular, we would like to introduce explicit semantics in the description of the offered services (not only in the description of some resources, such as multimedia objects), in the direction of a Semantic Peer-to-Peer Grid. This feature would allow peers to search for published services using ontological concepts.

The high-level architectural design of our services and service-oriented architectures should not be affected by the refactoring of OGSI which was proposed in January 2004 as *WS-Resourse Framework (WSRF)* [Czajkovski et al., 2004], aiming at exploiting new Web Services standards, specifically WS-Addressing. WSRF retains essentially all the functional capabilities present in OGSI, while changing some of the syntax and also partitioning OGSI functionality into five distinct, composable specifications (WS-ResourceProperties, WS-ResourceLifetime, WS-RenewableReferences, WS_ServiceGroup, WS-BaseFault). Major benefits of this new specification are the explicit distinction between the "service" and the stateful entities acted upon by the service, and the introduction of the forthcoming WSDL 2.0 capabilities.

References

[Amoretti et al., 2004a] Amoretti M., Conte G., Reggiani M., Zanichelli F., Designing Grid Services for Multimedia Streaming in an E-Learning Environment, in Proc. of the 13th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises (WETICE-2004), Workshop on Emerging Technologies for Next generation GRID (ETNGRID-2004), Modena, Italy, 2004.

[Amoretti et al., 2004b] Amoretti M., Conte G., Reggiani M., Zanichelli F., Service Discovery in a Grid-based Peer-to-Peer Architecture, in Proc. of the International Workshop on e-Business and Model Based IT Systems Design, St. Petersburg, Russia, 2004.

[Berners Lee et al., 2001] Berners Lee T., Hendler J., and Lassila O., The Semantic Web. Scientific American, 2001.

[Cui et al., 2004a] Cui G., Chen F., Chen H., Li S., A Grid Service Framework for Metadata Management in Self e-Learning Networks, in Proc. of the 2nd European Across Grids Conference, Nicosia, Cyprus, 2004.

[Cui et al., 2004b] Cui G., Chen F., Chen H., Li S., OntoEdu: Ontology Based Education Grid System for E-Learning, in Proc. of the GCCCE2004 International Conference, Hong Kong, 2004.

[Czajkovski et al., 2004] Czajkovski K., Ferguson D., Foster I., Frey J., Graham S., Maguire T., Snelling D., Tuecke S., From OGSI to WS-Resource Framework: Refactoring and Evolution, draft for The Globus Alliance, 2004.

[Foster et al., 2002] Foster I., Kesselman C., Nick J. M., Tuecke S., The Physiology of the Grid: an Open Grid Service Architecture for distributed systems integration, draft for The Globus Alliance, 2002.

[GGF] The Global Grid Forum home page: http://www.gridforum.org/

[Globus] The Globus Alliance home page http://www.globus.org/

[Jena] The Jena home page: http://jena.sourceforge.net/index.html

[Hunter, 2003] Hunter J., Enhancing the Semantic Interoperability of Multimedia through a Core Ontology. IEEE Transactions on Circuits and Systems for Video Technology, Special Issue on Conceptual and Dynamical Aspects of Multimedia Content Description, 2003.

[MPEG21] ISO/IEC JTC1/SC29/WG11, MPEG-21 Overview v.5, http://www.chiariglione.org/mpeg/standards/mpeg-21/mpeg-21.htm

[MPEG7] ISO/IEC JTC1/SC29/WG11, MPEG-7 Overview, http://www.chiariglione.org/mpeg/standards/mpeg-7/mpeg-7.htm

[CORBA] Object Management Group, Common Object Request Broker Architecture home page: http://www.corba.org/

[Pankratius and Gottfried, 2003] Pankratius V. and Gottfried V., Towards E-Learning Grids: Using Grid Computing in Electronic Learning, in Proc. of the IEEE Workshop on Knowledge Grid and Grid Intelligence (in conjunction with 2003 IEEE/WIC International Conference on Web Intelligence / Intelligent Agent Technology), Halifax, Canada, 2003.

[Tuecke et al., 2003] Tuecke S., Czajkowski K., Foster I., Frey J., Graham S., Kesselman C., Maquire T., Sandholm T., Snelling D., Vanderbilt P., Open Grid Services Infrastructure (OGSI), draft version 1.0, http://www.ggf.org/ogsiwg, 2003.