

Systems Support for Collaborative Learning

Colin Allison, Alan Ruddle
School of Computer Science, University of St Andrews, Scotland KY16 9SS
{ca,alanr@dcs.st-and.ac.uk}

Rosa Michaelson
Accountancy & Business Finance, University of Dundee, Scotland DD1 4HN
{r.michaelson@dundee.ac.uk}

Abstract

One of the distinguishing features of novel network based learning environments is their capability to support group work and collaboration. TAGS, the Tutor and Groups Support Scheme, is an inter-disciplinary, inter-institutional project, which brings together software systems builders, subject-specialists and educational content developers. Collaborative Learning is central to the pedagogical goals of TAGS, and this has lead to the concept of groups being used as a fundamental organising principle. Groups form the basis of (i) privileges and access control, (ii) information dissemination and event awareness, (iii) teamwork involving shared, multi-user educational resources, (iv) online management of group learning, (v) user-centric portal generation, and (vi) replicated servers. The technical implications of this heavy reliance on the group abstraction are described.

1. INTRODUCTION

From the perspective of good educational practice an online collaborative learning environment should provide certain features for learners:

1. group-work support [1]
2. interactive, engaging, responsive
3. real-world input
4. student-centred
5. anytime/anywhere

Those pedagogical goals imply certain technical requirements.

Support for group working means that online resources may be shared and used by multiple concurrent readers and writers. This requires concurrency control, and the usual associated concerns with liveness, safety and fairness. In addition, multi-user awareness is important for teamwork. Members of a group need to know what each other are doing, and have appropriate communication options to co-ordinate with each other. In particular, distributed real-time multimedia channels may be used to augment distributed resource coherence [2].

Interactivity implies that online working is much more than simply browsing lecture notes dumped on the Web. Regardless of the balance of computation between server and client the system must be geared to writers as well as readers. *Responsiveness* is essential for creating an interactive feel when working on the web. It means that the delay between a user making a request and a result being returned to them should be no longer than some period that is considered reasonable within a particular context. For example, a user working from home over a relatively low bandwidth modem line will have lower expectations than one attached to a high-speed campus network. If a network

service is perceived to be slow then it will be seen as an unproductive use of time, and may be abandoned.

Real world input is facilitated by direct Internet connection to real world data sets. For example, students of meteorology may gain access to real satellite weather image downloads within an hour of them happening and can learn from making their own weather predictions which can be later compared with established media channel predictions. The Finesse portfolio management facility [3] allows the management of a portfolio of shares by a student group. The groups maintain their portfolios by buying and selling shares chosen from a database of live data for companies quoted on the London Stock Exchange.

Student-centred means focusing on the user as an individual. It also implies strong monitoring capabilities, for manual inspection or automatic adaptation. In an interactive online environment it is of course possible to implement extensive activity logging (by careful choice of a meaningful activity), which can then be “mined” for useful feedback on a student’s use of a set of educational resources. This feedback can be used to inform the resource development and adaptation process and thereby enhance the learning environment.

Anytime/anywhere implies a resilient, highly available web-based service. Availability means coping with faults. Providing a service based on multiple servers can reduce periods of non-availability caused by server crashes. These servers may be distributed across a wide area, possibly globally. If they are also used to share the load then *replication* becomes an attractive strategy as it can, in principle, also support performance by providing an appropriate ratio of users to servers. This approach is further useful as a practical means of implementing *incremental scalability*, which may be necessary to cope with a dynamically changing user population.

2. THE TAGS FRAMEWORK FOR COLLABORATIVE LEARNING

TAGS [4] is a framework for the development, deployment and maintenance of web-based collaborative learning environments. At present TAGS is being used by tutors and students in the context of accredited degree programmes at six Scottish universities in ten departments, representing a wide range of subject areas: Physics and Astronomy; Medicine; Languages; Accountancy, Finance and Management; Information Technology; and Educational Studies. Further expansion of types of use is planned.




Figure 1: The TAGS Framework

The TAGS framework, see Figure 1, is concerned with supporting highly interactive multi-user, concurrent environments, accessible via standard Web browsers [5]. Requirements of the architecture include concurrency control, availability and responsiveness. When the latter two requirements are supported by replication local concurrency control mechanisms must be subsumed by global coherence. There can be a tension between strong replica synchronisation and maintaining interactive response times. TAGS accordingly provides a plug'n'play set of replication synchronisation mechanisms to support a range of resource types.

Tutors may be computer literate, but are rarely computing specialists with an understanding of distributed systems. Accordingly, a key role for TAGS is to provide useful abstractions for tutors to work with. For example, the three basic abstractions provided for tutors are users, groups, and resources. One can imagine providing tutors with further abstractions such as “course”, “class”, “assignment” and so on. In our experience systems, which, attempt to constrain the options of course, directors are not popular. It is more productive to provide basic building blocks. For example, the potential to group users readily supports the concept of a class, and the ability to group resources supports the notion of a course consisting of components.

TAGS also provides a developers and application programmers service interface (DAPSI), which allows educational resource developers to concentrate on content, by supplying commonly needed system services.

The project has addressed issues in usability, security, responsiveness, concurrency control, availability and infrastructure Quality of Service (QoS) that would have been unlikely to arise in a pure research environment. Much of the utility of TAGS comes from its strong support for group-based learning, which has in turn resulted in the group ethos permeating the system.

3. THE USE OF GROUPS AS A STRUCTURING PRINCIPLE

The concept of the *group* is central to TAGS, where it is used to support many functions, including the following:

- privileges and access control
- information dissemination and event awareness
- teamwork involving shared, multi-user educational resources
- the management of online collaborative learning,
- user-centric portal generation
- collaborative development

In practical terms, tutors construct a collaborative learning environment by using the Users, Groups and Resources management tool. This sets up arbitrary relationships between users and resources, using groups as the basis for the mapping. Users and groups are unique by name; resources are unique by name and type. Access rights can be specified when a resource is allocated to a group. A resource may simply be distinguished as Read-only or Read-Write, or it may export a more subtle set of access methods. Figure 2 illustrates an example set of relationships between users, groups and resources.




Figure 2: Basic Relations

John, Alice and Mark are members of the Students group. There are no resources shown allocated to the Students group. Jane is a Tutor and has created Groups A and B. John, Alice and Jane are in Group A and have Resources A and S allocated. S is allocated on a Read-only basis, whereas A is Read-Write. Mark and Jane are in Group B, which allows resource B as Read-Write and S as Read-only.

When a user is a member of multiple groups who have different access privileges to the same resource then they are credited with the highest level from their set of privileges.

It is hard to overstate the importance of access control and security in online environments. Privacy is essential to ensure that a student's online work is not corrupted or plagiarised (and is not in itself the result of plagiarism). Authentication is essential to the maintenance of *identity*, establishing that a user is who they say they are. Once authentication is established, authorisation needs to be performed each time a resource is accessed, to establish whether the user has the right to perform the requested operation on the resource.

The TAGS scheme for distributed authentication and authorisation is shown in Fig. 3. This allows for the secure remote access of distributed resources.




Figure 3: TAGS authentication and authorisation

When a user logs on to the system they are authenticated against a database (which may be replicated). Once a user is authenticated they may attempt to use a number of resources. Each access attempt is then authorised against the user database. Group membership allows access to resources. A simple web interface is provided to allow the allocation and revocation to groups of access rights to resources.

Note that a group does not explicitly have to be an access control group, or an information dissemination group, or any other type of group – it is the responsibility of the resource to use the groupings as it sees fit.

3.1 Support for Shared Resources

The concept of a learning resource in TAGS is deliberately loose. It can be a simple timetable, an automated assessment exercise or an interactive multi-user simulation.

In contrast to learners and tutors, developers are computing specialists. They are the people who develop and maintain resources. It is important to let people play to their strengths, and in the case of developers this means allowing them to focus on providing functionality without worrying about deployment, distribution, and management of a resource. These needs are generic and a collaborative learning environment can provide a useful set of services and interfaces for developers. Generic models for distribution allow a developer to choose the conditions, under which their resource can be made available, replicated for performance or fault tolerance, copied, re-used, and accounted for.

4. USER-CENTRIC PORTAL GENERATION

The purpose of a web portal in general is to provide a single, initial point of contact for a range of services. It is a technique widely used by web-based service providers. It reduces the time that individuals spend searching, although it only promotes options that are commercially sponsored. Portals sometimes offer facilities for personal customization; for example, an interest in specific sporting events or the local weather forecast. However, these options are very limited. They require identity and registration, which in the context of general web services must be user-driven as the service provider

has little idea of who is using a portal. Any type of user may use a search engine portal for example. By and large, web portals are public and cater for anonymous users.

User portals differ from public portals in that they are built entirely around the identity of the user. Their utility in a collaborative learning environment is that they can be generated for individuals and dynamically maintained using information that is already known about an individual's roles and responsibilities. For example, if a lecturer is (i) a 1st Year Advisor of Studies for the Science Faculty, (ii) teaching modules CS3013 and CS2001 and (iii) responsible for tutorial group CS1001/5, all the links to the relevant resources associated with these responsibilities can be aggregated onto a single page. Similarly, a student who is on courses CS3020, IS 1902 and EC4221 will find all their learning resources clearly presented as on their home page.

A users home page is updated dynamically with links to all the resources they have been allocated via their group memberships. This tends to mean that a student has a relatively clutter free, minimal home page; a tutor has more resources; and a course director may have an over-crowded view. Accordingly a view manager is being developed to help users who are members of a large number of groups to manage their home page.

5. RESOURCE REPLICATION ON CLUSTERS

Groups of servers are referred to as clusters. Replication and clustering can support high availability, good responsiveness and incremental scalability where appropriate middleware is provided.

A *resource* in TAGS may consist of a set of CGI programs, servlets, static HTML pages, applets, streaming media, multi-way conferencing channels and so on. The important feature of a resource in the clustering context is that it is the unit of replication management. We distinguish between resource-specific requirements and generic service needs. For example, if a resource contains continuous media then, in addition to the generic requirements of responsiveness and availability, it will also need appropriate network QoS for the satisfactory rendering of its media streams. If a resource is shared between many concurrent users then it will also require resource-specific concurrency control. If a resource is the focus of a collaborative effort then it will also need specific multi-user awareness features. Whenever a resource is replicated it will need *resource-specific replication coherence*.

The motivation for distinguishing between types of replication coherence is to avoid a "one-size-fits-all" approach. If such an approach is adopted then the synchronisation of replica updates must be able to satisfy the most stringent requirements, such as traditional atomic transactions. The undesirable consequence is that all resources, including those with relatively relaxed coherence requirements, must pay the same high price in terms of delay and protocol overhead. It can be argued that a uniform approach, such as distributed shared memory, is suitable for developers, but we believe that developers are capable of and willing to select from a range of coherence models in order to best meet the requirements of their applications. This means that the required response time can be traded off against the speed of convergence. The TAGS replicated resource architecture is more fully described in [6].

6. CONCLUSION

TAGS has evolved in line with the pedagogical needs of its users. The group ethos in TAGS is reflected from the application level all the way through the system design. The group mechanism provided by TAGS is part of a small and simple but powerful set of abstractions, which can be used to create a wide variety of collaborative learning environments.

ACKNOWLEDGMENTS

TAGS was funded under the SHEFC Use-of-the-MANS II and ScotCIT programmes (1997-2001).

REFERENCES

- [1] Michaelson, R. (1999) Web-based Group Work. *Proceedings of the 10th Annual CTI-AFM Conference*. Brighton, August, pp. 58-64. CTI-AFM Publications, East Anglia.
- [2] Allison, C., Huang, F and Livesey, M.J (1999) Object Coherence in Distributed Interaction. In Correia, N., Chambel, T. and Davenport, G. (eds), *Multimedia'99*, pp. 123-132. Springer-Verlag: New York.
- [3] Power, D.M., Michaelson R, and Allison, C. (1998) The Finesse Portfolio Management Facility. *Proceedings of the 9th CTI-AFM Conference*. York, October, pp. 119-125. CTI-AFM Publications, East Anglia.
- [4] Allison, C., Bramley M., Michaelson R., and Serrano, J. (1999) An Integrated Framework for Distributed Learning Environments. *Advances in Concurrent Engineering: the 6th ISPE International Conference on Concurrent Engineering*. Bath, September, pp. 345-353. Technomic publishing.
- [5] Allison, C., McKechnan, D., Lawson, H. and Michaelson, R. (2000) The TAGS Framework for Web-Based Learning Environments. *Web-Based Learning Environments 2000*, Portugal, 5-6 June. University of Portugal: FEUP Editions.
- [6] Allison, C., Bramley, M., Serrano, J., and McKechnan, D. (2000) Replicating the R in URL. *Proceedings of the 8th Euromicro Workshop on Parallel and Distributed Processing*. Rhodos, pp. 77-83. IEEE CS Press.