The university community: new services and technology

M. Amoretti G. Conte M. Reggiani F. Zanichelli

Dipartimento di Ingegneria dell'Informazione
Università di Parma - Italy
e-mail: {amoretti,conte,reggiani,mczane}@ce.unipr.it

Abstract. The advent of ubiquitous computing is going to change deeply the way higher education is delivered to university communities. Wireless and mobile technologies will surely influence the evolution of current e-learning services and push the development of totally new services enabling anytime, anywhere, anyhow learning. In this paper we describe our first steps towards providing advanced learning services and an effective infrastructure to our campus. Our approach is related to the Grid computing paradigm and takes advantage of the recent availability of support for e-services.

1 Introduction

With the increasing popularity of mobile technologies, high-speed wireless Internet access is now available in many locations around the world, including airports, cafes, corporate offices, universities, factories and homes. The widespread availability of portable terminals such as PDAs, laptops, and cellular phones is leading to the establishment of dynamic, distributed Virtual Organizations (VOs) where information should be provided with Quality of Service (QoS) compatible with the specific user terminal and network. In a university campus scenario the joint availability of portable devices and wireless technologies and their forthcoming functional integration will change the way in which students attend their education in the university, the way the institution and the professors communicate with the students and the way in which the students communicate each other. The scenario we are interested in also targets this additional space and time where learning does occur by means of alternative pathways, mainly including student-to-student interaction.

In this time devoted to consolidation of notions and self-confidence development, students will obtain great advantage by the ability to access to lecture notes, case studies, and multimedia contents, such as video recordings of previous classes, made available by teachers and instructors or by other students. The overall objective is to create an "ad-hoc" space where interaction and discussions among the participants can take place sharing experience and knowledge. Collaboration can be synchronous (a session is defined in advance and for a specific time) or asynchronous (collaborators send information that can be viewed and processed by other participants at a later time). The workspace (secure but very dinamic, topologically unconstrained) enables learners to organize and manage meetings, share applications and URLs, store documents, coordinate and communicate with each other.

A traditional e-learning model is often based on the access to unconnected or statically organized sources of contents (web pages, multimedia contents, etc.) that the user must look for, i.e. a client-server interaction mode. The next generation e-learning platforms must support cooperative use of geographically distributed computing and educational resources as an aggregated single e-learning environment. In this overall framework, our activity is addressing the definition and implementation of advanced services for the aggregated single e-learning environment, as well as the design and experimentation of a technical infrastructure enabling the effective exploitation of content distribution and multimedia streaming in light of edge device heterogeneity, mobility, content adaptation and scalability. First, in section 2 we will describe the levels of service we are focusing on. Finally, section 3 will discuss some technological aspects of the infrastructure we are devising for our university campus.

2 A hierarchy of levels for community-oriented services

The goal of a service enriched, technology empowered university communities can be approached from multiple directions, each with different levels of cost, complexity and managebility. Our intention is to explore a hierarchy of levels for community-oriented services which the university should provide in the era of ubiquitous and mobile computing. The main levels we envision are:

- *the first level* can be simply a generic connectivity service which can be provided to wired and wireless users. In both cases, suitable authentication, authorization and accounting (AAA) facilities should be made available to allow intranet and Internet access. In particular, we are currently investigating different techniques for secure, flexible wireless user authentication, ranging from web-based (https) to Virtual Private Network solutions. The network facilities of the University of Parma campus are currently being extended using Wi-Fi technology. Our Department has already implemented wireless access points within its premises. During the next months, wireless access will be extended to the whole open access areas (lounges, corridors) of the buildings of the School of Engineering.
- *the second level* can be defined as the availability of traditional web-based services directly or indirectly covering the many facets of e-learning in the university community. This level represents the current state of e-learning services offered by many universities around the world. A traditional Best-Effort intranet hosting a set of web-based services appears perfectly adequate to this scenario.
- *the third level* of service should build upon the previous ones to support distributed computing and educational resources as an aggregated single e-learning environment. Service abstractions (*e.g.* Web Services) and specific middleware support should take care of issues like service discovery, authorization, QoS and resource allocation and so forth as described in the next section. In this advanced scenario we regard as a key issue the ability to provide campus-wide multimedia streaming services exploiting an effective, QoS-enabled content delivery infrastructure. Caching and content-adapting proxies are being researched to obtain bandwidth savings and user terminal independence.
- *the fourth level* is related to an emerging scenario where resources and contents are not only provided by the official infrastructure but also by the whole university community. For example, several students own recent, capable notebooks. These resources can be partially and temporarily offered to the common infrastructure while personal data (lesson notes, exercises, ...) can be shared within the community, thus significantly increasing the size and the dynamics of the distributed system and possibly reducing the activity on official, persistent nodes. In fact, access to material located on a university server can be avoided if new requests are implicitly routed to copies still available on user computers which downloaded master files (as done by BitTorrent [3]). Computers belonging to a community of students often requiring the same information at the same time can possibly switch from the role of passive (destination) client to active source of information thus reducing the network and server workload. This communication model is very similar to the popular and controversial Peer-to-Peer (P2P) paradigm. Nevertheless, previously described services for authentication and QoS support should be incorporated into this new model of content distribution and delivery.

Such levels of service, along with the necessary enabling infrastructure, can also be seen as the desidered milestones for our current and future activity.

3 Technologies for sharing and interoperability of learning resources

Within the previously described framework for university community services, we are mainly focused on the exploitation and improvement of available technology to increase and extend the direct interaction between students and professors beyond the official class and office hours. The final goal is the development of an architecture providing a set of services to assist the virtual community in the use of university resources.

As an example consider tracing the activity of a student willing to book the access to a PC in laboratory for a given time slot to implement his/her assignment code. In an application like this, we rapidly identify several hot spots demanding advanced solutions. Firstly, we would like to support security across university organization boundaries, thus enabling a “single sign-on” mechanism. This should include delegation of credentials for probing booking databases of multiple sites. Then, a descriptions of the PC Laboratory facilities should be published to allow a discovery service to look for a computer matching the student request. This requires that all the university organizations agree on a standard format for the distribution of both static details about the machine and information on the state of the resource. Moreover, the organizations should also be able to accept asynchronous queries from the discovery service providing a common interface of diverse implementations.

In order to cope with the complex coordination task required by this kind of applications, a challenging solution can be based on the emerging Grid technology [1]. Previously described hot spots can be effectively implemented using the rich Grid platform, namely the Grid Security Infrastructure [6] (supporting single sign on, delegation and credential mapping), the Grid Resource Allocation and Management protocol [7] for service creation and management, and the Meta Directory Service [8], providing a uniform framework for discovering and accessing system configuration and status information.

The interest in Grid technology is also motivated by its continuous shifting from large scale computing to the world of robust and flexible e-services. The Open Grid Service Infrastructure (OGSI) [2] defines Grid services as Web services that provide well-defined interfaces and that follow specific conventions. In a service-oriented view, the definition of the interface and agreement on the communication protocols allows to shield users from the heterogeneity of the applications regarding user interfaces and techniques for access control and resource allocation.

Our first effort to encapsulate several implementations of similar services (running on multiple heterogeneous platforms) behind a common interface regards video streaming applications. An interesting and challenging service of educational organizations in the era of pervasive computing is to offer a 24/7 teaching resource access to their students, possibly including libraries of videolessons or additional multimedia information to revise or extend in-class learning sessions. In general, this would require the university to set up a number of video servers, organized in an architecture that allows a secure, standardized, adaptable, and flexible access to the available multimedia resources.

Our goal is to implement this architecture as a set of Grid Services connected through the Grid middleware. The definition of the Grid Service for video streaming application requires to identify the basic operation supported by video servers and to define a common interface for their basic operations using a standard format such as the Web Service Definition Language (WSDL). By wrapping existing systems to a consistent WSDL interface, we can not only shield users from the heterogeneity of video servers but also pave the way for stronger integration and interoperability among e-learning and other community applications, in a consistent way with the service-oriented perspective.

Recently, there has been some interest for a possible convergence of Grid and P2P technologies [4, 5]. While both distributed computing models share common objectives, *i.e.* the organization of resource sharing within virtual communities, they focus on different aspects heavily characterizing the two approaches. Grid computing mostly emphasizes security and remote resource access, as the main goal is to serve the needs of Virtual Organization, *i.e.* relatively small, closed communities. Conversely, the P2P approach targets highly dynamic, open communities whose aim is content sharing or massive parallel computing. In our view and application context, a possible convergence can be investigated by looking for a security model for P2P systems whose authentication, authorization and content validation techniques are able to fullfil the requirements of educational and corporate institutions. Specifically, new protocols for negotiating resource and service access (acquisition, reservation, ...) according to the specific user profiles (e.g. student curricula) and requirements (e.g. desired service level) must be considered to guarantee an effective and fair access to the community services.

4 Conclusions and future work

In this paper we have outlined our response to the challenges and opportunities of ubiquitous and mobile computing in the university community. Specifically, a hierarchy of levels of services has been presented along with some relevant technical issues emerging in our work. Thanks to Grid computing and the richness of its infrastructure, we are advancing in the development of advanced multimedia streaming services, offering access trasparency, security, QoS support and content adaptation.

Being also interested in the promising convergence of Grid and P2P models, we will investigate the very critical issue of providing security and QoS to P2P systems. The often massive scale and the high dynamics of peer arrival/departure seem to make traditional enforcing techniques quite inapplicable. However, the *campus-wide P2P computing grid* we envision appear to be less critical than an Internet-wide one and probably amenable to the application of Grid-like techniques for providing service level agreement among peers. Our future work will focus mainly on new methods for resource management, so as to be able to meet user requirements for performance, cost, security and other quality of service metrics.

References

1. I. Foster, C. Kesselman, S. Tuecke, Anatomy of the Grid, International J. Supercomputer Applications, 15(3), 2001.
2. I. Foster, et al., The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration, Tech. Report, Globus Project; http://www.globus.org/research/papers/ogsa.pdf (current June 2002).
3. BitTorrent - `http://bitconjurer.org/BitTorrent`.
4. I. Foster, A. Iamnitchi, On Death, Taxes, and the Convergence of Peer-to-Peer and Grid Computing. 2nd International Workshop on Peer-to-Peer Systems (IPTPS'03), February 2003, Berkeley, CA.
5. D. Talia, P. Trunfio, Toward a Synergy Between P2P and Grids, IEEE Distributed Systems Online, vol. 4, n. 7, July 2003.
6. I. Foster, C. Kesselman, J.M Tsudik, S. Tuecke, A Security Architecture for Computational Grids. ACM Conference on Computers and Security, 1998.
7. K. Czajkowski, I. Foster, N. Karonis, C. Kesselman, S. Martin, W. Smith, S. Tuecke, A Resource Management Architecture for Metacomputing Systems. 4th Workshop on Job Scheduling Strategies for Parallel Processing, 1998.
8. K. Czajkowski, S. Fitzgerald, I. Foster, C. Kesselman, Grid Information Services for Distributed Resource Sharing. 10th IEEE International Symposium on High Performance Distributed Computing, 2001.