Designing Grid Services for Multimedia Streaming in an E-learning Environment

Michele Amoretti, Gianni Conte, Monica Reggiani, Francesco Zanichelli
Dipartimento di Ingegneria dell'Informazione
Parco Area delle Scienze, 181A, I-43100, Parma, Italy
{amoretti, conte, reggiani, zanichelli}@ce.unipr.it

Abstract

Next generation e-learning platforms should support cooperative use of geographically distributed computing and educational resources as an aggregated environment to provide new levels of flexibility and extensibility. In this overall framework, our activity addresses the definition and implementation of advanced multimedia services for an aggregated Grid-based e-learning environment, as well as the design and experimentation of a content distribution and multimedia streaming infrastructure in light of edge device heterogeneity, mobility, content adaptation and scalability.

In this paper we initially present the general objectives and requirements that we are taking into account in the development of a multimedia access service for an e-learning platform. Then we describe a partial system prototype which capitalizes upon traditional features of Grid computing like providing access to heterogeneous resources and services of different administrative domains in a transparent and secure way. Moreover, our system takes advantage of recent proposals by the Global Grid Forum (GGF) aiming at a standard for service-oriented architectures based on the concept of Grid Service.

1. Introduction

In a global knowledge-based society, communities play a pivotal role and re-shape the process of learning and sharing knowledge in and among organizations. Knowledge sharing is the process where individuals mutually exchange their (implicit or explicit) knowledge and jointly create new knowledge. The role of technology, especially of Information and Communication Technology (ICT), is crucial. Communities are often defined in ICT terms as *Virtual Organizations* whose members exploit ICT tools to achieve some permanent or temporary common goals. ICT is often seen as a valuable means in bridging gaps of space and time between members of such communities.

Community support includes all methods for enabling communication and coordination in a group of people. It includes support for direct communication, support for indirect information and support for matchmaking. Community support systems (community platforms) that provide a rich communication medium for work or interest groups are gaining more and more attention in application areas ranging from leisure support and customer support to knowledge management. One of these application areas is the support of teaching and research activities in universities. ICT, in particular, is becoming of paramount importance for the so-called Informal Learning, which includes all the activities taking place outside classes and courses, that aim at understanding, creation of knowledge and skills acquisition [1].

In the specific context of a university campus the joint availability of portable devices and wireless technologies and their forthcoming functional integration will eventually change both the way in which students attend their education in the university and the way the institution and the professors communicate with the students as well as the way in which students communicate to each other. In this time devoted to consolidation of notions and self-confidence development, students will obtain great advantage by a 24/7 access to lecture notes, case studies, and multimedia contents, such as video recordings of previous classes, made available by teachers and instructors or even by other students.

In this paper we discuss our first results in the design of a campus-wide service 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.

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 [4]. Recent significant activities coordinated by the Global Grid Forum (GGF) [5] are currently defining a service-oriented architecture based on the concept of Grid Service [15]. A major goal of service Grids is the possibility to simplify the de-

sign of services geographically distributed 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.

Current research efforts are trying to integrate the two formerly distinct Web and Grid communities to create an integrated approach described in the Open Grid Service Architecture (OGSA) [4]. OGSA's Grid Services are essentially (potentially transient) stateful Web Services with an associated lifetime and whose behaviors must conforms to a set of interfaces defining the service semantic. Moreover, Grid Services should be compliant with standard mechanisms to support state management, referenceable handles and event notification. Grid Service standards are defined in the Open Grid Service Infrastructure (OGSI) specification [15]. In general, the service-oriented design methodology [4] is nowadays a major approach to cope with dynamic and flexible applications which need integration and interoperability among heterogeneous off-the-shelf software components (*e.g.* databases, service proxies, collaborative platforms, streaming servers, etc.).

The paper is organized as follows. Section 2 introduces the general issues and requirements for a campus-wide multimedia service and establishes our long-term goals. The system prototype under development is illustrated in Section 3. Finally, some related work and an outline of further work conclude the paper.

2. General requirements for a campus-wide multimedia service

Multimedia content distribution is rapidly becoming one of the most important network applications, both on the global Internet and on private intranets. Several universities around the world are already deploying or experimenting solutions for multimedia content distribution over their TCP/IP wired/wireless networks and the Internet. Recordings of class lectures, seminars, lab sessions are made available in digital form to be streamed and viewed either off- or on-campus, both for special e-learning programs as well as for blended learning services which aim at improving the experience of attending students. The availability of a wireless infrastructure adds further flexibility over time and space constraints in the learning activity, thus allowing students to exploit their own terminals to access teaching resources essentially when and where they intend to.

In the following we introduce a set of objectives and requirements to be taken into account in the design of a campus-wide service for multimedia content access and distribution. Several complex, often interacting issues contribute to the effectiveness of such a complex 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.

A great deal of variants on the concept of multimedia repositories have been introduced in the recent past. Digital archives, Video on Demand Systems, Streaming Servers often emphasize somewhat different aspects but they frequently combine cataloguing and search capabilities on their content with some delivery mechanisms to convey it to their users. Thus, they often integrate a database component to represent descriptive information about the available multimedia data, *i.e.* metadata, together with some degree of search capabilities built around it. Before analyzing the requirements for the search capabilities, it should be underlined that, owing to the ever increasing heterogeneity of user connectivity and terminals, it is becoming very important for these systems to constraint multimedia content search to what can be appropriately delivered and visualized at users premises. Network and computing resources currently available to users ultimately affect the quality of service they can experience.

The effective exploitation of a potentially huge amount of instructional multimedia material calls for powerful search capabilities which should allow users to perform highly selective queries in the service database. Within most video services in use today a rather limited set of metadata is exploited to describe the available multimedia contents. Although a set of standards is recently emerging regarding audio/video metadata [9, 16], widely used descriptions often only include a minimal set of common attributes (*e.g.* title, author, date, description, etc.) together with key-words based indexes. Further levels of search capabilities can be obtained by means of ontology-based, semantic descriptions [8] to automate search or to assembly course materials on account of user preferences.

Additional goals for designing a campus-wide video service can be related to system integration. On one hand, it should be designed so as to be able to integrate pre-existing, local services of the same type which, rather unsurprisingly to anyone familiar with universities, are often setup by different groups in different departments or even in the same

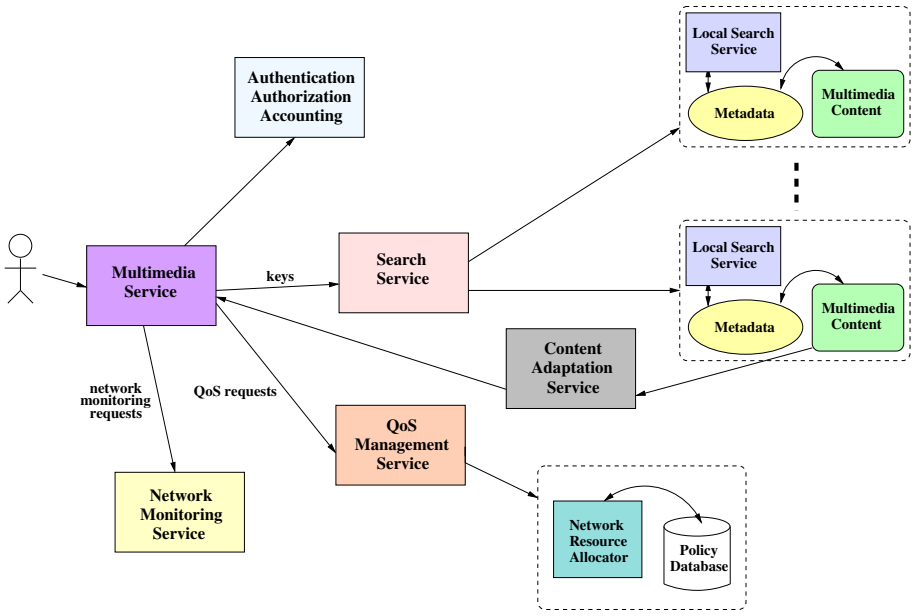

Figure 1. High-level organization of the campus-wide Multimedia Service.

one. The integration can be obtained with the goal of hiding the heterogeneity of such systems to users, by providing some unifying layer to search and access consistenly over all available services of the same kind.

On the other hand, a campus-wide video service should also be easily integrated into other larger e-learning or collaborative frameworks which can take advantage of distributed multimedia capabilities. From our viewpoint, both forms of integration of applications can be pursued by means of computing technologies relying on high level, standardized, self-describing interfaces.

One more system-level goal is represented by security management. The usual authentication, authorization and accounting capabilities should be available to the campus-wide video service. However, it would be very important to finely tune access control lists to multimedia material as well as to make this service part of a single sign-on solution which does not require users to authenticate multiple times at different services.

Finally, we envision that such a system could interact with resource managers on behalf of its users, possibly hiding the complexity of a feature-rich network infrastructure. For example, on best-effort IP networks, Content Delivery Networks (CDN) exploit caching proxies to move web and multimedia content towards network edges and closer to clients. These network intermediaries can also offer additional services such as content adaptation, personalization or location-aware data manipulation. QoS-enabled Diffserv networks allow for traffic classification and prioritization and can rely on a Bandwidth Broker component responsible for admission control and edge-router configuration (marking and shaping). Thus, in a Diffserv domain QoS-sensitive applications can negotiate Service Level Agreements leading to proper resource reservations to obtain guarantees regarding bandwidth, delay, jitter, and so forth.

The campus-wide video service should be designed to let users take advantage of the available infrastructure components and capabilities. This could be carried out transparently, that is directly determining all the information (without burdening the users) required to acquire and reserve the needed resources, (*e.g.* availability of proxies and/or bandwidth over the interested links), given the requested content, the state of the network and the characteristics of user terminal and connectivity. However, as it was already mentioned, QoS parameters can also be explicitly considered in the phase of content search, to the purpose of constraining the search only to material which can be adequately experienced by the end user.

Furthermore, coallocation should also be supported, *i.e.* the simultaneous allocation of multiple resources. The coallocation of resources enabling concurrent streaming tasks (*e.g.* for the distribution of some live or prerecorded content to several classrooms) poses significant additional challenges than the reservation of resources for a single streaming task.

Summing up, previous discussion of requirements for a campus-wide multimedia service motivates an high-level system organization as described in Figure 1.

3. Grid-based prototype of a Multimedia Service

In this section we describe the on-going development of a Multimedia Service prototype based on the Globus Toolkit 3 [6]. 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 [10], *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 Service* (one for each content provider), as illustrated in Figure 1. The user (in our e-learning application a student interested in a video lesson) 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.

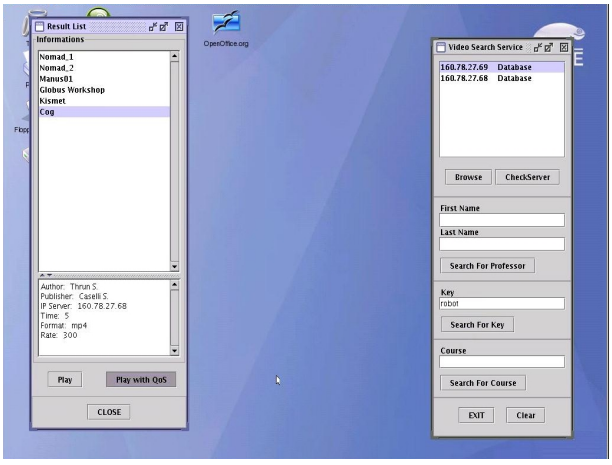

Figure 2. Screenshot of the user interface in the current prototype.

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 liveness. 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 informations 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 four aspects of our prototype, with the aim of illustrating key services and components of the implemented multimedia streaming system.

3.1. Authentication and Authorization

The Search Service can operate only if the user is identified by the system and his/her requests are compatible with resource access rights. In practice, a collection of multimedia content is usually made available only to a specific subset of users, e.g. a group of video lectures could be accessible only to students which have paid to attend the course.

One of the strongest motivations in using Grid technologies to build our multimedia streaming system, is that security functionalities can be naturally embedded in Grid Services which can be located and used as needed by applications. Specifically, security policies are published so that users can dinamically discover which credentials and mechanisms are needed to establish trust with the service. The security model that we used is based on the work of the Globus Project security team [14], as a contribution to the creation of a *standard* OGSA security model.

On the user side, the first step is to inspect the security policy of the target service, and to identify which mechanisms and credentials are required to submit a request. If the needed credentials are not present, a credential conversion service is contacted to convert the existing credentials to the needed format. Finally, a token processing and validation service is used to handle the formatting and processing of authentication tokens for exchange with the target service.

On the resource side, a token processing service is used to authenticate the user, by the determination of its identity and attributes. Only after this, the details of the request and user informations are presented to an authorization service for a policy decision. If also this step completes suc-

cessfully, the target service application receives the authorized request and starts processing it.

3.2. Searching for multimedia objects

If the user is authenticated and authorized to use the Search Service, a graphical user interface (GUI) is provided to initiate the search process. The Search Service instance can perform the following operations: browse, which returns the complete list of available multimedia objects; searchForProfessor, searchForKey and searchForCourese which return a partial list including only specific content.

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, on each active multimedia server is spawned an instance of Local Search Service. 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.

3.3. Streaming

If many providers match the query for a multimedia object, the choice of the one which will act as streaming source is left to the user. This is the only case of direct interaction between the user and the remote resources. Future work will consider a set of mediator Grid Services, e.g. for resource reservation (Bandwidth Broker), proxy caching, transcoding, etc.

Our current system considers only RTSP streaming servers and each user terminal must be equipped with a suitable multimedia player. The multimedia data flow is established between these applications. The stream is managed using the Real Time Streaming Protocol (RTSP) [13], which coordinates the delivery of multimedia objects acting as a ”VCR remote control”. At the lower transport level, the Real-Time Transport Protocol (RTP) [12] is used to provide support for applications with real-time properties such as continuous media (e.g., audio and video), including timing reconstruction, loss detection, security and content identification.

3.4. Updating multimedia content

Provided the user has the necessary authorization, the client application can be used not only for searching but also to add, delete or update multimedia content on the distributed servers. For example, to make a new video lecture available, the teacher should provide the path for the corresponding file and a useful description for keywords-based searches. In order to support that, the Search Service requires the activation of an istance of Update Service (not developed yet) on each destination server (chosen by the user or by the Search Service itself).

4. Related Work

A theoretical model that identifies the influence of Information and Communication Technology (ICT) on knowledge sharing is proposed in [7]. The authors tested the theoretical model within two ICT-facilitated communities for professionals in working conditions. The results of these case studies show that the positive contribution of ICT to knowledge sharing in communities is part of a complex set of influences and relationships. Easier exchanges between community members are facilitated by ICT, which helps create two public goods: a shared information base (communality) and communications independent of time and place (connectivity). The results also show that trust among members of a community, and their identification with the community, are important influences on knowledge sharing.

The novelty of our approach to the community support problem, and in particular to the multimedia streaming task, is the Grid Service-oriented design. In literature there are descriptions of projects based on the client-server paradigm and sometimes on Web Services, but nothing yet adopted the OGSI specification.

In [17] and [18] Zhang et al. define a component-based framework to support device-independent multimedia Web Services. Two intelligent agents are introduced into proxy server and service broker server respectively. Moreover, metadata are separated from multimedia content to enhance SOAP flexibility. Together with this, CC/PP-based user profile management provides an easy and flexible way to split and adapt multimedia services to appropriate composite devices, as well as increases the flexibility for users to manage multi-devices. The architecture is designed making two assumptions: (1) service requesters, with the help of service brokers, have already located the service providers that offer the requested services, and (2) users invoke Web Ser-

vices through their proxy server. The authors have adopted caching and replacements algorithms introduced in [11], but the architecture should be improved by the binding of SOAP to other more multimedia-oriented transportation protocols such as RTSP.

A framework for providing multimedia, multimodal, distributed services using Session Initiation Protocol (SIP) and Web Services technologies is proposed in [2]. Building blocks (from different vendors) are seen as services which can be registered in UDDI registry for dynamic binding at run time. There are a SIP applications server (used to build call routing/telephony and Instant Message and Presence), a Web application server, an IVR system, a media server, a TTS server and a conference server. SOAP is used to transfer control information among different components, which are both service providers and service requestors.

In [3] the authors introduce a robust mechanism for providing differentiated quality of service using application-specific characteristics of Web services. Transcoding is used to allow web servers to customize the size of objects constituting a web page, and hence the bandwidth consumed by that page, by dynamically varying the size of multimedia objects on a per-client basis.

5. Conclusions and future work

In this paper, we presented our research regarding a campus-wide system for multimedia content access and distribution. The Grid Service approach has been motivated with the need for integrability and scalability to enable future extensions, consistent navigational model for users, and support for heterogeneous devices.

We described the general requirements that have to be taken into account in the design of campus-wide services for multimedia content access and distribution. We also discussed how Grid Services technologies have provided an effective support to build a system prototype which copes with an initial subset of the requirements. Currently we are experimenting a testbed using several distributed multimedia resources available on our campus LAN.

As future work, we will extend the Multimedia Service to take advantage of QoS support for massive multimedia distribution. To this goal we plan to integrate caching proxies and bandwidth broker components as Grid Services into the existing system.

References

[1] M. Ahonen. Mobility, Accessibility and Learning. Keynote Speech in Mlearn 2003 Conference, London, May 2003.

[2] H. Cai, W. Lu, B. Yang, and L. Tang. Session Initiation Protocol and Web Services for Next Generation Multimedia Applications. In *IEEE Fourth International Symposium on Multimedia Software Engineering (MSE'02)*, Dec 2002.

[3] S. Chandra, C. Ellis, and A. Vahdat. Application-Level Differentiated Multimedia Web Services Using Quality Aware Transcoding. *IEEE Journal on Selected Areas in Communications*, 18(12), December 2000.

[4] I. Foster, C. Kesselman, J. Nick, and S. Tuecke. The Physiology of the Grid: an Open Grid Service Architecture for distributed systems integration. Draft, Globus Project, June 2002.

[5] Global Grid Forum. Home Page. http://www.gridforum.org/.

[6] Globus. Globus Toolkit 3. http://www.globus.org/.

[7] B. Hooff, W. Elving, J. Meeuwsen, and C. Dumoulin. Knowledge Sharing in Knowledge Communities. In *Communities and Technologies*, Sept 2003.

[8] J. Hunter. Enhancing the semantic interoperability of multimedia through a core ontology. *IEEE Transactions on Circtuits and Systems for Video Technology*, 2003.

[9] R. Koenen and F. Pereira. MPEG-7: a standardised description of audiovisual content. *Signal Processing: Image Communication*, 2000.

[10] Object Management Group. Common Object Request Broker Architecture. http://www.corba.org/.

[11] S. a. Paknikar. A Caching and Streaming Framework for Multimedia. In *8th ACM International Conference on Multimedia*, Oct 2000.

[12] H. Shulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A transport protocol for real-time applications. Request for Comments 1889, January 1996.

[13] H. Shulzrinne, A. Rao, and R. Lanphier. Real-time Streaming Protocol (RTSP). Request for Comments 2326, April 1998.

[14] F. Siebenlist. OGSA Security Roadmap. http://www.globus.org/ogsa/Security/ogsa-sec-roadmap-v13.pdf, April 2002.

[15] S. Tuecke and al. Open Grid Services Infrastructure (OGSI). http://www.ggf.org/ogsi-wg, Apr. 2003.

[16] S. Weibel, J. Kunze, C. Lagoze, and M. Wolf. Dublin core metadata for resource discovery. *RFC 2413, The Internet Society*, 1998.

[17] J. Zhang and J. Chung. A SOAP-Oriented Component-Based Framework Supporting Device-Independent Multimedia Web Services. In *IEEE Fourth International Symposium on Multimedia Software Engineering (MSE'02)*, Dec 2002.

[18] J. Zhang, J. Chung, and Z. Zhang. A Router Model for QoS-based Multimedia Web Services. In *IEEE International Conference on Multimedia & Expo, Baltimore, Maryland*, July 2003.